IMPLIED EXCHANGE RATE DISTRIBUTIONS: EVIDENCE FROM OTC OPTION MARKETS

José M. Campa
P.H. Kevin Chang
Robert L. Reider

Working Paper **6179**

NBER WORKING PAPER SERIES

IMPLIED EXCHANGE RATE DISTRIBUTIONS: EVIDENCE FROM OTC OPTION MARKETS

José M. Campa
P.H. Kevin Chang
Robert L. Reider

Working Paper 6179
http://www.nber.org/papers/w6179

NATIONAL BUREAU OF ECONOMIC RESEARCH
1050 Massachusetts Avenue
Cambridge, MA 02138
September 1997

For helpful comments, we are grateful to Geert Bekaert, Jens Carsten Jackwerth, Antonio Roma, and conference participants at the JIMF-LIFE Workshop on Developments in Exchange Rate Modeling (Maastricht, April 1997), the UCLA-USC-UCI Joint Conference (Rancho Bernardo, May 1997), the 1997 Western Finance Association Meetings (San Diego), and the 1997 French Finance Association Meetings (Grenoble). We also thank seminar participants at the Bank of England, Bocconi, ESSEC, INSEAD, Pompeu Fabra, and the Universities of Amsterdam, Arizona, and Salerno. This paper is part of NBER's research program in International Finance and Macroeconomics. Any opinions expressed are those of the authors and not those of the National Bureau of Economic Research.

Implied Exchange Rate Distributions: Evidence from OTC Option Markets
José M. Campa, P.H. Kevin Chang and Robert L. Reider
NBER Working Paper No. 6179
September 1997
International Finance and Macroeconomics

ABSTRACT

This paper uses a rich new data set of option prices on the dollar-mark, dollar-yen, and key EMS cross-rates to extract the entire risk-neutral probability density function (pdf) over horizons of one and three months. We compare three alternative smoothing methods---cubic splines, an implied binomial tree (trimmed and untrimmed), and a mixture of lognormals---for transforming option data into the pdf. Despite their methodological differences, the three approaches lead to a similar pdf clearly distinct from the lognormal benchmark, and typically characterized by skewness and leptokurtosis.

We then document a striking positive correlation between skewness in these pdfs and the spot rate. The stronger a currency the more expectations are skewed towards a further appreciation of that currency. We interpret this finding as a rejection that these exchange rates evolve as a martingale, or that they follow a credible target zone, explicit or implicit. Instead, this positive correlation is consistent with target zones with endogenous realignment risk. We discuss two interpretations of our results on skewness: when a currency is stronger, the actual probability of further large appreciation is higher, or because of risk, such states are valued more highly.

José M. Campa
Stern School of Business
New York University
44 West 4th Street
New York, NY 10012
and NBER
jcampa@stern.nyu.edu

P.H. Kevin Chang
Marshall School of Business
University of Southern California
Los Angeles, CA 90089-1421
kchang@sba.usc.edu

Robert L. Reider
J.P. Morgan
60 Wall Street
New York, NY 10260
reider_robert@jpmorgan.com

IMPLIED EXCHANGE RATE DISTRIBUTIONS: EVIDENCE FROM OTC OPTION MARKETS

This paper uses a rich new data set of options prices on the dollar-mark, dollar-yen, and key EMS cross-rates to extract the risk-neutral probability density function (pdf) over horizons of one and three months. Unlike most exchange rate forecasts, which provide only a point estimate of the future exchange rate, our options-based forecasts describe the entire range of realizations anticipated by the market, and the probability attributed to each range.[1] Knowledge of these exchange rate distributions can be helpful not only to the private sector for the management of currency risk but also to policymakers as a source of prompt market feedback to policy changes or other political and economic shocks.

This work contributes to a new and growing literature on options-based approaches to modelling expected asset returns. Interest in this topic has burgeoned in recent years, perhaps driven by the growth of real-world derivative markets, and a greater appreciation for the richness of information in option prices. Option prices have been used to describe the distribution of stock returns (Ait-Sahalia and Lo (1996), Jackwerth and Rubinstein (1996), Rubinstein (1994)), oil prices (Melick and Thomas (1997)), exchange rates (Campa, Chang, and Reider (1997); Malz (1996a); McCauley and Melick (1996a)) or interest rates (Abken(1995), McCauley and Melick (1996b)). For exchange rates in particular, recent empirical papers on the ERM (e.g. Campa and Chang (1996a), Malz (1996b); and Mizrach (1996)) have exploited the richer information structure embedded in options in analyzing credibility and realignment risk in the ERM around 1992. An important caveat, however, applies to this entire literature: options permit the estimation of the risk-neutral pdf, not the actual pdf. Any divergence, possibly based on risk, in these two distributions will simply not be reflected in option prices.

In all these cases, options, as forward-looking instruments, are able to incorporate information (such as risk, peso problems, or regime changes) that may not appear, or appear only much later, in exchange rate time series. [2] Even when stochastic volatility is explicitly modelled, as in a GARCH framework, parameters are often updated quite slowly even when major regime changes have occurred. In contrast, Campa, Chang, and Reider (1997), show that options (on exchange rates) are highly sensitive to news. In late July 1996, the pdf implied by options on the DEM-FRF rate widened sharply, even with relatively little change in the spot rate, in response to controversy and uncertainty regarding the French budget. Similarly, DEM-ITL expectations narrowed significantly with the victory of Prodi and the center-left party in April 1996, then again in November 1996 as the budget was approved and later the lira re-entered the ERM.

Another important advantage in describing exchange rate distributions using option prices is the generality of the methodology. Because an options-based approach does not presuppose any particular functional form, it can be applied equally to freely floating exchange rates, target zones, managed floats, or even contexts in which the regime is unknown or switching (as in Engel and Hamilton (1990)). Given the ambiguous degree of government influence (e.g. managed floating) on key exchange rates such as dollar-mark and dollar-yen, or the uncertainty surrounding Britain's potential participation in EMU, there is great value in using an approach that does not depend on a particular specification of exchange rate dynamics.

The implied probability density functions we derive enable us to document a striking empirical regularity across all the currency pairs we study: a consistently significant positive correlation between the level of the spot rate and the skewness of the risk-neutral distribution. The stronger a currency, the more exchange rate expectations are positively skewed, making large appreciations more likely than large depreciations. This suggests that exchange rate expectations are "extrapolative," or characterized by a kind of momentum in which a strong currency is associated with a relatively high probability of a large future appreciation.

Earlier studies that derive the pdf of exchange rates based on option prices include Leahy and Thomas (1996) for the Canadian dollar during the October 1995 Quebec referendum; Malz (1996b) for dollar-mark, dollar-yen, and pound-mark; and McCauley and Melick (1996a), which focuses primarily on methodology. Our work differs from these studies in a number of ways. First, we are fortunate to have a significantly more complete data set than Malz (1996a) or McCauley and Melick (1996a), in that we observe five strike prices for each exchange rate rather than three. As will be discussed later, even five points may not be enough to obtain the desired precision. By McCauley and Melick's own admission, three points are "the bare minimum." Second, we use three alternative techniques in mapping the option data into the pdf, comparing their advantages and disadvantages. We ultimately choose the binomial tree approach, which we believe to be somewhat more effective than the quadratic approach used by Malz (1996a) and more general than the mixture of lognormals approach used by Leahy and Thomas (1996) and McCauley and Melick (1996a). Finally, we not only derive the pdf's, but also demonstrate that skewness in these pdf's, by a number of different measures, is positively correlated with the spot rate for virtually all measures of skewness and all of our currency pairs at both the one-month and three-month horizons. While this correlation is briefly mentioned in McCauley and Melick (1996a), we document this phenomenon more completely, and evaluate its statistical significance for five currency pairs, five measures of skewness, and two time horizons.

The remainder of the paper is organized as follows. Section I describes the theoretical background behind the use of option prices to derive market expectations. Section II introduces our over-the-counter options data. Section III compares three derivations of the probability density functions using our option data. Section IV investigates the behavior of the pdf over time, and in particular the strong correlation between skewness in the pdf and the level of the spot rate. Section V concludes.

I. Market Expectations Embedded in Option Prices

Relative to other instruments, options have greater precision in describing market expectations of the future exchange rate. Our analysis will focus on the behavior of call options, financial securities giving the right but not the obligation to buy one unit of foreign currency at a pre-specified local currency strike price on a given expiration date. We will be working exclusively with European options, i.e. options that can be exercised only on their expiration date, and not earlier. With European options, there is no loss of information in using only calls and ignoring puts. By put-call parity, put options, financial securities giving the right but not the obligation to *sell* one unit of foreign currency at a pre-specified local currency strike price on a given expiration date, can always be replicated through a combination of the home currency, the foreign currency, and the call option with the same strike price and expiration date.

It was first shown in Breeden and Litzenberger (1978) that the decline in the value of a European call option for an infinitesimal increase in its strike price equals the (discounted) risk-neutral probability the option will, at expiration, finish in-the-money, i.e. with positive value. The option expires in-the-money whenever the exchange rate at expiration is above the strike price, i.e. at all points right of the strike price (assuming the exchange rate is on a graph's horizontal axis).

The mathematical derivation of this result is straightforward. Under risk neutrality, the price of a call option with strike price K and expiration date T, denoted $Call_{K,T}$ can be expressed as:

(1) $$Call_{K,T}=\frac{1}{1+i_T}\int_K^{\infty}(S_T-K)f(S_T)dS_T$$

where S_T is the exchange rate at time T, $f(S_T)$ is the risk-neutral probability density function for the exchange rate at time T, and i_T is the domestic interest rate prevailing between now and time T.

The partial derivative of this expression with respect to the strike price K can be expressed as:

(2) $$\frac{\partial Call_{K,T}}{\partial K} = -\frac{1}{1+i_T}[1-\int_0^K f(S_T)dS_T] = -\frac{1}{1+i_T}[1-F(K)]$$

where F(K) is the cumulative distribution function of the exchange rate at time T, evaluated at K.

If we take the partial derivative of equation (2) with respect to the strike price, we obtain the $1/(1+i_T)$ times the probability density function (pdf):

(3) $$\frac{\partial^2 Call_{K,T}}{\partial K^2} = \frac{1}{1+i_t}f(S_T)$$

These relations between the option value and the strike price will form the basis for our derivation of the implied probability density function.

One can also interpret this result more intuitively. For call options that will definitely finish in-the-money, a one-dollar increase in the strike price decreases the current value of the call by the present value of one full dollar: the holder of the option must at expiration pay one dollar more for an underlying security she would otherwise have acquired more cheaply. The partial derivative of the call with respect to the strike price equals minus (the present value of) one.

Conversely, for call options that will definitely finish out-of-the-money, i.e. worthless at expiration, a one-dollar increase in the strike price is irrelevant, since the holder of the option would not have purchased the security at either strike price. The partial derivative of the call with respect to the strike price equals (the present value of) zero.

For call options that may or may not finish in-the-money, a one-dollar increase in the strike price decreases the value of the call by (the present value of) an amount between zero and one dollar. The higher the probability[3] the option finishes in-the-money, the more likely a one-dollar increase in the strike price will ultimately matter to the option holder, and the greater the decrease in the option price. In fact, the first partial derivative equals minus (the present value of) the probability that the option will finish in-the-money. The option finishes in-the-money whenever S_T, the security price at expiration, is greater than K, the strike price. Therefore, the probability the option finishes in-the-money is the same as the probability $S_T > K$, which is 1-F(K).

Since the first partial derivative of the call with respect to strike price yields an expression with the risk-neutral cumulative distribution function, or cdf, we can take a second partial derivative with respect to strike price to obtain the probability density function, or pdf, denoted as f(K), our objective in fully describing exchange rate expectations.

Note that the probabilities derived from option prices will be the *risk-neutral* probabilities, since contingent claims are priced according to the risk-neutral distribution rather than the actual distribution. These two distributions will be equal if exchange-rate risk is not priced, or in the context of the capital-asset pricing model, unsystematic. For the remainder of this paper, the terms probability, cdf, pdf, and expected value will refer to these measures under the risk-neutral distribution. Similarly, since we are using the risk-neutral measure, present value is calculated using the risk-free interest rate.

II. Over-the-Counter Options Data: Volatility Smiles and Smirks

Our primary data consist of market quotes of over-the-counter options on five actively traded currency pairs: the dollar-mark, dollar-yen, mark-pound, mark-French franc, and mark-lira. In particular, we use closing prices on these options, as well as on exchange rates and Eurocurrency interest rates, recorded by currency option traders at a major money-center bank. Observations are weekly (every Wednesday) from 3 April 1996 to 5 March 1997, for a total of 48 weeks.

While exact quotes may differ slightly from bank to bank, the over-the-counter options market is relatively liquid and competitive, especially for the currency pairs we study.[4] Relative to exchange-traded options, our over-the-counter options are far more liquid. For reference, Table 1 provides information on daily turnover and outstanding notional amounts of over-the-counter and exchange-traded options on these currency pairs, as reported in the 1996 Bank for International Settlements triennial survey of foreign exchange activity. Relative to exchange-traded options, over-the-counter options typically have an order of magnitude more turnover and notional amount outstanding. For example, depending on the currency pair, daily turnover on exchange-traded options is 12% or less of its OTC counterpart, and the notional amount outstanding is 7% or less.

In the over-the-counter options market, price quotes are expressed as implied volatilities which traders by agreement substitute into a Garman-Kohlhagen formula (Black-Scholes adjusted for the foreign interest rate) to determine the option premium.[5] Since the volatility is the only unobservable parameter in the Black-Scholes formula, these volatilities---representing traders' subjective assessment of future movements in the underlying asset---uniquely determine the options' price. Traders do not necessarily believe that Black-Scholes holds. This method simply represents a one-to-one mapping from implied volatility quotes to the option's price, which is also referred to as the "option premium." Under this convention, option quotes do not necessarily require updating even as the spot rate evolves minute-by-minute. If prices were expressed as an option premium for a fixed strike, most short-run changes in the option premium would result from innovations in the spot rate, requiring far greater coordination between spot markets and option markets.

Our observations for a given maturity (1 and 3 months) consist of implied volatility quotes for at-the-money-forward options, i.e. whose strike price equals the forward rate of the same maturity, as well as four additional strike prices, two below and two above the forward rate. The two strike prices below the forward rate correspond to out-of-the-money put options with "delta" (partial derivative of the option price with respect to the spot exchange rate) equal to -0.10 and -0.25.[6] Similarly, the two strike prices above the forward rate correspond to out-of-the-money call options with delta equal to 0.10 and 0.25. The higher the level of implied volatility, the wider the spread in strike price between the -0.10 delta put, the lowest strike price, and the 0.10 delta call, the highest strike price.

Although Black-Scholes assumes a constant volatility, the implied volatility quoted by option traders will typically vary as a function of an option's strike price. This reflects a departure from the Black-Scholes assumptions, with the result that the probability distribution for the future exchange rate is not lognormal.

When volatility is stochastic but uncorrelated with changes in the spot rate, it can be shown that the Black-Scholes implied volatility is lowest at-the-money forward (strike price=forward rate), increasing for both in-the-money (call's strike < spot) and out-of-the-money (call's strike > spot) options. This pattern is referred to as the "volatility smile," so named for the appearance of a graph with implied volatility on the vertical axis and the option strike price on the horizontal axis. Different shapes of the volatility smile are consistent with different distributions of the underlying exchange rate. For instance, a symmetric volatility smile is consistent with leptokurtosis or "fat tails" in the distribution, i.e. higher probabilities than under the lognormal distribution of larger positive or negatives changes, as would result from returns with stochastic volatility. By put-call parity, the implied volatility at any strike is the same for call options and put options.

At times, the probability of future exchange rate realizations is not symmetrically distributed around the at-the-money strike price. For example, for the sterling-mark exchange rate, the market may perceive a greater probability of a large depreciation of the pound than a similar depreciation of the mark. When this asymmetry is present, the smile can be transformed into a "smirk," with the options' implied volatility rising more sharply for strike prices on one side of the forward rate than the other. Figure 1 depicts the volatility smile or smirk, for our five currency-pairs at the three-month horizon on a single day in our data sample, 11 September 1996.

As the graphs in Figure 1 indicate, these smiles are far from symmetric. Options are in fact capturing skewness of expectations---i.e. asymmetric market expectations in terms of the direction of potential large exchange rate changes. In this case, all five exhibit a one-sided rise in implied volatility in the direction of a major dollar depreciation vs. the mark or yen, or major mark appreciation vs. the three other European currencies, reflecting greater probability of this realization than a similar movement in the opposite direction.

We were also able to obtain, from another money center bank, supplementary data on "risk reversals," a traded instrument denoting the simultaneous purchase of an out-of-the-money call and sale of an equally out-of-the-money put. Risk reversals, which above all reflect the skewness of a distribution, are discussed further in Section IV, which focuses on skewness. Our supplementary data consist of daily observations from 1 April 1992 to 15 April 1997 of one-month 25Δ risk-reversals, i.e. where the call has 0.25 delta and the put -0.25 delta, for the dollar-mark, dollar-yen, mark-pound, mark-lira, and pound-dollar. This data source did not have historical risk reversal prices for the mark-French franc.

III. Alternative Derivations of the PDF

III.A. Three Alternative Approaches to the PDF

From these option prices, we derive the pdf using a few alternative techniques. If we could observe a continuum of option prices, there would be no need to use smoothing techniques, as the second partial derivative of the option price with respect to strike would lead directly to the pdf. Instead we have six data points to estimate each pdf: the five observed option prices plus the forward rate, which must equal the mean of the risk-neutral distribution. A different pdf is estimated for each week's observations, each currency pair, and each option maturity.[7]

We now derive the pdf using three alternative methods. The "cubic spline" approach fits the volatility smile---with Black-Scholes volatilities on the vertical axis and strike prices on the horizontal, while the other two ("implied binomial trees" and "mixture of lognormals") impose *a priori* restrictions on the probability distribution to be estimated and fit that distribution to the observed option prices.

(1) Cubic Splines Based on Volatility-Smoothing Approach of Shimko (1993).

Shimko (1993) introduces a method of fitting the volatility smile with a quadratic function to obtain a continuum of call prices as a function of strike price within the range covered by the data. The pdf is obtained by twice differentiating this result then discounting by the riskless interest rate. For strike prices outside the observed range, one can extrapolate by extending the quadratic function outwards.

Campa, Chang, and Reider (1997) use a modification of this approach providing greater flexibility in the shape of the smile, and hence the pdf. The method replaces the quadratic with cubic splines---polynomial functions of order three or lower. The polynomials between any two points are chosen such that where two polynomials meet at a single data point, the first derivatives of the two functions are equal and differentiable.[8] This assures that the first derivative is continuous throughout the range of strike prices, and everywhere differentiable. Outside the observed range, polynomials from the first and last regions are used to extend the "smile" left (by the same distance as between the first and second data points) and right (by the same distance as between the last and second-to-last data points). Beyond these points, the smile is treated as flat.

The cubic spline approach is attractive for its generality, as the third-order polynomial used to fit the volatility smile is permitted to change form over each interval. One drawback is the potential explosion of the implied volatility in the extrapolation of the smile outside the range of observed strike prices.

(2) "Implied Binomial Tree" Approach of Rubinstein (1994). Under the implied binomial tree approach, one solves directly for the risk-neutral probability p_i associated with each of the n+1 terminal nodes of an n-step binomial tree. Inputs to this approach are the m observed call options, where normally one chooses n >> m. For our five option observations, for example, we choose a 100-step tree. Under this method, probabilities are chosen to minimize the sum of squared difference between a prior probability p'_i and the new estimated probabilities p_i, subject to these constraints: each call must equal its expected value based on these probabilities, the probabilities must sum to one, the expected value of the future spot rate must equal the forward rate (denoted as $Forward_T$), and $p_i \geq 0$ for all i.

Mathematically, one minimizes a "loss function," in this case quadratic, solving the following constrained optimization:

subject to the constraints

(4) $$\underset{p_i}{Min} \sum_{i=0}^{n} (p_i - p'_i)^2$$

(5) $$Call_{K_j,T} = \frac{1}{1+i_T} \sum_{i=0}^{n} p_i \; Max(S_{i,T} - K_j, 0) \qquad for\ j = 1,2,3,...,m$$

(6) $$\sum_{i=0}^{n} p_i = 1$$

(7) $$\sum_{i=0}^{n} p_i \; S_{i,T} = Forward_T$$

(8) $$p_i \geq 0 \quad for\ all\ i = 0,1,2,3,...,n$$

Typically, the prior probabilities p'_i are based on the lognormal distribution, so that provided the constraints are satisfied, the distribution appears as close as possible to lognormal. Jackwerth and Rubinstein (1996) empirically compare, for the S&P 500, a number of loss functions other than quadratic (e.g. goodness-of-fit, absolute difference, maximum entropy, smoothness), and find similar estimated pdf's. When the loss function is altered, the effects on the resultant pdf are pronounced only at the tails, where, in any event, very little absolute probability is found. In general, the further from the observed strike prices, the more the derived probabilities depend on the chosen parameterization than on the actual data.

Empirically, one difficulty implementing the binomial approach is irregularities in the tails: outliers often have a probability orders of magnitude larger than under the lognormal benchmark. Probability sometimes rises inexplicably deep into the tails at nodes far from observed strike prices. Since the probability assigned to these outliers is low in absolute terms, one or more orders of magnitude in fact have little effect in constraints (5)-(7), where outliers enter only linearly. For higher moments such as kurtosis, these irregularities take on far greater importance: for example, the binomial approach applied to 48 weekly one-month dollar-mark options leads to pdf's with an average (excess) kurtosis of 19, as reported in Table 2c. One can reduce the excess kurtosis by "trimming" the distribution, or truncating the distribution by eliminating the ends in the tails (see Rosenberger and Gasko (1983) for details).[9] To accomplish this, we eliminate 20% of possible outcomes from the top and the bottom of the distribution, or the top and bottom 20 terminal nodes in a 100-step binomial tree. This typically leads to an exclusion of 0.5-1.0% of total probability.[10]

(3) "Mixture of Lognormals" Approach of Melick and Thomas (1997). Still another approach is to fit the distribution as a "weighted average" or mixture of lognormals, as in Leahy and Thomas (1996), Melick and Thomas (1997), and McCauley and Melick (1996a). This method has a natural economic interpretation---that of multiple alternative regimes. In Melick and Thomas (1997), for example, three different lognormal distributions for the price of oil correspond to various outcomes of the 1991 Gulf War.

Following this approach, we solve for a pdf $h(S_T)$ as a mixture of two lognormals, expressed as follows:

(9) $$h(S_T)=\pi_1 g(\mu_1,\sigma_1^2)+(1-\pi_1)g(\mu_2,\sigma_2^2)$$

where $h(S_T)$ represents the unknown composite pdf, and $g(\mu,\sigma^2)$ represents the pdf of a variable whose logarithm is normal with mean μ and variance σ^2. The five unknown parameters π_1, μ_1, μ_2, σ_1, and σ_2 are selected in order to satisfy as well as possible the constraints on the m observed call options and the forward rate:

(10) $$Call_{K_j,T}=\frac{1}{1+i_T}\int_{K_j}^{\infty}(S_T-K_j)h(S_T)dS_T \quad for\ j=1,2,3,...,m$$

and

(11) $$\int_0^{\infty} S_T\ h(S_T)dS_T = Forward_T$$

In solving this system, one typically defines a loss function expressed as the sum of squared percentage deviations from these constraints.

Statistically, one advantage of the mixture of lognormals approach is the smooth behavior of the tails. Probability in the tails declines monotonically and always decays quickly enough to prevent unreasonable kurtosis. On the other hand, this approach may impose too rigid a structure on the resultant pdf, unless one believes *a priori* that the underlying economics are well captured by a weighted sum of discrete outcomes, as in the 1991 Gulf War or perhaps certain periods of crisis in the ERM. Empirically, we find that the two lognormals identified in our derived pdf often make little economic sense. The parameters and probability associated with each distribution fluctuate widely from week to week for no apparent reason. This lack of economic meaning presumably arises from the sparseness of the data; with only five observations, small deviations in the price of any option can result in sizable parameter shifts.

As an alternative to these three statistically based methods of deriving the pdf, one could estimate a full stochastic volatility model (such as Heston (1993)) with potential nonzero correlation between shocks to the spot rate and shocks to volatility. The estimated parameters of the model could then be used to describe a full pdf.

III.B. Comparing Implied PDF's

We now compare these three alternative methods of deriving the pdf in order to choose a single approach to be used in our economic analysis of the pdf in Section IV. The three methods produce remarkably similar pdf's in that they more closely resemble each other than they do the lognormal benchmark, where at-the-money implied volatility is imposed for all strike prices. This is illustrated in Figure 2a for 3-month mark-lira options on 11 September 96. A comparison of the three methods indicates that this similarity holds more generally over other currencies and dates. All three capture key deviations from lognormal such as a fatter right-hand-side tail, and a mode shifted slightly to the left. Relative to the lognormal, these three options-based pdf's clearly reflect the additional information embedded in our volatility smile data.

As shown in Figure 2b, depicting the corresponding cdf's, the three options-based cdf's cross the lognormal cdf at nearly the same point near 1055 ITL/DEM, about 3% above the forward rate. To the right of this point, the derived cdf's lie below the lognormal benchmark, indicating greater probability attributed to realizations above this level. Though not visible upon casual inspection, this distribution is positively skewed, i.e. with a higher probability of a large mark appreciation than of a comparable mark depreciation.[11]

Tables 2a-2e report the mean, standard deviation, skewness, and kurtosis of exchange rate changes embedded in the pdf's computed from option prices for five currency pairs at the one-month and three-month horizons. For each currency pair and option horizon, we simply report the average, minimum, and maximum of the moments over the 48 weeks in our sample. These tables reveal the extreme kurtosis associated with the binomial tree approach, and how "trimming" sharply reduces this excess while excluding only about 1% of the distribution. For example, by excluding only a small fraction of the distribution, we greatly reduce excess kurtosis: for one-month dollar-mark options, excess kurtosis drops from over 19 to 0.237, a more reasonable figure. While reducing the kurtosis and standard deviation of returns, trimming has relatively little effect on returns' mean and skewness, the focus of the next section. In fact, trimming the binomial distribution typically makes the skewness more consistent with the skewness of the two other approaches (the cubic spline method and the mixture of lognormals). For the five currency pairs and two horizons, in 9 out of 10 instances the trimmed skewness is closer to the average of the two other approaches than is the untrimmed.

Overall, our comparison of the alternative methods indicates that *any* of the methods successfully incorporates the richness in volatility smile data that is absent under the simple Black-Scholes model. While each method has some strengths and weaknesses, for the analysis presented in Section IV, we have selected the trimmed binomial tree approach, which seemed to incorporate the key features of the data without imposing too rigid a parametric structure.

Clearly, the sparseness of our data limits the precision of any of these methods in fully identifying the implied pdf. Five option prices and the forward rate are enough to distinguish the underlying distribution from the lognormal, but not enough to determine precise probabilities, especially in ranges far from the actual data. Although we do not observe options outside these five strike prices, the observed options do incorporate all the probability from the strike to infinity (for calls) and the strike to zero (for puts). For any strike price, options prices indicate the product of the probability the option expires in-the-money and the expected value of the option conditional on expiring in-the-money. The technical difficulty lies in disentangling these two contributions to the options' value without observing more strike prices.

To improve precision, one might try to exploit information in the "term structure" in options prices, combining data on options with differing times-to-expiration. This requires making some assumptions about the structure of expectations---e.g. that, per unit time, the distribution of expected exchange rate changes is the same from time 0 to time 1 as from time 1 to time 3. Alternatively, one could impose more economic structure, possibly estimating a stochastic volatility model such as Heston (1993) with nonzero correlation between shocks to the spot rate and shocks to volatility. The estimated parameters of the model could then be used to describe a full pdf. One could also try to use the time series of option prices to improve precision. This approach, used by Ait-Sahalia and Lo (1996), requires some assumption about the stationarity of the distribution of exchange rate changes---e.g. that all options prices observed over a certain period were in fact noisy observations of the same distribution. Under such assumptions, these approaches are likely to generate better-defined pdf's and may be a useful area of future research.

IV. Patterns in the PDF: Skewness, Risk Reversals, and the Spot Price

Having identified these pdf's from options data, we now ask how these distributions and their characteristics evolve over time, and what economic factors may influence them. The mean of the risk-neutral distribution always equals the forward rate, so by covered interest parity, the first moment of exchange rate changes is strictly determined by interest rate differentials. The second moment, the variance, has already been the subject of extensive study, including work on the implied volatilities in option prices (e.g. Scott (1992), Bates (1996a) Jorion (1995)) as well as on the time series of exchange rate changes, often in a GARCH context (e.g. Hsieh (1988, 1989)).

In this paper, we focus on the third moment of the distribution, i.e. its skewness. As the volatility smiles on 11 September 1996 in Figure 1 indicate, these distributions can at times be significantly skewed. Tables 2a-2e further suggest that the skewness is often systematically one-sided. For example, for most smoothing methods, the distributions on this date are skewed towards a mark appreciation against the lira, French franc, dollar, and pound as well as a yen appreciation against the dollar. We will later investigate how this skewness changes sign or magnitude over time, and what factors---particularly the spot rate---may drive such changes in skewness.

For the stock market, skewness in expected returns has been documented extensively (see Bates (1991), Ait-Sahalia and Lo (1996), and Jackwerth and Rubinstein (1996), for example) using S&P options trading on the Chicago Board Options Exchange. For the S&P 500 index, negative skewness in the pdf appearing since the 1987 stock market crash reflects either a higher probability attached to a future crash, or higher valuation of returns in the event of a crash. In a stochastic volatility model, this skewness results from a negative correlation between shocks to price and shocks to volatility, as discussed in (among others) Hull and White (1987) and Heston (1993). Jackwerth and Rubinstein (1996) show that this negative skewness has been roughly constant in S&P options since about one year after 1987 crash.

In contrast, skewness in currency markets has been less systematic over time. A number of studies[12] have found negative skewness (expressed in foreign currency per dollar) in (primarily) dollar-mark options in the mid-1980's. Since the dollar at the time was unusually strong, this is consistent with mean-reversion: a large dollar depreciation was more probable than a large dollar appreciation. Since that period, however, Bates (1996b) reports considerable fluctuations in dollar-mark skewness, which has varied in sign but usually moved in the same direction as dollar-yen skewness (expressed in dollars per yen). Bates (1996b) finds that negative skewness (DEM/USD), implying higher probability of a large dollar depreciation, is associated with relatively higher trading volume in out-of-the-money mark calls; positive skewness (DEM/USD), or a higher probability of a dollar appreciation, is associated with relatively more volume in out-of-the-money mark puts. None of these studies, however, identify factors with a systematic influence on skewness---an important objective of this paper.

The third moment of a distribution, or its skewness, is defined mathematically as $E(x-\mu)^3/[E(x-\mu)^2]^{1.5}$, where μ is the mean of the distribution. The skewness coefficient is unitless, reflecting simply a distribution's asymmetry. Positive skewness indicates that the right tail is "fatter" (higher probability) and usually extends out further than the left tail; large positive realizations are more likely than large negative realizations. Note that such skewness does not imply that an appreciation is more likely than a depreciation. Regardless of its skewness, the mean of the risk-neutral distribution must always be the forward rate.

IV.A. Alternative Measures of Skewness

Given our over-the-counter options data, we have a number of approaches to measuring the skewness in the implied distribution:

(1) Skewness from pdf. Having estimated the pdf from options data, we can compute skewness directly. For reasons discussed earlier, we focus on the trimmed binomial tree approach to deriving the pdf from option prices, and thus use this pdf to compute skewness.

(2) Relative Intensities. Since skewness reflects asymmetry in the tails, we can explicitly compare the likelihood of large appreciations vs. large depreciations, where "large" is defined as at least a certain distance "x" above or below the forward rate, $Forward_T$. We compare the pdf in the tails below a critical level $\underline{S}$ (below the forward rate) and above another critical level $\bar{S}$ (above the forward rate) by examining their "intensity," a measure incorporating both probability and magnitude of realizations in these ranges.[13]

Mathematically the "intensity" of realizations above any exchange rate $\bar{S}$ can be defined as:

$$intensity_{\bar{S}+} = \int_{\bar{S}}^{\infty}(S_T-\bar{S})f(S_T)dS_T \tag{12}$$

By definition, $intensity_{S+}$ sums over all possible realizations of S_T the size of any appreciation beyond $\bar{S}$ multiplied by the probability that such a realization will take place. Since $f(S_T)$ is the risk-neutral distribution, the present value of

$\text{intensity}_{\bar{S}+}$ also equals the value of a call option on S with a strike price $\bar{S}$ and expiration date T, or $\text{Call}_{\bar{S},T}$.

Likewise, the "intensity" of realizations below any exchange rate $\underline{S}$ is defined as:

(13) $$intensity_{\underline{S}-} = \int_0^{\underline{S}} (\underline{S}-S_T) f(S_T) dS_T$$

Similarly, the present value of $\text{intensity}_{\underline{S}-}$ equals the value of a put option with strike $\underline{S}$ and expiration date T, or $\text{Put}_{\underline{S},T}$.

We can measure skewness by the difference of these two intensities, or the "relative intensity," which we define as $\text{intensity}_{\bar{S}+}$ - $\text{intensity}_{\underline{S}-}$. This is equivalent to $(1+i_t)(\text{Call}_{\bar{S},T} - \text{Put}_{\underline{S},T})$. By examining the *relative* intensity, we focus specifically on skewness. Clearly, we would not capture skewness by using only one tail, since a symmetric distribution with fat tails (leptokurtotic) or simply high volatility would exhibit higher intensity at both the left and right tails. Note that when $\bar{S}=\underline{S}=\text{Forward}_T$, no matter what the shape of the distribution, relative intensity will always equal zero, since the distribution's mean must equal the forward rate. Alternatively, by put-call parity, $\text{Put}_{F,T} = \text{Call}_{F,T}$, so relative intensity must be zero.

We choose $\bar{S}$ and $\underline{S}$ symmetrically around the forward rate. Recall that relative intensity, as a measure of skewness, is designed to capture the difference in intensity between "large" appreciations and "large" depreciations. Since the meaning of "large" will depend on the overall dispersion of the distribution, we want the distance between $\bar{S}$ and $\underline{S}$ to be a positive function of the volatility. In particular, we set $\ln \bar{S} \equiv \ln(\text{Forward}_T) + \sigma$ and $\ln \underline{S} \equiv \ln(\text{Forward}_T) - \sigma$, where $\sigma \equiv$ the standard deviation of expected returns in the pdf in question. To compare the more extreme part of the tails, we later define $\ln \bar{S} \equiv \ln(\text{Forward}_T)+1.5\sigma$ and $\ln \underline{S} \equiv \ln(\text{Forward}_T)-1.5\sigma$. We base the range $[\underline{S},\bar{S}]$ on the standard deviation of the distribution so that changes in "relative intensity" truly reflect changes in skewness without contamination from changes in volatility. If we had instead made $\underline{S}$ and $\bar{S}$ independent of volatility, then relative intensity would rise with volatility, even for constant skewness. This would occur because the size of the tails rises with volatility, so any difference in the left and right tails (for a fixed degree of skewness) would be proportionally magnified under higher volatility.

Alternatively, we could have ignored magnitude and simply measured the "total probability" right of $\bar{S}$ or left of $\underline{S}$, then computed the difference. Unlike intensity, which is always centered around the forward rate, probability is centered around the median rather than the mean. For example, in a positively skewed distribution, there must be greater total probability below the forward rate than above it. So in pure probability terms, the relative probability would be negative for $\bar{S}=\underline{S}=\text{Forward}_T$. Then, as $\underline{S}$ and $\bar{S}$ are brought further from Forward_T, the relative probability would at some point switch to positive. For these reasons, we use intensities rather than probabilities in measuring skewness.

(3) Risk Reversals. In over-the-counter FX option markets, there happens to exist an active market in an instrument that specifically reflects skewness. This instrument is called a "risk reversal," denoting the simultaneous purchase of an out-of-the-money call and sale of an equally out-of-the-money put. As the difference in the price of two options, the risk reversal closely resembles the earlier measure of "relative intensity." One distinction between these two measures is how the out-of-the-money strike prices are chosen. The degree of "out-of-the-moneyness" of risk reversals is expressed in terms of delta. Thus, the purchase of a 25Δ call accompanied by the sale of a 25Δ put would be referred to as a "25Δ risk reversal." In practice, the market is most liquid in 25Δ and, to a lesser degree, 10Δ risk reversals. As a rough guideline, the 25Δ call and 25Δ put have approximately a 25% probability of finishing in-the-money, i.e. the strike prices are such that approximately 25% of the distribution lies to the right of the call's strike and 25% to the left of the put's strike.[14] In comparison, 10Δ options have approximately a 10% probability of finishing in-the-money---and will therefore have strike prices further from the spot rate. As instruments traded in the market, risk reversals have the advantage of being directly observable. Hence, as one of our measures of skewness, they are less subject to error introduced by any smoothing technique used to derive the pdf.

The price of a risk reversal is the difference between the premium paid for the call option minus the premium received for the put option. Since prices of over-the-counter options are quoted in terms of Black-Scholes implied volatility, and these options are equally out-of-the-money, the price of a risk-reversal is by convention quoted as the volatility of the long option minus the volatility of the short option. Note that if Black-Scholes assumptions held, then volatility would the same across all strike prices and the price of a risk reversal would be approximately zero. Even with stochastic volatility uncorrelated with spot, the price of a risk reversal would be close to zero. For example, Heston (1993) notes "if volatility is uncorrelated with the spot return, then increasing the volatility of volatility increases the kurtosis of spot returns, not the skewness. [...] random volatility is associated with increases in the prices of far-from-the-money options relative to near-the-money options." Since the put and call in a risk reversal are equally "far-from-the-money," the risk reversal is nearly a pure measure of the skewness in the distribution.

By the Black-Scholes formula, equally out-of-the-money options, i.e. puts or calls with the same magnitude delta, have the same vega, or sensitivity to changes in volatility.[15] Therefore, risk reversals, the difference in two equally out-of-the-money options, are vega-neutral. Hence, option traders can buy and sell risk reversals without specifying the *level* of volatility of these out-of-the-money options. For most of the major currency pairs, risk reversals at 10 and 25 delta are the most liquid instruments after at-the-money straddles. To compute the actual level of volatilities for 10Δ and 25Δ options, one requires information not only on risk reversals but also on 10Δ and 25Δ "strangles."[16] A strangle is defined as an out-of-the-money call *plus* an equally out-of-the-money put, so obviously observing both the strangle and the risk reversal permits solution of the price of the call and put individually.

Risk reversals can have a positive or negative sign, depending on the direction of skewness in expected exchange rate changes. In the dollar-yen market, for example, if 25Δ out-of-the-money dollar calls are trading at a volatility 0.5 percentage points higher than 25Δ out-of-the-money dollar puts, then 25Δ risk reversals might be quoted as "0.5, dollar calls over," indicating a skewness biased towards large dollar appreciation. Conversely, if that distribution were skewed towards a large yen appreciation, the analogous quote would be "0.5, yen calls over."

Figures 3a-3e depict, for each currency pair, the price of 10Δ and 25Δ risk reversals in terms of the priced currency (i.e. the denominator or right-hand side currency) call minus the priced currency put. For example, since the exchange rate is quoted ITL/DEM, the risk reversal is plotted in terms of mark calls minus mark puts (vs. the lira). Similarly, since the exchange rate is quoted JPY/USD, the risk reversal indicates the price of dollar calls minus dollar puts (vs. the yen). Prices are expressed in volatility points, referred to by option traders as "vols." Note that we have also graphed the spot exchange rate on the same axes in order to provide a visual image of an empirical regularity to be discussed later: the positive correlation between the spot rate and risk reversals.

IV.B. Similarities of Three Skewness Measures

A comparison of these alternative measures of skewness is presented in the three-dimensional Figures 4a-4c, depicting the behavior of the skewness coefficient (trimmed) as well as the 1.0σ and 1.5σ relative intensities (untrimmed) as a function of the 10Δ risk reversal and the level of at-the-money volatility. The figures are based on simulated data in which other variables are held constant. In particular, both domestic and foreign interest rates are set at a fixed 5%, implying that the forward rate equals the spot rate. In these simulations, the relation between the 10Δ risk reversal (RR) and 25Δ risk reversal is captured by the following equation, expressed in volatility points (vols):

25Δ RR = -0.01586 + 0.4559 (10Δ RR).

These coefficients are the average of coefficients for all 5 currency-pairs in regressions of the 1-month 25Δ risk reversal on the 1-month 10Δ risk reversal. Since the regression coefficient is positive but less than one, and the constant term is negligible, the 25Δ risk reversal is positively correlated with the 10Δ risk reversal but of lower magnitude. This means that the smile or smirk becomes more pronounced further from the forward rate, as expected. To avoid introducing more deviations from the lognormal benchmark, we assume that, aside from the effects of the risk reversal, there is no "volatility smile." Furthermore, for simplicity we assume that the entire effect of the risk reversal appears in the put, with the smile always flat on the side of the call.[17] Thus, for a +1% risk reversal and 8% at-the-money volatility, the volatility is set at 8% for the 10Δ call, 8% at-the-money, and 7% for the 10Δ put.

The positive relation between the risk reversal and skewness, measured either by the skewness coefficient or the relative intensities, is clear in Figures 4a-4c. Starting from zero, any change in the risk reversal leads directly to a change of the same sign in any of these skewness measures. These graphs also indicate the relation between the risk reversal and these skewness measures is relatively independent of the level of volatility. Since the trimmed relative intensities (not shown) appeared to be less smooth than the untrimmed relative intensities, as a function of volatility and the risk reversal, we have used the untrimmed relative intensities as an indicator of skewness in the regressions below.

IV.C. Empirical Results: Positive Correlation Between Skewness and Spot

We run regressions to determine how skewness in exchange rate expectations may be related to the spot rate. Our five alternative left-hand side variables as measures of skewness include: 10Δ and 25Δ risk reversals, skewness (in the trimmed binomial tree), and "relative intensity" (untrimmed binomial tree) for 1 and 1.5 standard deviations from the forward rate. All regressions are performed in first-differences as a precaution against nonstationarity. A positive correlation between risk reversals and the spot rate was described qualitatively in McCauley and Melick (1996a), who also point out that dollar-mark and dollar-yen risk reversals appear to be positively correlated over time.

Note that exchange rates, skewness, and relative intensities are all calculated with exchange rates quoted according to market convention, e.g. JPY/USD ((*not* USD/JPY) terms. In this example, a positive skewness or relative intensity indicates a greater probability of a large dollar appreciation than an equivalent dollar depreciation. Risk reversals are expressed as right-hand-side currency (dollar in the example above) call minus right-hand-side currency put, so a positive risk-reversal also indicates a skewness towards a large appreciation of the right-hand-side currency.

The expected sign of the correlation between skewness and the spot rate depends on one's prior beliefs regarding exchange rate dynamics. For example, a credible target zone exchange rate regime is characterized by mean-reversion as well as bounded exchange rate realizations and therefore should exhibit a negative correlation between skewness and spot.[18] When the exchange rate is near its ceiling, the distribution is negatively skewed; large negative realizations have greater probability than large positive realizations, which are impossible if the regime is credible. Likewise, when the exchange rate is near its floor, the distribution is positively skewed. This pattern leads to a negative correlation between the spot rate and skewness. An interesting implication is that by testing for a negative correlation between skewness and the spot rate, we can test the hypothesis that exchange rates follow implicit bands---even if we cannot explicitly identify the floor and ceiling. Alternatively, if we believe that the exchange rate follows a random walk or a martingale process, there should be virtually no skewness in the distribution, and the correlation with the spot rate should be zero.

Empirically, we reject both the credible target zone and martingale framework for exchange rate dynamics. We find overwhelming evidence of a strong *positive* correlation between the spot exchange rate and skewness in the implied risk-neutral distribution. The stronger a currency, the more expectations are skewed towards a large appreciation of that currency. This result holds quite consistently across all 5 currency pairs and all 5 measures of skewness at the 1-month horizon, reported in Table 3a, and for the great majority of measures at the 3-month horizon, reported in Table 3b.

At the one-month horizon, coefficients on the spot rate are positive and significant for all currency pairs when we use risk reversals as a measure of skewness. When skewness is computed from the estimated (trimmed) distribution, the coefficient on the spot rate is significantly positive for 4 out of 5 currency pairs. Finally, when skewness is measured as relative intensity, half the coefficients are significantly positive and none are significantly negative.

At the three-month horizon, reported in Table 3b, results are qualitatively unchanged for the two risk reversals, which again indicate a significant positive effect of the spot rate on skewness for all five currencies. The (trimmed) skewness coefficient itself shows a positive effect for all currencies but the dollar-mark, and is significant for the dollar-yen and pound-mark. The relative intensity measures are positive and significant for the mark-lira and mark-French franc. None of the coefficients are significantly negative.

This positive relation between skewness and the spot rate means that expectations are in a sense "extrapolative": a currency that is relatively strong has a greater chance of becoming much stronger than much weaker. In a stochastic volatility framework, this would hold if the correlation between shocks to spot and shocks to volatility were a positive function of spot. When the dollar is strong, a large increase in volatility is more likely to be accompanied by a dollar appreciation than a dollar depreciation. Conversely, when the dollar is weak, a large increase in volatility is more likely to be accompanied by a dollar depreciation. This does not mean that on average the market expects appreciation rather than depreciation, or *vice versa*, since the mean of the distribution remains the forward rate.

The positive correlation between spot and skewness is consistent with the economic interpretation that exchange rates follow target zones (explicit or implicit) with endogenous realignment risk, as in Bertola and Caballero (1992). To generate the positive correlation, this endogenous realignment risk must vary positively with the location of the spot rate within the target zone: the closer the spot rate to the ceiling, the greater the probability of an upward realignment, and the closer to the floor, the greater the probability of a downward realignment. This positive correlation could also be consistent with "support" and "resistance" levels often mentioned in the context of technical analysis. Once an exchange rate crosses a certain critical threshold, it may be perceived to have entered a new "trading range."

To test the robustness of our results, we use supplementary data consisting of five years (April 1992 to April 1997) of daily 25Δ risk-reversal prices on the dollar-mark, dollar-yen, mark-pound, mark-lira, and pound-dollar. We are unable to estimate the full pdf because we do not know the level of implied volatility corresponding to 25Δ calls and puts, and furthermore are missing the 10Δ options.[19] Nonetheless, risk reversals alone suffice to determine the relation between skewness and the spot rate. We perform regressions of the change in the risk reversal against a constant and the change in the spot rate, adjusting standard errors for overlapping observations according to Newey and West (1987). As shown in Table 4, over this five-year period, the correlation between skewness and spot is positive for all five currency pairs, and significantly so for all but the pound-mark.

Since options reflect the risk-neutral distribution as opposed to the actual distribution, an alternative interpretation of the observed skewness in the risk-neutral distribution would be based on risk. For example, when the dollar is strong, the market may attach greater subjective valuation (per dollar) to states of the world in which the dollar becomes much stronger. The relatively high valuation of these states may be due to risk-aversion: producers and consumers might, for example, face high adjustment costs if an already-strong dollar were to appreciate significantly.

Under this risk-based explanation, markets do not necessarily view large appreciations of a strong currency as statistically more likely; they may simply attach a higher valuation to states of large appreciation.

V. Conclusion

In this paper, we have derived risk-neutral probability density functions for the dollar-mark, dollar-yen, pound-mark, mark-franc, and mark-lira from April 1996 to March 1997 using weekly over-the-counter options data. By using a relatively broad cross-section of strike prices we are able to provide a more precise estimate of the underlying pdf than most previous work on the subject. We also use three alternative methods of transforming option prices into the pdf---cubic splines, implied binomial trees, and a mixture of lognormals. We demonstrate that despite important methodological differences, all three approaches successfully capture how the underlying distribution differs from the lognormal (Black-Scholes) benchmark. We discuss advantages and disadvantages of each method, and ultimately opt to use the "trimmed binomial tree" approach. This method incorporates most of the advantages of the binomial tree method, but reduces the excess kurtosis by eliminating the extreme tails of the distribution.

There are a number of real-world implications of this options-based approach to deriving the risk-neutral pdf. First, identifying the exchange rate distributions could have important policy applications. Governments can almost instantaneously observe market reactions to their policies or to exogenous shocks. For example, option prices will immediately reflect the impact on market expectations of foreign exchange intervention, providing feedback to monetary authorities. Indeed, an understanding of expectations could help central banks time market intervention. One could also measure the impact of announced fiscal policies, changes in the government budget, or other news of economic or political importance. Central banks attempting to maintain an exchange rate within a given target zone, either through formal arrangement or informal exchange rate management, could use the distribution embedded in options to gauge the credibility of the target zone as perceived by the market. If there is international "contagion" in exchange rate crises, a country could gauge its own vulnerability by monitoring its currency's implied pdf in reaction to a speculative attack on another---even when there is little movement in the spot rate.

Second, knowledge of the perceived exchange rate distribution could also be of value to the private sector. In managing portfolios, investors could make decisions based on the market perception of exchange rate risk. Corporations could make hedging decisions based on the uncertainty perceived by the market. On an operational level, companies might use different exchange rate scenarios implied by a given probability distribution to make pricing, marketing, or capital-budgeting decisions. Of course, financial institutions could use these implied pdf's to price options whose payoff was any arbitrary function of the spot rate at expiration.[20]

In brief, deriving the exchange rate's probability distribution using an options-based methodology has a strong motivation from both a theoretical and practical perspective. Foreign exchange options, traded actively on both over-the-counter markets and exchanges worldwide, should in principle reflect the latest publicly available information as it appears on trading screens in the world's major financial centers. Options embody the mathematical principle of a probability distribution, even if forecasters themselves would not normally express forecasts in these terms. Since they trade at a range of strike prices, options permit the economist---with virtually no assumptions about investor preferences or exchange rate dynamics---to use market prices to determine this probability distribution.

After we have identified the pdf's, we proceed to investigate the relation between skewness in the pdf and the level of the spot exchange rate. We construct five different measures of skewness: the 10Δ and 25Δ risk reversals, the skewness coefficient from the estimated pdf, and two measures of "relative intensity," representing the relative size of the right and left tails. Our regression results indicate that for the overwhelming majority of currencies, measures of skewness, and time horizons, skewness in the distribution is a positive function of the spot rate. This means that the stronger a currency, the more expectations are skewed towards a large appreciation rather than a large depreciation. This can be viewed as "extrapolative" expectations: when a currency has already appreciated significantly, it becomes less difficult to imagine a further large appreciation. The consistently positive coefficients suggest that, at least over the April 1996 - March 1997 period, markets have not perceived these exchange rates to be operating within a binding target zone, implicit or explicit. Such findings are also inconsistent with a martingale characterization of exchange rates, favoring a target zone with realignment risk positively correlated with spot. These results were further corroborated by regressions using five years of daily data (April 1992 - April 1997) on 25Δ risk reversals only.

Extensions of this research could go in a number of directions, primarily related to a fuller understanding of exchange rate expectations. Interesting theoretical work might involve exploring alternative exchange rate models with differing implications for the correlation between skewness and the spot rate. One could test more explicitly whether a model of target zones with endogenous realignment risk indeed best fits the observed skewness pattern. Empirically, one could determine whether skewness indeed accurately predicted the correlation between shocks to spot and shocks to volatility, simultaneously testing rational expectations and the implications of a number of stochastic volatility option models. One could also investigate whether skewness in the risk-neutral distribution is correlated with the interest differential. Recall that such skewness may represent risk, i.e. higher subjective valuation of certain states of nature, as well as simple statistical probability. Under the interpretation of risk, a negative correlation between skewness and the interest differential may prove to be an identifiable source of a time-varying risk premium negatively correlated with the interest differential. This would be a step towards understanding more about the puzzling forward premium anomaly.

Another potentially useful direction would be to improve the precision of the estimated pdf by aggregating the data either over a cross-section of horizons (e.g., one-month, three-month, and six-month options) or over time. Each of these would require assumptions about the stationarity of the distribution of exchange rate changes, but could provide greater insights into the distribution, especially in the tail regions. If one assumed the distribution of exchange rate changes to be stationary, one could also compare the *ex ante* implied distribution in option prices with the series of *ex post* realized distributions. Ignoring small-sample issues such as peso problems, a comparison of the implied risk-neutral distribution with the realized statistical distribution could reveal interesting information regarding the "pricing kernel" mapping statistical probability into risk-neutral probability and ultimately used to price derivative assets. This could provide further understanding into the nature and pricing of exchange rate risk.

References

Abken, P.A. (1995) Using Eurodollar futures options: gauging the market's view of interest rate movements. *Economic Review* **16** (Federal Reserve Bank of Atlanta), 10-30.

Ait-Sahalia, Yacine, and Andrew W. Lo (1995) Nonparametric estimation of state-price densities implicit in financial asset prices. Working Paper.

Bank for International Settlements (1996) *Central Bank Survey of Foreign Exchange Market Activity in April 1995*, Basel: Bank for International Settlements.

Bates, David S. (1991) The Crash of '87: was it expected? *Journal of Finance* **46**, 1009-1044.

Bates, David S. (1996a) Dollar jump fears, 1984-1992: distributional abnormalities implicit in currency futures options. *Journal of International Money and Finance* **15**, 65-93.

Bates, David S. (1996b) Testing option pricing models. *Statistical Methods in Finance* **14567-611.**

Bertola, Giuseppe and Ricardo Caballero (1992) Stochastic devaluation risk and the empirical fit of target-zone models. *American Economic Review* **82**, 520-536.

Breeden, Douglas and Robert Litzenberger (1978) Prices of state-contingent claims implicit in option price. *Journal of Business* **51**, 621-651.

Campa, José M. and P.H. Kevin Chang (1995) Testing the expectations hypothesis of the term structure of volatilities in foreign exchange options. *Journal of Finance* **50**, 529-547.

Campa, José M. and P.H. Kevin Chang (1996a) Arbitrage-based tests of target zone credibility: evidence from ERM cross-rate options. *American Economic Review* **86**, 726-740.

Campa, José M. and P.H. Kevin Chang (1996b) The forecasting ability of correlations implied in foreign exchange options. Working Paper.

Campa, José M., P.H. Kevin Chang, and Robert L. Reider (1997) ERM bandwidths for EMU and after: evidence from foreign exchange options. *Economic Policy* **24**, 55-89 .

Chiras, Donald P. and Steven Manaster (1978) The information content of option prices and a test of market efficiency. *Journal of Financial Economics* **10**, 213-234.

Day, T. and C. Lewis (1992) Stock market volatility and the information content of stock index options. *Journal of Econometrics* **52**, 267-287.

Dooley, Michael P., and Jeffrey Shafer (1984) Analysis of short-run exchange rate behavior: March 1973 to November 1981. In: *Floating Exchange Rates and the State of World Trade and Payments*, eds Bigman and Taya, Cambridge, MA: Ballinger Publishing Company.

Engel, Charles M. and James D. Hamilton (1990) Long swings in the dollar: are they in the data and do the markets know it? *American Economic Review* **80**, 689-713.

Grundy, Bruce D. (1991) Options prices and the underlying asset's return distribution. *Journal of Finance* **46**, 1045-1069.

Heston, Steven L. (1993) A closed-form solution for options with Stochastic volatility with applications to bond and currency options. *Review of Financial Studies* **6(2)**, 327-343.

Hsieh, David A. (1988) Statistical properties of daily exchange rates: 1974-1983. *Journal of International Economics* **24**, 124-145.

Hsieh, David A. (1989) Modelling heteroscedasticity in daily foreign-exchange rates. *Journal of Business Economics and Statistics* **7**, 307-317.

Hull, John and Alan White (1987) The pricing of options on assets with Stochastic volatilities. *Journal of Finance* **42(2)**, 281-300.

Jackwerth, Jens Carsten and Mark Rubinstein (1996) Recovering probability distributions from contemporaneous security prices. *Journal of Finance* **51(5)**, 1611-1631.

Jorion, Philippe (1995) Predicting volatility in the foreign exchange market. *Journal of Finance* **50(2)**, 507-528.

Kroner, Kenneth R., Kevin P. Kneafsey, and Stijn Claessens (1993) Forecasting volatility in commodity markets. World Bank Policy Research Working Paper,no. 1226.

Lamoureux, Christopher G. and William D. Lastrapes (1993) Forecasting stock-return variance: toward an understanding of Stochastic implied volatilities. *Review of Financial Studies* **6**, 293-326.

Leahy, Michael P. and Charles P. Thomas (1996) The sovereignty option: market views on the Quebec referendum. International Finance Discussion Papers,no. 555, Federal Reserve Board.

Malz, Allan M. (1996a) Using option prices to estimate realignment probabilities in the European Monetary System: the case of Sterling-Mark. *Journal of International Money and Finance* **15**, 717-748.

Malz, Allan M. (1996b) Option-based estimates of the probability distribution of exchange rates and currency excess returns. Federal Reserve Bank of New York Working Paper.

McCauley, Robert and William R. Melick (1996a) Risk reversal risk. *Risk* **9(11)**, 54-57.

McCauley, Robert and William R. Melick (1996b) Propensity and density. *Risk* **9(12)**, 52-54.

Melick, William R. and Charles P. Thomas (1997) Recovering an asset's implied pdf from options prices: an application to crude oil during the Gulf crisis. *Journal of Financial and Quantitative Analysis* **32(1)**, 91-115.

Mizrach, Bruce (1996) Did option prices predict the ERM crisis? Rutgers University Department of Economics Working Paper 96-10.

Newey, Whitney K. and Kenneth D. West (1987) Hypothesis testing with efficient method of moments estimation. *International Economic Review* **28**, 777-787.

Rosenberger, James and Miriam Gasko (1983) Comparing location estimators: trimmed means, medians, and trimean. In: *Understanding Robust and Exploratory Data Analysis*, eds David Hoaglin, Frederick Mosteller, and John Tuckey, New York: John Wiley.

Rubinstein, Mark (1994) Implied binomial trees. *Journal of Finance* **49**, 771-818.

Scott, Louis O. (1992) The information content of prices in derivative security markets. *IMF Staff Papers* **39**, 596-625.

Shimko, David (1993) Bounds of probability. *Risk* **6**, 33-37.

Svensson, Lars E. O. (1992) An interpretation of recent research on exchange rate target zones. *Journal of Economic Perspectives* **6**, 119-144.

1. The risk-neutral distribution embedded in option prices may, of course, differ from the actual distribution. Grundy (1991) discusses bounds on how far these two distributions may diverge. The two distributions will be equal if investors are truly risk-neutral, or if risk in the underlying security (in this case the exchange rate) is not priced. This is much more likely to hold for exchange rates than for an underlying variable highly correlated with wealth, such as the S&P.

2. Indeed, many papers have shown that options contain information not found in the underlying time series that is useful for predicting future volatility (e.g. Chiras and Manaster (1978), Day and Lewis (1992) and Lamoureux and Lastrapes (1993) for stocks; Jorion (1995), Scott (1992) for exchange rates; Campa and Chang (1996b) for correlations between dollar-mark and dollar-yen exchange rates; and Kroner, Kneafsey, and Claessens (1993) for commodities)).

3. As discussed below, probabilities in this context in fact refer to "risk-neutral" probabilities.

4. In earlier research based on over-the-counter option quotes against the U.S. dollar (Campa and Chang (1995)), we found data from different banks to differ insignificantly. Also, the presence of currency option brokers, who use quotes from different banks to provide the best price to the market, contributes to some degree of price convergence.

5. The Garman-Kohlhagen formula is: $Call_{K,T}= S_0*\exp(-r_t^*T)N(d_1)-K*\exp(-r_TT)N(d_2)$, where $d_1=[\ln(S_0/K)+(r_T-r_t^*+\frac{1}{2}\sigma^2)T]/\sigma\sqrt{T}$, $d_2=d_1-\sigma\sqrt{T}$, σ is the implied volatility of the option, r and r^* are the continuously compounded domestic and foreign interest rates, and N is the cumulative normal distribution function.

6. The "delta" of a call equals $\exp(-r_t^*T)N(d_1)$, and the "delta" of a put equals $\exp(-r_t^*t)[N(d_1)-1]$ where $d_1 = [\ln(S_0/K)+(r_T-r_t^*+\frac{1}{2}\sigma^2)T]/\sigma\sqrt{T}$, σ is the implied volatility of the option, r and r^* are the continuously compounded domestic and foreign interest rates, and N is the cumulative normal density function. Given the implied volatilities, interest rates, time-to-expiration, and spot rate, a given delta for a put or call uniquely defines the strike price K.

7. Later in the paper we discuss possible aggregation across one or more of these dimensions as a means of increasing the number of observed data points used to estimate the pdf.

8. At the first and last data points, we assume that the second derivative equals zero, an approach referred to as the "natural spline."

9. Alternatively, one could use a loss function in (4) in which deviations from prior probabilities were penalized more heavily for realizations far from the forward rate, at the expense of a somewhat worse fit in the central part of the distribution.

10. An exception is the mark-franc, where up and down movements in the binomial tree are small because at-the-money volatility is low, and trimming leads to an exclusion of 1.7-1.8% of total probability.

11. Note that although the constraint was not explicitly imposed, the total probability in the cubic spline approach is close to one.

12. Bates (1996b) provides an excellent summary of these results.

13. The term "intensity" has been most commonly used in the literature on realignments in target zone regimes. The "intensity" of realignment is often defined as the sum, over all possible realignments, of the probability of that realignment times its size. See, for example, Svensson (1992) or Campa and Chang (1996).

14. This rule holds only approximately because, for both puts and calls, the delta is based on $N(d_1)$ while the probability of expiring in-the-money is based on $N(d_2)$, where $d_2 = d_1 - \sigma\sqrt{T}$.

15. The vega of a call or a put option equals $S_0 * \exp(-r_t^* T) N'(d_1)\sqrt{t}$. For a put and call of the same (absolute value) delta, d_1 will be of equal absolute value but opposite sign. Since $N'(x) = (1/\sqrt{2}\sqrt{\pi})\exp(-x^2)$ is an even function, vega of the put and call will be equal.

16. Malz (1996b) and McCauley and Melick (1996a) use at-the-money-forward straddles, 25Δ risk reversals, and 25Δ strangles to identify three points on the volatility smile for the one-month horizon.

17. Because of quoting conventions in over-the-counter markets, the exception is the pound-mark, where the effect of the risk-reversal appears in the call.

18. Mean-reversion may also be suggested by the evidence in Bates (1996a): negative skewness in the DEM/USD exchange rate disappeared or became less evident after the dollar declined from its record strengths in 1984-1985. Similarly, Abken (1995) notes that skewness increased sharply in Eurodollar futures options---indicating higher probability of a large rise in interest rates---over the 1993-1994 period when interest rates were unusually low. For the stock market, however, Jackwerth and Rubinstein (1996) find little evidence of time variation in skewness.

19. Although most banks today record option volatility data for a range of strike prices, with 25Δ options more commonly recorded than 10Δ options, the majority of banks have historical data only for at-the-money forward volatility.

20. Path-dependent options such as Asian options or barrier options could not, however, be priced from the pdf alone and would require numerical simulation or the construction of a full tree.

Table 1. Daily Turnover and Notional Amount Outstanding (in millions of dollars) in April 1995

CURRENCY	USD-DEM	USD-JPY	GBP-DEM	DEM-FRF	DEM-ITL
NOTIONAL AMOUNT OUTSTANDING (millions of U.S. dollars)					
OTC Markets	518,720	625,163	410,227 (DM vs. all currencies other than dollar + yen)		
Exchanges	36,395	21,749	2,225 (DM vs. all currencies other than dollar + yen)		
AVERAGE DAILY TURNOVER in April 1995 (millions of U.S. dollars)					
OTC Markets	10,241	13,266	1,366	1,878	1,876 (all other EMS vs. ITL)
Exchanges	1223	842	9	0	

Notional Amounts: BIS (1996), Tables 3-A-2
Daily Turnover: BIS (1996), Tables 9-B-2, 9-C-2

Table 2a. Moments of the Implied Distributions for ITL/DEM Returns

1 Month Options

MEAN	AVERAGE	MINIMUM	MAXIMUM
BINOMIAL	0.004006	0.002840	0.005086
MIXED LOG NORMALS	0.003490	0.001753	0.005263
TRIMMED BINOMIAL (20%)	0.003651	0.002399	0.005000
CUBIC SPLINE	0.001731	0.001094	0.002251

STANDARD DEVIATION	AVERAGE	MINIMUM	MAXIMUM
BINOMIAL	0.01993	0.01525	0.03048
MIXED LOG NORMALS	0.01934	0.01495	0.02828
TRIMMED BINOMIAL (20%)	0.01631	0.01228	0.02552
CUBIC SPLINE	0.007874	0.005845	0.01198

SKEWNESS	AVERAGE	MINIMUM	MAXIMUM
BINOMIAL	1.306	0.1444	2.365
MIXED LOG NORMALS	0.6832	-0.03914	1.442
TRIMMED BINOMIAL (20%)	0.9780	0.1507	1.537
CUBIC SPLINE	1.162	0.5932	1.793

KURTOSIS	AVERAGE	MINIMUM	MAXIMUM
BINOMIAL	16.23	4.515	21.18
MIXED LOG NORMALS	2.192	0.9810	4.930
TRIMMED BINOMIAL (20%)	1.346	0.03556	3.653
CUBIC SPLINE	2.379	1.078	4.807
LEFT OUT PROBABILITY	0.008365	0.004254	0.01581

3-Month Options

MEAN	AVERAGE	MINIMUM	MAXIMUM
BINOMIAL	0.01137	0.008756	0.01480
MIXED LOG NORMALS	0.01059	0.006538	0.01481
TRIMMED BINOMIAL (20%)	0.01084	0.007980	0.01472
CUBIC SPLINE	0.005034	0.003654	0.007970

STANDARD DEVIATION	AVERAGE	MINIMUM	MAXIMUM
BINOMIAL	0.03502	0.02660	0.05345
MIXED LOG NORMALS	0.03439	0.02395	0.05046
TRIMMED BINOMIAL (20%)	0.02990	0.02153	0.04639
CUBIC SPLINE	0.01421	0.01067	0.02132

SKEWNESS	AVERAGE	MINIMUM	MAXIMUM
BINOMIAL	1.401	0.4716	2.367
MIXED LOG NORMALS	1.076	0.2030	1.933
TRIMMED BINOMIAL (20%)	1.414	1.135	1.835
CUBIC SPLINE	1.509	1.125	2.456

KURTOSIS	AVERAGE	MINIMUM	MAXIMUM
BINOMIAL	12.35	5.333	17.37
MIXED LOG NORMALS	2.146	0.6242	4.379
TRIMMED BINOMIAL (20%)	1.823	0.4848	3.912
CUBIC SPLINE	2.408	0.9541	6.797
LEFT OUT PROBABILITY	0.008928	0.005315	0.01739

Table 2b. Moments of the Implied Distributions for FRF/DEM Returns

1-Month Options

MEAN	AVERAGE	MINIMUM	MAXIMUM
BINOMIAL	0.0002545	-5.590e-05	0.0005348
MIXED LOG NORMALS	-0.0003301	-0.002985	0.0006612
TRIMMED BINOMIAL (20%)	-6.211e-05	-0.0006367	0.0003709
CUBIC SPLINE	6.746e-05	-0.0002462	0.0005769

STANDARD DEVIATION	AVERAGE	MINIMUM	MAXIMUM
BINOMIAL	0.008146	0.004380	0.01156
MIXED LOG NORMALS	0.008976	0.004439	0.01653
TRIMMED BINOMIAL (20%)	0.006267	0.003820	0.009403
CUBIC SPLINE	0.003177	0.001897	0.005817

SKEWNESS	AVERAGE	MINIMUM	MAXIMUM
BINOMIAL	1.934	0.8481	3.375
MIXED LOG NORMALS	0.3946	-1.422	2.752
TRIMMED BINOMIAL (20%)	1.261	0.1360	2.156
CUBIC SPLINE	1.846	0.5949	3.587

KURTOSIS	AVERAGE	MINIMUM	MAXIMUM
BINOMIAL	17.00	9.477	22.15
MIXED LOG NORMALS	6.679	3.905	10.80
TRIMMED BINOMIAL (20%)	4.544	0.2851	8.254
CUBIC SPLINE	8.147	1.468	18.36
LEFT OUT PROBABILITY	0.01707	0.01002	0.02753

3-Month Options

MEAN	AVERAGE	MINIMUM	MAXIMUM
BINOMIAL	0.0009516	9.520e-06	0.002035
MIXED LOG NORMALS	0.0003432	-0.003599	0.001939
TRIMMED BINOMIAL (20%)	0.0004026	-0.001145	0.001955
CUBIC SPLINE	0.0003845	-0.0004452	0.001665

STANDARD DEVIATION	AVERAGE	MINIMUM	MAXIMUM
BINOMIAL	0.01505	0.01096	0.02122
MIXED LOG NORMALS	0.01558	0.009203	0.02404
TRIMMED BINOMIAL (20%)	0.01208	0.007535	0.01633
CUBIC SPLINE	0.005996	0.004063	0.01056

SKEWNESS	AVERAGE	MINIMUM	MAXIMUM
BINOMIAL	2.058	-0.1069	3.428
MIXED LOG NORMALS	1.266	-0.9991	3.210
TRIMMED BINOMIAL (20%)	1.666	1.035	2.314
CUBIC SPLINE	1.833	0.1964	3.473

KURTOSIS	AVERAGE	MINIMUM	MAXIMUM
BINOMIAL	15.10	8.204	22.41
MIXED LOG NORMALS	6.867	3.262	15.46
TRIMMED BINOMIAL (20%)	5.576	2.732	8.308
CUBIC SPLINE	7.061	-5.337	22.85
LEFT OUT PROBABILITY	0.01803	0.01115	0.02764

Table 2c. Moments of the Implied Distributions for DEM/USD Returns

1-Month Options

MEAN	AVERAGE	MINIMUM	MAXIMUM
BINOMIAL	-0.002237	-0.002508	-0.001858
MIXED LOG NORMALS	-0.001946	-0.004120	-0.0005112
TRIMMED BINOMIAL (20%)	-0.001951	-0.002809	-0.001047
CUBIC SPLINE	-0.0009376	-0.001091	-0.0007614

STANDARD DEVIATION	AVERAGE	MINIMUM	MAXIMUM
BINOMIAL	0.02806	0.02200	0.03548
MIXED LOG NORMALS	0.02460	0.01790	0.03475
TRIMMED BINOMIAL (20%)	0.02237	0.01791	0.02967
CUBIC SPLINE	0.01093	0.008653	0.01402

SKEWNESS	AVERAGE	MINIMUM	MAXIMUM
BINOMIAL	-0.7032	-2.081	0.4316
MIXED LOG NORMALS	-0.1977	-2.298	0.6805
TRIMMED BINOMIAL (20%)	-0.3053	-0.7198	-0.04676
CUBIC SPLINE	-0.3923	-1.109	0.1804

KURTOSIS	AVERAGE	MINIMUM	MAXIMUM
BINOMIAL	19.05	14.46	23.37
MIXED LOG NORMALS	1.165	-0.04524	9.994
TRIMMED BINOMIAL (20%)	0.2370	-0.07193	1.720
CUBIC SPLINE	1.413	0.7106	2.953
LEFT OUT PROBABILITY	0.006654	0.004602	0.009742

3-Month Options

MEAN	AVERAGE	MINIMUM	MAXIMUM
BINOMIAL	-0.006870	-0.01312	0.004110
MIXED LOG NORMALS	-0.006345	-0.01129	0.005838
TRIMMED BINOMIAL (20%)	-0.006102	-0.01200	0.005067
CUBIC SPLINE	-0.002878	-0.005579	0.001892

STANDARD DEVIATION	AVERAGE	MINIMUM	MAXIMUM
BINOMIAL	0.04979	0.03941	0.06024
MIXED LOG NORMALS	0.04527	0.03489	0.06221
TRIMMED BINOMIAL (20%)	0.04122	0.03181	0.04910
CUBIC SPLINE	0.01961	0.01532	0.02353

SKEWNESS	AVERAGE	MINIMUM	MAXIMUM
BINOMIAL	-1.119	-1.899	-0.5916
MIXED LOG NORMALS	-0.6325	-2.555	0.4226
TRIMMED BINOMIAL (20%)	-0.4883	-0.8413	0.1038
CUBIC SPLINE	-0.5800	-0.9750	-0.05455

KURTOSIS	AVERAGE	MINIMUM	MAXIMUM
BINOMIAL	17.35	13.73	20.07
MIXED LOG NORMALS	2.926	-0.1316	13.31
TRIMMED BINOMIAL (20%)	0.09261	-0.1017	0.8168
CUBIC SPLINE	0.9287	0.5378	1.404
LEFT OUT PROBABILITY	0.005374	0.004041	0.006852

Table 2d. Moments of the Implied Distributions for JPY/USD Returns

1-Month Options

MEAN	AVERAGE	MINIMUM	MAXIMUM
BINOMIAL	-0.004677	-0.005194	-0.004204
MIXED LOG NORMALS	-0.004175	-0.007408	-0.002988
TRIMMED BINOMIAL (20%)	-0.004184	-0.004941	-0.003563
CUBIC SPLINE	-0.001988	-0.002209	-0.001805

STANDARD DEVIATION	AVERAGE	MINIMUM	MAXIMUM
BINOMIAL	0.02829	0.02181	0.04163
MIXED LOG NORMALS	0.02522	0.01779	0.03906
TRIMMED BINOMIAL (20%)	0.02366	0.01775	0.03794
CUBIC SPLINE	0.01120	0.008580	0.01708

SKEWNESS	AVERAGE	MINIMUM	MAXIMUM
BINOMIAL	-1.248	-2.491	0.08228
MIXED LOG NORMALS	-0.6618	-2.737	0.1569
TRIMMED BINOMIAL (20%)	-0.7141	-1.159	-0.09873
CUBIC SPLINE	-0.8678	-1.591	-0.1052

KURTOSIS	AVERAGE	MINIMUM	MAXIMUM
BINOMIAL	15.51	9.858	17.95
MIXED LOG NORMALS	2.258	-0.1235	15.99
TRIMMED BINOMIAL (20%)	0.9304	0.1059	2.286
CUBIC SPLINE	1.673	0.4755	4.081
LEFT OUT PROBABILITY	0.006035	0.003799	0.007518

3-Month Options

MEAN	AVERAGE	MINIMUM	MAXIMUM
BINOMIAL	-0.01420	-0.01528	-0.01352
MIXED LOG NORMALS	-0.01389	-0.01896	-0.01157
TRIMMED BINOMIAL (20%)	-0.01306	-0.01485	-0.01218
CUBIC SPLINE	-0.006064	-0.006486	-0.005778

STANDARD DEVIATION	AVERAGE	MINIMUM	MAXIMUM
BINOMIAL	0.05231	0.04006	0.07129
MIXED LOG NORMALS	0.04827	0.03483	0.06931
TRIMMED BINOMIAL (20%)	0.04407	0.03262	0.06127
CUBIC SPLINE	0.02078	0.01567	0.02847

SKEWNESS	AVERAGE	MINIMUM	MAXIMUM
BINOMIAL	-1.635	-2.361	-0.6227
MIXED LOG NORMALS	-1.141	-2.821	-0.2915
TRIMMED BINOMIAL (20%)	-0.9649	-1.250	-0.6178
CUBIC SPLINE	-1.088	-1.451	-0.5220

KURTOSIS	AVERAGE	MINIMUM	MAXIMUM
BINOMIAL	14.16	9.996	17.37
MIXED LOG NORMALS	2.816	-0.2377	13.09
TRIMMED BINOMIAL (20%)	0.5958	-0.09561	1.631
CUBIC SPLINE	1.289	0.4661	2.142
LEFT OUT PROBABILITY	0.005482	0.004074	0.008482

Table 2e. Moments of the Implied Distributions for DEM/GBP Returns

1-Month Options

MEAN	AVERAGE	MINIMUM	MAXIMUM
BINOMIAL	-0.002573	-0.003246	-0.002148
MIXED LOG NORMALS	-0.002334	-0.003753	-0.001251
TRIMMED BINOMIAL (20%)	-0.003128	-0.004069	-0.002532
CUBIC SPLINE	-0.001103	-0.001408	-0.0009320

STANDARD DEVIATION	AVERAGE	MINIMUM	MAXIMUM
BINOMIAL	0.02281	0.01563	0.03514
MIXED LOG NORMALS	0.02082	0.01377	0.03292
TRIMMED BINOMIAL (20%)	0.02018	0.01357	0.03245
CUBIC SPLINE	0.009248	0.006325	0.01456

SKEWNESS	AVERAGE	MINIMUM	MAXIMUM
BINOMIAL	1.239	0.4068	1.745
MIXED LOG NORMALS	-0.2229	-0.9431	0.1874
TRIMMED BINOMIAL (20%)	-0.8210	-1.498	-0.4588
CUBIC SPLINE	-0.3988	-0.9255	-0.04594

KURTOSIS	AVERAGE	MINIMUM	MAXIMUM
BINOMIAL	14.90	11.04	17.85
MIXED LOG NORMALS	0.4684	-0.2567	2.793
TRIMMED BINOMIAL (20%)	0.7211	-0.1399	3.390
CUBIC SPLINE	0.8865	0.3963	1.821
LEFT OUT PROBABILITY	0.002896	0.001794	0.003941

3-Month Options

MEAN	AVERAGE	MINIMUM	MAXIMUM
BINOMIAL	-0.007805	-0.009662	-0.006343
MIXED LOG NORMALS	-0.008269	-0.01094	-0.005568
TRIMMED BINOMIAL (20%)	-0.008757	-0.01111	-0.007177
CUBIC SPLINE	-0.003320	-0.004111	-0.002720

STANDARD DEVIATION	AVERAGE	MINIMUM	MAXIMUM
BINOMIAL	0.03900	0.02894	0.05667
MIXED LOG NORMALS	0.03786	0.02469	0.05426
TRIMMED BINOMIAL (20%)	0.03432	0.02443	0.05144
CUBIC SPLINE	0.01586	0.01149	0.02338

SKEWNESS	AVERAGE	MINIMUM	MAXIMUM
BINOMIAL	1.024	0.1515	1.550
MIXED LOG NORMALS	-0.8600	-2.659	-0.3410
TRIMMED BINOMIAL (20%)	-1.057	-1.392	-0.7283
CUBIC SPLINE	-0.6647	-1.133	-0.3074

KURTOSIS	AVERAGE	MINIMUM	MAXIMUM
BINOMIAL	14.53	10.72	17.52
MIXED LOG NORMALS	1.755	-0.3263	11.36
TRIMMED BINOMIAL (20%)	0.5465	-0.1550	1.474
CUBIC SPLINE	0.7983	0.3485	1.568
LEFT OUT PROBABILITY	0.002794	0.001709	0.003701

Table 3a. Skewness (1-Month) Measures as a Function of $\ln(S_t)$ (in first differences)

ITL/DEM 1-Month	Constant	Spot	Rbar-sq.	DW
10Δ Risk Reversal	0.0845 (0.0520)	52.60 (7.71)	0.45	2.04
25Δ Risk Reversal	0.0364 (0.0251)	28.25 (3.72)	0.51	2.20
Skewness (trimmed)	0.0147 (0.0388)	11.96 (5.76)	0.06	2.42
Relative Intensity (±σ)	0.0001 (0.0001)	0.12 (0.02)	0.43	2.32
Relative Intensity (±1.5σ)	0.0000 (0.0001)	0.04 (0.01)	0.19	1.94
FRF/DEM 1-Month				
10Δ Risk Reversal	0.0525 (0.0540)	89.54 (26.48)	0.18	1.81
25Δ Risk Reversal	0.0205 (0.0238)	48.37 (11.69)	0.26	2.01
Skewness (trimmed)	0.0267 (0.0441)	34.18 (21.65)	0.03	2.14
Relative Intensity (±σ)	0.0000 (0.0001)	0.13 (0.03)	0.23	2.36
Relative Intensity (±1.5σ)	0.0000 (0.0000)	0.05 (0.01)	0.16	2.48
DEM/USD 1-Month				
10Δ Risk Reversal	-0.1337 (0.0664)	47.61 (5.91)	0.58	2.18
25Δ Risk Reversal	-0.0633 (0.0309)	22.80 (2.75)	0.59	2.09
Skewness (trimmed)	-0.0087 (0.0157)	3.28 (1.39)	0.09	3.29
Relative Intensity (±σ)	0.0000 (0.0001)	0.02 (0.01)	0.03	2.44
Relative Intensity (±1.5σ)	0.0000 (0.0001)	-0.00 (0.01)	-0.02	2.47

Weekly observations 3 April 1996 - 5 March 1997

Standard errors in parentheses.

Table 3a (continued). Skewness (1-Month) Measures as a Function of $\ln(S_t)$ (in first differences)

	Constant	Spot	Rbar-sq.	DW
JPY/USD 1-Month				
10Δ Risk Reversal	-0.1543 (0.0631)	62.52 (5.96)	0.70	2.41
25Δ Risk Reversal	-0.0734 (0.0288)	29.61 (2.72)	0.71	2.37
Skewness (trimmed)	-0.0227 (0.0181)	10.11 (1.71)	0.42	2.45
Relative Intensity (±σ)	0.0000 (0.0002)	0.03 (0.02)	0.00	2.40
Relative Intensity (±1.5σ)	0.0000 (0.0001)	0.00 (0.01)	-0.02	2.31
DEM/GBP 1-Month				
10Δ Risk Reversal	-0.1320 (0.0527)	36.49 (5.20)	0.51	1.72
25Δ Risk Reversal	-0.0684 (0.0304)	18.80 (3.00)	0.45	1.67
Skewness (trimmed)	-0.0153 (0.0193)	7.10 (1.90)	0.22	2.31
Relative Intensity (±σ)	0.0003 (0.0002)	-0.03 (0.02)	0.03	2.35
Relative Intensity (±1.5σ)	0.0001 (0.0001)	-0.01 (0.01)	0.04	2.29

Weekly observations 3 April 1996 - 5 March 1997

Standard errors in parentheses.

Table 3b. Skewness (3-Month) Measures as a Function of $\ln(S_t)$ (in first differences)

	Constant	Spot	Rbar-sq.	DW
ITL/DEM 3-Month				
10Δ Risk Reversal	0.0735 (0.0442)	36.82 (6.55)	0.36	2.30
25Δ Risk Reversal	0.0332 (0.0204)	21.71 (3.02)	0.48	2.45
Skewness (trimmed)	0.0082 (0.0166)	2.42 (2.47)	-0.00	2.42
Relative Intensity (±σ)	0.0002 (0.0002)	0.22 (0.03)	0.45	2.76
Relative Intensity (±1.5σ)	0.0001 (0.0001)	0.09 (0.02)	0.40	2.19
FRF/DEM 3-Month				
10Δ Risk Reversal	0.0583 (0.0480)	62.12 (23.53)	0.11	2.19
25Δ Risk Reversal	0.0193 (0.0204)	37.83 (10.01)	0.22	2.25
Skewness (trimmed)	0.0070 (0.0282)	17.60 (13.83)	0.01	2.02
Relative Intensity (±σ)	0.0000 (0.0001)	0.22 (0.04)	0.35	2.39
Relative Intensity (±1.5σ)	0.0000 (0.0000)	0.08 (0.02)	0.20	2.32
DEM/USD 3-Month				
10Δ Risk Reversal	-0.0452 (0.0328)	19.48 (2.92)	0.48	2.35
25Δ Risk Reversal	-0.0232 (0.0160)	9.85 (1.42)	0.50	2.26
Skewness (trimmed)	0.0039 (0.0271)	-1.37 (2.42)	-0.02	2.98
Relative Intensity (±σ)	-0.0000 (0.0002)	0.01 (0.02)	-0.01	2.26
Relative Intensity (±1.5σ)	0.0000 (0.0001)	-0.01 (0.01)	-0.01	2.73

Weekly observations 3 April 1996 - 5 March 1997

Standard errors in parentheses.

Table 3b (continued). Skewness (3-Month) Measures as a Function of $\ln(S_t)$ (in first differences)

	Constant	Spot	Rbar-sq.	DW
JPY/USD 3-Month				
10Δ Risk Reversal	-0.0791 (0.0339)	34.25 (3.21)	0.71	2.50
25Δ Risk Reversal	-0.0383 (0.0161)	16.71 (1.52)	0.72	2.49
Skewness (trimmed)	-0.0098 (0.0127)	3.82 (1.20)	0.16	2.36
Relative Intensity (±σ)	-0.0000 (0.0003)	0.03 (0.03)	0.00	1.96
Relative Intensity (±1.5σ)	-0.0000 (0.0001)	0.01 (0.01)	-0.02	2.47
DEM/GBP 3-Month				
10Δ Risk Reversal	-0.0972 (0.0433)	27.93 (4.27)	0.47	1.72
25Δ Risk Reversal	-0.0506 (0.0242)	14.31 (2.39)	0.47	1.58
Skewness (trimmed)	-0.0079 (0.0101)	4.65 (1.00)	0.31	2.32
Relative Intensity (±σ)	0.0003 (0.0002)	0.00 (0.02)	-0.02	2.30
Relative Intensity (±1.5σ)	0.0002 (0.0001)	-0.02 (0.01)	0.11	2.28

Weekly observations 3 April 1996 - 5 March 1997

Standard errors in parentheses.

Table 4. Daily 25Δ 1-month Risk Reversal as a Function of $\ln(S_t)$ (in first differences)

	Constant	Spot	Rbar-sq.	DW
ITL/DEM	-0.00134 (0.00297)	0.00498 (0.00175)	0.04	2.34
DEM/USD	-0.00012 (0.00238)	3.18590 (0.42703)	0.06	2.38
JPY/USD	0.00043 (0.00355)	0.09217 (0.01191)	0.13	2.38
DEM/GBP	0.00058 (0.00219)	0.65701 (0.49231)	0.00	2.46
USD/GBP	0.00024 (0.00223)	6.91730 (2.16970)	0.02	2.18

These regressions use supplementary data (1 April 1992 - 15 April 1997).

Standard errors, in parentheses, are corrected for overlapping observations using the technique of Newey and West (1987).

Figures

1. Volatility Smile

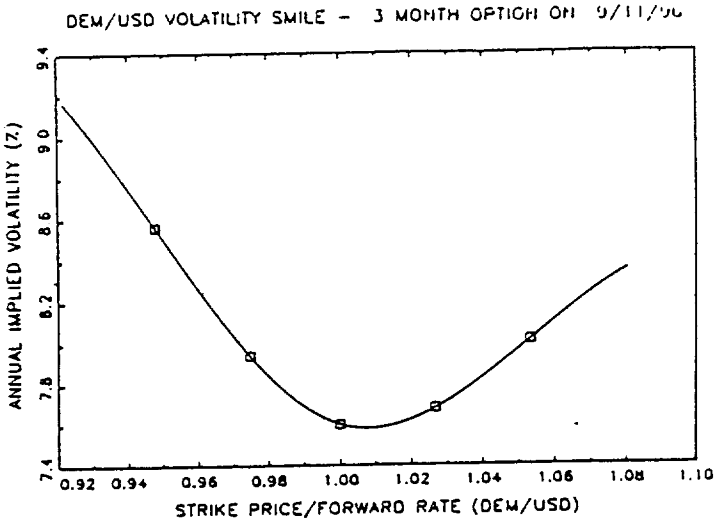

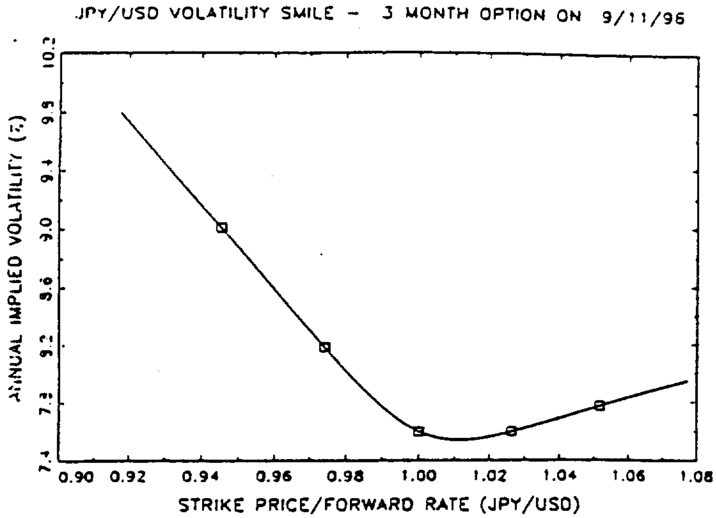

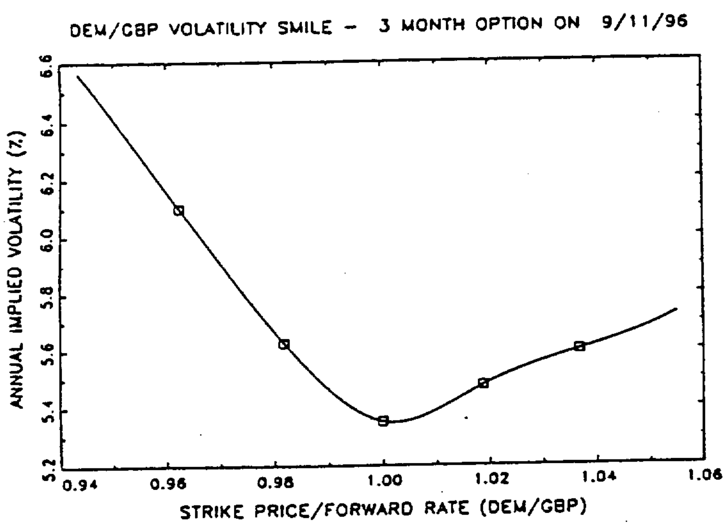

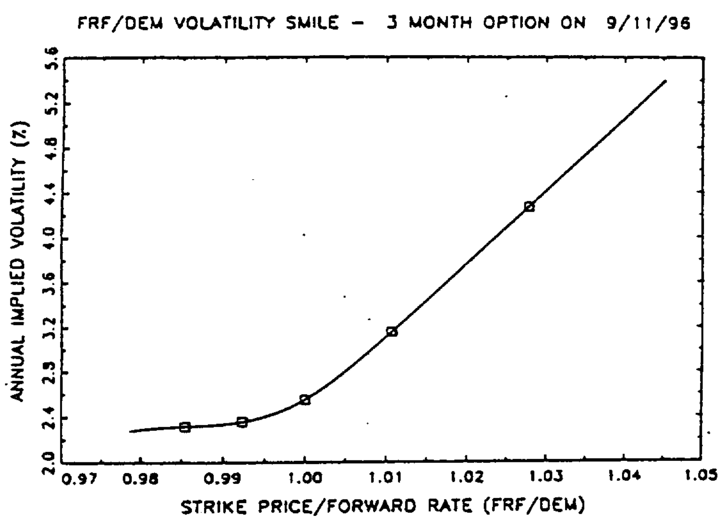

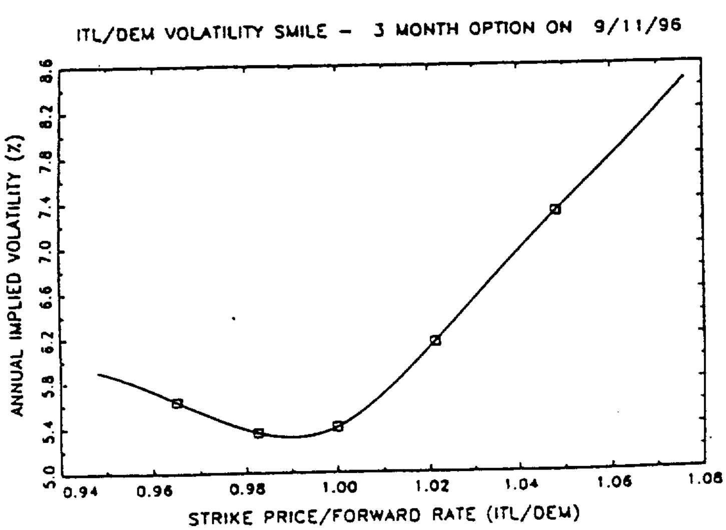

Figure 2a. A Comparison of PDF's: Cubic Spline/Binomial Tree/Mixture of Lognormals vs. Black-Scholes (log normal)

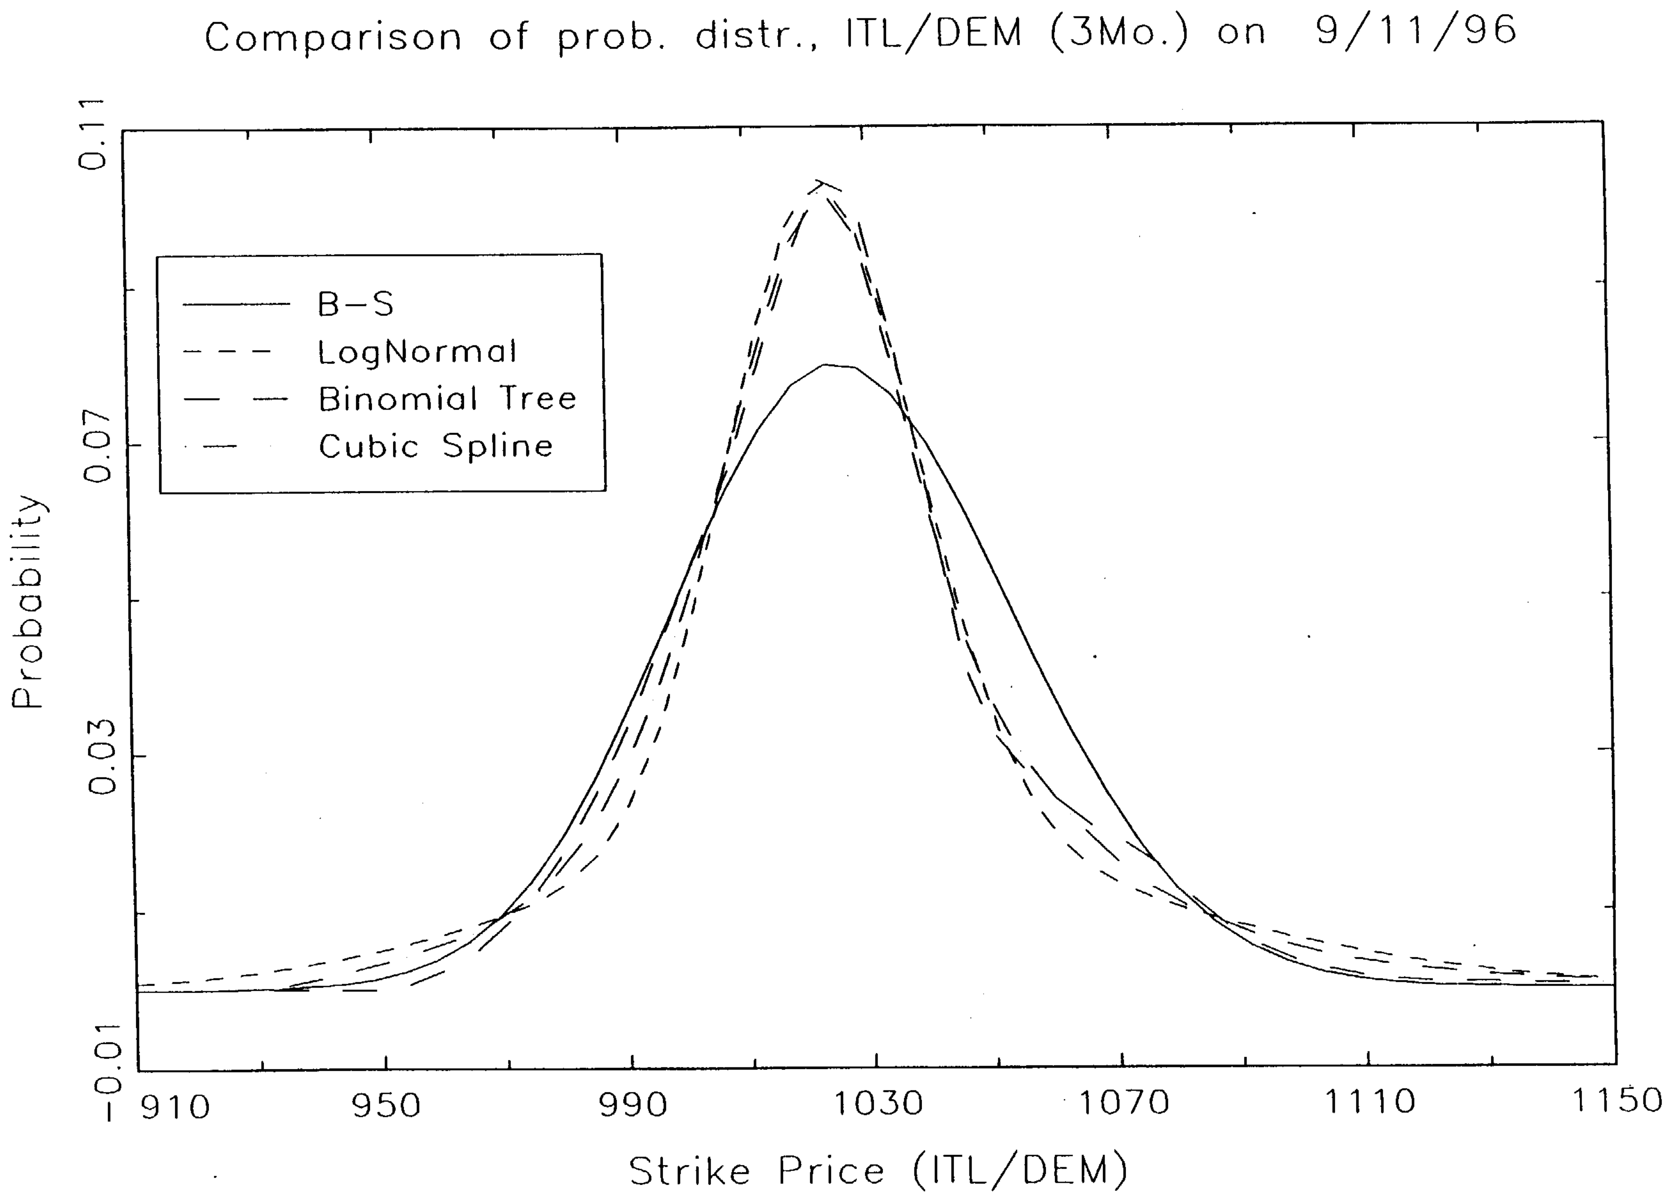

Figure 2b. A Comparison of CDF's: Cubic Spline/Binomial Tree/Mixture of Lognormals vs. Black-Scholes (log normal)

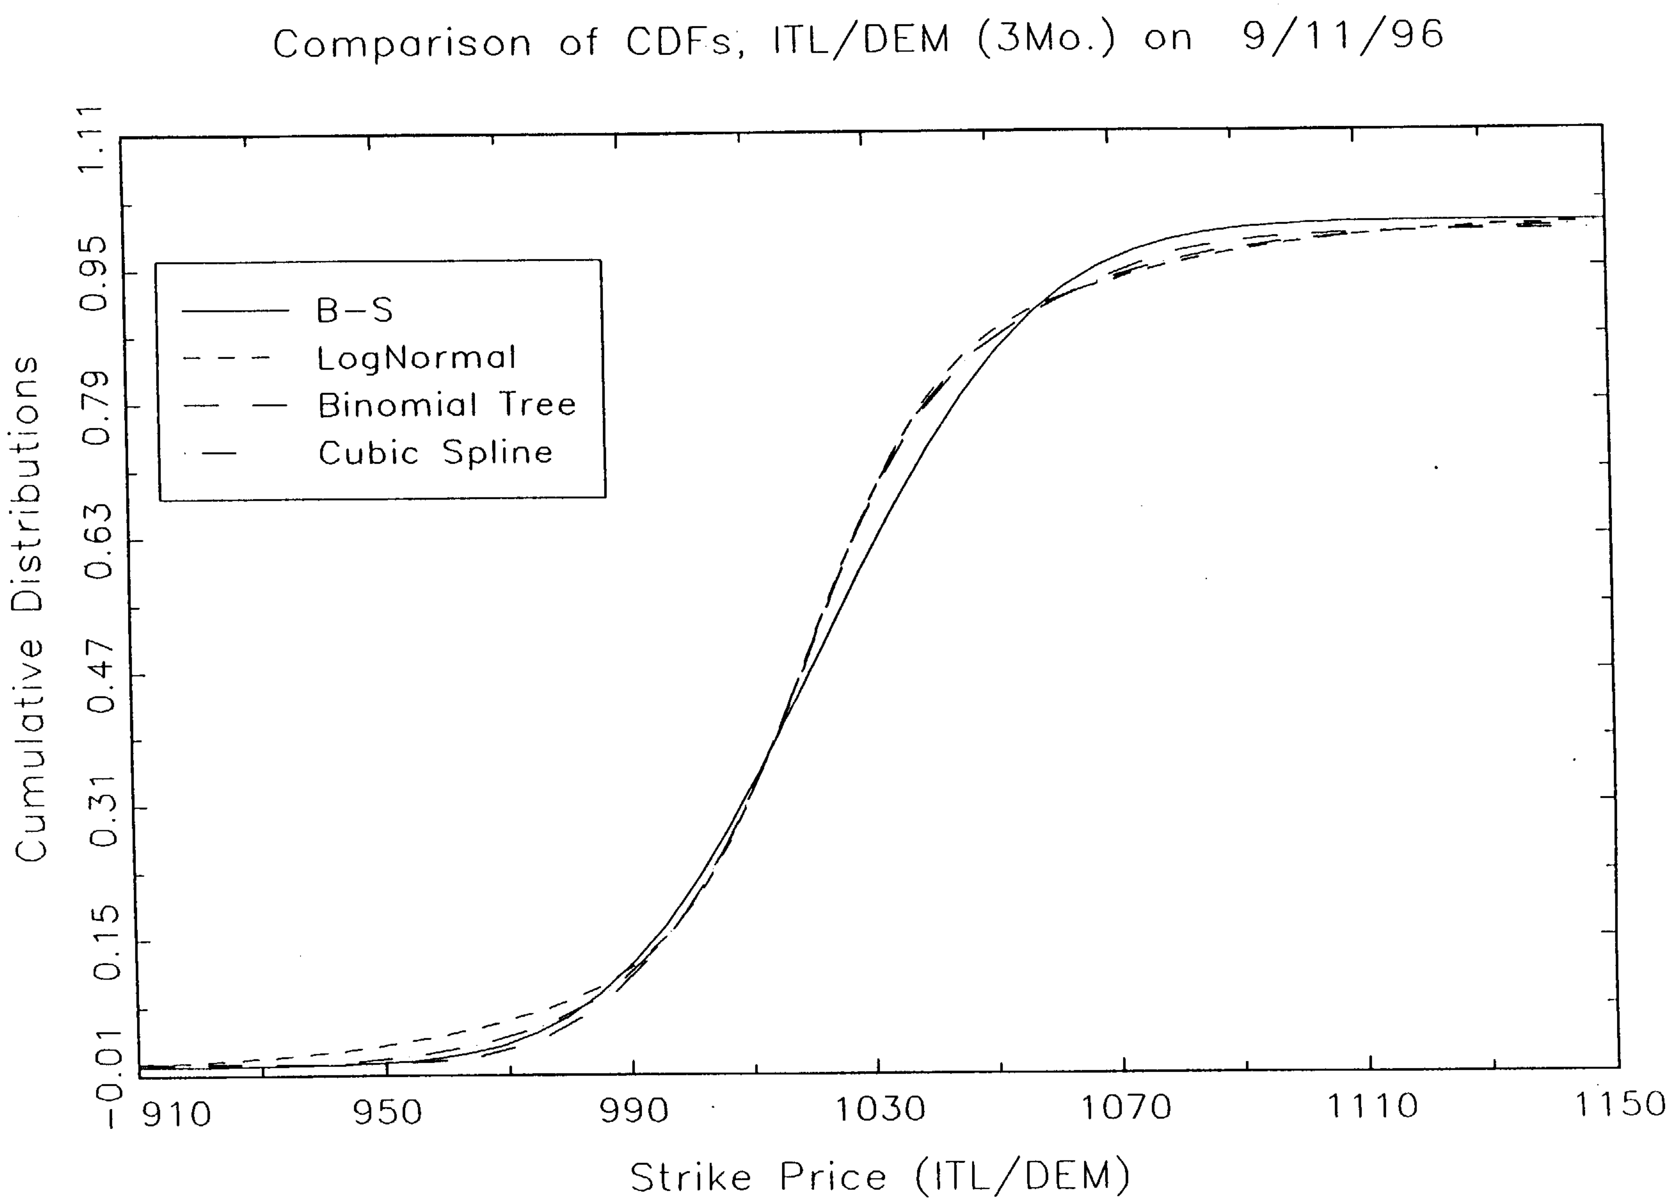

Figure 3a. 10 and 25 Delta Risk Reversals vs. Spot Rate (ITL/DEM)

Risk reversal ≡ (volatility of mark call) minus (volatility of mark put)

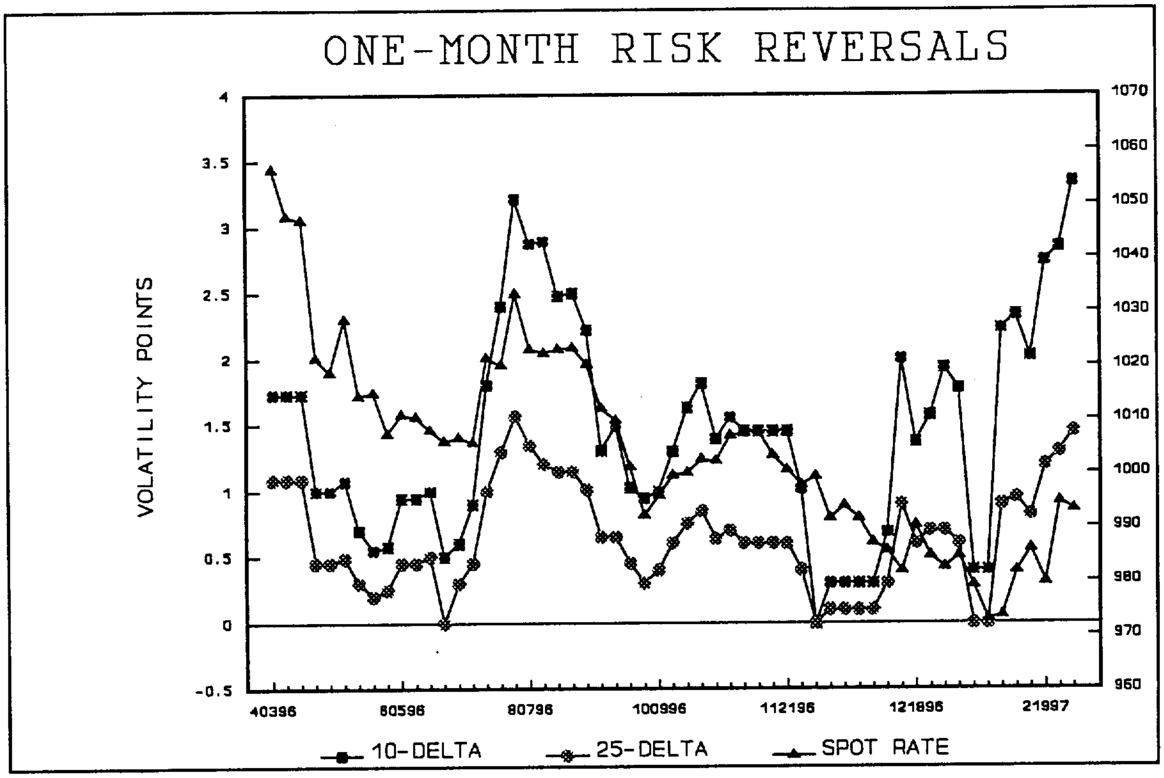

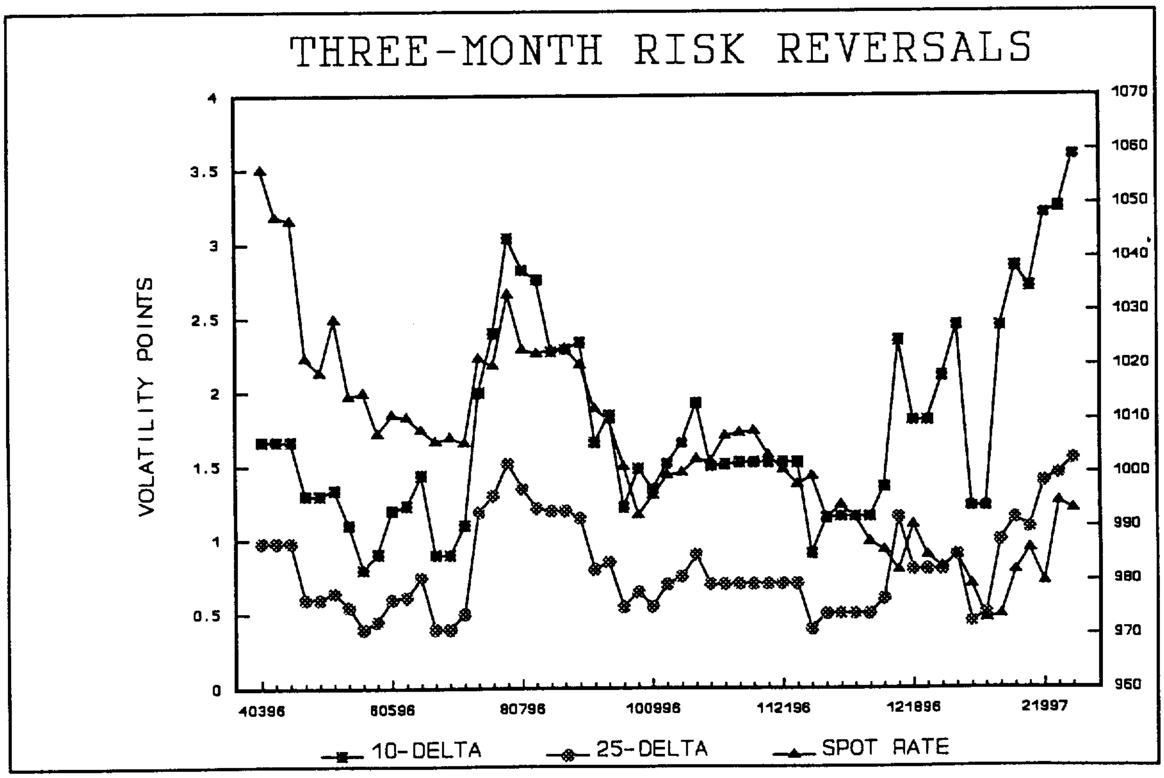

Figure 3b. 10 and 25 Delta Risk Reversals vs. Spot Rate (FRF/DEM)

Risk reversal ≡ (volatility of mark call) minus (volatility of mark put)

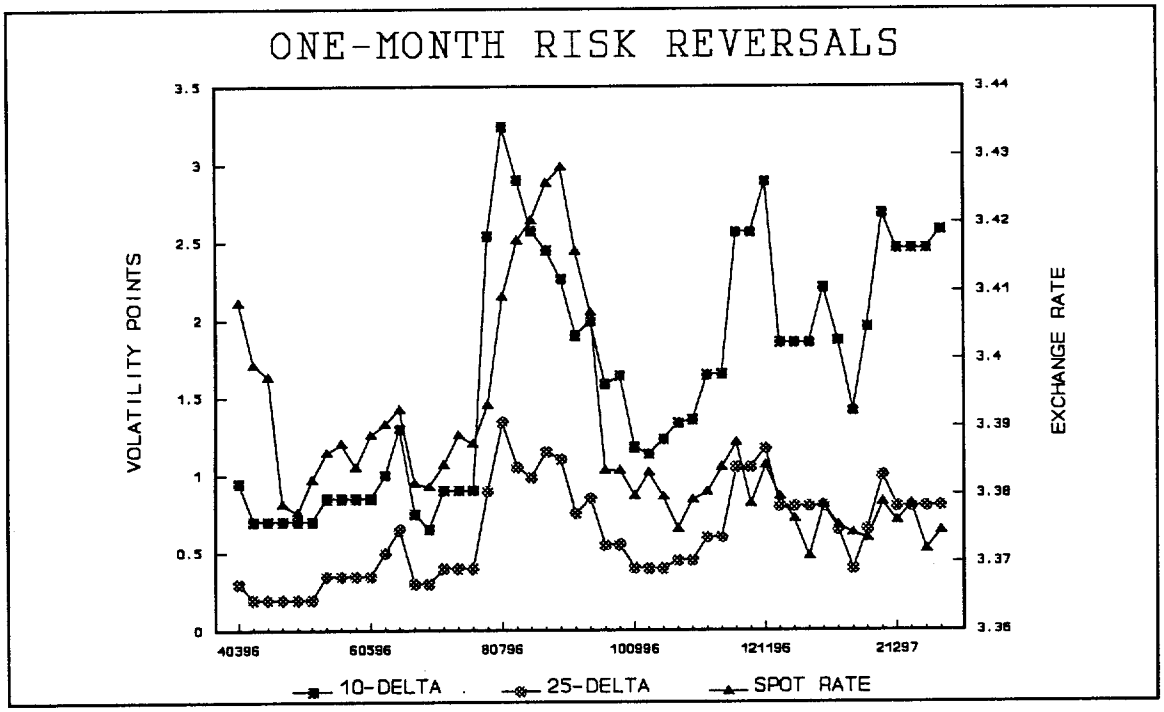

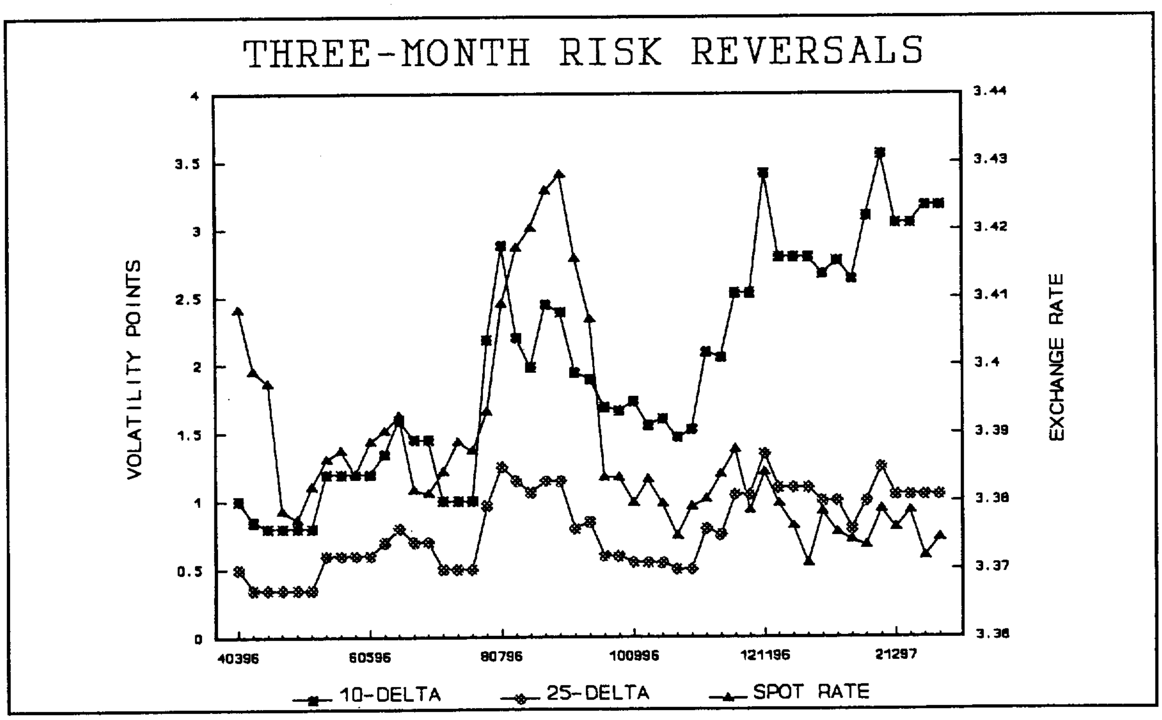

Figure 3c. 10 and 25 Delta Risk Reversals vs. Spot Rate (DEM/USD)

Risk reversal ≡ (volatility of dollar call) minus (volatility of dollar put)

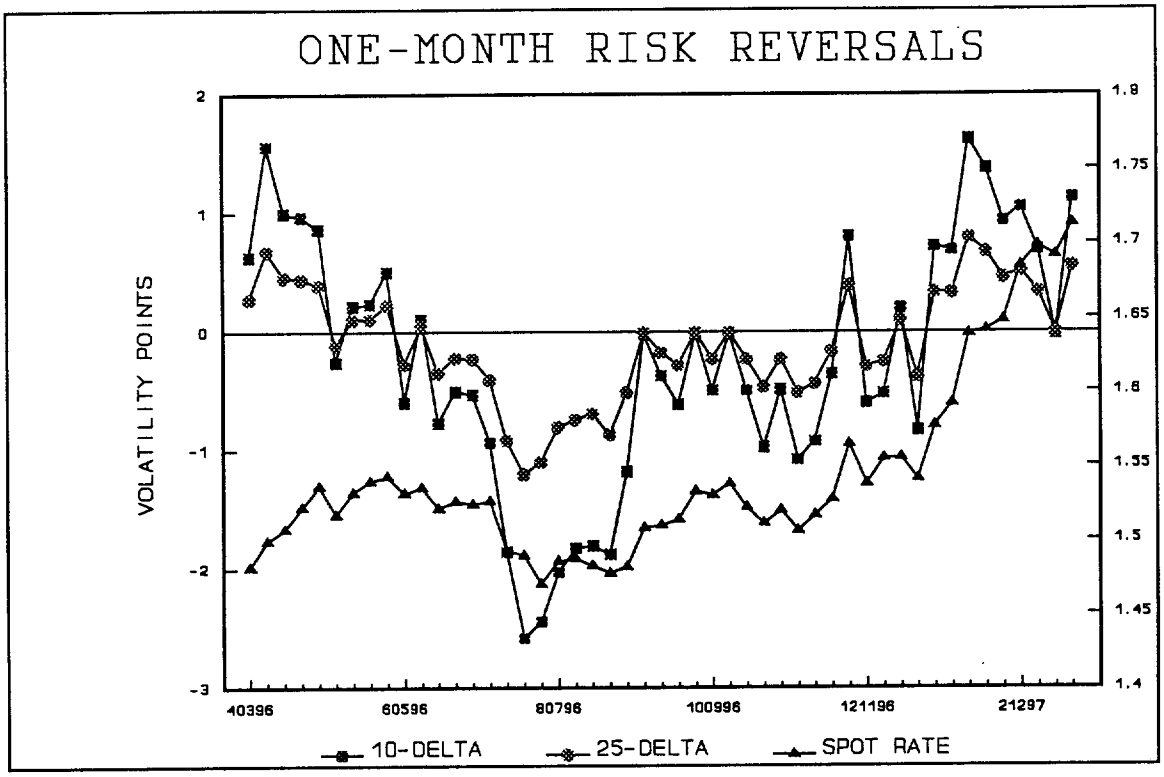

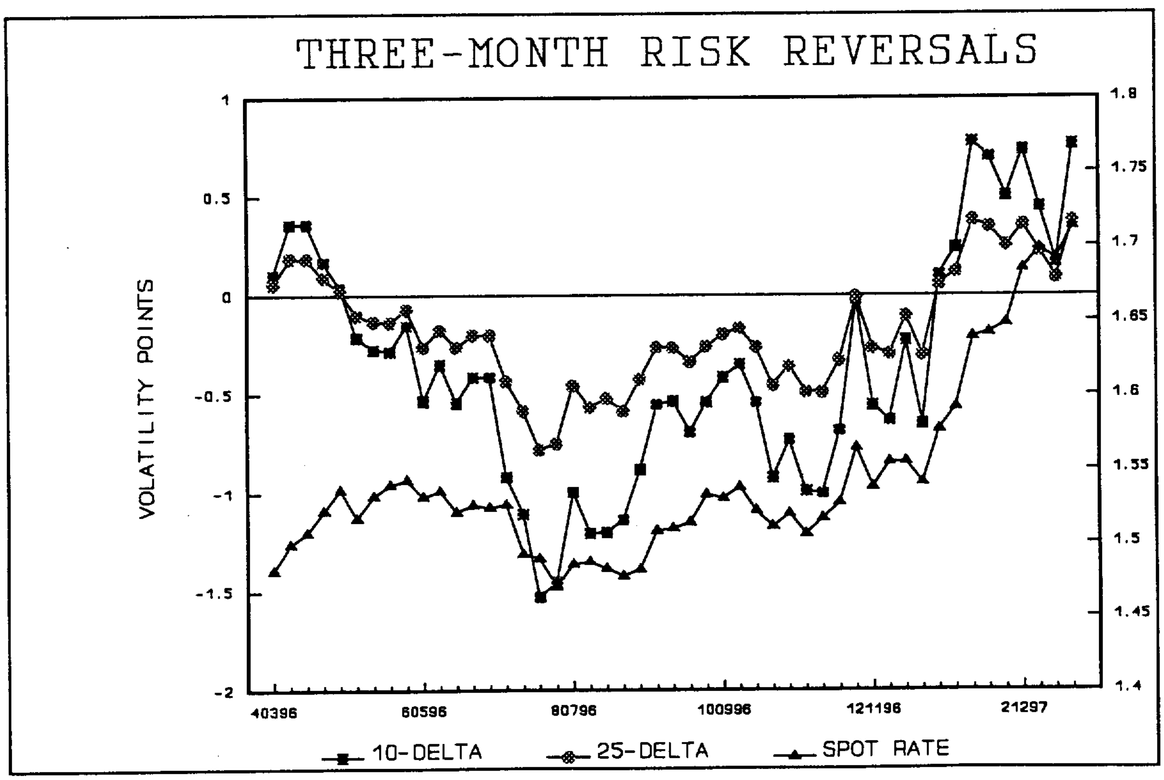

Figure 3d. 10 and 25 Delta Risk Reversals vs. Spot Rate (JPY/USD)

Risk reversal ≡ (volatility of yen call) minus (volatility of yen put)

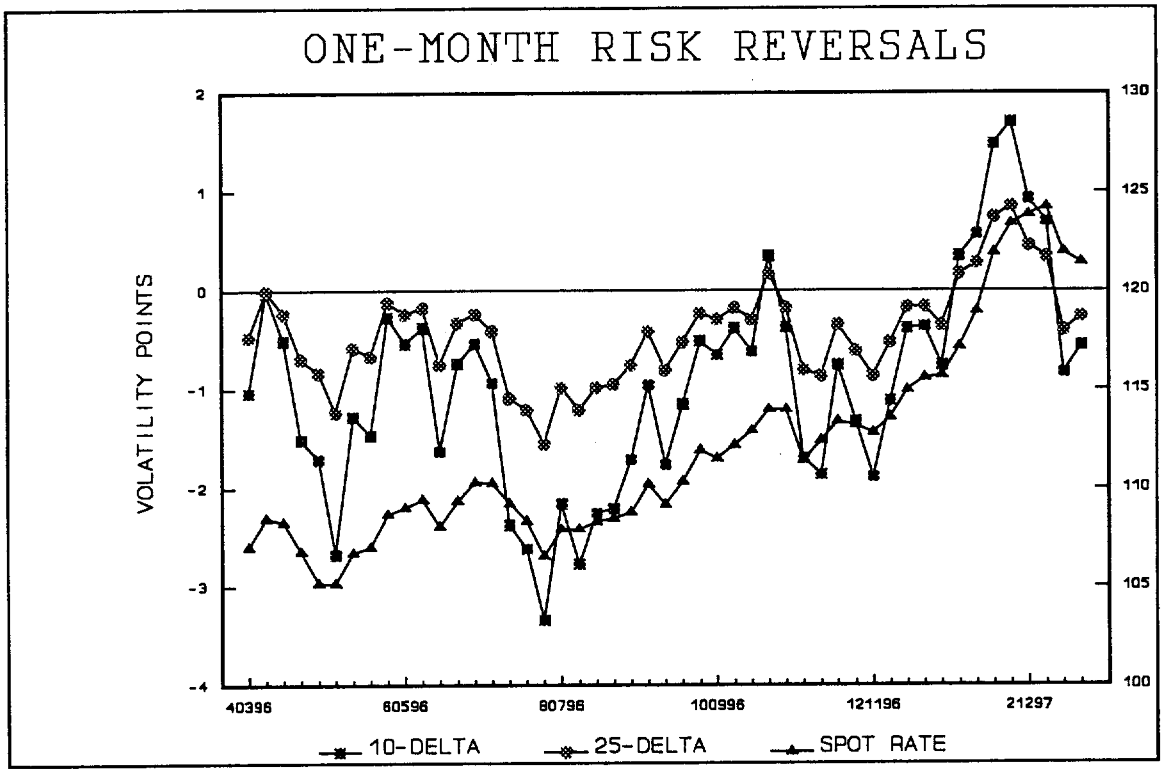

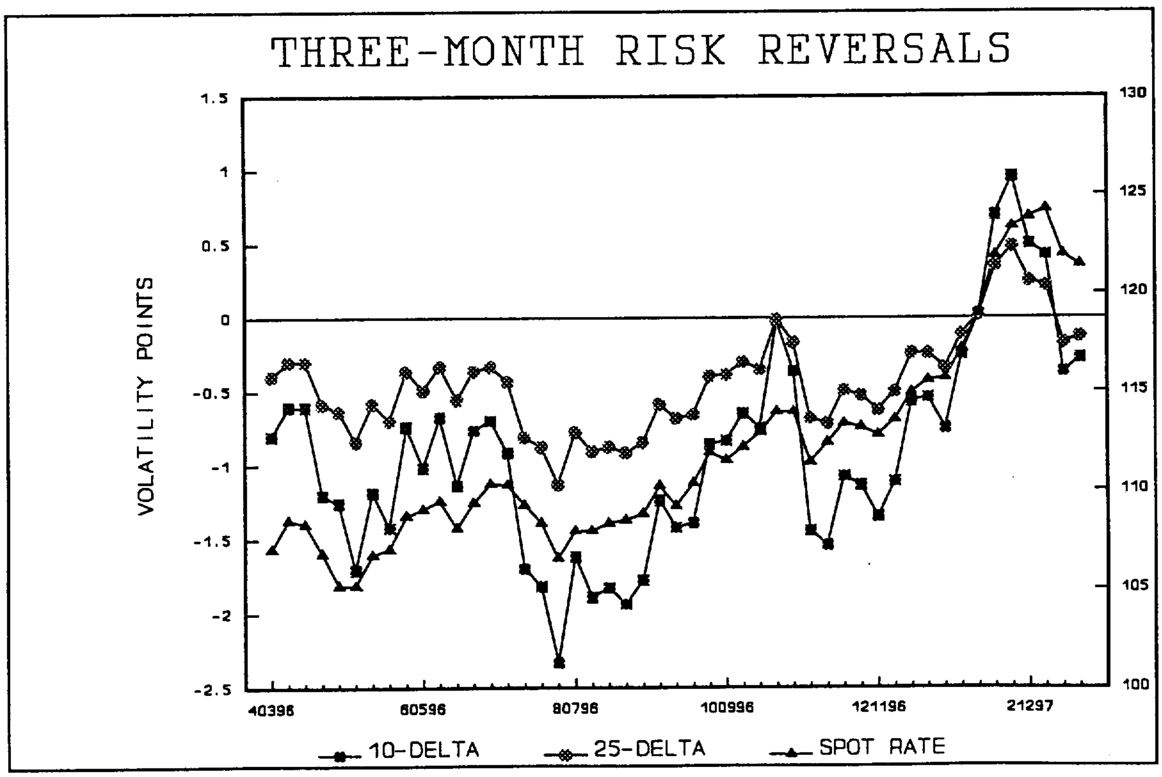

Figure 3e. 10 and 25 Delta Risk Reversals vs. Spot Rate (DEM/GBP)

Risk reversal ≡ (volatility of pound call) minus (volatility of pound put)

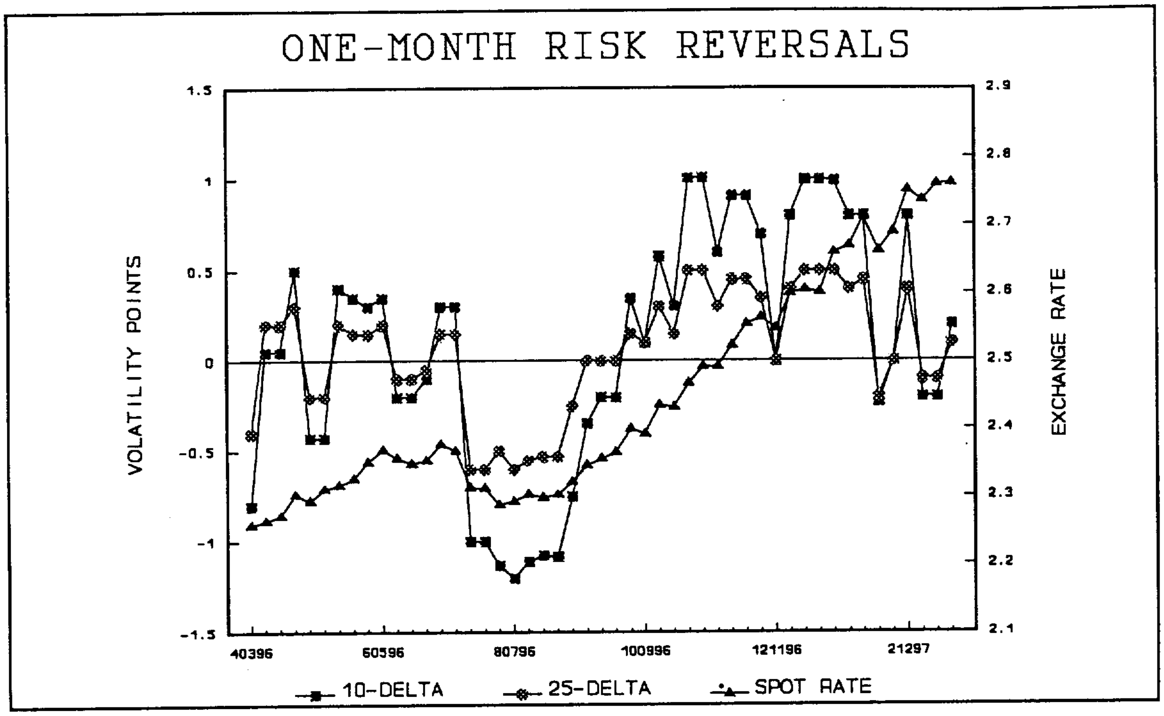

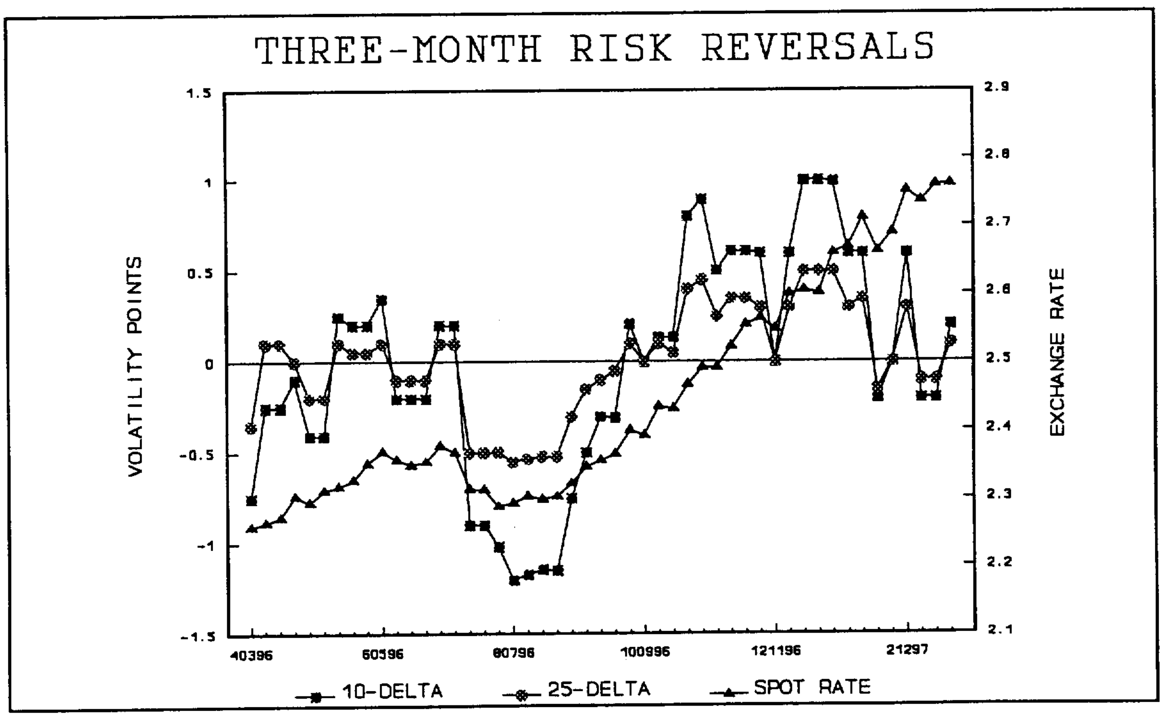

Figure 4a. Skewness As a Function of At-the-Money Volatility and 10Δ Risk Reversal (1-Month Horizon)

1-Month Risk reversal ≡ (volatility of call) minus (volatility of put)

Assumptions used in simulation:

$i=i^*=5\%$

25Δ RR = -0.01586 + 0.4559 (10Δ RR), expressed in volatility points.

Smile flat on side of the call

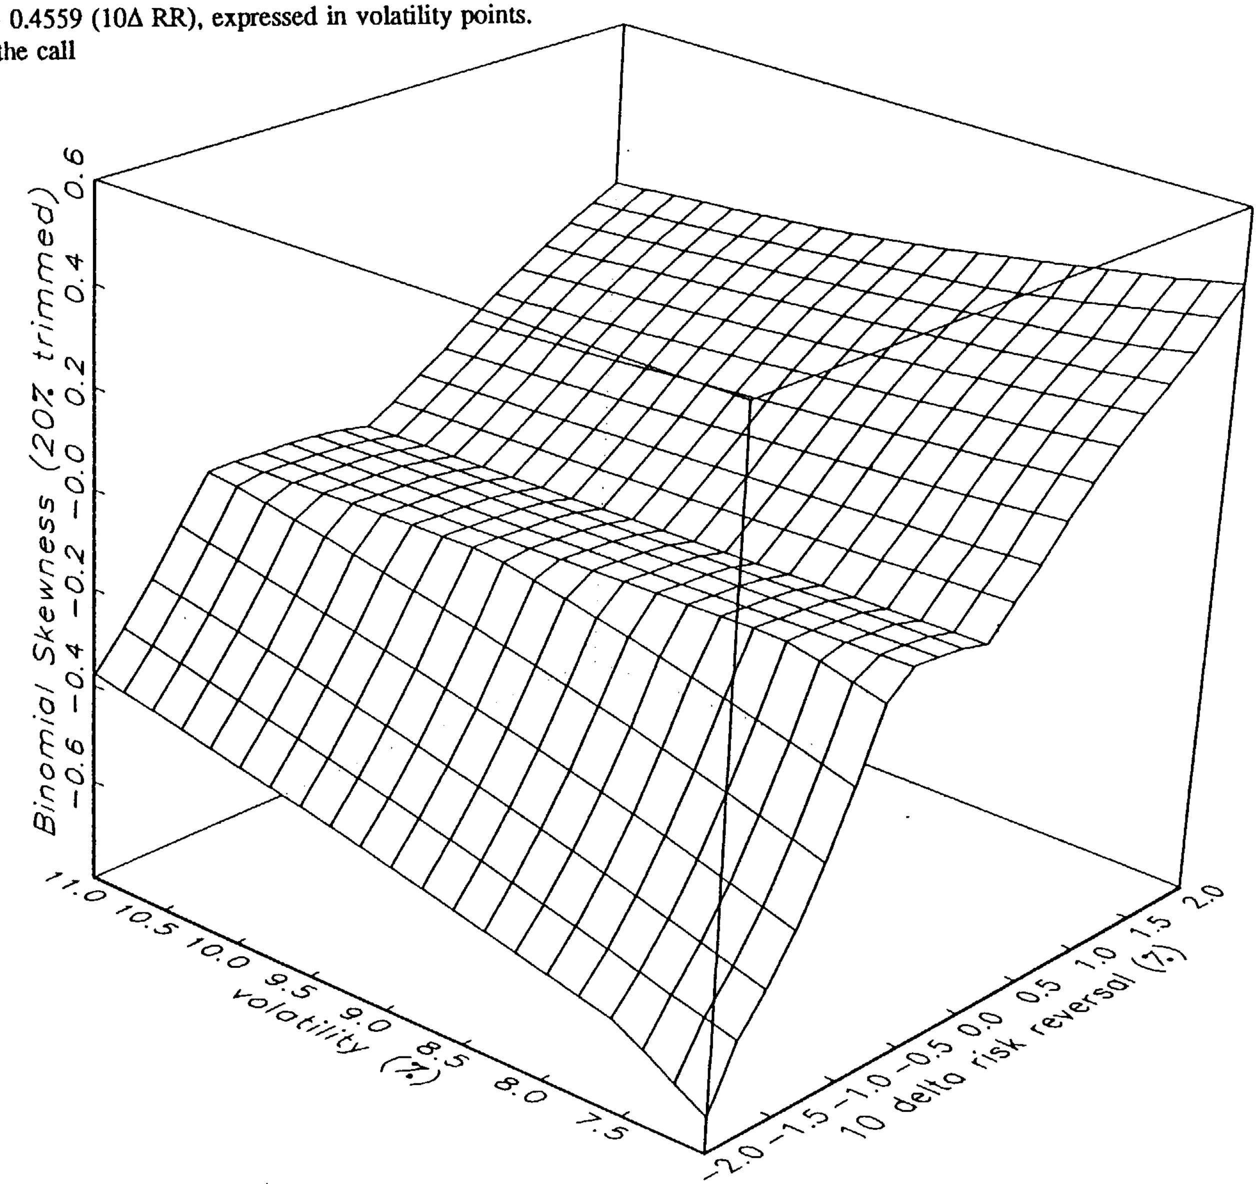

Figure 4b. Relative Intensity (±σ) As a Function of At-the-Money Volatility and 10Δ Risk Reversal (1-Month Horizon)

1-Month Risk reversal ≡ (volatility of call) minus (volatility of put)

Assumptions used in simulation:

$i=i^*=5\%$

25Δ RR = -0.01586 + 0.4559 (10Δ RR), expressed in volatility points.

Smile flat on side of the call

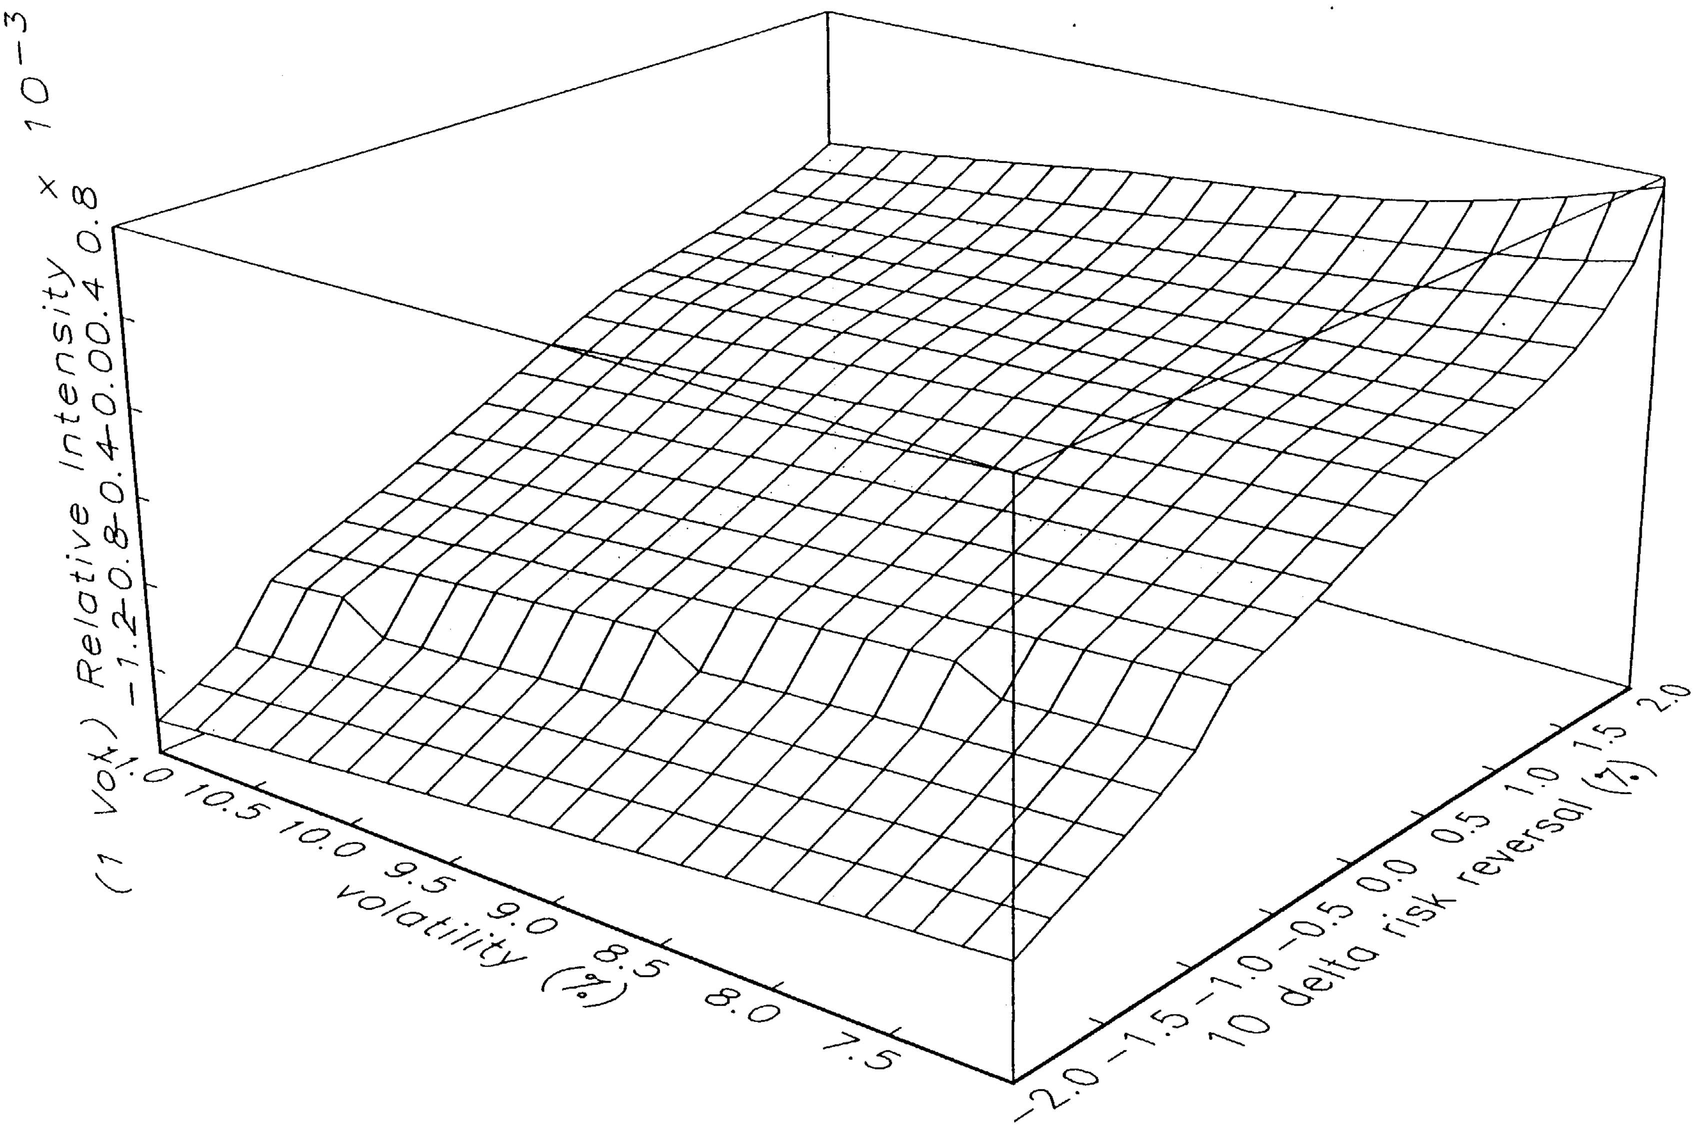

Figure 4c. Relative Intensity (±1.5σ) As a Function of At-the-Money Volatility and 10Δ Risk Reversal (1-Month Horizon)

1-Month Risk reversal ≡ (volatility of call) minus (volatility of put)

Assumptions used in simulation:

$i=i^{*}=5\%$

25Δ RR = -0.01586 + 0.4559 (10Δ RR), expressed in volatility points.

Smile flat on side of the call