Second Draft: July 1991

Estimating the Dynamics of Foreign Currency Volatility

by
William K.H. Fung and David A. Hsieh

1. Introduction

The volatility of financial markets has long been a favored subject of study for participants and academics alike. Although the many and varied approaches to the subject over the past decades have yet to reach a consensus on how 'volatility' should be modelled, one indisputable conclusion appeared to have emerged – namely, that volatility is volatile. Here the measurement problem is somewhat complicated. A number of studies on this subject have used standard deviations of daily price changes to study volatility changes in the stock market and its equilibrium implications; see for example French, Schwert, and Stambaugh (1987), Schwert (1990), and Schwert and Seguin (1990). The requirement of using daily data to compute standard deviations have limited these studies to monthly changes in volatilities.

More recent studies have begun to use implied volatilities derived from observed option prices to study daily changes in volatilities; see for example, Merville and Pieptea (1989), Franks and Schwartz (1990). The attraction lies in the notion of being able to use an ex ante measure of volatility while finessing the sample size problem at the same time. Unfortunately, these advantages are not obtained without costs. Firstly, there is a problem of internal consistency. On the one hand, implied volatilities are used to measure time variations in volatilities of the underlying asset price. On the other, the implied volatilities themselves are derived from option models that assume a stationary process for the underlying asset price. Hence it is difficult to infer from the nonstationary behavior of implied volatilities properties of the underlying asset's price dynamics. Explicit attempts to adjust option models for stochastic volatilities would not help, since all known option models of this type require further assumptions on the future evolution of volatility as well as the market price for volatility risk; see for example Hull and White (1987a, 1987b, 1988), Johnson and Shanno (1987), Scott (1987), Wiggins (1987), and Chesney and Scott (1989). In other words, using implied volatilities from a particular option model requires some prior knowledge about the behavior of volatility, which is the object of analysis to begin with. Secondly, it is not clear that implied volatility necessarily has superior information content than other volatility measures based solely on time series of prices. At least, the recent work by Canina and Figlewski (1990) cast doubt on this popular assumption. Thirdly, even if one accepts implied volatility to be a reasonable proxy for the market's expectation of future volatility, it is unclear as to the horizon over which this 'forecast' represents.

In this paper we present a method for assessing these alternative measures of volatility based on a variation of Hsieh (1991) and provide empirical evidence on their stochastic behavior on a daily frequency. Using tick-by-tick data, we are able to compute daily volatility observations from intra-day price movements. In addition, matching tick-by-tick data of options prices to the underlying asset prices, we calculate the implied volatility at the end of each trading day. This allows us to compare alternative measures of daily volatility over a reasonable sample size. While this methodology is applicable to most financial markets, we limit our application to the foreign currency futures market, due to the volume of tick-by-tick data required to conduct a study of this type.

The paper is organized as follows. Section 2 describes some of the popular beliefs on the statistical properties of implied volatilities which motivate the subsequent test design. Section 3 defines the alternative volatility measures and contrasts their information content. Section 4 compares their ability to forecast future volatility. Section 5 deals with the relation between volatility and market level. Section 6 summarizes the results and discusses some implications.

2. Implied Volatility: Some Accepted Wisdom

Since the 1960s, finance has been blessed with the growing availability of security price databases, and there is now a well documented literature on the statistical distribution of security prices. Curiously, much less is known about the statistical properties of option prices. The majority of the work in this area has centered on tests of alternative option pricing models. One could argue that since an option is just a derived asset to its "underlying," its time series behavior should, afortiori, be driven by the underlying asset's price dynamics. Running this argument in reverse, one may be able to infer properties of the underlying asset prices by studying the dynamics of option prices. This would be particularly useful if the option prices depend on an aspect of the underlying asset prices, such as its standard deviation, which is not directly observable. In fact, the connection between option prices and expected volatilities is much more explicit in practice.

Typically, over-the-counter (OTC) currency options are quoted in annualized percentage standard deviation rather than price. Figure 1 graphs the volatility quotes of two major currency options brokers, Tradition (Reuters, page TRAD) and the Bierbaum-Martin Group (Telerate, page 9042) on Jan. 18, 1991, as an illustration. These volatility quotes are for European style at-the-money spot, as distinct from forward, currency options of varying maturities. The accepted practice for converting a volatility quote to price is the Garman-Kohlhagen (1982, GK for short) variation of the standard Black-Scholes (1973, BS for short) model. The advantage of this format is that volatility levels are much slower moving than spot prices, and therefore, provided there is no ambiguity in converting volatility quotes to prices, one can avoid the need to continuously update option prices as the spot exchange rate moves. Here the GK model is being used as a conversion factor between volatilities and option prices, similar to the conversion factor used by the Chicago Board of Trade between the Treasury bond futures contract and deliverable cash bonds. Our point is that the notion of price and asset volatility is very interchangeable in the market place.

It is therefore natural to think of studying implied volatility derived from standard option models without necessarily assuming the "correctness" of the model itself. This is analogous to computing the basis of a futures contract to its underlying cash instrument without implying that the basis is at "fair value." From an empirical perspective, we think of the process of computing implied volatilities as a method for forecasting future volatility. The question then boils down to this: can implied volatility forecast future volatility better than standard measures based on historical prices? One of our goals here is to design a framework for assessing this issue. Obviously, if the resultant implied volatilities display stochastic behavior incompatible with the assumptions of the option model used to create them, we will need to address some difficult theoretical issues in option pricing. But we must first document the facts.

One stylized fact is the existence of a maturity or term structure in the volatility quotes, which is illustrated in Figure 1. This is certainly consistent with the mean reversion behavior of implied volatilities reported in Merville and Pieptea (1989) for stock options. A simple way to see this is to apply a risk neutral valuation argument. In such a world, one would be willing to form expectations of future volatility levels, incorporate them directly into option prices without further adjustment for risk. Therefore, in much the same way as in the case of an interest rate term structure, the expected average volatility may differ with maturity point. With risk aversion, the additional complication is how to determine the size of the risk premium embedded in these volatility quotes.

A related question is how asset price volatility behavior affect implied volatility and, in turn, option prices. Once again, only sufficient conditions can be provided. Suppose the volatility of asset prices mean reverts. Unless this phenomenon is totally disregarded in forming expectations of future volatility or price dynamics, options with different maturities would be priced to reflect this. After all, delta hedging an option position is ultimately governed by the amount of price volatility realizable. In the next section, we provide some empirical observations on this issue.

A second common observation concerns the negative association between implied volatility and market level which is frequently posited as accepted wisdom. Whilst a theoretical argument linking volatility to equity prices can be constructed, as in Christie (1982), we know of no analogous theory for currencies. In fact, the notion of market level is hard to define in the foreign exchange market. Figure 2 provides some tentative clues on this point. The price-to-implied volatility relation is routinely monitored by major investing institutions. As a result, brokerage houses regularly report on this. Figure 2 is a typical graph published by the Discount Corporation of New York Futures, which plots the implied volatility of the nearest to at-the-money options on the lead month futures contract along side the futures prices. The precise model and interest rate used to compute the implied volatilities is of minor importance here, since the early exercise value of at-the-money options is low relative to its value and since short term interest rate plays a small role on options on futures contracts. Therefore, even if there are biases in the computation due to the early exercise feature of these options, they are likely to be small.

Although the plots themselves do not constitute a formal test, the fact that market participants are interested in them provide a prima facie case for further academic investigation. Pause for a moment on the currency plot in Figure 2. The solid line represent an index of futures prices in British Pound, Deutschemark, Japanese Yen, and Swiss Franc. Since the individual futures prices are expressed in number of units of US currency per unit of foreign currency, a decline in the index level corresponds to a rising US dollar and vice versa. It is interesting to note that during the dollar rally through 1988 ending in the summer of 1989, implied volatility was high. Subsequently to that, during the second half of the dollar decline throughout the second half of 1990 and into the early part of 1991, implied volatility was also rising. Although the second rise in implied volatility is more likely to be related to the Gulf War, nonetheless, we have seen two dollar moves in opposite directions coinciding with high levels of implied volatility. Therefore, if there is an association between implied volatility and market level, how this can be measured for currency options is far from clear. We provide some empirical observations of this phenomenon in the next section.

The third stylized fact posits a "strike price bias" in option implied volatilities. Figure 3 illustrates this. The horizontal axis expresses the closing futures price as a percent of various strike prices for the March DM option on March DM futures in the International Monetary Market on January 2, 1991. Implied volatilities of closing option prices at various strikes are expressed as a percent of the implied volatility of the nearest to at-the-money option. The model used is based on the Barone-Adesi and Whaley (1987) extension to the BS model. The Quantitative Strategies Group in the Institutional Futures Division of Lehman Brothers monitors these statistics on a daily basis. Clearly, the implied volatilities of out-of-the-money options skew away from that of the at-the-money options. This is sometimes referred to as the "smile" across strikes. There are a number of ad hoc institutional reasons often put forward as explanations; few, if any, lend themselves to formal tests. A more systematic cause of this can be found in the work on option models admitting stochastic volatilities with non-zero correlation between volatility and price.

Consider the structure analyzed in Wiggins (1987) where the underlying asse t price S(t) follows a diffusion process given by:

$$dS(t) = \mu_s S(t)\ dt + \sigma(t)\ S(t)\ dz_s,$$

where the volatility diffusion process is:

$$d\sigma(t) = f(\sigma(t))\ dt + \theta\ \sigma(t)\ dz_\sigma,$$

with the instantaneous correlation

$$\rho\ dt = dz_s\ dz_\sigma.$$

Variations to the function $f(\sigma(t))$ were used in Scott (1987), Hull and White (1987a, 1987b, 1988), and Chesney and Scott (1989). All of the above papers used a mean reversion process for $f(\sigma(t))$, with a stationary variance θ. In a loose way, we can interpret the stylized facts where volatility is stochastic, displays a maturity effect, and appears to be correlated to market levels, as lending support to the above structure. With simulated results, most of the above authors showed that even with ρ set to zero, the presence of stochastic volatility alone is sufficient to induce implied volatilities for out-of-the-money options computed using standard models to be higher than their at-the-money counterparts. However, with ρ set to be negative, it will have an asymmetric effect on puts and calls. We prefer to adopt this as the motivation for testing this "smile" across strikes. Results are reported in the next section.

3. Volatility Measures and Their Information Content

In this section, we define various measures of volatility and contrast their information content. This methodology applies to analysis of volatility in financial markets in general. In this paper, we limit our scope to the currency markets. Like the US government bond market, the foreign exchange (FX for short) market remains an over-the-counter market where transactions are generally conducted through interbank networks. Liquidity is perhaps the highest of all financial markets. In 1986, the average volume of worldwide trading is over $200 billion per day. (See Krugman and Obstfeld, 1988,

p. 314, footnote 1.) Of the various currencies, the most liquid currency is the Deutschemark/US Dollar exchange rate. Due to the nature of the market, transactions data history simply does not exist in any meaningful way. Quotation prices are mainly through the normal information agencies such as Reuters and Telerate, which pool currency quotes from contributing banks. As a result, although official sources on daily exchange rates exist, such as the Federal Reserve's data base, higher frequency data such as tick-by-tick quotes remain the private property of a limited number of financial institutions. Both the availability and the data quality of such databases are only becoming known through recent work by Goodhart and Demos (1990) and Müller et al (1990). The result thus far leads us to believe that more work is required before the data can be subjected to serious tests. In relation to this, option prices on spot FX rates are even harder to obtain. Therefore, in this paper, we use the Deutschemark (DM for short) futures contract in the Chicago Mercantile Exchange which also has available the associated options data.

The tick-by-tick (also called "quote capture" or "time-and-sales") data contain the time and price of every transaction in which the price has changed from the previous transaction. In addition, a bid price is recorded if it is above the previous transaction, and an ask price is recorded if it is below the previous transaction. Since these bid and ask prices do not represent actual transactions, we eliminated them from our sample. Note that there is no information on the number and volume of transactions at any given price. In addition to the tick-by-tick data, the Chicago Mercantile Exchange also provides daily data on the open/high/low/settlement prices of each futures contract, as well as its daily volume and open interest. Our data begins on Feb. 25, 1985, when daily price limits were removed on currency futures, and ends on Jul. 30, 1989, totaling 1121 daily observations.

We begin our analysis with a description of the daily settlement prices. Let F_t denote the futures settlement price at date t. Let $x_t = \log[F_t/F_{t-1}]$ be the continuously compounded rate of change, where "log" means natural logarithm. Figure 4 is a plot of x_t, and Table 1 presents some statistical description of the data. The mean is close to zero (the t-statistic for this test is 1.62). There is little autocorrelation in x_t, but there is some autocorrelation in $|x_t|$. The latter finding is consistent with the predictability volatility changes.

Our goal is to provide a description of volatility changes. Some further notation is needed. Define $E_t[x_{t+1}] = \mu_t$, and $V_{t-1}[x_t] = \sigma_t^2$. Based on the information in Table 1, we assume that $\mu_t = 0$, and that σ_t changes over time.

Since σ_t is not observable, we need to proxy it. To do so, we use the tick-by-tick data on the DM futures contract in the CME. For each trading day, we calculate σ_t as the standard deviation of the 15 minute rates of change of the nearby futures contract. We call this "realized volatility."

It is appropriate to discuss the rationale for using the 15 minute interval, rather than a shorter time span. In tick-by-tick data, as in most transactions data, there are bid-ask bounces, causing a large and negative first order serial correlation in the data. We need a sufficiently long time interval to remove the effects of any bid-ask bounce. This is achieved using a 15 minute interval.

We should also discuss the method used to convert the standard deviation of 15-minute data to a daily volatility. Essentially, we multiply the standard deviation of the 15-minute data by the factor $(96/M)^{½}$, where 96 is the number of 15-minute intervals in a 24 hour day, and M is the number of 15-minute intervals in the trading day. This "scaling up" is based on the observation that currency markets are open around the clock, and therefore the intraday open-to-close price volatility is proportional to the daily close-to-close price volatility. [The appropriateness of this method is illustrated in Table 4, and will be discussed later.]

We examine the time series behavior of σ_t in terms of its natural logarithm. There are three reasons for this transformation. In the first place, the standard deviation is always a positive number. If we had graphed σ_t, it would be easy to pick out sudden increases but difficult to see sudden decreases, which may be responsible for the popular belief that volatilities are more likely to jump upwards than downwards. In the second place, we would like to predict σ_t. If we had worked with σ_t itself, our prediction may have been negative! This is avoided by working with $\log[\sigma_t]$. In the third place, $\log[\sigma_t]$ is a much better behaved series than σ_t itself, because the logarithmic transformation "pulls" the outliers in. This makes statistical analysis much simpler and nicer.

Figure 5 is a plot of $\log[\sigma_t]$. Some statistics of $\log[\sigma_t]$ are provided in Table 2. Volatility changes over time. Its smallest and largest values are -6.29 and -2.96, or 2.9% and 83.1%, respectively, on an annualized basis. Volatility is also serially correlated. Based on the Akaike (1974) information criterion, $\log[\sigma_t]$ is best described as an AR(7). The regression is given in the lower panel of Table 2. The sum of the autoregressive coefficients is 0.794, which indicates that $\log[\sigma_t]$ is mean reverting. The fitted values, $\log[\sigma 1_t]$, are plotted in Figure 6.

This characterization of DM futures prices is similar to that of spot exchange rates. For example, Hsieh (1989) found a significant amount of nonlinear dependence in daily rates of change of spot exchange rates, which is consistent with Engle's (1982) autoregressive conditional heteroskedasticity (ARCH) model, or Bollerslev's (1986) generalized autoregressive conditional heteroskedasticity (GARCH) model. We have decided not to pursue that method of extracting volatility, in view of the fact that the tick-by-tick data provide a natural way to obtain daily measures of volatility, and that Hsieh (1991) finds that GARCH models for currency futures suffer from misspecification.

4. Forecasting Volatility

We next examine some of the methods used to forecast volatility. First, we consider historical volatility, which is the standard deviation of past observations of x_t. We use a 20-day rolling measure:

$$\sigma 2_t = [\quad \Sigma_{i=1}^{20} \ (x_{t-i} - \Sigma_{j=1}^{20} x_{t-j}/20 \ )^2 \ /20 \ ]^{\frac{1}{2}}.$$

Figure 7 is a plot of $\log[\sigma 2_t]$. Since $\sigma 2_t$ is a rolling measure of volatility, it is not surprising that $\sigma 2_t$ is highly autocorrelated, as shown in Table 3.

Second, we use the implied volatility of at-the-money (ATM) call and put options on the DM futures, $\sigma 3_t$ and $\sigma 4_t$. This is obtained as follows. For each day, we choose the nearby DM futures contract and the options on that contract which mature in the same month. For example, in January 1989, we use the March 1989 option and futures. Contracts are rolled over when the option has less than 10 days to maturity. We match futures and options prices using the tick-by-tick data from the CME. The option with the strike price closest to the futures price at the end of the day is chosen. The interest rate is taken to be the treasury bill rate for the bill which matures nearest to the options expiration date. The implied volatility is that volatility which equates the options price and the model price, using the Barone-Adesi and Whaley (1987) approximate solution to an American option.

Figures 8 and 9 are plots of $\log[\sigma 3_t]$ and $\log[\sigma 4_t]$. It is interesting to note that ATM call and put implied volatilities are almost identical. This is quite different from the behavior of stock index options. Statistics of the series are given in Table 3. They indicate that there is strong autocorrelation in both series. This should not be surprising, since the forecast of an AR(7) is likely to be much more correlated than the AR(7) itself, as demonstrated in Table 3.

Third, we use the lagged value of the 15-minute volatility, σ_{t-1}. This can be interpreted as a naive forecast.

Fourth, we use the fitted values of the AR(7) model. This can be interpreted as a "rational expectations" forecast.

We perform two comparisons of the ability of these series to forecast future price volatility. In the first comparison, we test whether these measures of volatility can standardize the daily rates of change, as follows. For each day, we divide the rate of change of the settlement price (x_t) by the previous day's measure of volatility, i.e.:

$$z_t = x_t/\sigma_{t-1};\ z2_t = x_t/\sigma 2_{t-1};\ z3_t = x_t/\sigma 3_{t-1};\ z4_t = x_t/\sigma 4_{t-1}.$$

If σi_{t-1} is the appropriate forecast of volatility, then zi_t should have mean zero, variance one, and no evidence of heteroskedasticity. Table 4 presents the statistics of the zi's for the 1121 observations (Feb. 26, 1985 to Jul. 30, 1989). The means of the zi's are not statistically different from zero (the t-statistics are 0.737, 1.434, 1.486, and 1.454, respectively, for z_t, $z2_t$, $z3_t$, and $z4_t$). The standard deviation of $z2_t$ is statistically different from unity at the 5% significance level, while those of z_t, $z3_t$ and $z4_t$ are not statistically different from unity (the t-statistics are 0.00, 2.063, -0.602, and -0.636, respectively, for z_t, $z2_t$, $z3_t$, and $z4_t$). In addition, the autocorrelation of the absolute values of the zi's show that the standardization has removed all evidence of heteroskedasticity. These tests, although crude, show that the 15-minute and implied volatilities can predict the next day's price volatility, but historical volatility tends to underpredict the next day's price volatility.

In the second comparison, we calculate the root mean squared errors and mean absolute errors of the 1-day ahead forecasts of price volatility, σ_t, using the various measures of volatility from the previous day. The results are reported in Table 5. It shows that the historical volatility and the lagged 15-minute volatility are the worse predictors. The put and call implied volatilities can forecast realized volatility better, by about 17% in root mean squared error, and by about 8% in mean absolute error. Not surprisingly, the fitted values of the AR(7) is the best predictor of future volatility, beating the historical volatility by 20% in root mean squared error and 10% in mean absolute error.

The impressive performance of implied volatility in forecasting 1-day ahead realized volatility, at least relative to the fitted volatility, brings us to the following issue. Is the implied volatility a forecast of the average realized volatility over the remaining life of the option, or is it merely forecasting the next day's realized volatility? This is a difficult question to answer. The problem is that, even though we have over 1000 observations, there are only 17 quarterly expiration cycles in our data. In other words, the options are forecasting overlapping periods, so that the forecast horizon shortens as the options age. This creates substantial statistical problems. The overlapping forecast horizon was addressed in Canina and Figlewski (1990) in the context of stock index options. But the asymptotic distribution used in their work is most likely a poor approximation of the finite sample distribution, especially in our case.

One way to answer this question is to use the AR(7) model to forecast the average future volatility for the life of the option. We denote this series as σf_t. The logarithm of σf_t is plotted in Figure 10. Due to the mean reversion of the AR(7), σf_t is a smoother series than $\sigma 1_t$, as indicated by its smaller standard deviation in Table 3. In addition, the autocorrelation coefficients of σf_t are uniformly smaller than those of $\sigma 1_t$. Relying on the principle of matching moments and autocorrelation coefficients, the evidence favors the view that implied volatilities are 1-day ahead forecasts of future volatility.

5. Volatility and Levels

We now deal with the issue of whether volatility is related to asset levels. Figure 11 is a plot of the levels of the DM futures contract. Superimposing Figure 11 on realized volatility (Figure 5), historical volatility (Figure 6), fitted volatility (Figure 7), call and put implied volatilities (Figures 8 and 9), there does not appear to be any relation between the level of DM futures and any measures of volatility.

In trying to test the relation between volatility and level, we encountered some important statistical problems. It is natural to test the relation between volatility and level by regressing the former on the latter, and testing whether the regression coefficient is statistically negative.

This procedure runs into the problem of "spurious" correlation, as defined in Granger and Newbold (1974). Whenever one regresses a highly autocorrelated series (such as the logarithm of volatility) on another higher autocorrelated series (such as the logarithm of DM futures price), the standard errors of the estimated coefficients are incorrect, leading to erroneous inference. To avoid this problem, we ran the regression in first differences. We found that there is no relation between volatility and levels. This is reported in Table 6. None of the regression coefficients of $\log[F_{t-1}]$ are statistically different from zero.

An implication of this finding is that the observed behavior that out-of-the-money options tend to have different implied volatilities than at-the-money options (as illustrated in Figure 3) cannot be due to a negative relation between the level of the asset and volatility. It is more reasonable to attribute this behavior to the incorrect assumption of the Black (1976) model that futures prices follow a geometric brownian motion. In particular, the distribution of the rates of change has fatter tails than the normal, which means that the Black (1976) model underprices out-of-the-money options, i.e. their implied volatilities are higher than at-the-money options.

6. Concluding Remarks

In summary, we have found the following empirical facts regarding the volatility of the DM futures contract:

a) Realized volatility changes over time and is mean reverting; it can be modelled as an AR(7).

b) The historical volatility is a poor predictor of future volatility, relative to at-the-money put and call volatilities.

c) At-the-money put and call volatilities are almost identical, and are good predictors of future volatility, almost as good as the in-sample predicted values of the AR(7) model.

d) The evidence slightly favors the view that implied volatilities are 1-day ahead forecasts of future volatilities.

e) None of the measures of volatility (realized, fitted, implied) are related to the level of the DM futures. This means that the "mispricing" of the Black (1976) option model on out-of-the-money options relative to at-the-money options cannot be attributed to this explanation.

There are some interesting implications for these findings. The fact that an AR(7) process is needed to describe volatility indicates that it has a complicated structure. The simple stochastic volatility models used in recent option pricing models may not be adequate. Nevertheless, the AR(7) process can easily be used in simulations. This can help in describing the distribution of asset price movements, which is critically important in determining the capitalization needs of leverage positions. See Hsieh (1991) for details. Lastly, an accurate model of the dynamics of volatility will help in asset allocation decisions. This will be left for future research.

References:

Akaike, H., 1974, A New Look at the Statistical Model Identification, *IEEE Transactions on Automatic Control* 19, 716-723.

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

Black, F., and M. Scholes, 1973, The Pricing of Options and Corporate Liabilities, *Journal of Political Economy*, 81, 637-654.

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

Bollerslev, T., 1986, Generalized Autoregressive Conditional Heteroskedicity, *Journal of Econometrics* 31, 307-327.

Canina, L., and S. Figlewski, 1990, The Informational Content of Implied Volatility, unpublished manuscript, New York University.

Chesney, M., and L. Scott, 1989, Pricing European Currency Options: A Comparison of the Modified Black-Scholes Model and a Random Variance Model, *Journal of Financial and Quantitative Analysis* 24, 267-284.

Engle, R., 1982, Autoregressive Conditional Heteroscedasticity With Estimates of The Variance of U. K. Inflations, *Econometrica* 50, 987-1007.

Franks, R., and E. Schwartz, 1988, The Stochastic Behavior of Market Variance Implied in the Prices of Index Options, unpublished manuscript, University of California at Los Angeles and London School of Economics.

French, K., G.W. Schwert and R. Stambaugh, 1987, Expected Stock Returns and Volatility, *Journal of Financial Economics* 19, 3-29.

Garman, M., and S. Kohlhagen, 1983, Foreign Currency Option Values, *Journal of International Money and Finance*, 2, 231-238.

Goodhart, C., and A. Demos, 1990, Reuters Screen Images of the Foreign Exchange Market: The Deutschemark/Dollar Spot Rate, *Journal of International Securities Markets* 4, 333-348.

Granger, C., and P. Newbold, 1974, Spurious Regressions in Econometrics, *Journal of Econometrics* 2, 111-120.

Hsieh, D., 1989, Testing for Nonlinear Dependence in Daily Foreign Exchange Rates, *Journal of Business* 62, 339-368.

Hsieh, D., 1991, Implications of Nonlinear Dynamics for Financial Risk Management, FORCE Working Paper No. 91-107, Fuqua School of Business, Duke University.

Hull, J., and A. White, 1987a, The Pricing of Options on Assets with Stochastic Volatilities, *Journal of Finance*, 42, 281-300.

Hull, J., and A. White, 1987b, Hedging the Risks from Writing Foreign Currency Options, *Journal of International Money and Finance* 6, 131-152.

Hull, J., and A. White, 1988, An Analysis of the Bias in Option Pricing Caused by a Stochastic Volatility, *Advances in Futures and Options Research* 3, 29-61.

Johnson, H., and D. Shanno, 1987, Stock Price Volatility, Mean-reverting Diffusion, and Noise, *Journal of Financial Economics* 24, 143-151.

Krugman, P., and M. Obstfeld, 1988, *International Economics: Theory and Practice*, Glenview, IL: Scott, Foresman/Little, Brown.

Merville, L., and D. Pieptea, 1989, Stock-Price Volatility, Mean-Reverting Diffusion, and Noise, *Journal of Financial Economics* 24, 193-214.

Müller, U., M. Dacorogna, R. Olsen, O. Pictet, M. Schwarz, and C. Morgenegg, 1990, Statistical Study of Foreign Exchange Rates, Empirical Evidence of a Price Change Scaling Law, and Intraday Analysis, *Journal of Banking and Finance* 14, 1189-1208.

Schwert, G.W., 1990, Why Does Stock Market Volatility Change Over Time? *Journal of Finance* 44, 1115-1153.

Schwert, G.W., and P. Seguin, 1990, Heteroskedasticity in Stock Returns, *Journal of Finance* 44, 1129-1155.

Scott, L., 1987, Option Pricing When the Variance Changes Randomly: Theory, Estimation, and an Application, *Journal of Financial and Quantitative Analysis* 22, 419-438.

Wiggins, J., 1987, Option Values Under Stochastic Volatility: Theory and Empirical Estimates, *Journal of Financial Economics* 19, 351-372.

Table 1
Statistics of Rates of Change of DM Futures
$x_t = \log[F_t/F_{t-1}]$

Summary statistics of x_t:

Mean	0.376510E-03
Median	0.000000
Std dev	0.777268E-02
Skewness	0.212509
Kurtosis	5.51004
Maximum	0.483210E-01
Minimum	-0.326300E-01

Autocorrelation Coefficients:

Lags	x_t	$\lvert x_t \rvert$
1	-0.011	0.055
2	0.007	0.035
3	0.045	0.087
4	-0.030	0.064
5	0.001	0.079
6	0.000	0.107
7	-0.008	0.104
8	0.029	0.077
9	-0.002	0.049
10	-0.006	0.129
11	-0.012	0.033
12	-0.005	0.057
13	0.022	0.097
14	-0.021	0.056
15	0.051	0.098
16	-0.050	0.054
17	-0.035	0.015
18	-0.007	0.055
19	-0.013	0.013
20	0.031	0.033

Table 2
Statistics of Log Volatility (15 Minute Rates of Change)

Mean	-4.74154
Median	-4.76615
Std dev	0.469094
Skewness	0.639847E-01
Kurtosis	3.00770
Maximum	-2.95990
Minimum	-6.29327

Autocorrelation Coefficients:

Lags	
1	0.465
2	0.418
3	0.392
4	0.417
5	0.412
6	0.388
7	0.388
8	0.358
9	0.343
10	0.324
11	0.313
12	0.314
13	0.294
14	0.298
15	0.308
16	0.285
17	0.271
18	0.201
19	0.283
20	0.252

AR(7) regression: $y_t = \log[\sigma_t]$

$$y_t = \underset{(.164)}{-.9004} + \underset{(.030)}{.215}\, y_{t-1} + \underset{(.031)}{.116}\, y_{t-2} + \underset{(.030)}{.077}\, y_{t-3} + \underset{(.030)}{.133}\, y_{t-4}$$

$$+ \underset{(.031)}{.114}\, y_{t-5} + \underset{(.031)}{.062}\, y_{t-6} + \underset{(.030)}{.094}\, y_{t-7}$$

$\bar{R}^2 = .336$, SEE $= .379$

Heteroskedasticity-consistent standard errors in parentheses.

Table 3
Statistics of Log Alternative Volatility Measures

	$\sigma 2$	$\sigma 3$	$\sigma 4$	$\sigma 1$	σf
Mean	-4.9844	-4.8671	-4.8644	-4.7446	-4.7411
Median	-4.9910	-4.8877	-4.8877	-4.7417	-4.7455
Std Dev	0.3332	0.2133	0.2140	0.2740	0.1344
Skewness	0.032	-0.030	-0.025	-0.198	-0.254
Kurtosis	2.901	2.658	2.654	2.613	4.268
Maximum	-4.2097	-4.2988	-4.3319	-4.0982	-4.3180
Minimum	-6.0921	-5.7153	-5.6557	-5.5830	-5.4107

Autocorrelation Coefficients:

Lag		$\sigma 2$	$\sigma 3$	$\sigma 4$	$\sigma 1$	σf
Lag	1	0.973	0.966	0.965	0.943	0.939
	2	0.943	0.944	0.946	0.910	0.887
	3	0.914	0.925	0.926	0.895	0.833
	4	0.885	0.906	0.908	0.852	0.785
	5	0.857	0.891	0.894	0.809	0.741
	6	0.827	0.878	0.879	0.775	0.703
	7	0.793	0.863	0.864	0.721	0.665
	8	0.757	0.847	0.848	0.695	0.635
	9	0.721	0.834	0.832	0.672	0.607
	10	0.683	0.818	0.817	0.650	0.579
	11	0.644	0.802	0.802	0.628	0.552
	12	0.604	0.787	0.788	0.605	0.526
	13	0.563	0.770	0.772	0.586	0.503
	14	0.522	0.757	0.756	0.567	0.483
	15	0.480	0.741	0.741	0.551	0.465
	16	0.437	0.726	0.726	0.533	0.445
	17	0.394	0.710	0.711	0.514	0.424
	18	0.354	0.695	0.697	0.493	0.401
	19	0.312	0.681	0.681	0.482	0.388
	20	0.272	0.667	0.668	0.462	0.366

Table 4
Statistics of Standardized Data

		z_t	$z2_t$	$z3_t$	$z4_t$
Mean		0.022	0.049	0.043	0.042
Std dev		1.000	1.144	0.969	0.967
Skewness		-0.401	-0.098	0.062	0.057
Kurtosis		7.972	5.175	4.167	4.226
Maximum		4.035	4.893	4.647	4.659
Minimum		-7.507	-5.800	-3.347	-3.405
T-test:					
Mean=0		0.737	1.434	1.486	1.454
SD=1		0.000	2.063	-0.602	-0.636
Autocorrelation coefficients of absolute values:					
Lag	1	-0.066	0.034	-0.002	-0.001
	2	-0.080	-0.022	-0.037	-0.036
	3	-0.053	0.013	0.012	0.012
	4	-0.035	0.021	0.016	0.014
	5	-0.043	0.038	0.030	0.029
	6	0.008	0.034	0.043	0.042
	7	0.037	0.042	0.044	0.045
	8	0.008	0.005	0.022	0.022
	9	0.026	0.019	0.022	0.022
	10	0.053	0.039	0.066	0.066
	11	-0.024	-0.029	-0.015	-0.017
	12	-0.002	-0.021	0.004	0.003
	13	-0.013	0.013	0.033	0.033
	14	-0.002	-0.005	0.018	0.015
	15	0.058	0.019	0.048	0.048
	16	0.025	-0.007	0.042	0.039
	17	-0.024	-0.063	-0.026	-0.025
	18	-0.017	-0.039	0.011	0.011
	19	-0.013	-0.045	-0.017	-0.017
	20	0.022	-0.070	-0.014	-0.013

Table 5
Forecast Errors of $\log[\sigma_t]$

Log of	σ_{t-1}	$\sigma 2_{t-1}$	$\sigma 3_{t-1}$	$\sigma 4_{t-1}$	$\sigma 1_{t-1}$
Root Mean Squared Errors	0.482	0.487	0.406	0.405	0.390
Mean Absolute Errors	0.601	0.613	0.562	0.562	0.550

Table 6
Regression of Volatility on Asset Levels

First differences:

Call implied volatility on futures price level:

$$\log[\sigma 3_t] = \underset{(0.0015)}{-0.0024} + \underset{(0.163)}{0.128} \log[F_{t-1}] - \underset{(0.056)}{0.130} \log[\sigma 3_{t-1}]$$

$\bar{R}^2 = .015$

Put implied volatility on future price level:

$$\log[\sigma 4_t] = \underset{(0.0015)}{-0.0025} + \underset{(0.166)}{0.289} \log[F_{t-1}] - \underset{(0.057)}{0.189} \log[\sigma 4_{t-1}]$$

$\bar{R}^2 = .034$

Realized volatility on futures price level:

$$\log[\sigma_t] = \underset{(0.0115)}{-0.0008} - \underset{(1.343)}{1.201} \log[F_{t-1}] - \underset{(0.032)}{0.764} \log[\sigma_{t-1}]$$

$$- \underset{(0.038)}{0.628} \log[\sigma_{t-2}] - \underset{(0.044)}{0.536} \log[\sigma_{t-3}] - \underset{(0.044)}{0.386} \log[\sigma_{t-4}]$$

$$- \underset{(0.043)}{0.253} \log[\sigma_{t-5}] - \underset{(0.040)}{0.174} \log[\sigma_{t-6}] - \underset{(0.030)}{0.066} \log[\sigma_{t-7}]$$

$\bar{R}^2 = .370$

Heteroskedasticity-consistent standard errors in parentheses.

Discount Corporation of New York Futures

Futures Price vs. Implied Volatility*

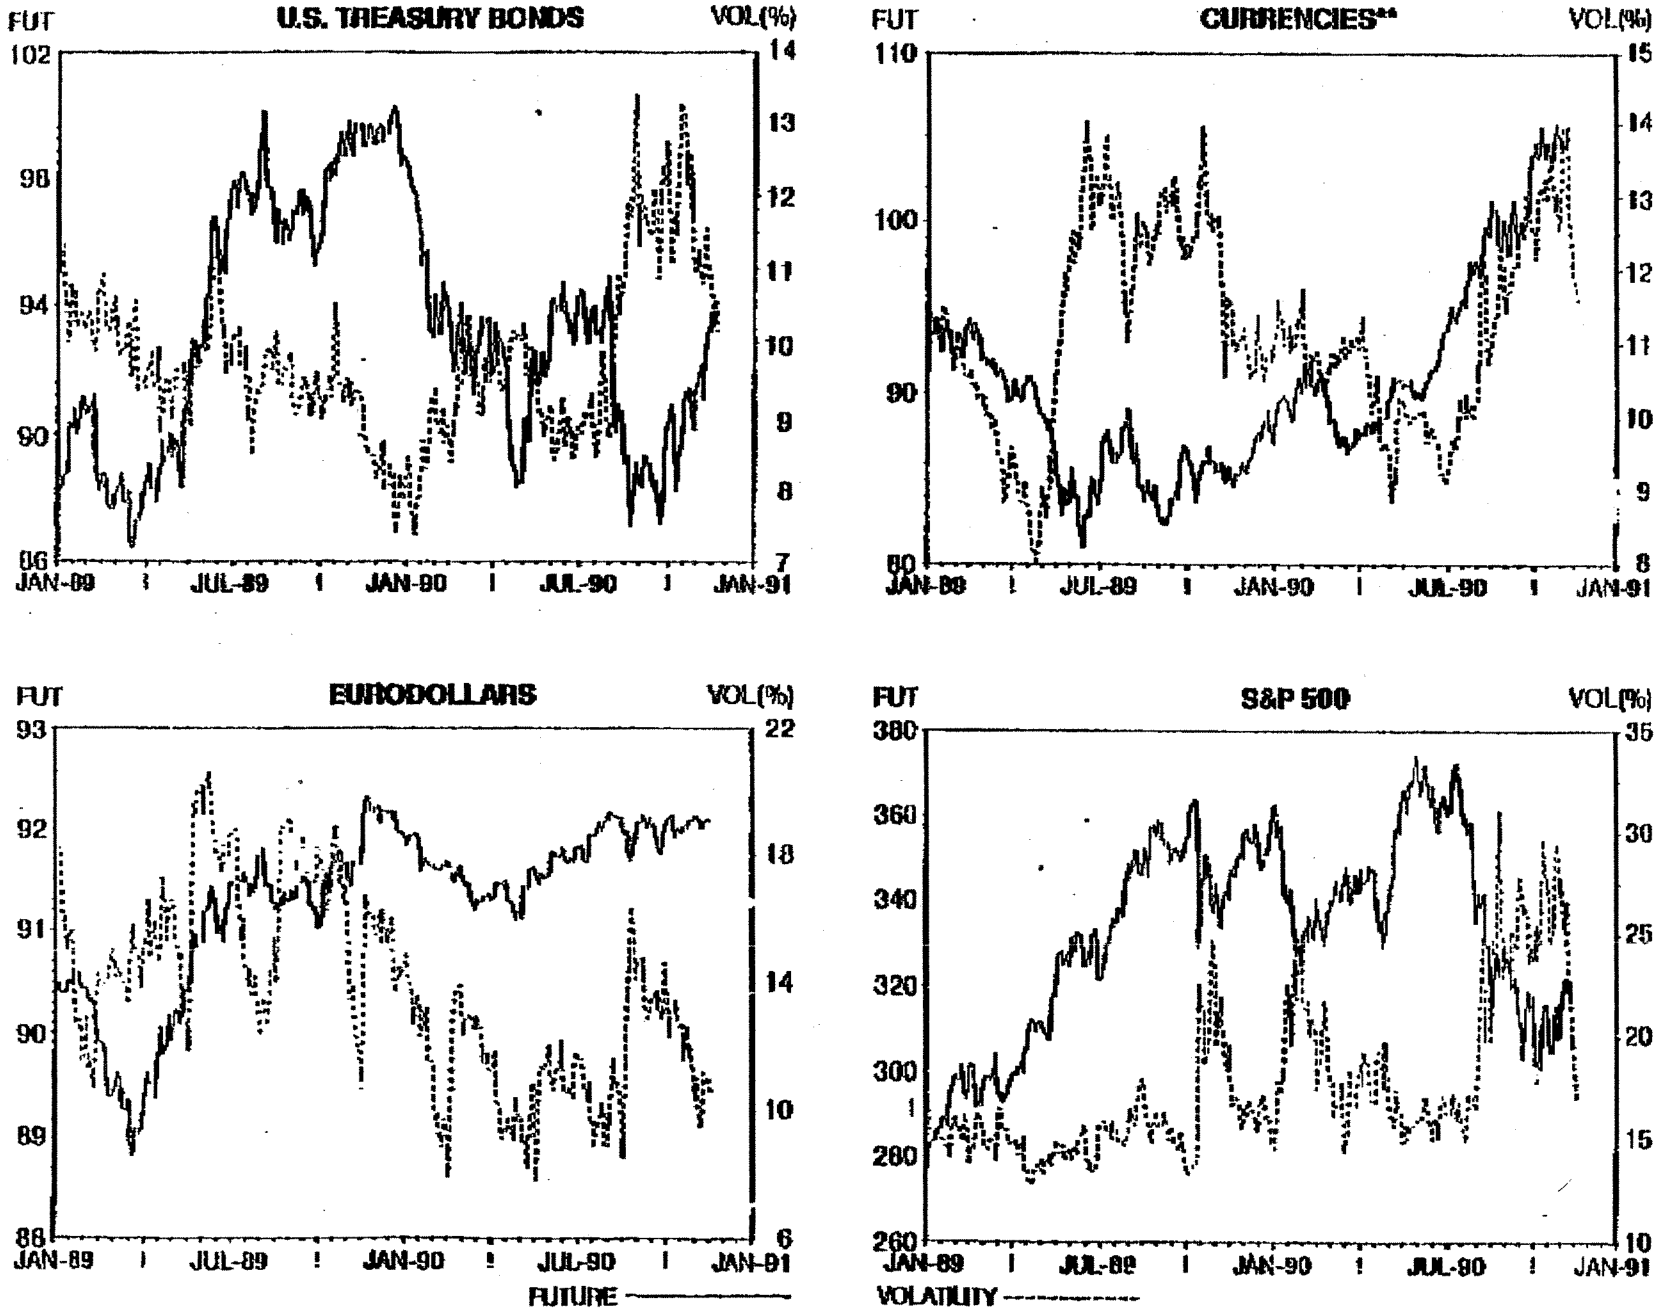

*Lead-month switch 20 days before option expiration

**Index of British pound, Deutsche mark, Japanese yen, Swiss franc

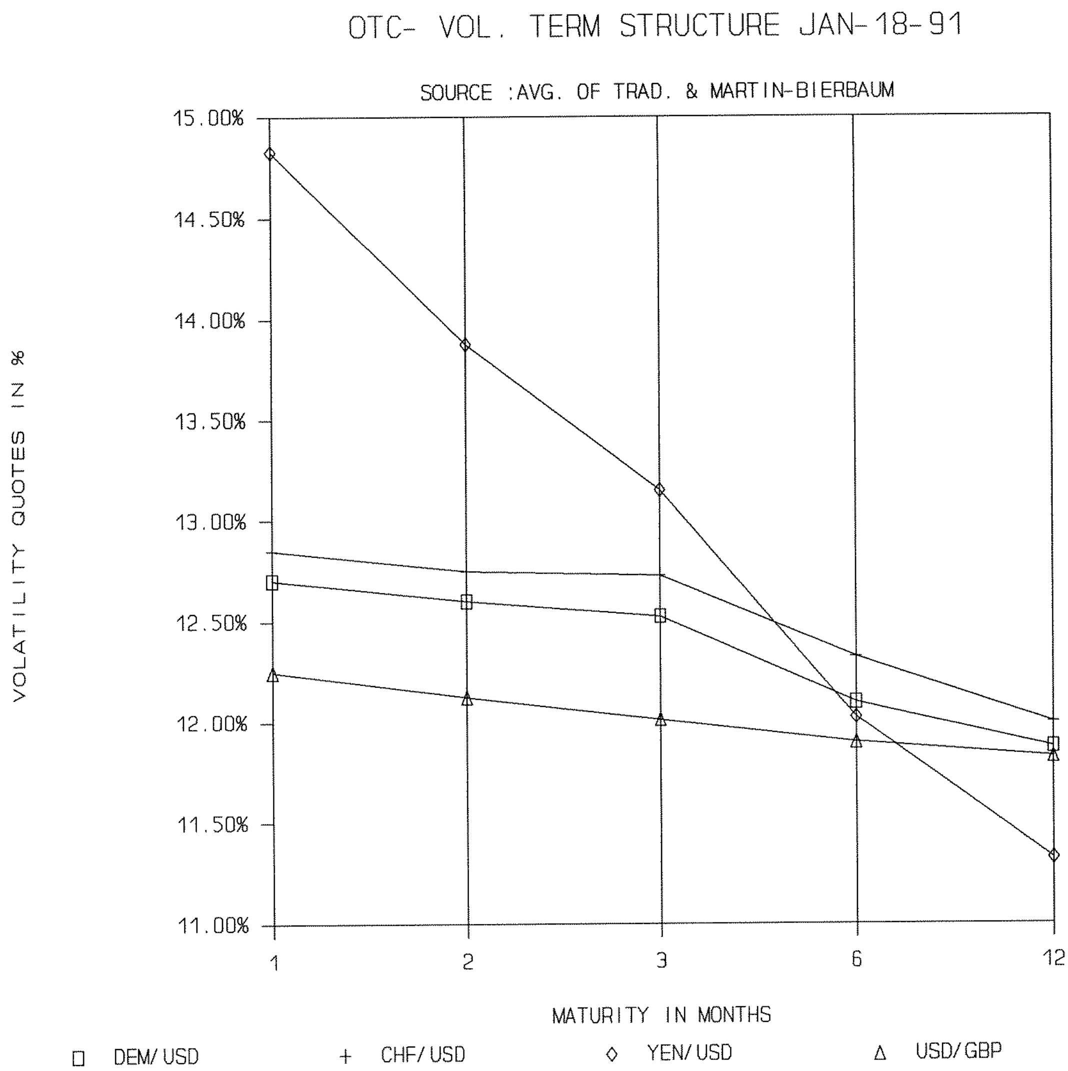
OTC- VOL. TERM STRUCTURE JAN-18-91
SOURCE :AVG. OF TRAD. & MARTIN-BIERBAUM
VOLATILITY QUOTES IN %
15.00%
14.50%
14.00%
13.50%
13.00%
12.50%
12.00%
11.50%
11.00%
1
2
3
6
12
MATURITY IN MONTHS
DEM/USD
CHF/USD
YEN/USD
USD/GBP

Figure 1

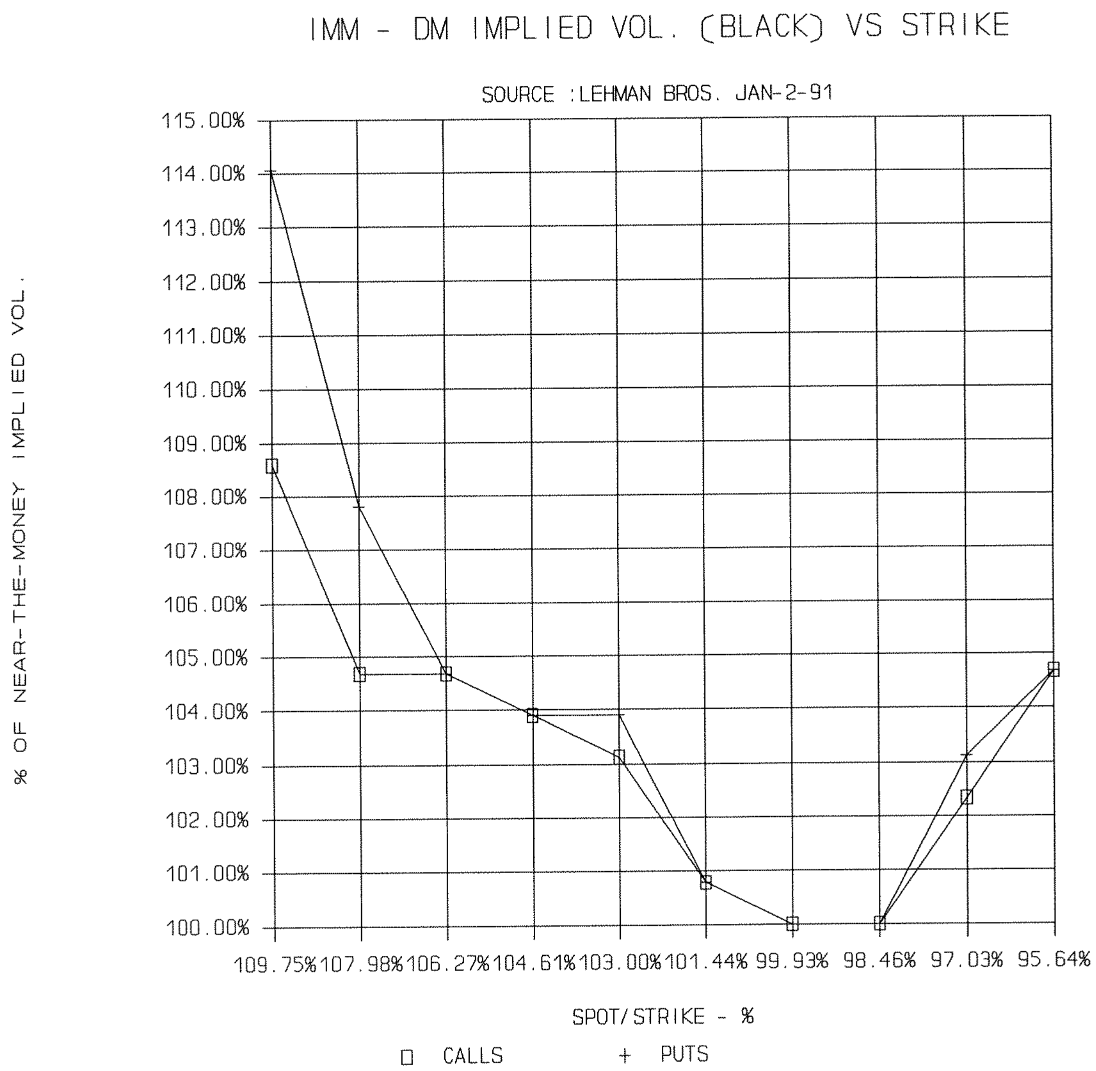
IMM - DM IMPLIED VOL. (BLACK) VS STRIKE
SOURCE :LEHMAN BROS. JAN-2-91
% OF NEAR-THE-MONEY IMPLIED VOL.
115.00%
114.00%
113.00%
112.00%
111.00%
110.00%
109.00%
108.00%
107.00%
106.00%
105.00%
104.00%
103.00%
102.00%
101.00%
100.00%
109.75%
107.98%
106.27%
104.61%
103.00%
101.44%
99.93%
98.46%
97.03%
95.64%
SPOT/STRIKE - %
□ CALLS
+ PUTS

Figure 3

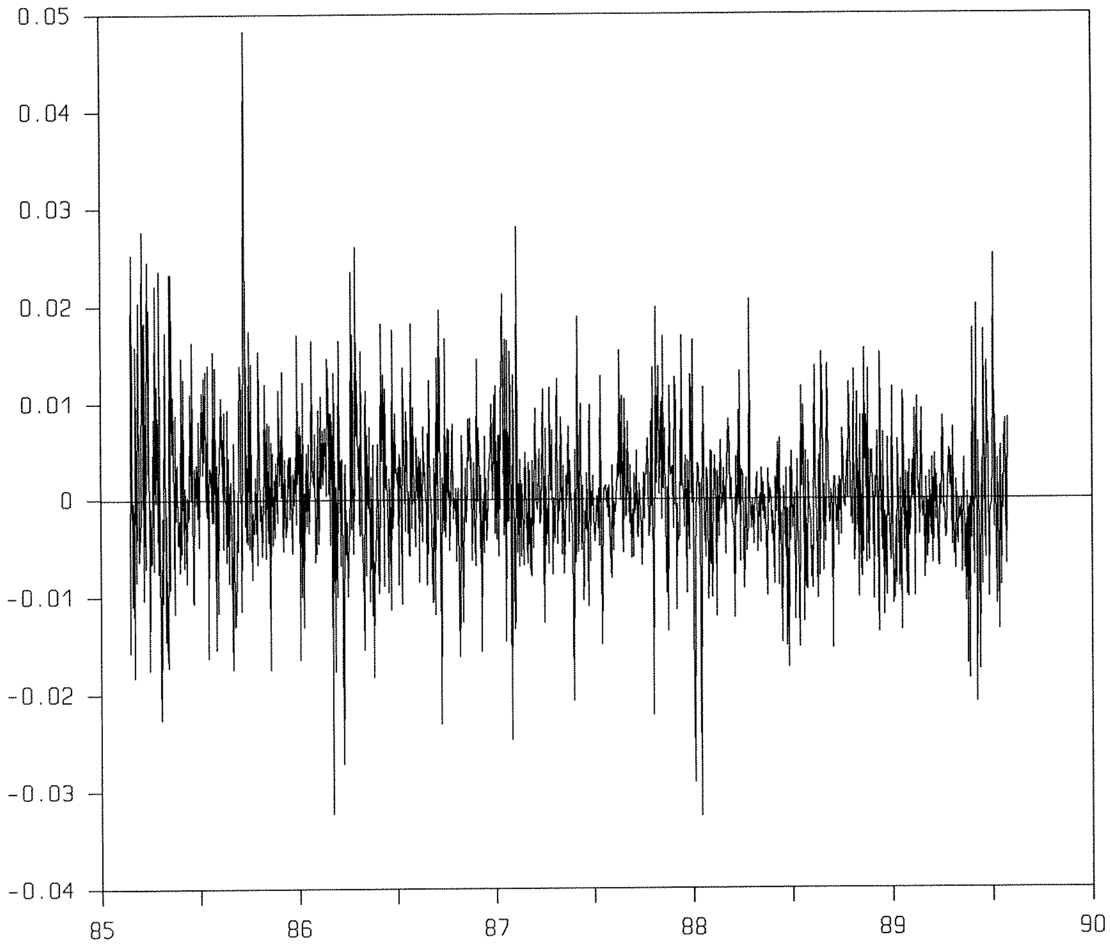

Figure 4 DM Futures Rates of Change

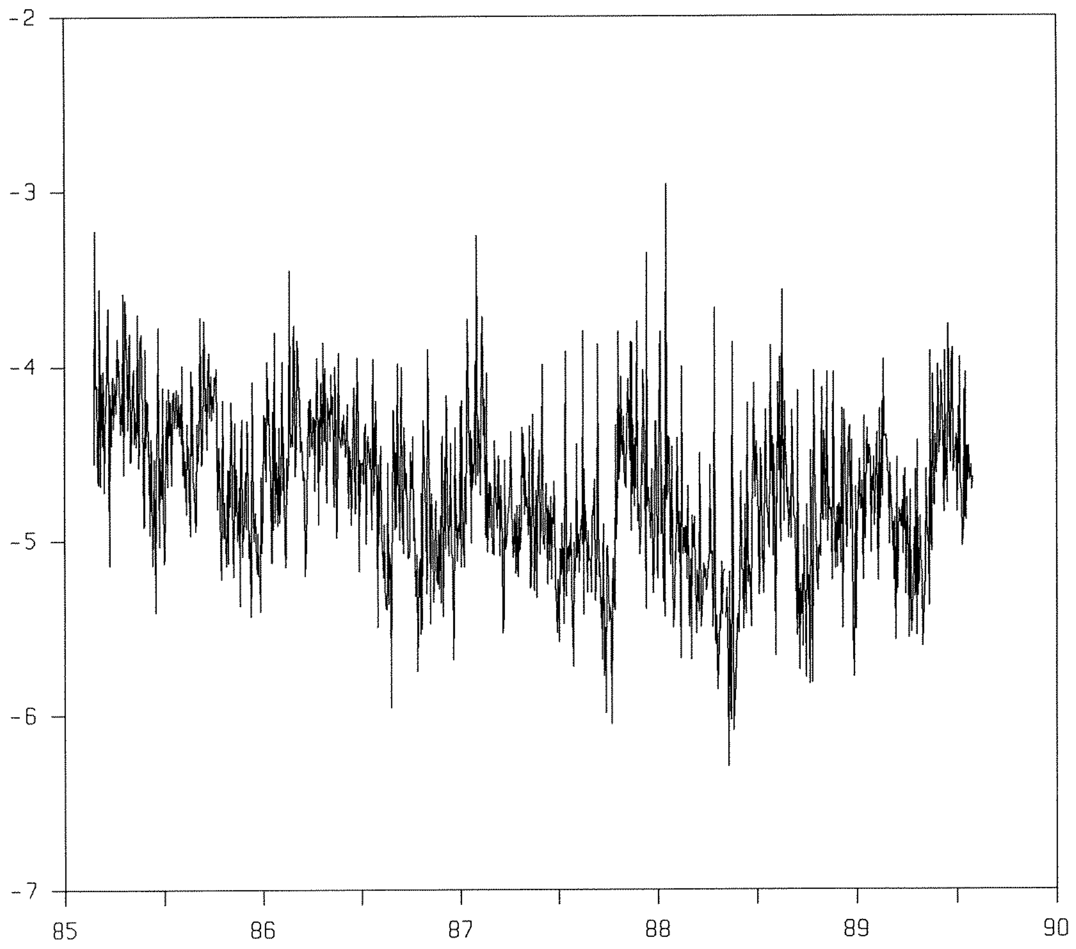

Figure 5 Realized Volatility (15 minutes DM Futures Prices)

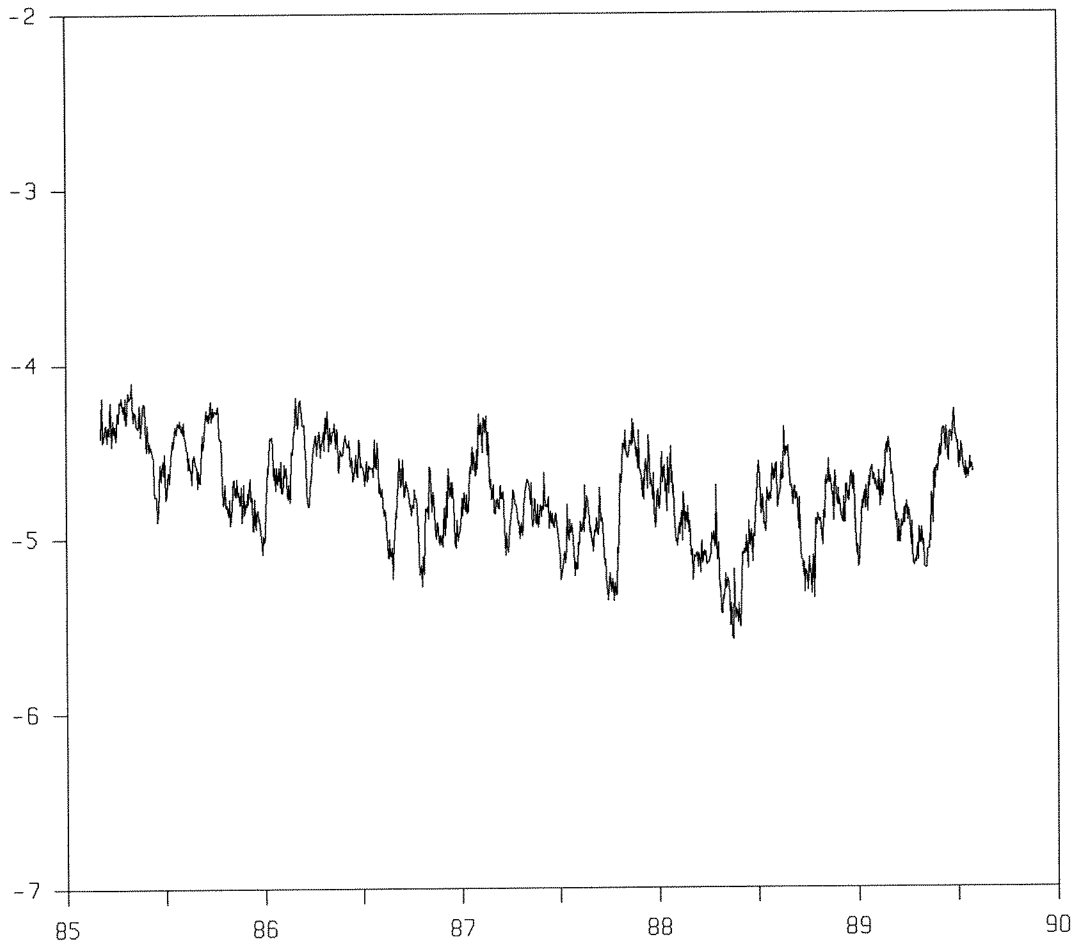

Figure 6 Fitted Volatility

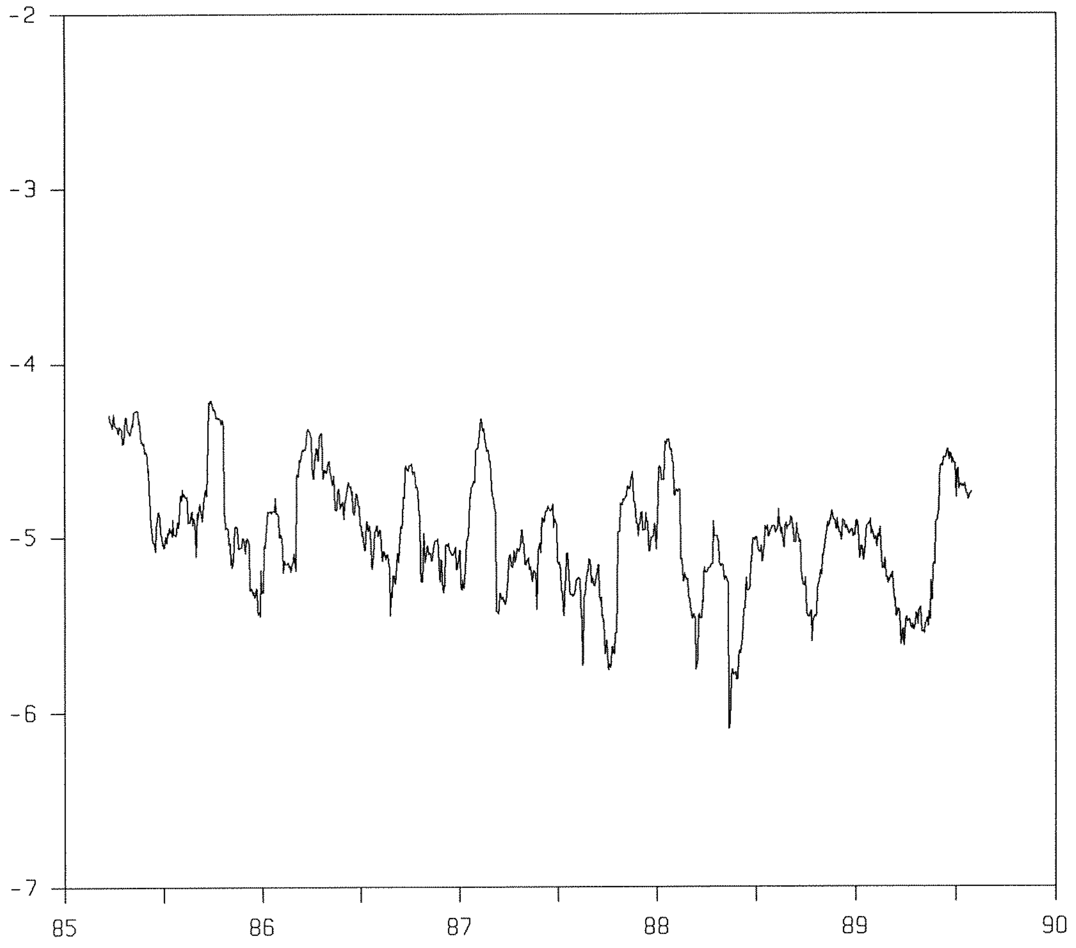

Figure 7 Historical Volatility (20 days)

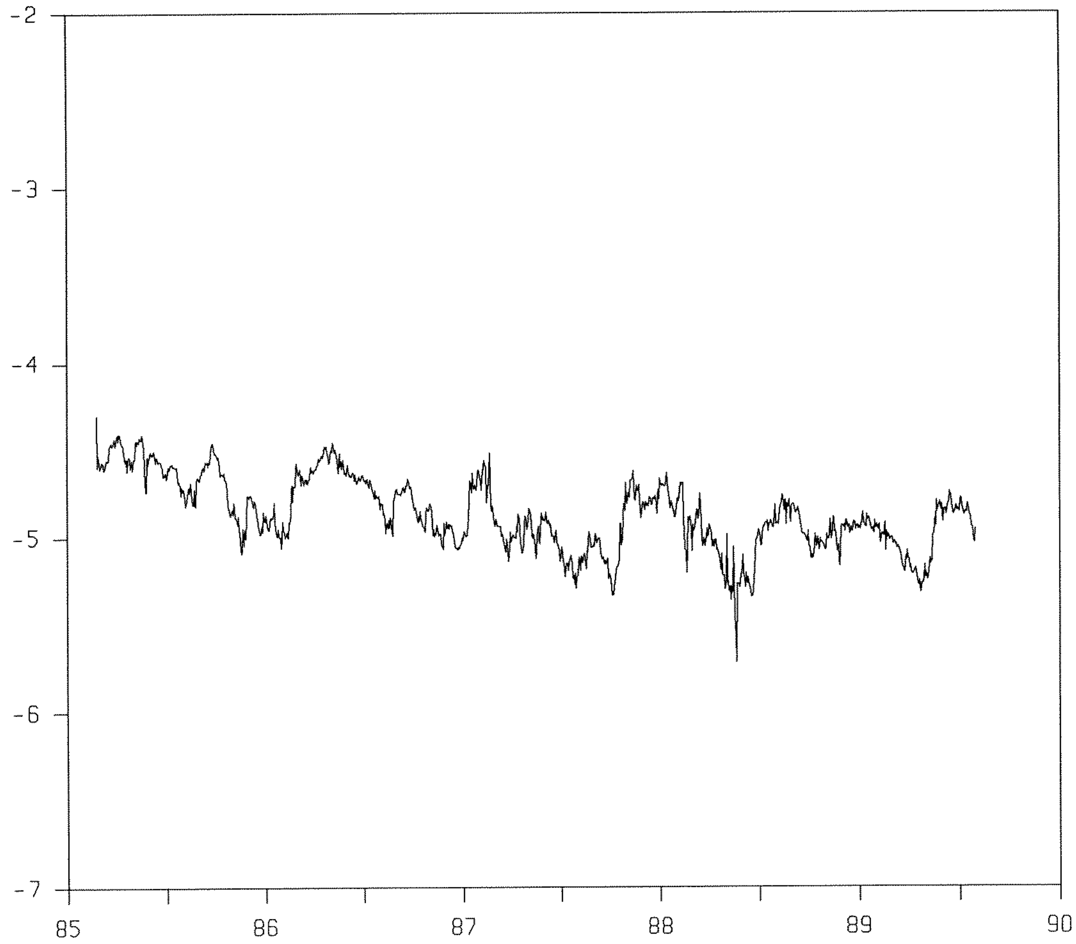

Figure 8 Call Volatility

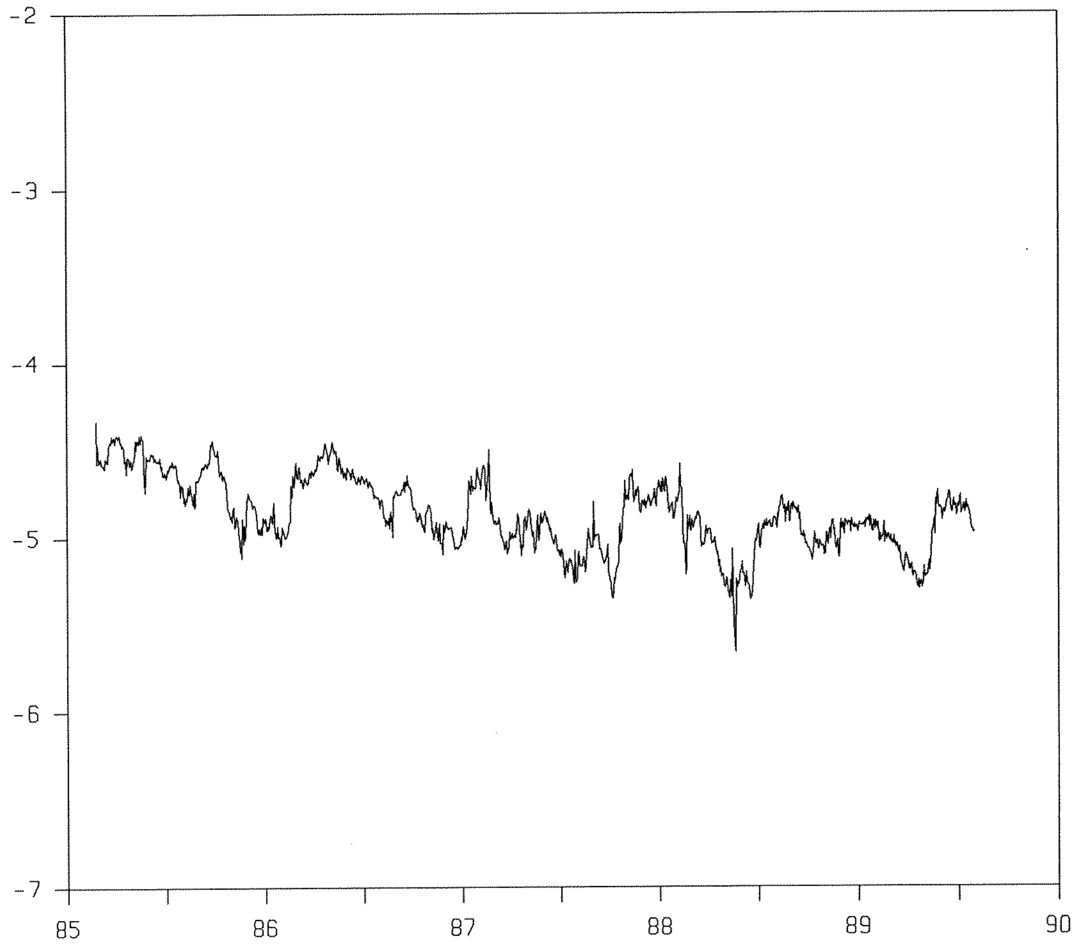

Figure 9 Put Volatility

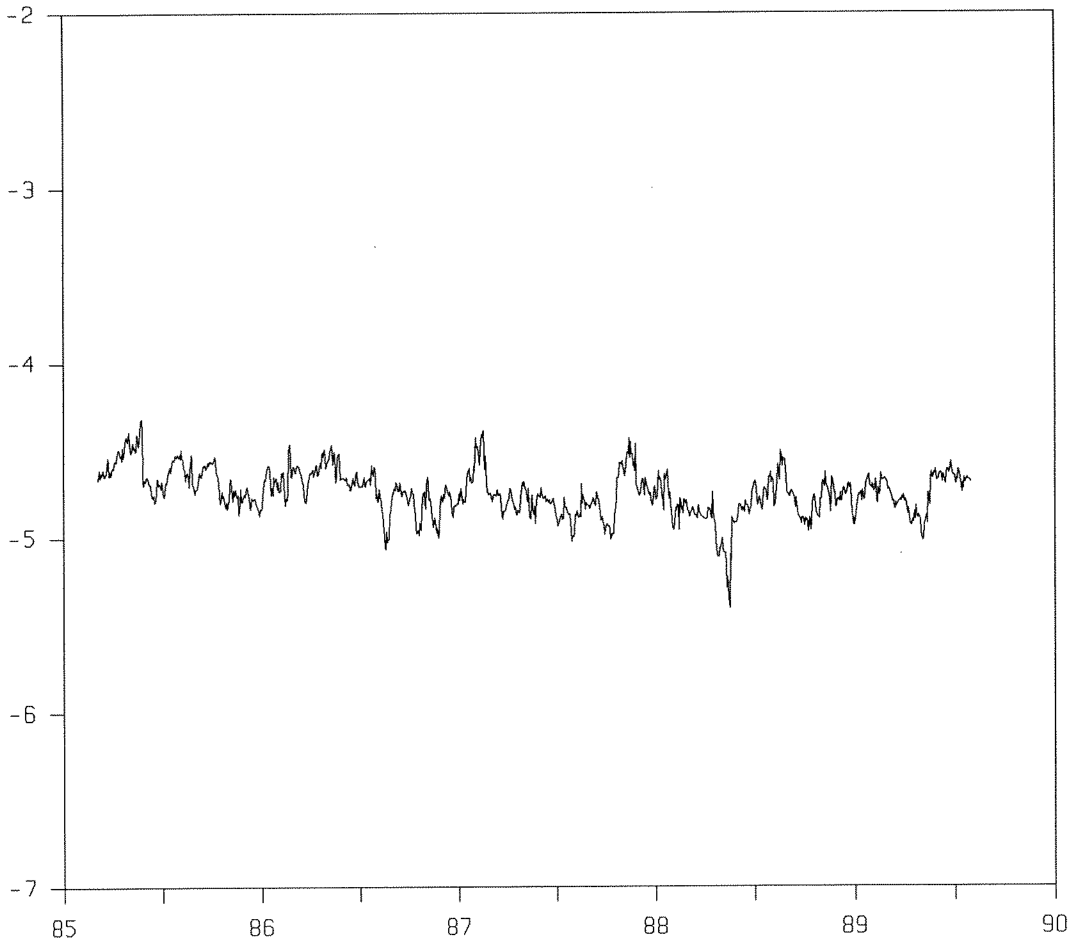

Figure 10 Average Future Volatility

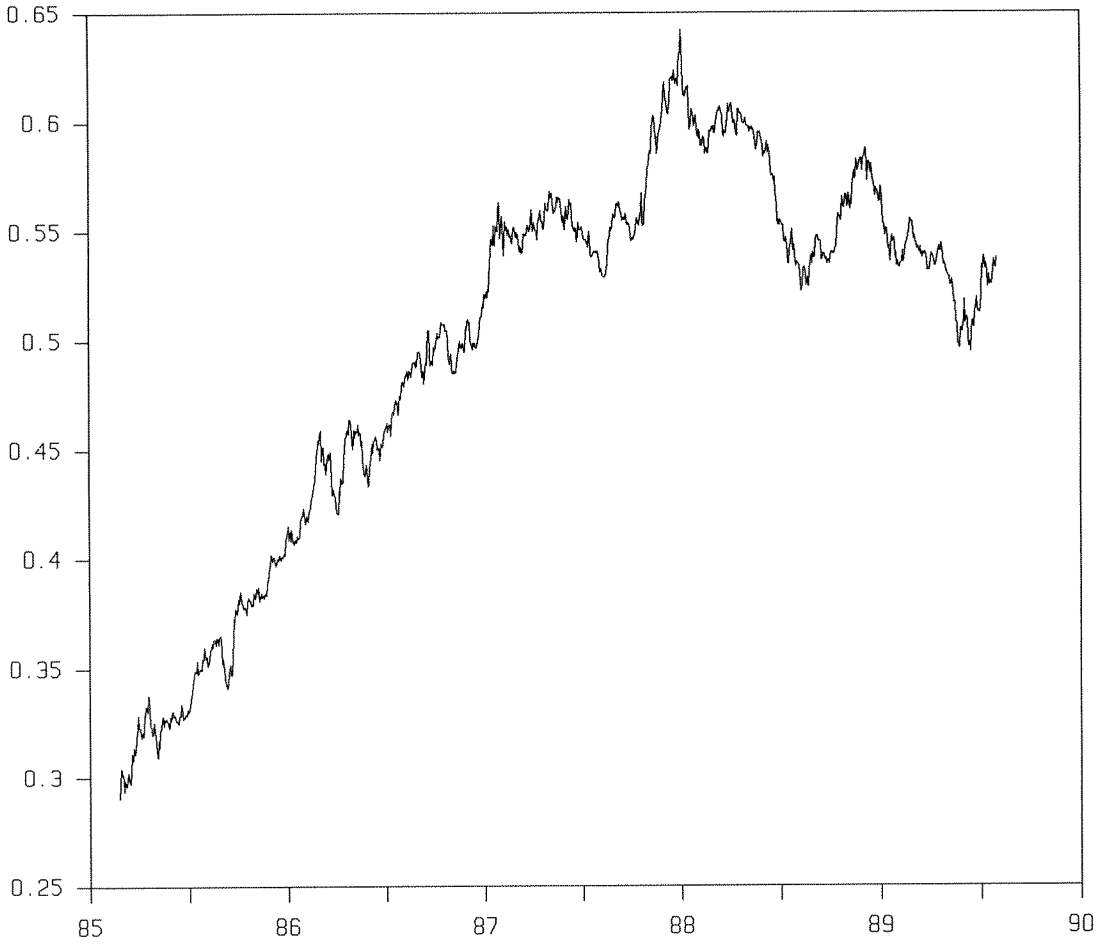

Figure 11 Levels of the DM Futures