Volatility Scaling in Foreign Exchange Markets

Jonathan Batten
Department of Banking and Finance
Nanyang Technological University, Singapore
email: ajabatten@ntu.edu.sg

and

Craig Ellis
School of Finance and Economics
University of Technology, Sydney
email: craig.ellis@uts.edu.au

Volatility Scaling in Foreign Exchange Markets

Abstract

When distributions are non-Gaussian or display linear dependence it may not be appropriate to annualise the risk coefficient determined by the linear rescaling of the variance from other time intervals. This paper investigates the scaling relationships for daily spot foreign currency returns: the Deutsche mark- U.S. dollar (DMK/USD), the Swiss franc-USD (SWF/USD), the Japanese yen-USD (JPY/USD), and the English pound-USD (GBP/USD), from February 1985 to May 1998. We find that all four series were non-Gaussian and displayed similar scaling properties with the estimated variance, based upon a scaling at the square root of time, significantly underestimating the actual level of risk predicted from a normal distribution. The economic implications of these results were then established by estimating the premiums on a series of foreign currency options using a variation of a Black-Scholes model with varying strike prices. Based on a 180 day maturity, the results suggested that the inappropriate scaling of risk underestimated the price of call and put options for DMK/USD, SWF/USD and JPY/USD, but not for GBP/USD.

Volatility Scaling in Foreign Exchange Markets

Non-normality is a feature of many financial time-series. However many market valuation models assume normality or attempt to adjust for the non-normality of price distributions. Many of these valuation models also require an annualised risk coefficient. However, under the random walk model, the temporal dimension of risk is irrelevant with the risk of the asset at any time interval being estimated by the linear rescaling of the risk from other time periods. That is, the convention is for risk to be scaled at the square root of time. This crucial assumption has rarely been tested and may be inappropriate for many valuation techniques.

Researchers generally have begun their valuation analysis by investigating the distributional qualities of financial time-series data. Evidence of non-normality in the distributions of daily and weekly returns for spot and forward currencies has been extensive (Hodrick, 1987). More recently the literature has focused upon the low and high order correlation structures of time-series. Low order correlation structures suggest short-term memory, or short-term dependence, which tend to decay rapidly over 2 to 3 days. Alternately high order structures suggest long-term memory, or long-term dependence, which persists providing positive economic benefits to investors with trend following forecasting systems.

The presence of long-term memory effects may be investigated by calculating an exponent named after the original developer Hurst (1951); the Hurst exponent, or using the modification suggested by Lo (1991) to eliminate low order persistence. Findings have generally been in favour of the presence of fractal structures in spot currency (Cheung, 1993; Batten and Ellis, 1996), stock indicies (Mills, 1993; Nawrocki, 1995) and futures markets (Fang, Lai and Lai, 1994; Corazza, Malliaris and Nardelli, 1997). While these researchers noted the presence of long and short-term dependence in the lag structure of time-series data, the implication for the estimation and modeling of risk (the standard deviation of the log of returns) has not investigated.

Statistical long-term dependence has particular implications for modelling the behaviour of financial market asset prices. Modelled as Brownian or fractional Brownian line-to-line functions[1], the principle of scale invariance suggests an observable relationship between asset returns across different time frequencies. Series exhibiting long-term dependence (of the type associated with fractional Brownian motion) should scale by a factor equivalent to their Hurst exponent. By contrast, Gaussian series should scale by the factor H = 0.5.

An example of a common application of scaling in financial time-series involves the annualising of risk in the Black-Scholes Option Pricing Model. Under the assumption of a normal Gaussian distribution the Hurst exponent H = 0.5 implies for time-series data that mean annual increments (or returns) should be 12 times the equivalent monthly mean and 52 times the equivalent weekly mean. Similarly, the standard deviation of annual increments should be $\sqrt{12}$ times that of the monthly increments. However, series generally exhibit higher peaks and fatter tails than are associated with the normal distribution and become more leptokurtic as the frequency at which the returns are measured increases. One implication of this result is that rescaling returns by the square root of time necessarily incorporates the use of an inappropriate scaling factor. By way of example, Peters (1994) demonstrated that empirical estimates of the standard deviation of daily returns for spot USD/JPY scaled at a rate H = 0.59. In terms of the annualising of risk in the Black-Scholes model, this result suggests that volatility would be underestimated when scaled using H = 0.5.

The presence of fractal structures in time-series data also has implications for investors and risk managers. Analysing mean absolute logarithmic price changes for five foreign exchange rate pairs, Muller, Dacorogna, Olsen, Pictet, Schwarz and Morgenegg (1990) found that intraday price changes scaled at approximately H = 0.59. Measuring the standard deviation of interday returns as the proxy for risk, Holton (1992) and Estrada (1997) demonstrated that when asset prices did not follow a random walk, annualising risk by the square root of time ($T^{½}$) either overestimated or underestimated the true level of risk associated with an investment. Empirical analysis by both authors confirmed the tendency for the risk of financial asset returns to scale at a faster rate than the square root of time leaving investors with incorrect investment horizons and risk managers with incorrectly priced financial assets. Neither Holton, nor Estrada, attempted a solution to the problem of scaling risk for non-random (non-Gaussian) time-series.

---

1 The definition of a line-to-line function is one where (as in the case of time-series observations) increments in the series are constrained not to move backward. A Brownian line-to-line function is

To reconcile these findings, initially we investigate the distributional qualities of the daily returns of long-term currency data and determine the dependence relationships. The second objective was to establish the relationship between the data's distributional qualities, and the appropriate scaling relationships. Four long-term daily spot foreign currency series were investigated: the Deutsche mark-U.S. dollar (DMK/USD), the Swiss franc-USD (SWF/USD), the Japanese yen-USD (JPY/USD), and the English pound-USD (GBP/USD) from February 1985 to May 1998. These markets constitute the largest foreign exchange markets in the world measured in terms of turnover, are highly liquid, and have low transaction costs. The BIS (1996) survey on foreign exchange market activity estimated these currency pairs accounted for 61% of total daily turnover in the spot foreign exchange markets in 1995 (DMK/USD US$253.9 billion (29%), JPY/USD US$242.0 billion (18%), GBP/USD US$77.6 billion (9%) and SWF/USD US$60.5 billion (5%)). Trading also occurs on a 24 hour basis, with almost instantaneous transmission of news items to market participants using computerised technology and on-line broking services. Consequently these markets are as close to the efficient market ideal as is currently possible.

Initially, the distributional qualities of the data were investigated with all four series exhibiting features of non-normality (leptokurtosis), although the lagged correlation structure did not suggest the presence of long-term dependence. The four series were subsequently tested for the presence of long-term dependence using the Hurst technique. Then, the volatility of these four distributions was estimated using two hundred and fifty-two different scaling intervals (t–(t-1)) to (t –(t-252)) with all four scaled series underestimating the level of risk in the series, although to different degrees. The economic implications of these results were also significant. Employing a simple Black model for pricing currency options, we find that the incorrect scaling of volatility led to significant incorrect pricing decisions. Some questions still remain including what is the relevant and appropriate time horizon implied by the lag structure?

The next section discusses the scaling and self-similarity of Brownian and fractional Brownian motion systems in more detail. Then Section II establishes the theoretical basis and research methodology for the empirical research. Section III describes the currency data and provides evidence on its distributional qualities. Section IV presents our empirical results on

therefore analogous to a Gaussian random walk through time.

the scaling relationships. The final section summarises the results and allows for some concluding remarks.

I. Scaling and Self-similarity in Brownian and Fractional Brownian Motion.

The concepts of standard Brownian motion (sBm) and fractional Brownian motion (fBm) may be defined in terms of the relative level of dependence between increments. One characteristic of these processes is their self-similar behaviour. Developing the principles of fractal geometry, Mandelbrot (1977, 1982) proposed the concept of fractional Brownian line-to-line functions to describe time-series that exhibited an underlying fractal distribution. Encompassing values of the Hurst exponent $0 \le H \le 1$, the distributions of these functions were defined to be similar at different scales. Ordinary Brownian line-to-line motion B(t) in time t was characterised by small and mutually independent increments in B such that

$$F(x) = Pr\,(X < x) = Pr\left(\frac{B(t+\Delta t) - B(t)}{|\Delta t|^{H}} < x\right) \qquad (1)$$

where F(x) was the probability distribution of x. For real values of x and t, and H = 0.5, Equation (1) conformed to standard Brownian motion for all t and Δt. Distinguishing this from fractional Brownian line-to-line motion, $B_H(t)$, Mandelbrot relaxed the restriction of H = 0.5 necessary for a Gaussian process, thus allowing any $0 \le H \le 1$. For $0 \le H < 0.5$ the series $B_H(t)$ was characterised as being anti-persistent. Anti-persistent series would diffuse more slowly than ordinary Brownian line-to-line functions and would therefore appear “rougher” than a graph corresponding to a Gaussian process. Alternatively, persistent series with $0.5 < H \le 1$ would appear “smoother” than a graph of a Gaussian process of the same function.

Defining the fractal dimension (D) of a line-to-line function as a measure of the degree of irregularity of the graph Mandelbrot (1987) proposed an alternative interpretation of the value of the Hurst exponent was the amount of space the graph fills in two-dimensional space. Measured as

$$D = 2 - H \qquad (2)$$

the fractal dimension of an ordinary Brownian line-to-line function is $D = 1.5$. For the limiting values $H = 0$ and $H = 1$, the value of the fractal dimension is $D = 2$ and $D = 1$ respectively. The intuition of this result is that the graph corresponding to $H = 0$ would fill its space entirely. For the upper limit $H = 1$, the resulting graph would be linear, and hence only fill one dimension (length). In so far that Brownian motion, as originally defined by Bachelier (1900), implies an equal probability of incremental movement in *all* directions, the definition of the line-to-line function may be seen as a more precise description of time-series behaviour. These functions are explained in the following section, then evidence and implications of scaling relationships are provided.

A. Characteristics of fractional Brownian motion and line-to-line functions.

The related concepts of self-similarity, self-affinity and scale invariance can be used to describe the relationship between the parts and the whole of any function. For standard Brownian and fractional Brownian functions, the relationship is expressed in terms of self-similarity. For line-to-line functions (time-series) the correct expression is self-affine.

Consider a function (S) made up of the points $X = [X_0, X_1,..., X_n]$, where the probability of incremental movement is unrestricted with respect to the direction of the movement. Changing the length of the function by a common factor $r<1$, such that $rX = [rX_0, rX_1,..., rX_n]$, will yield a new function rS, whose geometric length is less than that of the original function. For the appropriate value of r, self-similarity implies the original function S can be recovered by N times contiguous replications of the self-similar rescaled function rS. In other words, the function S is scale invariant, ie. it is invariant to the change in scale by the factor r.

For line-to-line functions measured with respect to time, Mandelbrot (1977) showed such functions would instead be self-affine. Consider the same function (S), measured now as a line-to-line function comprising the points $X(t) = [X(t_0), X(t_1),..., X(t_n)]$, in time t. Changing the time scale of the function by the ratio $r<1$, the required change in scale of the amplitude was shown to be r^H for a self-affine function.

Given the function S was a Brownian line-to-line function, the distance from $X(t_0)$ to a point $X(t_0 + t)$ was shown by Mandelbrot (1977) to be a random multiple of $\sqrt{t}$. Setting $t_0 = 0$ it followed for $t>t_0$ that

$$X(t_0 + t) - X(t_0) \approx e\,|(t_0 + t) - t_0|^{0.5} \approx e\,t^{0.5} \tag{3}$$

where e was a random variable with zero mean and unit variance. Properly rescaled in time by r, and in amplitude by √r, the increments of the self-affine rescaled function (rS)/√r would be

$$\frac{X(t_0 + rt) - X(t_0)}{\sqrt{r}} \tag{4}$$

For the correct choice of scaling factor, the two functions S and (rS)/√r are statistically indistinguishable, such that they have the same finite dimensional distribution functions for all t_0 and all r>0. Consistent with the known value of the Hurst exponent for Gaussian series (H = 0.5), the scaling factor $\sqrt{r} = r^H$ is characteristic of all self-affine Brownian line-to-line functions. Allowing for $0 \leq H \leq 1$, $H \neq 0.5$ it follows for fractional Brownian line-to-line functions that Equation (3) can be generalised by

$$X(t_0 + t) - X(t_0) \approx e\,|(t_0 + t) - t_0|^{H} \approx e\,t^{H} \tag{5}$$

Equation (4) can similarly be generalised by

$$\frac{X(t_0 + rt) - X(t_0)}{r^H} \tag{6}$$

for a fractional line-to-line function. Self-affine line-to-line functions also appear the same graphically when properly rescaled with respect to H. Mandelbrot (1977) noted that fractional Brownian line-to-line functions should exhibit statistical self-affinity at all time scales. Independent of the incremental length (or frequency of observation) of $X(t_0 + t) - X(t_0)$, the relative level of persistence or anti-persistence should remain consistent. Peitgen, Jürgens and Saupe (1992), also revealed that the limit formula for the fractal dimension of Brownian and fractional Brownian line-to-line functions in Equation (2) could be proved to follow directly from the principles of scale invariance.

B. Evidence and implications of scale invariance for foreign currency time-series.

Statistical long-term dependence has particular implications for modelling the behaviour of financial market asset prices. Modelled as Brownian or fractional Brownian line-to-line functions, the principle of scale invariance suggests an observable relationship between asset returns across different time frequencies. Series exhibiting long-term dependence (of the type associated with fBm's) should scale by a factor that is equivalent to their Hurst exponent. Therefore Gaussian series should scale by the factor H = 0.5.

One objective of this research paper was to determine whether the returns from holding a spot currency position scaled in a manner that was consistent with their underlying structures. The first part of this Section will briefly recount empirical literature relating to evidence of statistical long-term dependence in financial asset markets. This will provide the motivation for the current analysis of the economic implications of long-term dependence in the various foreign currency time-series. It should be noted that the failure to identify long-term dependent effects suggests support for the nominated currency conforming to normally distributed standard Brownian motion.

The fractal structure of financial price data has been widely investigated by Peters (1991, 1994) whose results indicated that a number of financial time-series displayed a significant degree of long-term dependence. Ambrose, Ancel and Griffiths (1992) disputed these results on methodological grounds. In currency futures markets Kao and Ma (1992) described short-term price dependence in GBP/USD and USD/DEM contracts. However, in spot exchange markets Cheung (1993) and Batten and Ellis (1996) have also observed evidence of long-term dependence. Investigating the economic implications arising from price dependence, Batten and Ellis found evidence of arbitrage profits to speculators holding long USD/JPY positions when the market was characterised by positive persistence. Recent research into the distribution of financial asset returns has also provided evidence of long-term dependent effects. In US agricultural futures markets Corazza, Malliaris and Nardelli (1997) suggested evidence of long-term dependence using the Hurst exponent was consistent with the distributions of the contracts being Pareto stable.

II. Research Methodology

Modelling spot currency time-series as a line-to-line function measured with respect to time, scaling invariance can be used to describe the relationship between the moments of the distribution of the time-series at different time intervals. While the formal definition of self-affinity provided by Mandelbrot (1982) required that *every* mathematical and statistical characteristic of the time-series under observation and its self-affine rescaled function should be examined, in empirical practice this proof can be "inferred from a single test that is only concerned with one facet of sameness" (Mandelbrot, 1982: 254). The test conducted in this study is the self-affine relation described by Equation (5) and Equation (6) and is an examination of scaling relations between the volatility of returns measured over four different time intervals; k = 1, 5, 22 and 252 days. The values of k used corresponded to daily, weekly, monthly and annual return intervals. Continuously compounded returns for four spot currency pairs, DMK/USD, SWF/USD, JPY/USD and GBP/USD, were measured using $X_q = \ln(P_t - P_{t-k})$ for each above nominated value of k. The volatility of returns was estimated by the second moment of the distribution of each returns series.

Using the principles of scaling invariance, the volatility of returns at any time interval can be estimated from the volatility at any *other* interval such that for any combination of k and n, $\infty \geq k \geq n \geq 1$, where n is any interval length less than k

$$s_q = [\, s^2 (P_t - P_{t-k})\,]^{0.5} = (k/n)^H \, [\, s^2 (P_t - P_{t-n})\,]^{0.5} \tag{7}$$

Where the time-series under observation conforms to a standard Brownian line-to-line function the value of the exponent H in Equation (7) is H = 0.5. For fractional Brownian line-to-line functions, the value of H should be $0 \leq H \leq 1$; $H \neq 0.5$.

Based on the above principle, the Hypotheses tested in this study were:

H_0 : $\sigma_q = (k/n)^{0.5}\, \sigma(P_t - P_{t-n})$

H_1 : $\sigma_q = (k/n)^H\, \sigma(P_t - P_{t-n})$, where $1 \geq H \geq 0$, $H \neq 0.5$

Implied standard deviations for each interval (k = 5, 22 and 252) were estimated from the standard deviation of n interval returns (n = 1, 5 and 22) for all n<k, and the results compared to the observed standard deviations. The imputed value of the scale exponent, for which the

standard deviation of k interval returns could be exactly estimated from the n interval standard deviations, was then calculated for each currency pair. The significance of the imputed scale exponent was that this represents the value of H for which the implied k interval standard deviations equals exactly their observed values. The acceptance of the null hypothesis that the appropriate scale exponent (H) is H = 0.5, will imply the series under observation conforms to a random Gaussian distribution. Observed values of H which are significantly different from H = 0.5, will imply the rejection of the null hypothesis.

In order to test for the economic significance of our findings, implied annual standard deviations (k = 252) for each currency pair were estimated from the standard deviation of returns over daily (n = 1), weekly (n = 5) and monthly (n = 22) intervals. The results were then used to calculate the values of a series of in-the-money, at-the-money and out-of–the-money European call and put foreign currency options.

III. Foreign Currency Returns

The time-series properties of the natural logarithm (loge) of interday returns (P_t - P_{t-1}), for the four exchange rates (DMK, SWF, JPY and GBP against the US dollar) from 22 February 1985 to 27 May 1998 (n = 3327 observations) are presented in Tables 1 to 4. These are discussed in turn.

(insert Table 1 about here)

The first table (Table 1) provides information on the moments of the four time-series. For all series, the sample means were slightly negative but not significantly different to zero. However, the sum of the loge returns over the sample period (e.g -0.6416 for the DMK/USD) shows the impact of the devaluation of the USD against all four currencies. For example, the DMK/USD loge return (-0.6416) is equivalent to a nominal positive return of 52.64% from holding DMK instead of USD, excluding the effect of the interest rate differentials over the sample period. Consequently, technical trading strategies that exploited the long-term depreciation of the dollar (such as moving average systems) would be expected to show positive returns. Which of the many technical trading strategies was best, given holding and transaction costs, provides scope for further analysis.

The four series were also non-Gaussian, due largely to their leptokurtic nature than due to skewness, which was close to zero. Three normality tests are reported (Anderson-Darling, Jacque-Bera and Shapiro-Wilk) and each failed to accept normality in the four series. This result is not surprising since leptokurtic, non-normal distributions are common in financial time-series. Also, a note of caution should be exercised in interpreting these three test results since they are very sensitive to small deviations from normality when the sample size is large. However these small differences may impact on pricing and risk management practice due to the higher probability of observations being away from the mean when the series displays leptokurtosis. Overall the four series displayed different degrees of departure from normality, with the JPY/USD and the GBP/USD being more normal than the DMK/USD and SWF/USD. These phenomena may be partly explained by the correlation structure of the four currencies pairs presented in Table 2.

(insert Table 2 about here)

While the four currency pairs suggest significant positive correlation (consistent with the DMK, JPY, SWF and GBP all appreciating against the USD), the co-movement of returns differed between the various currency pairs. Not surprising given the currency interdependence of Europe, the highest positive correlations were between the DMK and SWF (0.916), DMK and GBP (0.731), and SWF and GBP (0.698), while the lowest was between the GBP and JPY (0.472). These results highlighted the limited ability for portfolio managers to diversify currency portfolios during the sample period.

(insert Table 3 about here)

The four time-series were also tested for autocorrelation up to 252 lags, with the results for lags 1, 2, 5, 22 and 252 presented in Table 3. The results suggest slight positive short-term autocorrelation that dissipated after lag one. The significance of the autocorrelations was tested using t-statistics and Ljung-Box Q statistics (reported as LBQ in Table 3). The t-statistics were all high (ie. $t > 1.25$) for all currency pairs at lag one, but not high thereafter. However the LBQ statistic was only significant for SWF/USD (p-value = 0.0908) and GBP/USD (p-value = 0.0026) at lag 1. Over the longer lag structure three of the currency pairs (DMK/USD, JPY/USD and GBP/USD) had autocorrelations which could be judged as being significantly different to zero (at the 90% level of confidence) at lag 22 (LBQ p-values were 0.0640, 0.0680 and 0.0524 respectively), while two of the currency pairs (JPY/USD and GBP/USD) had autocorrelations which could be judged as being significantly different to zero (at the 90% level of confidence) at lag 252 (LBQ p-values were 0.0899 and 0.0514 respectively).

The next table, Table 4, reports tests for mean stationarity in the four series. Values for the Augmented Dickey-Fuller (ADF) test at lags of 9 and 28, and for the Dickey-Fuller (DF) test (lag 0) are presented for each currency pair. Calculated for the loge returns series, the significantly high negative values for each currency indicated that all the returns series were stationary over the sample period. Critical values for the ADF and DF tests at the 1% and 5% levels were –3.96 and –3.41 respectively.

(insert Table 4 about here)

Though not reported in Table 3, the four time-series were also tested for partial autocorrelations up to 252 lags. Only the t-statistics for the DMK/USD at lag 1 (correlation = 0.0255, t-statistic = 1.4714), and the GBP/USD at lag 1 (correlation = 0.052141, t-statistic = 3.0075) were significant. Given the series were stationary, and had significant partial and auto-correlations at lag 1, the 4 series were then tested as AR(1) and ARMA(1,1) processes. None of the four series had significant coefficients when tested as ARMA(1,1) processes, however the GBP/USD and SWF/USD both had significant coefficients when tested as AR(1) processes (i.e. GBP/USD coefficient was 0.0256, t-statistic 3.01, p-value 0.003, and the SWF/USD coefficient was 0.0293, t-statistic 1.69, p-value 0.091). This result is consistent with the significant autocorrelations at lag 1 for both series as provided by the LBQ statistic.

Overall, although the four series provided evidence of non-normality in the form of leptokurtic returns with a slight positive autocorrelation structure. Examination of the autocorrelation structure of each series showed no compelling evidence that the appropriate scale exponent was other than H = 0.5 (ie a Gaussian or random walk series). This result alone would imply that the standard deviation of interday return's should scale at $\sqrt{t}$. The next section investigates this issue in more detail by determining the scaling properties of the four currency series. Then, the implications for the pricing of currency options when volatility is incorrectly scaled are determined using simple currency option pricing models.

IV. Scaling Properties of Foreign Currency Returns

The scaling properties between the standard deviation of k and n interval returns ($\infty \geq k \geq n \geq 1$), for each of the four foreign currency pairs are presented in Tables 5 and 6. Table 5 results present statistical evidence of volatility scaling, while results presented in Table 6 demonstrate the economic significance of the empirical findings. Each set of results is discussed in turn.

(insert Table 5 about here)

Results presented in Table 5 show implied values for the scale exponent H in Equation (7) above. Using observed values of the standard deviation of k and n interval returns, the implied scale exponent is the value of H in Equation (7) for which rescaled values of the standard deviation of n interval returns exactly equal the standard deviation of k interval returns. Under the null hypothesis H_0, the value of the scale exponent should be H = 0.5. Rejection of the null hypothesis will imply that the currency series tested did not conform to a Gaussian random walk over the sample period.

Implied values of the scale exponent in Table 5 are typically closer to H = 0.5 when the difference between the return intervals k and n is small. That is, when the standard deviation of weekly returns (k = 5) are estimated by rescaling daily return standard deviations (n = 1), implied H values are closer to H = 0.5, than when daily return standard deviations are rescaled to estimate the standard deviation of annual returns (k = 252). For example, the implied value of H for estimating the annual return standard deviation from the standard deviation of daily returns for the JPY/USD is H = 0.54573. Rescaling the daily return standard deviation to estimate the standard deviation of weekly JPY/USD returns however yields H = 0.49948. This result is consistent across all pairs of k and n and all foreign currency pairs, except for annual GBP/USD returns, which show the opposite result.

Considering the size of implied values of the scale exponents in Table 5, it is not possible to conclude absolutely that the DMK/USD, JPY/USD and SWF/USD foreign currency series are long-term dependent. While values of H for the interval length n = 22 are generally higher than those for n = 5 and n = 1, none of the H values are significantly greater than H = 0.5. However, using the classical rescaled range techniques, Peters (1994) found that daily JPY/USD and DMK/USD returns scaled at H = 0.642 and H = 0.624 respectively. Overall, lower values of the scale exponent for the GBP/USD at *all* interval lengths provides support for acceptance of the null hypothesis for the GBP/USD exchange rate.

(insert Table 6 about here)

Evidence pertaining to the economic significance of the scale exponent values is provided in Table 6. Using a variation of the Black-Scholes Option Pricing Model for foreign currency options[2], annualised standard deviation estimates for intervals of n = 1, 5, and 22 lags were used to price in-the-money, at-the-money and out-of-the-money call and put options for each of the four currency pairs. Annual standard deviations were estimated using Equation (7) above. Under the assumption that each of the four currency pairs conformed to a Gaussian random walk, the value of the scale exponent in Equation (7) was H = 0.5. For the calculation of the Black-Scholes model, spot exchange rates for each currency were their actual value as at 27 May 1998. In-the-money and out-of-the-money exchange rates were set at -/+ 10.0% of the spot exchange rate. The domestic (US) and foreign risk-free interest rates were arbitrarily set at 8.0% and 10.0% each, and the time to maturity was one-half year (180 days based on a 360 day year). Results for each currency pair using rescaled standard deviations for interval lengths n = 1, 5, and 22 are reported in Table 6. Results using observed annual standard deviations (k = 252) are also presented. Under the null hypothesis H_0, option values derived from the rescaled standard deviations should not be different from those using the observed annual standard deviations.

Except for options written on the GBP/USD, option values derived from the rescaled standard deviations consistently underestimated their real value, based on the observed annual standard deviation (reported as n=252 last column in Table 6). Values for out-of-the-money call contracts were underestimated by a maximum of 84.0% (JPY/USD n= 1 call price of 0.0013) to a minimum of 40.0% (SWF/USD n= 22 call price of 0.0003). For out-of-the-money puts, it

2 Option prices were for European calls and puts. The model employed is analogous to the Black-Scholes continuous dividend option pricing model where the dividend yield (q) is replaced by the foreign risk-

was a maximum of 89.2% (JPY/USD n=1,5,22 put price of 0.0004). With respect to at-the-money calls and puts, call contract values were underestimated by a maximum of 25.5% (JPY/USD n=1 and 5 call price of 0.1825) and a minimum of 16.9% (SWF/USD n=22 call price of 0.0196). Equivalent put option values were undervalued by a maximum of 19.7% (JPY/USD n=1 put price of 0.2520) and a minimum of 13% (SWF/USD n=22 put price of 0.261). Finally, prices for in-the-money calls and puts derived from the rescaled standard deviations were undervalued by less than 1.0%. The result for in-the-money options is consistent with the option value being determined mostly by its intrinsic value than the level of volatility, with the reverse being the case for out-of-the money options. The results are also consistent with lower (higher) option vegas for in-the-money (out-of-the-money) positions. Across all contracts and for all currency pairs, option values derived from rescaled daily standard deviations were more highly undervalued than those based on rescaled monthly standard deviations.

V. Summary and Conclusions

The correct estimation of financial asset risk has important implications for investors using standard asset pricing models. Under the usual assumptions of independent and Gaussian distributed increments, traditional methods of estimating risk have required using an annualised risk coefficient which is calculated by linear rescaling of the variance from shorter time intervals. Only when the returns series under observation is independent, will rescaling provide correct estimates of the underlying level of risk associated with an investment. Dependence between increments in the returns series will conversely lead the investor to underestimate or overestimate their exposure to risk. The higher the underlying level of dependence, the greater the possibility of error in the estimation process.

Examining long-term returns for four spot foreign currencies (DMK/USD, GBP/USD, SWF/USD and JPY/USD) from February 1985 to May 1998, the objective of this study was to determine the statistical and economic implications for investors of rescaling financial asset risk. Using standard statistical tests, the distribution of returns for each of the four currency series were shown to be non-normal. However no significant evidence of dependence between series increments for the four currencies was found.

---

free interest rate.

Estimating scaling relationships between the volatility of returns at different time intervals produced some evidence of dependence not found using traditional techniques, with three of the four series tested having scale exponents (H) greater than 0.5 for all time intervals (H = 0.5 being the expected exponent value for independent increments). While the exponent values were not significantly large enough to conclude in favour of statistical long-term dependence in the currency returns series, the economic implications of the exponent values *were* significant. Using a simple Black-Scholes foreign currency option pricing model, linearly rescaled volatility estimates were shown to misprice the option value by as much as 25.5% for at-the-money contracts. These results are significant, since they demonstrate that even small deviations from independence in asset returns can result in significant economic benefits or costs. Investors should therefore exercise caution when using short-term returns to estimate longer-term risk, so as to avoid underestimating their real exposure to risk.

VI. References

Ambrose, B.W., Ancel, E.W. and Griffiths, M.D. 1992. The Fractal Structure of Real Estate Investment Trust Returns: The Search for Evidence of Market Segmentation and Non-Linear Dependency. Journal of the American Real Estate and Urban Economics Association. 20 (1). 25-54.

Bachelier, L. 1900. Teorie de la Speculation. Annales de l'Ecole Normale Superieure. 21-86.

Bank for International Settlements (BIS) 1996. Central Bank Survey of Foreign Exchange and Derivatives Market Activity 1995. Basle May.

Batten, J. and Ellis, C. 1996. Fractal Structures and Naive Trading Systems: Evidence from the Spot US Dollar/Japanese Yen. Japan and the World Economy. 8. 411-421.

Cheung, Y. 1993. Long Memory in Foreign-Exchange Rates. Journal of Business & Economic Statistics. 11 (1). 93-101.

Corazza, M., Malliaris, A.G., and Nardelli, C. 1997. Searching for Fractal Sstructures in Agricultural Futures Markets. Journal of Futures Markets. 17 (4). 433-473.

Estrada, J. 1997. Random Walks and the Temporal Dimension of Risk. Universidad Carlos III de Madrid. Departamento de Economia de la Empresa. Working Paper 97-24. April.

Fang, H., Lai, K.S. and Lai, M. 1994. Fractal Structure in Currency Futures Price Dynamics. Journal of Futures Markets. 14 (2). 169-181.

Hodrick, R.J. 1987. The Empirical Evidence on the Efficiency of Forward and Futures Foreign Exchange Markets. Harwood Academic Publishers: London.

Holton, G.A. 1992. Time: The Second Dimension of Risk. Financial Analysts Journal. November-December. 38-45.

Hurst, H.E. 1951. Long-term Storage Capacity of Reservoirs. Transactions of the American Society of Civil Engineers. 116. 770-799.

Kao, G.W. and Ma, C.K. 1992. Memories, Heteroscedasticity, and Price Limit in Currency Futures Markets. Journal of Futures Markets. 12 (6). 679-692.

Lo, A.W. 1991. Long-Term Memory in Stock Market Prices. Econometrika. 59 (5). 1279-1313.

Mandelbrot, B.B. 1977. Fractals. Form, Chance, and Dimension. W.H. Freeman and Company: New York.

Mandelbrot, B.B. 1982. The Fractal Geometry of Nature. W.H. Freeman and Company: New York.

Mandelbrot, B.B. 1987. Fractals. Encyclopedia of Physical Science and Technology. Volume 5. 579-593.

Mills, T.C. 1993. Is There Long-Memory in UK Stock Returns? Applied Financial Economics. 3. 303-306.

Muller, U.A., Dacorogna, M.M., Olsen, R.B., Pictet, O.V., Schwarz, M. and Morgenegg, C. 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.

Nawrocki, D. 1995. R/S Analysis and Long-term Dependence in Stock Market Indices. Managerial Finance. 21 (7). 78-91.

Peitgen, H., Jurgens, H. and Saupe, D. 1992. Chaos and Fractals. New Frontiers of Science. Springer-Verlag: New York.

Peters, E.E. 1991. Chaos and Order in the Capital Markets. John Wiley & Sons: New York.

Peters, E.E. 1994. Fractal Market Analysis. John Wiley & Sons: New York.

TABLE 1
Descriptive Statistics of the
Spot Foreign Currency Interday Returns
22 February 1985 to 27 May 1998

Currency n = 3327	DMK/USD	SWF/USD	JPY/USD	GBP/USD
Mean	-1.80 E-04	-1.80 E04	-2.10 E-04	-1.22 E-04
Standard deviation	7.08 E-03	7.72 E-03	6.86 E-03	6.92 E-03
Skewness	-9.60 E-02	-2.10 E-01	-4.10 E-01	1.70 E-03
Excess kurtosis	2.3395	1.8663	3.8328	5.0939
Sum of loge returns	-0.6416	-0.6617	-0.6476	-0.4156
Nominal return from short dollar strategy	52.646%	51.598%	52.331%	65.99%
Anderson-Darling	20.344 p=0.000	14.623 p=0.000	30.064 p=0.000	35.661 p=0.000
Jacque-Bera	763.84	507.30	2129.66	3597.05
Shapiro-Wilk	0.9860 p<0.01	0.9901 p<0.01	0.9777 p<0.01	0.9720 p<0.01

TABLE 2
Correlations of the
Spot Foreign Currency Interday Returns
22 February 1985 to 27 May 1998

	DMK/USD	SWF/USD	JPY/USD
SWF/USD	0.916		
p-value	0.000		
JPY/USD	0.608	0.612	
p-value	0.000	0.000	
GBP/USD	0.731	0.698	0.472
p-value	0.000	0.000	0.000

TABLE 3

Autocorrelations of the
Spot Foreign Currency Interday Returns
22 February 1985 to 27 May 1998

	DMK/USD	SWF/USD	JPY/USD	GBP/USD
Lag 1	0.03	0.03	0.02	0.05
t-statistic	1.47	1.69	1.21	3.01
LBQ	2.17	2.86	1.46	9.05
p-value	0.1407	0.0908	0.2269	0.0026
Lag 2	-0.02	0.00	0.02	0.01
t-statistic	-1.26	-0.26	1.23	0.82
LBQ	3.76	2.93	2.97	9.73
p-value	0.1526	0.2311	0.2265	0.0077
Lag 5	0.03	0.02	-0.01	0.02
t-statistic	1.61	1.17	-0.54	1.39
LBQ	7.87	4.75	5.14	19.60
p-value	0.1636	0.4472	0.3990	0.0015
Lag 22	0.00	-0.02	0.01	0.03
t-statistic	0.23	-0.87	0.71	1.78
LBQ	30.04	32.85	32.58	33.73
p-value	0.1175	0.0640	0.0680	0.0524
Lag 252	-0.016	-0.004	0.003	0.001
t-statistic	-0.90	-0.22	0.18	0.02
LBQ	273.71	258.22	282.61	289.71
p-value	0.1660	0.3804	0.0899	0.0514

TABLE 4

Stationarity Tests of the
Spot Foreign Currency Interday Returns
22 February 1985 to 27 May 1998

	DMK/USD	SWF/USD	JPY/USD	GBP/USD
DF (0)	-56.592	-56.094	-56.906	-55.056
ADF(9)	-17.875	-17.837	-17.044	-17.942
ADF(28)	-10.305	-10.030	-10.077	-10.368

TABLE 5

Implied Scale Exponent (H) for Estimation of k Interval Standard Deviation from Standard Deviation of n Interval Returns

Currency	Interval Length k =	Implied Scale Exponent (H)		
		n = 1	n = 5	n = 22
DMK/USD	5	0.49519		
	22	0.50723	0.52073	
	252	0.53400	0.54993	0.56680
SWF/USD	5	0.50167		
	22	0.51080	0.52104	
	252	0.53510	0.54882	0.56487
JPY/USD	5	0.49948		
	22	0.52123	0.54562	
	252	0.54573	0.56472	0.57576
GBP/USD	5	0.50887		
	22	0.51965	0.53173	
	252	0.49887	0.49476	0.47341

TABLE 6

Foreign Currency Option Values
using Implied Annual Standard Deviations

(A) In-the-money

Currency	Exchange Rate	Option Exercise	1[a]	5[b]	22[c]	252[d]
			Call			
DMK/USD	0.5607	0.4486	0.1024	0.1024	0.1024	0.1026
SWF/USD	0.6785	0.5428	0.1240	0.1240	0.1240	0.1243
JPY/USD	7.2659	5.8127	1.3272	1.3272	1.3276	1.3305
GBP/USD	1.6307	1.3046	0.2978	0.2978	0.2979	0.2978
			Put			
DMK/USD	0.5607	0.6729	0.1132	0.1132	0.1132	0.1136
SWF/USD	0.6785	0.8142	0.1371	0.1371	0.1372	0.1374
JPY/USD	7.2659	8.7190	1.4669	1.4669	1.4679	1.4737
GBP/USD	1.6307	1.9568	0.3291	0.3292	0.3293	0.3291

(B) At-the-money

Currency	Exchange Rate	Option Exercise	1	5	22	252
			Call			
DMK/USD	0.5607	0.5607	0.0141	0.0140	0.0145	0.0175
SWF/USD	0.6785	0.6785	0.0189	0.0189	0.0196	0.0236
JPY/USD	7.2659	7.2659	0.1825	0.1823	0.1969	0.2442
GBP/USD	1.6307	1.6307	0.0385	0.0392	0.0414	0.0383
			Put			
DMK/USD	0.5607	0.5607	0.0195	0.0193	0.0198	0.0229
SWF/USD	0.6785	0.6785	0.0254	0.0254	0.0261	0.0300
JPY/USD	7.2659	7.2659	0.2520	0.2518	0.2664	0.3137
GBP/USD	1.6307	1.6307	0.0541	0.0548	0.0570	0.0539

(C) Out-of-the-money

Currency	Exchange Rate	Option Exercise	1	5	22	252
			Call			
DMK/USD	0.5607	0.6729	0.0001	0.0001	0.0001	0.0004
SWF/USD	0.6785	0.8142	0.0002	0.0003	0.0003	0.0005
JPY/USD	7.2659	8.7190	0.0013	0.0013	0.0022	0.0080
GBP/USD	1.6307	1.9568	0.0002	0.0002	0.0003	0.0002
			Put			
DMK/USD	0.5607	0.4486	0.0000	0.0000	0.0000	0.0002
SWF/USD	0.6785	0.5428	0.0001	0.0001	0.0001	0.0004
JPY/USD	7.2659	5.8127	0.0004	0.0004	0.0008	0.0037
GBP/USD	1.6307	1.3046	0.0001	0.0001	0.0001	0.0001

Notes:
[a] Option value using an implied annual standard deviation from the standard deviation of daily returns.
[b] Option value using an implied annual standard deviation from the standard deviation of weekly returns.
[c] Option value using an implied annual standard deviation from the standard deviation of monthly returns.
[d] Option value using observed standard deviation of annual returns.