An Investigation of Long Range Dependence in Intra-Day Foreign Exchange Rate Volatility

By

Marc Henry and Richard Payne

DISCUSSION PAPER 264

March 1997

FINANCIAL MARKETS GROUP

AN ESRC RESEARCH CENTRE

LONDON SCHOOL OF ECONOMICS

ISSN 0956-8549-264

An Investigation of Long Range Dependence in Intra-Day Foreign Exchange Rate Volatility

Marc Henry and Richard Payne[1]

Financial Markets Group
London School of Economics
Houghton Street
London
WC2A 2AE

First Draft: March 1, 1996

This draft: May 20, 1997

Abstract

A comprehensive set of estimates of long memory in the volatility of three intra-day foreign exchange data series is presented. Robust semiparametric methods are used. Deseasonalizing procedures are proposed and permit the use of fully parametric methods which provide efficient tests of long memory. The hypothesis of long range dependence in the raw returns is rejected. In the volatility series, however, there is evidence of a long range dependent component, a finding which is significant and consistent across currencies. Furthermore, the hypothesis of I(1) volatility is strongly rejected in favour of a covariance stationary alternative, with evidence that previous findings of near-integrated volatility are due to the omission of long-range dependent components.

[1]Both authors are affiliated to the Department of Economics and Financial Markets Group, London School of Economics, London WC2A 2AE. Both gratefully acknowledge the financial assistance of the E.S.R.C., grant R000235892, and the Financial Markets Group. The first author also acknowledges the financial assistance of the Ministère de L'Education Nationale in the form of an A.M.X. grant. We would like to thank seminar particpants at CREST and the London School of Economics and Christian Gouriéroux, Wolfgang Härdle, Andrew Harvey, Bob Nobay and Peter Robinson for helpful comments.

1 Introduction

In recent years a vast amount of empirical work has been devoted to the characterisation of the temporal dependence in financial time series. Many authors have examined the time-series structure in asset returns, trading volumes and, perhaps most extensively, return volatility. Such studies are valuable in that they yield insights into issues such as the discrimination between regular and irregular market activity, the nature of information flows into financial markets, the way in which this information is assimilated into asset prices and the manner in which information is transmitted between markets. The current study extends the research in this area. We examine the relative performance of two classes of econometric models, ARMA-type representations and long range dependent specifications, in explaining the time-variation in financial time-series. Our empirical analysis concentrates on modelling the volatility process associated with a year long intra-daily sample of three major exchange rates.

The standard methodologies used to assess the degree of temporal dependence in a given time-series are the ARIMA class of models for series such as returns and volumes and GARCH or Stochastic Volatility specifications for modelling conditional heteroskedasticity.[1] Both of the above classes of models imply similar restrictions on the nature of the temporal dependence in the series in question. Their covariance stationary representations imply that the process is completely mean reverting and that this mean reversion is exponential in nature. In terms of the properties of the Wold representation, the series of Wold coefficients is hence both absolutely and square summable. The usual alternative to covariance stationarity is to assume the series is first order integrated i.e. an ARIMA(p,1,q) in the conditional mean or an IGARCH conditional variance process. This representation implies vastly different behaviour for the time-series in question. Shocks are completely persistent and hence the series of Wold coefficients are neither absolutely nor square summable.

Clearly, then, there is a very large distinction between the behaviours implied by I(0) and I(1) processes. On the one hand, one has a process which mean reverts at an exponential rate whilst on the other hand, the process in completely persistent. A far more general description of temporal dependence can be achieved by using more recent models which allow for long range dependence or long memory. One such model is the ARFIMA representation put forward by Granger and Joyeux (1980). Similar to the ARIMA class they permit non-zero orders of integration but, importantly, the order of integration is allowed to be non-integer. This generalisation allows one to model covariance stationary processes which have non-summable autocovariances.

[1] Bollerslev, Chou, and Kroner (1992) and Ghysels, Harvey, and Renault (1996) give excellent reviews of the GARCH and Stochastic Volatility modelling appraoches respectively.

Hence the effect of a shock on the process is completely mean reverting although very long lasting (the rate of decay in the Wold coefficients in hyperbolic rather than exponential as in the I(0) case.) The model also permits non-stationary processes which also mean revert. Hence, the ARFIMA model can be viewed as a more flexible alternative to the standard ARIMA, permitting a far a more general characterisation of the temporal dependencies in a given time-series.

We propose a comprehensive methodology for assessing the nature of temporal dependence. The methodology entails the following steps. The first step is to test the order of integration in the process using the methodology presented in Robinson (1996). The test is based on an underlying ARFIMA structure for the series in question and permits any degree of integration (integer or fraction) as a null hypothesis. Next, we gain a precise estimate of the degree of integration using two robust semiparametric long memory estimators. The first is the Log Periodogram regression (due to Geweke and Porter-Hudak (1983)) and the second is the Gaussian estimate proposed by Künsch (1987). Both of these estimates have been proven, in Robinson (1995b) and Robinson (1995a), to have standard asymptotics and are robust to any short range dependent features in the data. At this point, and analogous to the I(1) case, one could filter the long range dependence from the series and fit a covariance stationary ARMA to the residuals using traditional model selection procedures. We, however, go on to fit a fully parametric model to the series. The model is an extended version of the Long Memory in Stochastic Volatility (LMSV) model given in Harvey (1993), allowing for both short and long range dependence, and is sensitive to any mis-specification due to its fully parametric nature. The reason for fitting the fully parametric model is to permit one to assess the contributions of the short and long memory components to the overvall dependence in the series. This is accomplished via a set of quasi-Likelihood Ratio statistics for the fully parametric model. Hence, ultimately, we can discriminate between the long and short range dependent features of the process.

The data we examine in this work are the volatilities associated with three intraday foreign exchange (FX) returns series (the exchange rates in question being the DEM/USD, JPY/USD and JPY/DEM.) A pervasive result from previous work on this type of data is that the volatility process can be characterised as non-stationary (see, inter alia, Andersen and Bollerslev (1994), DeGennaro and Shrieves (1995) and Guillaume (1995).) There are however, some indications that this result may be due to mis-specification of the volatility models employed. First, the temporal aggregation results for GARCH processes do not hold when applied to FX data. The degree of persistence one identifies in daily data, for example, far exceeds that which would be implied by the results of estimations from data sampled at 1 hour intervals.[2] Second, other authors (e.g. Dacorogna, Müller, Nagler, Olsen, and Pictet (1993)) have noted that the correlograms of these intra-day volatility series decay far more slowly than the exponential decay which is associated with conventional GARCH or SV models. The combination of these two points serves as the motivation for our investigation of long-memory in volatility.

A theoretical motivation for the presence of long memory in asset price volatility can be generated by combining the simple mixture of distributions model in Tauchen and Pitts (1983) and the results on aggregation in Granger (1980). The former demonstrate, in a highly stylized framework, that both the volume and volatility in asset markets inherit the temporal dependencies associated with the latent flow of information into the market. Now assume that information flows are heterogeneous. Specifically, assume that there are an infinity of information arrival processes, each of which follows a stationary autoregression. The heterogeneity is modelled by variation in the AR parameters, which we assume follow a beta distribution. As Granger (1980) demonstrates, the aggregate information flow process will then exhibit long range dependence and, hence, so will volatility.

The paper is set out as follows. In Section 2 we present a more detailed account of our empirical methodology. Section 3 introduces the data employed in our study. As previously mentioned we focus on intra-daily FX volatilities, which are sampled at a ten minute calendar interval. Section 4 presents our estimation and testing results. We find that the FX returns process is well characterised by an I(0) process, in line with the efficient markets hypothesis. Results for the three volatility series demonstrate that all are covariance stationary and exhibit significant long memory. Further, estimation and testing of the fully parametric model demonstrates that the finding of non-stationarity in intra-day volatility is due to mis-specification. When one permits the possibility of long memory in volatility all specifications strongly indicate covariance stationarity. Section 5 concludes the paper and points to some directions for further work.

[2]See Andersen and Bollerslev (1994), for example.

2 Methodology

2.1 Semiparametric initial estimate

Let $\{r_t\}^T$ be our time series of raw returns and define $y_t \equiv \log r_t^2$ as an estimate of the volatility of the process.[3] Assume y_t to be covariance stationary with lag-j autocovariance γ_j. We require the autocovariances to satisfy the basic long range dependent specification

$$\gamma_j \sim cj^{2d-1} \text{ as } j \to \infty \tag{2.1}$$

where c is a finite positive constant and d is the slope parameter. Processes following 2.1 are generally labelled I(d) processes. An I(d) process is invertible when $d > -1/2$, in which case it can be reparametrized in an infinite moving average or Wold representation. I(0) denotes the short range dependent processes (including covariance stationary ARMA). In that case, the autocovariances and the weights in the Wold decomposition decay exponentially and are absolutely summable. When $0 < d < 1/2$, the process is still covariance stationary but autocorrelations decay hyperbolically and so slowly as to be non summable, even though they remain square summable. The Wold innovations share the same properties which imply mean reversion of the process. When $1/2 \leq d < 1$, the process becomes nonstationary while retaining its mean reverting characteristic. I(1) denotes the standard unit root case and implies complete persistence. The main feature of specification 2.1 is the smooth description of temporal dependence from the short range dependent I(0) case through the boundary of stationarity I(1/2) to the unit root I(1) case. The parameter d provides a measure of temporal dependence in the process. The focus of our attention is the range $d \in (0, 1/2)$ within which the process is termed covariance stationary, long range dependent.

Specification 2.1 translates to a related specification in the frequency domain.[4] We assume that y_t has spectral density $f(\lambda)$ satisfying

$$\gamma_j = \int_{-\pi}^{\pi} f(\lambda)\cos(j\lambda)d\lambda.$$

The basic long range dependent specification in the frequency domain is

$$f(\lambda) \sim g\lambda^{-2d} \text{ as } \lambda \to 0^+ \tag{2.2}$$

[3]The use of this proxy for volatility is motivated by the assumption that returns follow the following process, $r_t = \sigma e^{y_t/2}\epsilon_t$ where $\epsilon_t \sim N(0,1)$. This follows Harvey (1993) and Harvey, Ruiz, and Shephard (1994).

[4]See Robinson (1994) for details on the correspondence between the frequency domain specification 2.2 and the time domain specification 2.1.

where g is a finite positive constant and d is the slope parameter.

The two leading estimates of the long memory parameter d in the specification 2.2 are the Gaussian semiparametric estimate proposed by Künsch (1987) and the log-periodogram estimate proposed by Geweke and Porter-Hudak (1983). The asymptotic properties of both these estimates were derived in Robinson (1995a) and Robinson (1995b) respectively. Both estimation procedures are based on the periodogram

$$I(\lambda) = \frac{1}{2\pi T}\left|\sum_{t=1}^{T} y_t \exp(it\lambda)\right|^2$$

computed at the harmonic frequencies $\lambda_j = 2\pi j/T$ for $j \in [1, T/2)$. Although the periodogram is not a consistent estimate of the spectral density and is highly volatile, it is nevertheless an excellent tool in consistent frequency domain estimation.

The log-periodogram estimate is often preferred for its intuitive appeal. The long memory property is characterized, as mentioned above, by the slow hyperbolic decay of the autocorrelations. In the frequency domain, this property translates approximately to the existence of a hyperbolic pole around zero frequency as in 2.2. The idea of the log-periodogram is to linearize this local functional relation and obtain $\hat{d}$ as a least squares estimate of the slope parameter from the regression of a spectral estimate (in this case the periodogram) against a simple function of frequency. $\hat{d}$ is obtained from the regression

$$\log\left(I(\lambda_j)\right) = c - 2d\log(\lambda_j) + U_j \quad \text{for } j = l, \ldots, m \tag{2.3}$$

where m is the bandwidth and l is a trimming parameter such that $\frac{1}{l} + \frac{l}{m} + \frac{m}{T} \to 0$ as $T \to \infty$. A major drawback of the log-periodogram estimate for the type of data investigated in the current paper is the assumption of unconditional Gaussianity of the regression errors (U_j) which is difficult to relax due to the high degree of nonlinearity involved (the issue of consistency when Gaussianity is relaxed has been recently addressed by Velasco (1995)). Another drawback is the pathological behaviour of the very low periodogram ordinates (see Comte and Hardouin (1995)) which need to be trimmed out thereby adding another user-chosen parameter to the estimation when no automatic selection procedure has yet been proposed for the bandwidth itself.

The Gaussian semiparametric estimate is preferred here for its efficiency and robustness properties.[5] Like the log-periodogram, it is semiparametric in the sense that it relies on the low harmonics of the periodogram, but in this case, the lowest harmonics need not be trimmed out, so that it is consistent under the minimal bandwidth requirement

$$\frac{1}{m} + \frac{m}{T} \to 0 \quad \text{as } T \to \infty, \tag{2.4}$$

[5] Log-periodogram estimation results will be reported for comparison.

and an automatic bandwidth selection procedure is available in Henry and Robinson (1996). Estimation is based on the maximization of a local form of frequency domain log-likelihood

$$\frac{1}{m}\sum_{j=1}^{m}\left\{\log g\lambda_j^{2d} + \frac{\lambda_j^{2d}}{g}I(\lambda_j)\right\}$$

for $0 < m < [T/2]$ which, after concentrating out the constant g, is equivalent to minimizing

$$R(d) = \log\left\{\frac{1}{m}\sum_{j=1}^{m}\lambda_j^{2d}I(\lambda_j)\right\} - \frac{2d}{m}\sum_{j=1}^{m}\log\lambda_j.$$

The most appealing features of this estimate are its asymptotic normality (under a slightly stronger requirement than 2.4) with asymptotic variance $1/4m$ independent of d and g (see Robinson (1995a)) and the robustness of this asymptotic normality result to conditional heteroskedasticity of general form in the Wold innovations of the process under investigation (including long memory GARCH and other specifications introduced in Robinson (1991b)) . This robustness property, proved in Henry and Robinson (1997), does not include the case when the estimate is applied to the squares of a conditionally heteroskedastic martingale difference process but it is conjectured by the authors that it may continue to hold possibly with a different asymptotic variance.

2.2 Fully parametric specification

The semiparametric techniques advocated above rely only on the specification of the spectral density on a degenerate band of frequencies. They are therefore based on the concentration of the variance of the process in a neighbourhood of frequency zero and are insensitive to any short memory behaviour of the series. Such short range dependent behaviour in the series, if misspecified in a fully parametric model, will bias the estimation of the long range dependent parameter itself. Thus, the robust method advocated above could serve as a pre-estimation technique and enable us to create a fractionally differenced series $\Delta^{\hat{d}}y_t$ which, as noted in Henry and Robinson (1997) is an asymptotically valid approximation to an I(0) series without any parametric assumption on the autocorrelations of the underlying I(0) process $\Delta^d y_t$. On this differenced series, traditional model selection methods (using the AIC for instance) may be carried out to identify the order of a covariance stationary ARMA model for instance.

Our purpose, however, in using the semiparametric estimates is to yield robust pre-estimates of d which can be compared with those obtained from a fully parametric model which permits both long and short range dependence. If then there are no signs of systematic bias in the estimates obtained from the fully parametric model, we can proceed to compare the contributions made to overall temporal dependence by each of the long and short memory components.

The parametric model we adopt is an extension of the LMSV model of Harvey (1993):

$$\begin{cases} y_t = c + h_t + \xi_t \\ \\ (1-L)^d(1-\phi L)h_t = \eta_t \end{cases} \tag{2.5}$$

where c is a constant, ξ has mean zero and variance $\pi^2/2$, $\eta \sim N(0, \sigma_\eta^2)$ and d lies within the stationarity and invertibility range $(-1/2, 1/2)$.[6] This framework is consistent with the semiparametric specification described above insofar as the spectral density of a process thus specified follows 2.2.

The estimation procedure is a frequency domain log likelihood maximization[7]. Exact maximum likelihood is proven (Fox and Taqqu (1986)) to be asymptotically most efficient in the case of a correctly specified normal error distribution η_t. The estimates of the fractional differencing parameter $\hat{d}$ and of the autoregressive parameter $\hat{\phi}$ maximize

$$\mathcal{L}(d,\phi) = -\sum_{j=1}^{T-1} \log(g_j) - \sum_{j=1}^{T-1} \frac{I(\lambda_j)}{g_j} \tag{2.6}$$

where

$$g_j = \frac{\sigma_\eta^2(1-\phi^2)|4\sin^2(\frac{\lambda_j}{2})|^{-d}}{2\pi(1-2\phi\cos\lambda_j+\phi^2)} + \sigma_\xi^2.$$

The first order autoregressive short range dependent specification in volatility is chosen for comparison with the traditional stochastic variance specification. Both specifications naturally suffer from the ignored nonlinearity in the ξ_t which is likely to affect h_t. But this new framework allows us to discriminate long memory and strong autoregressive effects with a simple quasi-likelihood ratio test. The AR(1)-LMSV

[6] This specification for the log squared returns is derived from the formulation in footnote 3, hence $\xi_t = \log(\epsilon_t^2)$ is distributed as a $\log\chi^2$ variate.

[7] Note that in this case, the maximization is performed over the whole range of harmonic frequencies. In the semiparametric case, only a degenerate band of harmonic frequencies was used. The present estimate is therefore sensitive to any short range dependent misspecification. It is likely to be sensitive to the seasonal component in the series discussed in Section 2.4, but it is nonetheless reported before as well as after deseasonalization for completeness.

model (equation(2.5)) is compared to two nested alternatives

$$(LM) \begin{cases} y_t = c + h_t + \xi_t \\ \\ (1-L)^d h_t = \eta_t \end{cases} \quad \text{and} \quad (AR) \begin{cases} y_t = c + h_t + \xi_t \\ \\ (1-\phi L) h_t = \eta_t \end{cases}$$

obtained for $\phi = 0$ and $d = 0$ respectively. The frequency domain likelihood is computed as in 2.6 for each of these nested models and the following tests are performed using the likelihood ratio principle:

$$H_0 : d = 0 \quad against \quad H_a : d > 0.$$

$$H_0 : \phi = 0 \quad against \quad H_a : 0 < |\phi| < 1.$$

If the subscript $._u$ denotes the unconstrained estimates, the likelihood ratio statistic is $2\left[\mathcal{L}(\hat{d}_u, \hat{\phi}_u) - \mathcal{L}(\hat{d}, \hat{\phi})\right]$ where $\mathcal{L}$ is the concentrated form of the quasi-likelihood in 2.6.

2.3 Testing for persistence, long range dependence and stationarity

Bearing in mind that the Gaussian semiparametric estimate is justified for all values of d in a compact subset of $(-\frac{1}{2}, \frac{1}{2})$, a preliminary test of stationary and invertibility of the log squared returns y_t is required. The testing procedure presented here is fully parametric, and sensitive, as indicated above, to mispecified short range dynamics. The conclusions of the test need therefore to be confirmed after the model selection stage. This testing procedure relies on efficients tests of long range dependence (Robinson (1996) and Gil-Alaña and Robinson (1995)) which permit a wide class of null hypotheses. The object is the test of the hypothesis of persistence in foreign exchange volatility. As was mentioned above, most of the available methods for testing for unit roots (see Diebold and Nerlove (1989) for a review on the subject) have non standard limiting distributions and lack Pitman efficiency.[8] Unit root tests against autoregressive alternatives, in particular, are based on the Wald, Likelihood ratio and Lagrange multiplier principles, but they lack the sufficient degree of smoothness across the parameter of interest that would yield null χ^2 limiting distributions and Pitman efficiency. A process following the autoregression $x_t = \rho x_{t-1} + \epsilon_t$ is weakly dependent for $|\rho| < 1$, non-stationary for $\rho = 1$ (the unit root case) and explosive for $|\rho| > 1$.

Moreover, these tests give only one possible persistence null hypothesis. The testing procedure used here, on the other hand, allows one to postulate any value of d

[8]These tests are improved in Elliot, Stock, and Rothenberg (1994)

(integer or fraction) as a null hypothesis and possesses efficiency and a null χ^2 limiting distribution. In the fully parametric LMSV model, the volatility satisfies $(1-L)^d(1-\phi L)h_t = \eta_t$ which can be rewritten as $(1-L)^d h_t = u_t$ where u_t is a stationary AR(1), therefore I(0), process. u_t has spectral density

$$\frac{\sigma_\eta^2}{2\pi} f_u(\lambda;\phi) = \frac{\sigma_\eta^2}{2\pi}\left[\frac{1-\phi^2}{1-2\phi\cos\lambda+\phi^2}\right].$$

Suppose we want to test the hypothesis $\mathrm{H}_0 : d = d_0$. Let $I_u(\lambda)$ be the periodogram of the residuals $\tilde{u}_t = (1-L)^{d_0} h_t$. The frequency domain quasi-likelihood is

$$\mathcal{L}(\sigma_\eta^2,\phi) = -\sum_{j=1}^{T-1}\log\left(f_{uj}\right) - \sum_{j=1}^{T-1}\frac{I_u(\lambda_j)}{\sigma_\eta^2 f_{uj}} \tag{2.7}$$

where $f_{uj} = f_u(\lambda_j;\phi)$. Concentration of this likelihood yields $\sqrt{n}$-consistent estimates $\hat{\phi} = arg\min_\phi \sigma_\eta^2(\phi)$ and $\hat{\sigma_\eta^2} = \sigma_\eta^2(\hat{\phi})$ where $\sigma_\eta^2(\phi) = \frac{2\pi}{T}\sum_{j=1}^{T-1}\frac{I_u(\lambda_j)}{f_{uj}}$.

The test statistic is constructed on the score principle. Let ω be the $(T-1)\times 1$ vector with j-th element $\log|4\sin^2(\frac{\lambda_j}{2})|$, let $\hat{f}_u$ be the $(T-1)\times 1$ vector with j-th element $f_u(\lambda_j;\hat{\phi})$, and let M be the projector on the space orthogonal to the $(T-1)\times 1$ vector with j-th element $\frac{\partial}{\partial\phi}\log f_u(\lambda_j;\hat{\phi})$. The test statistic is

$$\hat{\mathcal{S}} = -\frac{\pi}{\hat{\sigma_\eta^2}}\,\frac{\omega'\hat{f}_u}{\|M\omega\|}.$$

Under suitable regularity conditions (Robinson (1996)), $\hat{\mathcal{S}} \xrightarrow{\mathcal{D}} N(0,1)$ as $T\to\infty$. The resulting testing rules for H_0 are summarized in the table below:

Alternative Hypothesis	Reject H_0 when
$\mathrm{H}_1 : d > d_0$	$\hat{\mathcal{S}} > z_\alpha$
$\mathrm{H}_1 : d < d_0$	$\hat{\mathcal{S}} < -z_\alpha$
$\mathrm{H}_1 : d \neq d_0$	$\hat{\mathcal{S}} > z_{\alpha/2}$

Note: Rules for α-level tests of $H_0 : d = d_0$ against various alternatives. z_α is the quantile of a standard normal variate.

This testing procedure provides us with two efficient tests of persistence: the null $H_0 : d = 1$ against the alternative $H_1 : d < 1$, which is the unit root test, and the null $H_0 : d = 1/2$ against the alternative $H_1 : d < 1/2$, which is a non-stationarity test.[9]

[9] I(d) processes can be seen as increasingly nonstationary as d increases from 1/2 to 1.

2.4 Deseasonalization

A salient feature of the intra-day volatility data considered in this work is the seasonality in volatility described in more detail in Andersen and Bollerslev (1994) and Payne (1996). Seasonal components appear in the periodogram as peaks at certain harmonic frequencies. These peaks affect all periodogram based estimation and the parametric estimation becomes invalid.

If one thinks of the spectrum of the process with strong seasonal components as a mixed spectrum, there is a need for spectral estimation methods which remove the Dirac mass points at the seasonal frequencies and smooth out the leakage from these peaks into the neighbouring frequencies.[10] The spectral estimate used here is a Double-Window smoother proposed by Priestley (1981) which is designed to remove seasonal components and the leakage around the seasonal frequency. Suppose the volatility series is decomposed into two uncorrelated components $y_t = z_t + \varsigma_t$ where z_t has a continuous spectral density and $\varsigma_t = \sum_{r=1}^{K} A_r \cos(\omega_r t + \phi_r)$. The examination of this mixed spectrum is greatly simplified by the knowledge of the seasonal harmonics ω_r, which correspond to the weekly frequency and multiples of the daily frequency.[11] The amplitudes of the seasonal harmonics can be estimated through a regression of y_t against $(\cos(\omega_r t + \phi_r))_{r=1}^{K}$, and the spectrum of z_t is consistently estimated with a Double-Window smoother. The spectral window adopted is the Bartlett-Priestley window

$$W(\theta; M) = \begin{cases} \frac{3M}{4\pi}\left\{1 - \left(\frac{M\theta}{\pi}\right)^2\right\}, & |\theta| \leq \frac{\pi}{M}, \\ \\ 0, & |\theta| \geq \frac{\pi}{M}, \end{cases}$$

where M is the bandwidth.[12] Call $\hat{f}_M(\omega) = \int_{-\pi}^{\pi} I(\theta) W(\omega - \theta; M) d\theta$ the spectral estimate using $W(\theta; M)$. The Double Window spectral estimate is constructed as

[10] von Sachs (1994) proposes a peak insensitive non-parametric procedure to estimate the continuous part of the spectrum, treating the periodic components as outliers (so that it does not permit the estimation of the discrete component in the spectral density). Kooperberg, Stone, and Truong (1995a) propose a fully integrated estimation procedure for both the continuous and the discrete parts of the spectrum. Only related asymptotic results are proposed (see Kooperberg, Stone, and Truong (1995b)), the method is computationally very expensive and there is no indication that it deals with leakage efficiently.

[11] Hence we can avoid employing tests to detect harmonic components (Whittle, Bartlett, Hannan or Priestley, in Priestley (1981)).

[12] This spectral window is a smoothed version of the Daniell (or rectangular) window and it is chosen for its compact support.

follows:

$$\hat{f}_{DW}(\omega) = \begin{cases} \hat{f}_n(\omega), & |\omega - \omega_r| > \frac{\pi}{n}, \\ \\ (\hat{f}_m(\omega) - c\hat{f}_n(\omega))/(1-c), & |\omega - \omega_r| \le \frac{\pi}{n}, \end{cases}$$

where $n > m$, $c = W(0;m)/W(0;n)$ and the ω_r's are the harmonics of the seasonal components defined above. A cross-validated likelihood maximizing procedure for the determination of both bandwidths (see Hurvich (1985), Beltrão and Bloomfield (1987) and Robinson (1991a) for the asymptotics) proved computationally too expensive and gave poor results. An ad hoc choice of bandwidths $n = \sqrt{T}$ and $m = n/10$ was preferred.

3 The Data

As indicated in the Introduction, the focus of this work is the behaviour of volatility in the intra-day Foreign Exchange (FX) market. We study three sets of FX returns, on the DEM/USD, JPY/USD and JPY/DEM, covering the period from the beginning of October 1992 to the end of September 1993.[13] These return series are filtered transcriptions of the tick-by-tick quotation series which appear on the Reuters FXFX page. Each quote encompasses a timestamp, bid and ask quotation pair, plus identifiers which allow one to determine the inputting bank and its location. In this study we ignore the identification of the inputting institution, using the tick-by-tick data solely to construct a homogenous time-series in calendar time.

The basic horizon over which we calculate returns is 10 minutes.[14] This yields, for each currency, a time-series with 37583 observations. The basic summary statistics of the returns are shown in Table 1.

TABLE 1: Summary Statistics for Exchange Rate Returns.

Rate	Mean	s.d.	Skew	Kurtosis	ρ_1	ρ_2	ρ_3	$Q(10)$
DEM/USD	6×10^{-6}	0.001	0.16	9.61	-0.076	-0.040	-0.005	306.8
JPY/USD	-4×10^{-4}	0.074	-0.06	13.85	-0.09	-0.015	0.0035	334.8
JPY/DEM	-5×10^{-4}	0.045	-0.25	7.93	0.0066	-0.0004	0.0009	13.5

Notes: the coefficients ρ_1, ρ_2 and ρ_3 represent the first through third sample autocorrelations respectively. The $Q(10)$ statistic is the Box-Ljung test statistic for up to tenth order serial correlation. The Box-Ljung statistic is distributed χ^2_{10} and has critical value 23.2 at 1%.

The above table illustrates the following facts. First, all three return series have a mean which is insignificantly different from zero. A point which conforms with many earlier studies is that there is pronounced excess kurtosis in the returns distribution. This, as pointed out by Bollerslev and Domowitz (1993), is a natural feature of time-series which display conditional heteroskedasticity, although their analysis shows that after correcting for the conditional heteroskedasticity much of the kurtosis remains.

[13]These data were supplied by Olsen and Associates (Zurich), to whom both authors are most grateful.

[14]Returns are determined as follows: at each 10-minute observation point the last mid-quote entered into the system is taken as the market price. We then first difference this quote series to obtain returns. At points when no quote is entered in a 10 minute interval, an artificial quote is calculated by linear interpolation between the nearest preceding and succeeding quotes. Finally, all weekend quotes are eliminated from the analysis due to the lack of FX market activity at these times. We define weekends as 21:00 GMT Friday to 21:00 GMT Sunday. Note also that the results presented in this paper carry over to the analysis of percentage returns.

Finally, the autocorrelation coefficients show that there is some temporal dependence in the return series, the DEM/USD and JPY/USD demonstrating negative autocorrelation whilst the JPY/DEM displays positive first-order autocorrelation. The significance of these autocorrelation coefficients is confirmed in the Box-Ljung statistics, which demonstrate that one cannot reject the hypothesis of up to tenth order serial correlation.

TABLE 2: Summary Statistics for the Logarithm of Squared Returns.

Rate	Mean	s.d.	Skew	Kurtosis	ρ_1	ρ_2	ρ_3	$Q(10)$
DEM/USD	-14.67	1.28	0.828	0.072	0.281	0.244	0.220	15948.03
JPY/USD	-7.68	3.36	-1.313	1.24	0.24	0.183	0.166	8433.88
JPY/DEM	-8.66	3.16	-0.99	0.45	0.353	0.29	0.264	20463.21

Notes: the coefficients ρ_1, ρ_2 and ρ_3 represent the first through third sample autocorrelations respectively. The $Q(10)$ statistic is the Box-Ljung test statistic for up to tenth order serial correlation. The Box-Ljung statistic is distributed χ^2_{10} and has critical value 23.2 at 1%.

In Table 2 we present identical sets of statistics for our volatility proxy. We employ the logarithm of squared returns as our volatility measure, a choice which is motivated by the Long Memory in Stochastic Volatility model which was presented in Section 2. The main feature of these results lies in the correlation structure of volatility. As is visible from comparing Tables 1 and 2, there is far larger dependence in volatility than in returns. The first-order autocorrelations are between 3 and 5 times greater for volatility than for returns, whilst the Box-Ljung statistics are, at least an order of magnitude greater. The characterisation of this temporal dependence is the focus of this work.

In order to clarify the nature of the dependencies in volatility in Figures 1 to 3 we present the first 1000 periodogram and logged periodogram ordinates for the YEN/USD volatility plus the first 1000 sample autocorrelations.[15]

Examining first the correlogram, one feature which is immediately apparent is the existence of a pronounced daily seasonal in volatility. This seasonal has recently been the subject of many papers, including Dacorogna, Müller, Nagler, Olsen, and Pictet (1993), Andersen and Bollerslev (1994) and Payne (1996). It is generated by the 24 hour activity in the foreign exchange market and the alterations in market activity which occur as trading shifts from the Far East to Europe to North America and so on. There is also evidence of seasonality at the weekly frequency. In the

[15]Throughout the work we present graphical examples for this currency only as those for the other two currencies are qualitatively similar.

current context, however, this component is of no intrinsic interest and simply masks the underlying temporal structure of volatility. Hence, when estimating our long memory specifications we filter this component.

In the periodogram of the data this seasonal component is represented by peaks at integer multiples of the fundamental seasonal frequency.[16] A feature of the periodogram which is more relevant to the current study is the behaviour of the periodogram in a neighbourhood of zero frequency, where the peak (visible on Figures 1 and 2) can be viewed as tentative evidence for the presence of long memory in our volatility series.

[16]As there are 144 ten minute intervals in one day, the seasonal frequency is $\frac{2\pi}{144}$, corresponding, approximately, to harmonic 228.

4 Results

Our results are presented in two stages. In the first sub-section we document the results of pre-testing for long range dependence, using the procedure developed in Robinson (1996). We then go on to present the estimations of the fractional differencing parameter, for both returns and volatility, using the two semiparametric procedures described in Section 2 and the fully parametric AR(1)-LMSV model. Finally, a set of specification tests of the AR(1)-LMSV model is presented.

4.1 Testing for Long Range Dependence

In Table 3 we present the test results for raw exchange rate returns for the DEM/USD, JPY/USD and JPY/DEM. As indicated in Section 2 the test statistic proposed by Robinson (1996) ($\hat{\mathcal{S}}$) has a limiting standard normal distribution under the specified null hypothesis, implying a one-sided rejection region of 2.32 at 1%. Our null hypotheses are formulated as follows. First, standard efficient markets theory indicates that asset prices should follow a random walk, implying that returns should be I(0). This defines one hypothesis as $H_0 : d = 0$. Second, we employ the theoretical bounds for stationarity and invertibility of the fractionally integrated representation for returns as hypotheses, yielding $H_0 : d = -\frac{1}{2}$ and $H_0 : d = \frac{1}{2}$.

TABLE 3: Tests of the significance of the Fractional Differencing Parameter for Raw Exchange Rate Returns

Rate	d=-0.5	d=0	d=0.5
DEM/USD	113.48	-5.02	-18.51
JPY/USD	112.55	-4.68	-19.71
JPY/DEM	122.09	0.46	-18.02

Note: Testing of the value of the fractional differencing parameter is carried out via the test procedure developed in Robinson (1996).

Results yield the following observations. For all three currencies one can strongly reject the hypothesis that $d = 0.5$ in favour of $d < 0.5$ implying that returns are covariance-stationary. Similarly one can conclude that the return processes are invertible given the sign and magnitude of the test statistics corresponding to the hypothesis that $d = -0.5$. A more interesting result appears when examining the second column of Table 3. Whereas for the JPY/DEM the test statistic indicates that one cannot reject the hypothesis of FX *quotations* following an I(1) process, for the DEM/USD and JPY/USD there is evidence that the degree of fractional integration in returns is negative. This then implies that the degree of integration for the *quotation* series of these two currencies is between one half and unity, such that these rates are non-stationary but not I(1).

TABLE 4: Tests of the significance of the Fractional Differencing Parameter for Exchange Rate Volatility ($\log(r^2)$).

Rate	d=0	d=0.5	d=1
DEM/USD	33.23	-13.15	-20.53
JPY/USD	21.12	-14.05	-20.70
JPY/DEM	31.06	-12.40	-20.38

Note: Testing of the value of the fractional differencing parameter is carried out via the test procedure developed in Robinson (1996).

Tables 4 and 5 present the results from the same testing framework on raw and deseasonalised FX volatility (computed, as indicated in Sections 2 and 3 as $\log(r_t^2)$.) As an indication of the efficacy of our deseasonalisation procedure, in Figure 4 we present the periodogram of our deseasonalised volatility. Comparison with Figure 1 demonstrates that the amplitudes at the seasonal frequencies are greatly reduced, although not completely eliminated.[17]

TABLE 5: Tests of the significance of the Fractional Differencing Parameter for Deseasonalised Exchange Rate Volatility.

Rate	d=0	d=0.5	d=1
DEM/USD	22.38	-14.26	-20.62
JPY/USD	18.54	-14.63	-20.28
JPY/DEM	24.26	-13.65	-20.55

Note: Testing of the value of the fractional differencing parameter is carried out via the test procedure developed in Robinson (1996).

The hypotheses of interest in our examination of volatility are as follows. First, given the many previous studies which have demonstrated that intra-day FX volatility has an IGARCH or almost-integrated SV representation, is there a random walk in

[17]As an alternative to our deseasonalisation procedure, we also computed all estimations for the volatility of a time-scale transformed series of midquotes. We used the theta-time scale proposed by Dacorogna, Müller, Nagler, Olsen, and Pictet (1993). Results from these estimations were very similar to those for deseasonalised volatility and are available upon request from the authors.

volatility? Second, can one characterise volatility as being long range dependent? The former dictates examination of $H_0 : d = 1$ whilst the hypotheses pertinent to the latter are $H_0 : d = 0$ and $H_0 : d = \frac{1}{2}$.

The final column of Tables 4 and 5 presents the evidence pertinent to the hypothesis of a random walk in volatility. There is strong evidence that, for all currencies, this hypothesis can be strongly refuted in favour of a degree of integration in volatility of less than unity. Further, this conclusion is stable across both raw and deseasonalised volatility. There is still, however, the possibility of non-stationarity in volatility if $d \in [0.5, 1]$. Column 2 of the tables indicates that the non-stationarity hypothesis can be rejected also, with the test statistics indicating that $d < 0.5$ for all three currencies. Finally, evidence on the long range dependence of volatility is shown in column 1. From both tables one can draw the conclusion that the true value of d lies between zero and one half, evidence of long memory in the volatility of all currencies.

Hence, the testing procedure indicates the following. Returns can be characterised as short range dependent, covariance stationary processes, with some indication of negative degrees of fractional integration for the DEM/USD and JPY/USD. The volatility results indicate covariance-stationarity also, with the non-stationarity and I(1) hypotheses convincingly refuted, although there is consistent evidence of long range dependence. The volatility processes are therefore completely mean reverting, a result which is more comfortable than that of I(1) volatility from a theoretical point of view.

4.2 Semiparametric Estimations

We first present point estimates for the fractional differencing parameter for the three returns series. The semiparametric procedures of Robinson (1995a) and Geweke and Porter-Hudak (1983) are employed in estimation. Table 6 gives the results.

Examining first the GPH estimates it is quite clear that the negativity of d indicated in the previous subsection is a very minor economic phenomenon, indicating that returns display very small anti-persistent tendencies. For no currency does $\hat{d}$ exceed 0.05 in absolute value. This implies that the *quotation* series may be regarded as following I(1) processes to more-or-less any degree of precision, in line with the efficient markets hypothesis.

This conclusion becomes less clear when one examines the Gaussian estimates. Whilst the results for the JPY/USD and JPY/DEM are very similar to their GPH counterparts, the value of $\hat{d}$ derived for the DEM/USD is now greatly negative. This implies a covariance-stationary and invertible representation for DEM/USD returns which displays non-negligible anti-persistence. Given the confluence between the testing and GPH results, however, we are inclined to treat this feature as an anomaly and describe the return generating process as approximately I(0)

TABLE 6: Estimation of the Fractional Differencing Parameter ($\hat{d}$) for Exchange Rate Returns.

Rate	$\hat{d}$	
	Rob95	GPH92
DEM/USD	-0.38 (0.02)	-0.01 (0.05)
JPY/USD	-0.02 (0.02)	-0.03 (0.05)
JPY/DEM	-0.01 (0.02)	-0.01 (0.05)

Note: Estimation of the long memory models is carried out via the semiparametric procedures of Robinson (1995) (Rob95) and Geweke and Porter-Hudak (1992) (GPH92). Standard errors in parentheses.

The estimation results for the volatility series are presented (for raw and deseasonalised volatility) in Tables 7 and 8. Here we complement the semiparametric procedures used in the analysis of returns with the fully parametric AR(1)-LMSV model.

TABLE 7: Estimation of the Fractional Differencing Parameter ($\hat{d}$) for Exchange Rate Volatility.

Rate	$\hat{d}$		
	Rob95	GPH92	LMSV
DEM/USD	0.29 (0.02)	0.21 (0.06)	0.37 (0.01)
JPY/USD	0.27 (0.02)	0.19 (0.06)	0.24 (0.01)
JPY/DEM	0.27 (0.02)	0.28 (0.06)	0.32 (0.01)

Note: Estimation of the long memory models is carried out via the semiparametric procedures of Robinson (1995) (Rob95) and Geweke and Porter-Hudak (1992) (GPH92), plus the fully parametric Long Memory in Stochastic Volatility model of Harvey (1993) (LMSV). Standard errors in parentheses.

The results of both tables demonstrate that the testing procedures contained in the previous subsection deliver the correct inferences. Across currencies and estimators there is consistent evidence that the value of $\hat{d}$ for the volatility series is between 0.2 and 0.3. This indicates that volatility can be characterised as covariance-stationary, invertible and long range dependent. The only real difference in estimation results for raw and deseasonalised volatility is that the AR(1)-LMSV estimates tend to be slightly greater for the former and greater than the results delivered by the semi-parametric estimators. This is likely to be due to mis-specfication of the short-range dependence in the series i.e. omission of an explicit seasonal in the fully parametric model. As one might expect, the difference in estimated $\hat{d}$ between the semi and fully parametric procedures is far smaller for deseasonalised volatility.

TABLE 8: Estimation of the Fractional Differencing Parameter ($\hat{d}$) for Deseasonalised Exchange Rate Volatility.

	$\hat{d}$		
Rate	Rob95	GPH92	LMSV
DEM/USD	0.29	0.19	0.26
	(0.02)	(0.04)	(0.01)
JPY/USD	0.30	0.18	0.22
	(0.02)	(0.04)	(0.01)
JPY/DEM	0.30	0.25	0.26
	(0.02)	(0.04)	(0.01)

Note: Estimation of the long memory models is carried out via the semiparametric procedures of Robinson (1995) (Rob95) and Geweke and Porter-Hudak (1992) (GPH92), plus the fully parametric Long Memory in Stochastic Volatility model of Harvey (1993) (LMSV). Deseasonalisation is carried out via a frequency domain Double-Window Smoother (DWin.). Standard errors in parentheses.

4.3 Specification Tests on the Fully Parametric Model

The final step in our empirical methodology involves a series of estimations and specification tests on the fully parametric model outlined in Section 2.2. These tests allow us to examine the relative contributions of short memory and long memory components to the temporal dependence in the volatility process. As indicated in Section 2.2 our fully parametric model nests a pure LMSV model (obtained by setting $\phi = 0$ in equation 2.5) and a standard AR(1)-SV model (obtained by restricting $d = 0$ in 2.5.) By estimating the unrestricted model and these two restricted alternatives we can employ a frequency-domain Likelihood Ratio test to gauge the significance of the long memory and AR(1) components. Results of these estimations and the associated tests, for deseasonalised volatility only, are given in Table 9.

Examining first the results which correspond to the simple AR(1)-SV models one can note the appearance of the common result that the 'underlying' volatility process has an autoregressive parameter very close to unity in all cases. This conforms with the results of many earlier studies. Moving on to the pure LMSV models it is quite clear that the conclusions of the previous estimations still hold and further that the long memory specification gives a far better fit than does the autoregression (as evinced by the lower minimised log likelihood.)

TABLE 9: Extensions of various LMSV specifications and Likelihood Ratio Testing using Deseasonalised Volatility

Rate	DEM/USD			JPY/USD			JPY/DEM		
	LM-AR	LM	AR	LM-AR	LM	AR	LM-AR	LM	AR
d	0.26 (0.01)	0.25 (0.01)	-	0.22 (0.01)	0.24 (0.01)	-	0.26 (0.01)	0.24 (0.01)	-
ϕ	0.12 (0.01)	-	0.96 (0.01)	0.15 (0.01)	-	0.96 (0.01)	0.15 (0.01)	-	0.93 (0.01)
LogL	3097.8	3098.2	3143.2	17228.9	17231.0	17304.8	15645.6	15646.1	15726.4
LR	-	0.85	90.90**	-	4.10*	151.75**	-	0.88	161.50**

Note: Estimation of the long memory models is carried out via the the fully parametric Long Memory in Stochastic Volatility model of Harvey (1993) (LMSV). Columns headed LM-AR present results from the model presented in equation 2.5. Columns headed LM are estimated with the restriction that the autoregressive parameter is zero and columns headed AR are estimates from the model where the fractional differencing parameter is set to zero. The final row of the table gives [illegible] significant at 5%, ** denotes significance at 1%. Deseasonalisation is carried out via a frequency domain Double-Window Smoother . Standard errors in parentheses.

A comparison with the unrestricted model yields the following observations. First, the improvement in fit of the combined model over the pure LMSV model is marginal, as the comparison of Log Likelihoods displays. Second, and more importantly, the value of the autoregressive parameter is far lower in the unrestricted model, in all cases between 0.1 and 0.2. Given the pure AR(1)-SV results, this may be taken as evidence that the previous findings of near unit-root behaviour in intra-day FX volatility are caused by model mis-specification through the omission of long-range dependent components. This result confirms the evidence from daily data found in Baillie, Bollerslev, and Mikkelsen (1993). These conclusions are reinforced by the quasi-LR test statstics. In the unrestricted model one can convincingly reject the hypothesis that the degree of fractional integration in volatility is zero for all currencies, whereas the hypothesis that the autoregressive parameter is zero cannot be rejected in two of the three cases and is only marginally rejected in the third.

A clear comparison of the AR-SV and LM-SV models, in terms of how well they fit the data, is shown in Figure 5. The figure graphs the actual correlogram of the volatility process for the DEM/USD alongside those implied by the estimated AR-SV and LM-SV models. It is immediately apparent that the long memory specification gives a far better approximation of the true volatility dynamics than does the autoregressive model. The autoregressive model greatly overstates the low order autocorrelations but dies out too quickly to mimic the persistently positive high-order autocorrelations in the DEM/USD data. Intuitively, the estimated autoregressive parameter is driven very close to unity in order to try to approximate the long memory in volatility, but this only results in the low order autocorrelations being far too high whilst the exponential decay in the correlogram ensures that, even if the estimated autoregressive parameter is arbitrarily close to unity, the autoregressive specification cannot match the persistence in volatility exhibited by the data.

Hence, to restate the main findings; returns can be characterised as I(0) processes to a great degree of precision; volatility, on the other hand, is best represented by a covariance stationary, long range dependent process; furthermore, there is evidence that the common finding of near-integrated volatility processes is driven by the misspecfication of traditional models which do not permit long range dependence.

5 Conclusions

We have presented in this work a study of the long range dependent properties of intra-day exchange rate returns for three major currencies. Evidence on the returns series tends to imply that they are I(0) processes, a result which supports the I(1) hypothesis for exchange rates. We however do not reject strongly the hypothesis of antipersistence (I(d) characteristics with $d < 0$) in the returns, which would imply mean reversion (however slow it may be) in the exchange rate series. In contrast, testing and estimation results for volatility are consistent across currencies and give a clear indication that the volatility of exchange rates series follows a covariance stationary long range dependent process. Further results suggest that the commonly found near unit-root behaviour of intra-day FX volatility is driven by the mis-specification of traditional models which do not allow for long range dependence. The existence of long range dependence in intra-day volatility is in line with similar findings for daily exchange rate volatility found in Baillie, Bollerslev, and Mikkelsen (1993) and Harvey (1993). A point to note, however, is that whereas our results indicate a covariance stationary volatility process ($0 < \hat{d} < 0.5$), those of both Baillie, Bollerslev, and Mikkelsen (1993) and Harvey (1993) indicate that $\hat{d}$ exceeds 0.5 and, therefore, volatility is non-stationary.

As far as methodology is concerned, the present paper has emphasized the performance of several individual econometric methods in the treatment of a stationary long memory component in a time series. It needs to be stressed that a practitioner should also rely on the strength of the sequential methodology. The pretesting procedure enabled us to identify the range in which the dependence parameter lies, in this case $0 < d < 1/2$, and therefore to identify the basic characteristics of the process within the chosen framework, namely its long range dependence, covariance stationarity and mean reversion. Robust estimation of d is then carried out to identify the main feature of interest in the data, namely the degree of temporal dependence. One may then proceed to a prefiltering of the volatility series y_t, yielding the differenced series $\Delta^{\hat{d}} y_t$ where $\hat{d}$ is the pre-estimate of the degree of persistence, prior to identifying the short memory structure (as we have consistently identified seasonal components) which leads to the selection of a fully parametric model on the basis of which one may then test relevant hypotheses and perform a full estimation possibly including posited volatility determinants.

For efficient forecasting purposes it would be desirable to perform this sequential procedure in a multivariate framework, preferably with a parsimonious factor decomposition for volatility . This entails a fractional cointegration analysis of our vector of volatilities and is the focus of further study (in Henry and Payne (1997)).

References

ANDERSEN, T., and T. BOLLERSLEV (1994): "Intraday seaonality and volatility persistence in financial markets," Unpublished Paper.

BAILLIE, R., T. BOLLERSLEV, and H. MIKKELSEN (1993): "Fractionally Integrated Generalized Autoregressive Conditional Heteroskedasticity," mimeo, Northwestern University.

BELTRÃO, K., and P. BLOOMFIELD (1987): "Determining the bandwidth of a kernel spectral estimate," *Journal of Time Series Analysis*, 8, 21–38.

BOLLERSLEV, T., R. CHOU, and K. KRONER (1992): "ARCH modeling in finance," *Journal of Econometrics*, 52, 5–59.

BOLLERSLEV, T., and I. DOMOWITZ (1993): "Trading patterns and prices in the interbank foreign exchange market," *Journal of Finance*, 48, 1421–1444.

COMTE, F., and C. HARDOUIN (1995): "Refression on log-regularized periodogram for fractional models at low frequencies," Discussion Paper, CREST.

DACOROGNA, M., U. MÜLLER, R. NAGLER, R. OLSEN, and O. PICTET (1993): "A geographical model for the daily and weekly seasonal volatility in the foreign exchange market," *Journal of International Money and Finance*, 12, 413–438.

DEGENNARO, R., and R. SHRIEVES (1995): "Public Information Releases, Private Information Arrival, and Volatility in the Foreign Exchange Market," Preprint.

DIEBOLD, F., and M. NERLOVE (1989): "Unit roots in economic time series: a selective survey," *Advances in Econometrics*, 8, 3–69.

ELLIOT, G., J. STOCK, and T. ROTHENBERG (1994): "Efficient tests of an autoregressive unit root," preprint.

FOX, R., and M. TAQQU (1986): "Large sample properties of parameter estimates for strongly dependent stationary Gaussian time series," *Annals of Statistics*, 14, 517–132.

GEWEKE, J., and S. PORTER-HUDAK (1983): "The estimation and application of long memory time series models," *Journal of Time Series Analysis*, 4, 221–238.

GHYSELS, E., A. HARVEY, and E. RENAULT (1996): "Stochastic Volatility," in *Handbook of Statistics*, vol. 14.

GIL-ALAÑA, L., and P. ROBINSON (1995): "Testing of unit root and other nonstationary hypotheses in macroeconomic time series," mimeo, LSE.

GRANGER, C. (1980): "Long memory relationships and the aggregation of dynamic models," *Journal of Econometrics*, 14, 227–238.

GRANGER, C., and R. JOYEUX (1980): "An introduction to long memory time series models," *Journal of Time Series Analysis*, 1, 15–29.

GUILLAUME, D. (1995): "On the intra-day performance of GARCH processes," mimeo., Financial Markets Group, London School of Economics.

HARVEY, A. (1993): "Long Memory in Stochastic volatility," mimeo, London School of Economics.

HARVEY, A., E. RUIZ, and N. SHEPHARD (1994): "Multivariate stochastic variance models," *Review of economic studies*, 61, 247–264.

HENRY, M., and R. PAYNE (1997): "Fractional cointegration and common factors in Foreign Exchange rates volatility.," mimeo, LSE.

HENRY, M., and P. ROBINSON (1996): "Bandwidth choice in Gaussian semiparametric estimation of long range dependence," in *Athens Conference on Applied Probability and Time Series Analysis*, ed. by P.M.Robinson, AND M.Rosenblatt, vol. 2. Springer Verlag, In memory of E.J. Hannan.

HENRY, M., and P. ROBINSON (1997): "Long memory conditional heteroskedasticity in semiparametric estimation of long memory," mimeo, LSE.

HURVICH, C. (1985): "Data-driven choice of a spectrum estimate: extending the applicability of cross-validation methods," *Journal of the American Statistical Association*, 80, 933–940.

KOOPERBERG, C., C. STONE, and Y. TRUONG (1995a): "Logspline estimation of a possibly mixed spectral distribution," *Journal of Time Series Analysis*, 16, 359–388.

KOOPERBERG, C., C. STONE, and Y. TRUONG (1995b): "Rate of convergence for Logspline spectral density estimation," *Journal of Time Series Analysis*, 16, 389–401.

KÜNSCH, H. (1987): "Statistical aspects of self-similar processes," in *Proceedings of the First World Congress of the Bernouilli Society*, pp. 67–74. VNU Science Press.

PAYNE, R. (1996): "Announcement effects and seasonality in the intra-day foreign exchange market," Discussion Paper 238, Financial Markets Group, London School of Economics.

PRIESTLEY, M. (1981): *Spectral analysis and time series*, Probability and Mathematical Statistics. Academic Press.

Robinson, P. (1991a): "Automatic frequency domain inference on semiparametric and nonparametric models," *Econometrica*, 89.

Robinson, P. (1991b): "Testing for strong serial correlation and dynamic conditional heteroskedasticity in multiple regression," *Journal of Econometrics*, 47.

Robinson, P. (1994): "Time series with strong dependence," in *Advances in Econometrics*, ed. by C. Sims, vol. 1. Cambridge University Press.

Robinson, P. (1995a): "Gaussian semiparametric estimation of long range dependence," *Annals of Statistics*, forthcoming.

Robinson, P. (1995b): "Log periodogram regression of time series with long range dependence," *Annals of Statistics*, 23, 1048–1072.

Robinson, P. (1996): "Efficient tests of nonstationary hypotheses," *Journal of the American Statistical Association*, forthcoming.

Tauchen, G., and M. Pitts (1983): "The price variability-volume relationship on speculative markets," *Econometrica*, 51, 485–505.

Velasco, C. (1995): "Consistent log periodogram regression," mimeo, LSE.

von Sachs, R. (1994): "Peak-insensitive non-parametric spectrum estimation," *Journal of Time Series Analysis*, 15, 429–452.