A Reexamination of Fractional Integrating Dynamics in Foreign Currency Markets

Hyun J. Jin
Research Assistant Professor
Department of Ag-Business & Applied Economics
North Dakota State University
Fargo, ND 58105-5636
Phone: (701) 231-8993
Fax: (701) 231-7400
Email: hjin@ndsuext.nodak.edu

John Elder
Assistant Professor
Department of Management, Marketing and Finance
North Dakota State University
Fargo, ND 58105-5137

Won W. Koo,
Professor and Director
Center for Agricultural Policy and Trade Studies
Department of Ag-Business & Applied Economics
North Dakota State University
Fargo, ND 58105-5636

Selected Paper prepared for presentation at the American Agricultural Economics Association Annual Meeting, Denver, Colorado, August 1-4, 2004

A Reexamination of Fractional Integrating Dynamics in Foreign Currency Markets

Abstract: This paper reexamines foreign currency markets for evidence of fractional integration, and extends the extant literature in several important dimensions. First, we utilize a new semiparametric wavelet-based estimator, which is far superior to the more prevalent GPH estimator on the basis of mean squared error. Second, we utilize a broader and longer sample, which better facilitates the detection of long memory dynamics. Our analysis yields interesting empirical results that contrast with other recent studies. In particular, we find new evidence that a large proportion (fourteen out of nineteen) of exchange rate series display evidence of long memory, with little variation over alternative sample periods.

Key Words: exchange rates, weak-form market efficiency, long memory, fractional integration; wavelets.

1. Introduction

Exchange rates and their volatility are important determinants of international capital flows, relative prices, foreign direct investment, trade in goods and services, and macroeconomic performance (particularly of small open economies). Although some exchange risk can be hedged in derivatives markets, longer-term fluctuations may be more difficult, or expensive, to hedge, distorting relative prices and the optimal allocation of resources. As such, researchers have long been interested in the univariate properties of exchange rates, with particular recent interest in whether shocks to the return process tend to decay rapidly, as with non-integrated processes, or whether they tend to decay much more slowly, as with fractionally integrated processes. In the latter case, exchange rates are said to display evidence of "long memory". Long memory is a particularly interesting feature because its presence indicates evidence of nonlinear dependence in the first and second moments, and hence, evidence of a predictable component. Predictability would cast doubt on the weak-form efficiency of foreign exchange markets, since long memory in exchange rate returns would imply that it may be possible to exploit the dependency in order to consistently earn speculative profits.

Estimates of the degree of integration are also relevant in constructing more robust models of long-term currency dynamics (c.f., Breidt, Crato, and de Lima (1998)). For example, fractional integration in exchange rates suggests that any possible cointegrating relationships may also be fractionally integrated, implying that equilibrium adjustments are long memory processes as well (c.f., Baillie and Bollerslev 1989 and 1994). Similarly, the long memory implied by such processes suggests that the speed of convergence to purchasing power parity may be quite slow, consistent with findings of, for example, Papell (2002).

Initially, research indicated considerable empirical evidence in support of long memory dynamics in exchange rates, while more recent research has suggested that the empirical evidence is inconclusive, with results sensitive to the particular sample examined. For example, Booth, Kaen, and Koveos (1982), applying classical rescaled-range (R/S) analysis to daily exchange rates, find evidence of negative dependence during the fixed exchange rate period 1965-1971, and positive long-term dependence during the post-1974 flexible exchange rate period in three major exchange rates. Cheung (1993), employing the popular Geweke and Porter-Hudak (GPH, 1983) estimator on five major exchange rates from 1974-1987, also finds evidence of long memory dynamics. Pan, Liu, and Bastin (1996) provide additional evidence of long memory dynamics in this same sample, utilizing modified R/S and variance-ratio tests. Baillie and Bollerslev (1989 and 1994) also find evidence of fractional integration in exchange rates. In contrast, Barkoulas, Baum, Caglayan, and Chakraborty (hereafter BBCC, 2003) find little convincing evidence of long memory in eighteen currency return series over the 1974-1995 period. They conclude that exchanges rates are best characterized as a martingale, rather than fractionally integrated.

BBCC differentiate their results from previous studies, such as Cheung's, on two dimensions. First, BBCC utilize the Gaussian semiparametric estimator developed by Robinson (1995a), rather than the more popular GPH estimator utilized by Cheung. The two estimators, however, are conceptually similar in the sense that they both exploit the low-frequency spectral behavior of fractionally integrated processes, and BBCC indicate that they obtain qualitatively similar results when using the GPH estimator. This suggests that the more important difference is the longer span of their data. Indeed, they report that Cheung's results are sensitive to the 1974-1987 sample, and, on basis of longer 1974-1995 span, they find little evidence of long memory dynamics.

The objective of this study is to reexamine exchange rates for evidence of long memory, utilizing a newly developed semi-parametric estimator, developed by Jensen (1999), as well as a longer, broader and more recent sample, which includes nineteen exchange rate series with observations from 1975 through 2002. This new estimator utilizes functional transforms known as wavelets to estimate the fractional integration parameter, and is denoted the wavelet OLS estimator.

The wavelet OLS estimator is a considerable improvement over the popular GPH estimator on the basis of mean squared error (MSE). For example, Jensen shows that, for relevant sample sizes (T = 128 through 1024), the MSE of the wavelet OLS estimator is approximately four to six times smaller than that of the GPH estimator. We also extend the sample period beyond that of previous studies, which, as suggested by BBCC, may better facilitate detection of the low-frequency spectral behavior associated with long memory processes.

Anticipating the principal results, we find that fourteen out of nineteen exchange rate series display evidence of long memory dynamics based on the wavelet OLS estimator. This contrasts sharply with results obtained by the popular GPH estimator, which indicate little or no evidence of long memory in recent samples. Our results are statistically significant, and the evidence of long-memory is relatively robust to alternative sample periods and alternative frequencies. Finally, we find that a surprisingly large fraction of return series display evidence of negative fractional integration (i.e., with the fractional integration parameter between −0.5 and 0), indicating that the associated exchange rate displays evidence of long memory, but that the return series displays evidence of antipersistence.

To reiterate, this study extends the extant literature in several important dimensions. First, we use Jensen's wavelet OLS estimator, which is superior to the more popular GPH estimator on the basis of MSE, to estimate the fractional integration parameter. Second, this study uses a broader and longer sample, covering 19 exchange rate series from 1975 through 2002. The longer time span improves the ability of statistical tests to detect long-memory dynamics. Third, our results contrast sharply with those obtained by recent studies based on the GPH estimator.

The remainder of the paper is organized as follows. The second section provides background on fractional integration and spectral-based estimators, such as the GPH and GSE estimators, and presents the relatively new wavelet OLS estimator. The third section details the data and presents the results. The final section concludes.

2. Methodology

2.1. Fractional Integration

Fractional integration is the main conceptual framework for describing long memory processes in financial time series. Fractional integration is a generalization of integer integration, under which time series are usually presumed to be integrated of order zero or one. For example, an autoregressive moving-average process integrated of order d (denoted ARMA(p,d,q)) can be represented as

(1) $$(1-L)^d\Phi(\mathrm{L})x_t = \varphi(L)\,u_t\,,$$

where u_t is an i.i.d. random variable with zero mean and variance σ^2_u; L denotes the lag operator; and $\Phi(\mathrm{L})$ and $\varphi(L)$ denote finite polynomials in the lag operator with roots outside the unit circle. For $d = 0$, the process is stationary, and the effect of a shock to u_t on x_{t+j} decays geometrically as j increases. For $d = 1$, the process is said to have a unit root, and the effect of a shock to u_t on x_{t+j} persists into the infinite future.

In contrast, fractional integration defines the function $(1-L)^d$ for non-integer values of d. Following Hosking (1981) and Granger and Joyeaux (1980), a more general definition of $(1-L)^d$ can be derived from a power series expansion as follows

(2) $$(1-L)^d = \sum_{k=0}^{\infty}\binom{d}{k}(-L)^k = 1 - dL - \tfrac{1}{2}d(1-d)L^2 - \tfrac{1}{6}d(1-d)(2-d)L^3 - \ldots\,.$$

This power expansion can be re-expressed in terms of the gamma function as

(3) $$(1-L)^d = \sum_{k=0}^{\infty}\frac{\Gamma(k-d)L^k}{\Gamma(-d)\Gamma(k+1)}\,.$$

It turns out that for $-0.5 < d < 0.5$ the process x_t is stationary and invertible. For such processes, the effect of a shock u_t on x_{t+j} decays as j increases, but the rate of decay is much slower than for a process integrated of order zero. More precisely, the autocovariance function for zero-integrated processes decays geometrically, while the autocovariance function for a fractionally integrated process decays hyperbolically,[1] with the sign of the autocovariances being the same as the sign of *d*. In this sense, fractional integration captures long memory dynamics more parsimoniously than non-integrated ARMA processes.

Important distinguishing features of fractionally integrated processes can be expressed in the frequency domain, as represented by the power spectrum. Recall that the power spectrum is a plot of the spectral density versus frequency, and can be interpreted as representing the proportion of the variance due to each frequency. For example, consider as a benchmark the power spectrum of a white noise process. A MA(1) process with a $\varphi_1 = 0.75$ is somewhat "smoother" (i.e., it changes more slowly) than white noise. Such a process has slowly decaying, positive autocovariances, and its power spectrum would indicate that a larger fraction of its variance is due to low frequency components. In contrast, an MA(1) process with $\varphi_1 = -0.75$ is "choppier" than white noise. It has slowly decaying, negative autocovariances, and its power spectrum would indicate that a larger fraction of its variance is due to high frequency components.

Similarly, a fractionally integrated process with $0 < d < 0.5$ is "smoother" than white noise. The extent of the long range dependence, however, is such that a *very* large portion of its variance is explained by *very* low frequency components. More precisely, the spectral density at frequency zero is infinite. In contrast, a fractionally integrated process with $-0.5 < d < 0$ has negative autocovariances, with the spectral density at frequency zero equal to zero. For this reason, a fractionally integrated process with $-0.5 < d < 0$ is sometimes referred to antipersistence.

2.2. Estimators of the Fractional Integration Parameter

The most prevalent method for estimating the fractional integration parameter *d* is the GPH method, although other methods and variations have been proposed (e.g., Sowell, 1992; Robinson, 1994; Robinson, 1995; and Breidt, Crato, and De Lima, 1998). The GPH estimator is known as a semi-parametric estimator because it yields an estimate of the fractional integration parameter, *d*, without specification of the short-term dynamics, i.e., the autoregressive and moving-average parameters. This is a considerable advantage over maximum likelihood based methods because a) misspecification of the short-term dynamics does not bias or otherwise impact the estimate of the fractional integration parameter; and b) a closed-form solution for the fractional integration parameter can be obtained, avoiding numerical optimization of a likelihood function with a relatively dense Hessian.

The GPH estimator is based on the low-frequency spectral behavior of a series. As described previously, a fractionally integrated process with $0 < d < 0.5$ has a very large portion of its variance explained by very low frequency components. As such, the periodogram (i.e., the sample spectral density) should indicate an inverse relationship between the level of the periodogram and the frequency at which the level is evaluated. The GPH estimator captures this relationship through a simple OLS regression, in logs, of the level of the periodogram on the frequencies. Since the (log) relationship holds only for low frequencies, including higher frequencies in the regression induces a tradeoff between precision (reducing the standard error of the regression estimate), and bias.

More precisely, the GPH estimator is as follows. Let $I(w_j)$ denote the sample periodogram at the j^{th} Fourier frequency, evaluated at $w_j = 2\pi j/T$, $j = 1, 2, ..., T/2$. The GPH estimator of *d* is then based on the OLS regression of the log periodogram on the log frequency

(4) $$\ln [I(w_j)] = \beta_0 + \beta_1 * \log(w_j) + \varepsilon_j,$$

where $j = 1, 2, ..., J$ and $\hat{d} = -\frac{\hat{\beta}}{2}$.

The wavelet OLS estimator of the fractional integration parameter *d* was introduced into econometrics by Jensen (1999). The intuition behind the wavelet estimator is similar to that of the

GPH estimator, in that wavelet analysis can be used to decompose time series into "low frequency" and "high frequency" components, although the terminology and methodology differ.

In particular, wavelets are functional transforms, which are analogous to the Fourier transforms utilized in spectral analysis. That is, utilizing Fourier transforms, spectral analysis can be used to decompose a series into low frequency and high frequency components by writing the series as a sum of sine and cosine functions of different frequencies and amplitude. A short-coming of spectral analysis is that, because the underlying sine and cosine functions are "smooth", the extracted low frequency component of the time series must also be smooth. For example, spectral analysis may be applied to an observed time series on aggregate output, such as GDP, to extract variations occurring at "business cycle" frequencies. This removes variations due to high frequency, or *noise*, components, as well as very low, or *trend*, components. Invariably, a plot of the extracted "business cycle" frequency component will appear "smooth", representing slow transitions from periods of rising to falling output.

In contrast, the analogous decomposition with wavelets is achieved by utilizing functional transforms, known as wavelets, which do not necessarily impose such smoothness on the underlying series. Thus, wavelets can localize a process in time and scale -- either "zooming in" on the behavior at a particular point in time or "zooming out" to reveal long-run features. Wavelet transforms can be simple step functions, or other more sophisticated functions, designed to have specific properties, such as forming a basis in the space of all square-integrable functions along the [0,1] interval.

As an example of a wavelet decomposition, consider the following. Let $f(x)$ be any real-valued function on the interval [0,1]. The function $f(x)$ can be expressed in the wavelet domain as:

$$f(x) = a_0 + \sum_{i=0}^{\infty} \sum_{j=0}^{2^i - 1} a_{ij} \Psi(2^i x - j) , \tag{5}$$

where $\{x_t\}$ is a covariance-stationary discrete time process, a_0 and a_{ij} are scalars and the wavelet function $\Psi(z)$ is defined as

(6) $\Psi(z) =$ 1 if $0 \le z < 0.5$,

-1 if $0.5 \le z < 1$

0 otherwise

The index i is scales the wavelet $\Psi(z)$, compressing or dilating the wavelet to capture low and high "scale" features of the underlying process. The index j is the transition index, or window, that shifts the wavelet function $\Psi(z)$. The particular wavelet function $\Psi(z)$ defined in (6) is known as the Haar wavelet. It is probably the simplest wavelet, and is often used in heuristic applications. Many other, more sophisticated wavelets, have been proposed. For example, the very popular Daubechies (1988) wavelet is commonly used in applications outside economics, such as signal processing (e.g., denoising digitized "signals" such as a nuclear magnetic resonance signals and digital images). Jensen (1999) and Tkacz (2001) suggest that the Daubechies wavelet has desirable properties for time series analysis, and so we use the Daubechies wavelet in our analysis.

Just as the GPH estimator captures low-frequency spectral behavior, the wavelet estimator captures the low-"scale" wavelet behavior. That is, the wavelet "scales", or resolutions, which contribute the most to a series' variance are associated with the wavelet coefficients with the largest variance. Hence, the wavelet coefficients sample variance can be utilized to provide a parametric estimate of the fractional integration parameter d. More precisely, Jensen (1999) proves that, as the scaling coefficient $i \to 0$, the wavelet coefficients a_{ij} in equation (5), associated with a mean zero I(d) process with $|d| < 0.5$, are distributed N(0, $\sigma^2 2^{-2id}$). This can be rewritten as a log-linear relationship in d by letting $R(i) = \sigma^2 2^{-2id}$ denote the wavelet coefficient's variance at scale i and taking logs

(7) $$\ln [\mathrm{R}(i)] = \ln[\sigma^2] - d \ln[2^{2i}].$$

Hence, an estimate of the fractional integration parameter d may be obtained by performing OLS on equation (7), substituting the sample variance of the wavelet coefficient $\hat{R}(i) = \frac{1}{2^i} \sum_{k=0}^{2^j - 1} a_{i,j}{}^2$

for the population variance $R(i)$. Jensen (1999) shows that the wavelet OLS estimator $\hat{d}$ is a consistent estimator of the fractional integration parameter d. Jensen also shows, via a Monte Carlo analysis, that the mean squared error (MSE) of the wavelet estimator is approximately four to six times smaller than the MSE the GPH estimator at relevant sample sizes (T = 128, 256, 512, 1024), although the bias tends to be larger (but negative), while the bias of the GPH estimator is most often positive.

In practice, the wavelet defined in equation (5) is evaluated over a moving window whose width is a power of two. In order to avoid boundary effects associated with the evaluation of equation (5), the sample size be a power of 2, or, if not, the sample must padded with zeros. As discussed below, we trim our sample so the actual number of observations is 2^8 or 2^7.

3. Data and Results

3.1. Data

Our data set consists of U.S. dollar nominal exchange rates for 19 industrial countries: Australian dollar, Austrian schilling, Belgian franc, British pound, Canadian dollar, Danish krone, Deutsche mark, Finnish markka, French franc, Greek drachma, Italian lire, Japanese yen, Netherlands guilder, New Zealand dollar, Norwegian krone, Portuguese escudo, Spanish peseta, Swedish krona, Swiss franc. Non-European exchange rates are obtained from the *International Financial Statistics* of the International Monetary Fund and exchange rates for countries participating in the European Monetary Union are obtained from the *Financial Statistics* of the Federal Reserve Board.

The frequency of observation is monthly, and the full sample covers the post-Bretton Woods period, beginning in January 1974 through December 2002, yielding 347 total observations.[2] For Australian dollar and Greek drachma, the full sample period begins when the peg to the U.S. was lifted, in November 1974 and April 1975, respectively, yielding 337 and 332 return observations. As discussed previously, in order to facilitate application of the wavelet OLS estimator, the full sample is trimmed so that the effective number of observations is a power of 2. In particular, our effective full sample runs from September 1981 though December 2002, yielding 256 usable observations.

The exchange rate data are expressed as continuously compounded returns, that is $x_t \equiv \log(S_t/S_{t-1})$, where S_t is the nominal exchange rate. This transformation is necessary in order to induce covariance-stationarity, which required for the GPH and Wavelet estimators to yield sensible results. Standard unit root tests, such as the Augmented Dickey-Fully, provide evidence this transformation is sufficient.[3] Summary statistics are reported in Table 1 and five major currencies' returns are plotted in Figure 1 through 5. The summary statistics indicate that most of the return series have positive skewness, with five of the return series exhibiting both relatively large skewness and large kurtosis. These are: the Australian dollar, the Greek drachma, the New Zealand dollar, the Portuguese escudo and the Swedish krona. It therefore appears as though these features are limited to less major international currencies.

3.2. Results

We estimate the fractional integration parameter d by the wavelet OLS estimator described previously, using the Daubechies (1988) wavelet, with two different "smoothing" parameters (Daubechies-4 and 12) in order to ensure robustness.[4,5] Since exchange rate returns are stationary, the null hypothesis of interest is whether the returns series are integrated of order zero (H_0: $d = 0$), versus the alternative of fractional integration (H_A: $d \neq 0$).

Estimates for the fractional integration parameter d over the 1981-2002 sample are reported in the third and fourth columns of Table 2, along with t-statistics for the null hypothesis $d = 0$. Table 2 includes some surprising results. First, fourteen out of nineteen exchange rate return series display statistically significant evidence of fractional integration across both wavelet transforms. These include some major currency rates, such as the British pound, Deutsch mark and Japanese yen, as well as minor currency rates. For five return series (Australian dollar, British pound, Greek drachma, New Zealand dollar, and Norwegian krone), the estimate of the fractional integration parameter is negative ($-0.5 < \hat{d} < 0$) and significant, while for the other nine series (Austrian schilling, Belgian franc, Danish krone, Deutsch mark, Italian lire, Japanese yen, Netherlands guilder, Portuguese escudo, and Spanish peseta) the estimate is positive ($0 < \hat{d} < 0.5$) and significant.

Negative estimates for d, or antipersistence, suggest that the return series is choppier than white noise, with negative autocovariances and unbounded variance at high spectral frequencies. It also implies that $0.5 < \hat{d} < 1$ for the series in levels, so that the exchange rate series still displays evidence of long memory with positive autocovariances, but with infinite variance. This relatively high incidence of antipersistence in the return series contrasts sharply with BBCC and Cheung (1993), which indicate no evidence of antipersistence. The disparate results, however, may be a manifestation of small sample bias, as the GPH estimator usually has positive bias while the wavelet OLS estimate has negative bias (c.f., Jensen 1999). Given the much lower MSE for the wavelet estimator, however, our results suggest that antipersistence in exchange rate returns may be more prevalent than previously acknowledged.

Interestingly, the five return series with negative fractional integration include the three series with the greatest kurtosis (Australian dollar, Greek drachma, and New Zealand dollar). We suspect that these features of the data may have a common source, such as frequent and substantial fundamental events that disproportionately impact these less prominent series. The volatility induced by such events could account for both the excess kurtosis and the high frequency spectral behavior associated with antipersistence.

In sum, five return series display evidence of antipersistence and nine return series display evidence of conventional long memory, for a total of fourteen exchange rate series exhibiting evidence of fractional integration. For the remaining five return series the evidence of fractional integration is inconclusive, with significance depending upon which wavelet is employed to estimate the fractional integration parameter. For three of these five series, (Canadian dollar, French franc, and Swiss franc), $\hat{d}$ is significant for the Daubechies-4 wavelet, which, as noted by Jensen (1999), is less likely to be impacted by boundary effects than wavelet transforms with higher order smoothing parameters.

In order to ensure that our results are robust to alternative frequencies, we also performed the wavelet OLS estimation for weekly observations and obtaine similar results. For example, on a sample corresponding to table 2, adjusted for boundary effects, the only substantial change is for the Italian lire and Spanish peseta, for which the evidence of fractional integration with weekly observations is inconclusive, with significance and sign depending upon which wavelet is employed to estimate the fractional integration parameter.

Due to their prominence in international transactions, the dynamic behavior of the Deutschemark and the Franc are potentially interesting. Indeed, our Wavelet-based analysis suggests that the two exchange return series have very similar properties, with both displaying strong evidence of long memory over the 1974-1995 sample, and both displaying evidence of long memory over the 1981-2002 for the Daubechies-4 wavelet. In addition, for both return series the estimate of the integration parameter is $0<d<0.5$, suggesting persistence with positive autocorrelations. Intuitively, our estimates indicate that the effects of innovations in the univariate return series due to, for example, macroeconomic events, tend to persist for many months, but eventually die out. Of course, our empirical analysis is strictly descriptive of the univariate data series. Identifying the source and effects of a particular macroeconomic shock (due to, for example, innovations in monetary policy, inflation differential etc..) would require a simultaneous equations model, which is beyond the scope of this paper.

3.3. Reconciliation with Previous Research

Cheung (1993) originally reported evidence that major exchange rate series, except for the British pound, displayed evidence of long memory. Several subsequent studies supported those findings, until BBCC reported substantial evidence of temporal instability. In particular, BBCC find only weak and sporadic evidence of long memory in their updated sample (1974-1995 versus 1974-1989).

In order to reconcile our results with previous findings, we undertake the following three exercises. First, to evaluate the impact of the wavelet OLS estimator, we also estimate $\hat{d}$ by GPH method for our 1981-2002 sample. Second, to better contrast our results with BBCC, we estimate $\hat{d}$ by the wavelet OLS method and the GPH method over their 1974-1995 sample. Third, we estimate $\hat{d}$ over Cheung's sample period for the relevant exchange return series. These are discussed, in turn below.

The estimates for $\hat{d}$ based on the GPH estimator, with two common bandwidths, for the most current (1981-2002) sample are reported in the final two columns of Table 2. We use the GPH estimator as a benchmark because it is the most popular method for estimating $\hat{d}$; because it is conceptually similar to the GSE estimator employed by BBCC; and because BBCC indicate their results are robust with respect to the GPH estimator. The disparate results obtained from the GPH estimator versus the wavelet OLS estimator are striking. By the GPH estimator, none of the series display persuasive evidence of long memory, while fourteen out of nineteen display evidence on the basis of the wavelet estimator. Although there are no extant Monte-Carlo studies on the relative power of the wavelet versus the GPH estimator, Jensen has shown that the wavelet estimator has considerably smaller MSE error than the GPH estimator, although somewhat greater bias. This implies substantially lower sampling variability, which may translate into greater power against the null of no fractional integration. The estimates presented here certainly seem to suggest that the wavelet estimator may have greater power.

To better discern whether the disparate results of the wavelet versus the GPH estimator are an artifact of the 1981-2002 sample period, we also provide wavelet estimates over the 1974-1995 sample, comparable to BBCC, with a small adjustment to the starting month (September rather than January) so that there are 256 monthly observations on the return series. The results are reported in Table 3. Again, the results are striking. Fourteen out of nineteen series display persuasive evidence of long memory, versus five in BBCC. Each of the five series reported by Barkoulas to have evidence of long memory (French franc, Danish krone, Luxembourg franc, Portuguese escudo, Spanish peseta) is also found to display evidence of long memory by the wavelet estimator, except for the Luxembourg franc which is not reported here. In addition, there appears to be considerable temporal stability in the wavelet OLS results, as only the Austrian schilling, Australian dollar and the Greek drachma exchange rate series display substantial diminished evidence of long memory, while the French franc, Swedish krona and Swiss franc display increased evidence of long-memory. The evidence on long memory in remaining thirteen exchange rate series is unchanged from Table 2, although the sign on the return series is different for three of the currency returns (Japanese yen, New Zealand dollar, Norwegian krone).

Finally, we apply the wavelet estimator to the series examined by Cheung, using the 1977-1987 sample, which is slightly modified from Cheung's, to accommodate 128 observations. The results are not reported here, but they are obtainable from the authors upon request. All five major currencies have statistically significant long memory in both degrees of smoothing, reinforcing the findings of Cheung, and providing further evidence of the temporal stability of the wavelet estimator. Indicative of the small-sample biases, the wavelet point estimates are again smaller in magnitude that the GPH estimates reported by Cheung.

4. Summary and Conclusion

This study reexamines exchange rates for evidence of long memory, utilizing a newly developed semi-parametric estimator, developed by Jensen (1999), as well as a longer, broader and more recent sample. This new estimator we utilize, denoted the wavelet OLS estimator, is a considerable improvement over the popular GPH estimator on the basis of mean squared error, with a MSE of approximately four to six times smaller than that of the GPH estimator for relevant sample sizes.

We find that fourteen of nineteen exchange rate series display persuasive evidence of long-memory dynamics based on the wavelet OLS estimator. This contrasts sharply with results obtained by the popular GPH estimator, which indicate little or no evidence of long-memory in recent samples. Our results are statistically significant and the general evidence of long-memory is relatively robust to alternative sample periods and alternative frequencies. Finally, we find that a surprisingly large fraction of return series display evidence of negative fractional integration (i.e., with the fractional integration parameter between –0.5 and 0), indicating that the associated exchange rate displays evidence of long memory, but that the return series displays evidence of antipersistence.

Our empirical results suggest several areas for future research. First, our results suggest that the GPH estimator may suffer from low power relative to the wavelet OLS estimator, at least for some relevant alternative hypotheses. A Monte-Carlo study on the relative power of the two estimators could confirm this. If the wavelet OLS estimator does have greater power, then it will be an extremely useful in re-assessing the prevalence of long memory in other financial series, such as stock returns and commodity prices.

References

Baillie, R. T. and T. Bollerslev. "Common Stochastic Trends in a System of Exchange Rates." *Journal of Finance* 44 (1989):167-181.

Baillie, R. T. and T. Bollerslev. "Cointegration, Fractional Cointegration, and Exchange Rate Dynamics." *Journal of Finance* 49-2 (1994): 737-745.

Barkoulas, J., C. Baum, M. Caglayan, and A. Chakraborty. "Persistent Dependence in Foreign Exchange Rates? A Reexamination." Forthcoming in *Global Financial Markets: Issues and Policies*, Greenwood Press, Westport CT, 2003.

Booth, G. G., F. R. Kaen, and P. E. Koveos. "R/S Analysis of Foreign Exchange Markets Under Two International Monetary Regimes." *Journal of Monetary Economics* 10 (1982): 407-415.

Breidt, F. J., N. Crato, and P. De Lima. "The Detection and Estimation of Long Memory in Stochastic Volatility." *Journal of Econometrics* 83-2 (1998): 325-348.

Brockwell, P. J. and R. A. Davis. *Time Series: Theory and Methods*. 2nd ed. Springer, New York, 1991.

Cheung, Y. "Long Memory in Foreign-Exchange Rates." *Journal of Business and Economic Statistics* 11 (1993): 93-101.

Daubechies, I. "Orthonormal Bases of Compactly Supported Wavelets." *Communications on Pure and Applied Mathematics* 41 (1988): 909-996.

Dickey, D.A. and W. Fuller. "Likelihood Ratio Tests for Autoregressive Time Series with a Unit Root." *Econometrica* 49 (1981): 1057-1072.

Geweke, J. and S. Porter-Hudak. "The Estimation and Application of Long Memory Time Series Models." *The Journal of Time Series Analysis* 4-4 (1983): 221-238.

Granger, C. W. and R. Joyeaux. "An Introduction to Long-Memory Time Series Models and Fractional Differencing." *The Journal of Time Series Analysis* 1-1 (1980): 15-29.

Hosking, J. R. "Fractional Differencing." *Biometrika* 68 (1981): 165-176.

Jensen, M. J. "Using Wavelets to Obtain a Consistent Ordinary Least Squares Estimator of the Long-memory Parameter." *Journal of Forecasting* 18 (1999): 17-32.

Kwiatkowski, D., P. Phillips, P. Schmidt, and Y. Shin. "Testing the Null of Stationarity Against the Alternative of a Unit Root." *Journal of Econometrics* 54 (1992): 159-178.

Pan, M., Y. A. Liu, and H. Bastin. "An Examination of the Short-Term and Long-Term Behavior of Foreign Exchange Rates." *The Financial Review* 31-3 (1996): 603-622.

Papell, D.H (2002)," The Great Appreciation, the Great Depreciation, and the Purchasing Power Parity Hypothesis." *Journal of International Economics*, 57: 51-82.

Robinson, P. "Efficient tests of nontationary hypotheses." *Journal of the American Statistical Association* 89 (1994): 1420-1437.

Robinson, P. "Gaussian Semiparametric Estimation of Long Range Dependence." *Annals of Statistics* 13 (1995): 1630-1661.

Sowell, F. "Maximum Likelihood Estimation of Stationary Univariate Fractionally Integrated Time Series." *Journal of Econometrics* 53 (1992): 165-188.

Tkacz, G. "Estimating the Fractional Order of Integration of Interest Rates Using a Wavelet OLS Estimator." *Studies in Nonlinear Dynamics & Econometrics* 5-1 (2001): 19-32.

Table 1.
Summary Statistics for Nominal Exchange Rate Returns, January 1974 through December 2002.

Currency	Mean	Median	Standard Deviation	Maximum	Minimum	Skewness	Kurtosis
Australian dollar	0.00121	0.00000	0.0128	0.0832	-0.0384	1.5676	7.5229
Austrian schilling	-0.00055	-0.00035	0.0138	0.0526	-0.0434	0.1341	0.9439
Belgian franc	-0.00012	0.00058	0.0115	0.0343	-0.0308	-0.0235	0.0540
British pound	0.00040	0.00012	0.0132	0.0554	-0.0570	0.0837	1.9604
Canadian dollar	0.00056	0.00050	0.0058	0.0271	-0.0151	0.5026	1.6634
Danish krone	0.00009	0.00143	0.0112	0.0365	-0.0310	-0.0979	0.0143
Deutsche mark	-0.00050	0.00010	0.0115	0.0373	-0.0310	-0.0960	-0.0185
Finnish markka	0.00046	-0.00037	0.0128	0.0656	-0.0328	0.6586	2.6416
French franc	0.00028	0.00108	0.0112	0.0390	-0.0311	0.0484	0.2714
Greek drachma	0.00300	0.00153	0.0126	0.0756	-0.0261	1.3219	5.5220
Italian lire	0.00133	0.00107	0.0111	0.0473	-0.0275	0.3520	0.8645
Japanese yen	-0.00127	0.00049	0.0145	0.0500	-0.0651	-0.3434	1.4250
Netherlands guilder	-0.00040	0.00061	0.0114	0.0373	-0.0308	-0.0872	0.0507
New Zealand dollar.	0.00130	0.00021	0.0142	0.1081	-0.0471	1.7303	12.1892
Norwegian krone	0.00015	-0.00069	0.0123	0.0498	-0.0419	0.3398	1.5346
Portuguese escudo	0.00247	0.00210	0.0115	0.0716	-0.0221	0.9872	4.4481
Spanish peseta	0.00127	0.00096	0.0113	0.0651	-0.0278	0.6051	2.6998
Swedish krona	0.00073	-0.00070	0.0132	0.0747	-0.0335	1.1748	5.1681
Swiss franc	-0.00110	-0.00035	0.0155	0.0671	-0.0505	0.1885	0.9081

Table 2.
Wavelet OLS and GPH Estimates of the Fractional Integration Parameter on Nominal Exchange Rate Returns: Sep 1981–Dec 2002.

Exchange Rate	Signif	$\hat{d}$ D-4	$\hat{d}$ D-12	$\hat{d}$ GPH(0.45)	$\hat{d}$ GPH(0.5)
Australian dollar	– –	–0.1186**	–0.0470**	0.2916	0.2226
		(–14.29)	(–5.61)	(1.14)	(1.06)
Austrian schilling	++	0.0764**	0.0734**	0.0829	–0.03951
		(9.80)	(6.76)	(0.32)	(–0.19)
Belgian franc	++	0.1112**	0.0780**	0.2070	0.2443
		(11.48)	(3.99)	(0.80)	(1.16)
British pound	– –	–0.2969**	–0.1452**	–0.1563	–0.1532
		(–14.79)	(–9.52)	(–0.61)	(–0.73)
Canadian dollar		–0.0658**	–0.0067	0.2732	0.4501**
		(–7.71)	(–0.68)	(1.07)	(2.14)
Danish krone	++	0.0942**	0.0771**	0.1556	0.2345
		(9.99)	(4.60)	(0.61)	(1.12)
Deutsche mark	++	0.1100**	0.1181**	0.2974	0.2373
		(13.38)	(10.84)	(1.16)	(1.28)
Finnish markka		–0.0645**	0.0502**	0.3289	0.1994
		(–4.23)	(6.60)	(1.28)	(0.95)
French franc		0.0679**	0.0145	0.2543	0.2106
		(5.46)	(0.57)	(0.99)	(1.00)
Greek drachma	– –	–0.0905**	–0.3372**	0.1003	0.1780
		(–6.27)	(–11.00)	(0.30)	(0.85)
Italian lire	++	0.0308**	0.0621**	0.3575	0.2802
		(2.79)	(4.97)	(1.40)	(1.33)
Japanese yen	++	0.0617**	0.0557**	0.1016	0.0679
		(10.21)	(5.72)	(0.40)	(0.32)
Netherlands guilder	++	0.1117**	0.1183**	0.2750	0.2094
		(13.11)	(10.02)	(1.07)	(1.00)
New Zealand dollar	– –	–0.2353**	–0.0448**	–0.2213	–0.1626
		(–10.22)	(–3.51)	(–0.86)	(–0.77)
Norwegian krone	– –	–0.1719**	–0.0448**	–0.2266	–0.0130
		(–7.88)	(–3.27)	(–0.86)	(–0.62)
Portuguese escudo	++	0.0688**	0.1290**	0.3807	0.3621
		(5.33)	(12.50)	(1.48)	(1.72)
Spanish peseta	++	0.0835**	0.0630**	0.3325	0.3884
		(7.70)	(3.07)	(1.30)	(1.85)
Swedish krona		–0.0140	0.0301**	0.0178	0.1631
		(–1.46)	(2.67)	(0.69)	(0.78)
Swiss franc		–0.0576**	0.0168	–0.0272	0.1240
		(–3.86)	(1.66)	(–0.11)	(0.59)

This table reports estimates of the fractional integration parameter *d* obtained by the wavelet OLS estimator with the Daubechies-4 and Daubechies-12 wavelets on each nominal exchange rate return series. The values in parentheses are *t*-statistics for the null hypothesis $d = 0$. The superscript ** indicates that the null is rejected at the 5% level of significance. The column label "Signif" includes indicators "++" and "– –" for whether the fractional integration parameter is significant for both wavelets, with the indicated sign. The GPH estimator is reported for two estimation windows, $T^{0.45}$ and $T^{0.50}$.

Table 3.
Wavelet OLS and GPH Estimates of the Fractional Integration Parameter on Nominal Exchange Rate Returns: Sep 1974–Dec 1995.

Exchange Rate	Signif	$\hat{d}$ D-4	$\hat{d}$ D-12	$\hat{d}$ GPH(0.45)	$\hat{d}$ GPH(0.5)
Australian dollar		0.0356**	–0.0224**	0.1625	–0.0176
		(6.53)	(–4.91)	(0.63)	(–0.08)
Austrian schilling		0.0044	–0.0152**	0.002	0.1542
		(0.76)	(–2.34)	(0.01)	(0.73)
Belgian franc	++	0.1600**	0.1605**	0.2352	0.3452
		(30.05)	(26.28)	(0.92)	(1.64)
British pound	– –	–0.0741**	–0.1069**	0.0449	–0.0079
		(–5.09)	(–5.80)	(0.18)	(–0.03)
Canadian dollar		–0.0114	–0.0617**	0.3252	0.3046
		(–0.85)	(–3.76)	(1.27)	(1.45)
Danish krone	++	0.1392**	0.1388**	0.1815	0.3149
		(18.81)	(17.47)	(0.71)	(1.50)
Deutsche mark	++	0.0816**	0.0639**	0.0257	0.1513
		(12.73)	(7.83)	(0.10)	(0.72)
Finnish markka		0.0090	–0.0337**	0.0851	0.0835
		(0.83)	(–2.85)	(0.33)	(0.40)
French franc	++	0.1683**	0.1614**	0.2714	0.3560
		(22.63)	(21.04)	(1.06)	(1.69)
Greek drachma		0.0097	0.0756**	0.3869	0.3557
		(0.55)	(7.77)	(1.51)	(1.69)
Italian lire	++	0.1353**	0.1285**	0.2623	0.3880
		(15.73)	(12.37)	(1.02)	(1.85)
Japanese yen	– –	–0.2632**	–0.1835**	–0.0606	0.0513
		(–11.79)	(–10.22)	(–0.24)	(0.24)
Netherlands guilder	++	0.0938**	0.0805**	0.0290	0.1449
		(14.46)	(9.94)	(0.11)	(0.69)
New Zealand dollar	++	0.0355**	0.0235**	0.000	0.0469
		(4.19)	(3.27)	(0.00)	(0.22)
Norwegian krone	++	0.0372**	0.0236**	–0.0175	0.0056
		(4.59)	(3.70)	(–0.07)	(0.03)
Portuguese escudo	++	0.2504**	0.2343**	0.3966	0.3647
		(31.61)	(28.61)	(1.55)	(1.73)
Spanish peseta	++	0.1936**	0.1948**	0.2697	0.4065
		(35.12)	(31.28)	(1.05)	(1.93)
Swedish krona	++	0.0604**	0.0666**	0.1200	0.2301
		(5.53)	(8.43)	(0.47)	(1.09)
Swiss franc	– –	–0.0101	–0.0485**	0.0412	0.1261
		(–1.13)	(–5.61)	(0.16)	(0.60)

This table reports estimates of the fractional integration parameter *d* obtained by the wavelet OLS estimator with the Daubechies-4 and Daubechies-12 wavelets on each nominal exchange rate return series. The values in parentheses are *t*-statistics for the null hypothesis $d = 0$. The superscript ** indicates that the null is rejected at the 5% level of significance. The column label "Signif" includes indicators "++" and "– –" for whether the fractional integration parameter is significant for both wavelets, with the indicated sign. The GPH estimator is reported for two estimation windows, $T^{0.45}$ and $T^{0.50}$.

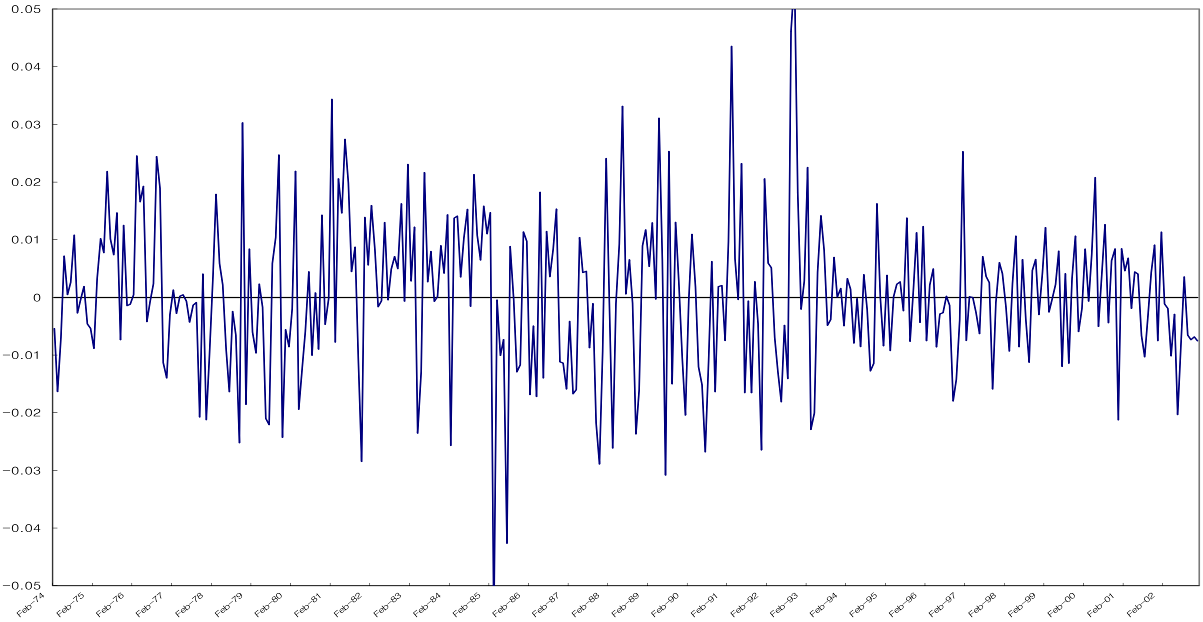

Figure 1
British Pound Returns, from Jan. 1974 through Dec. 2002.

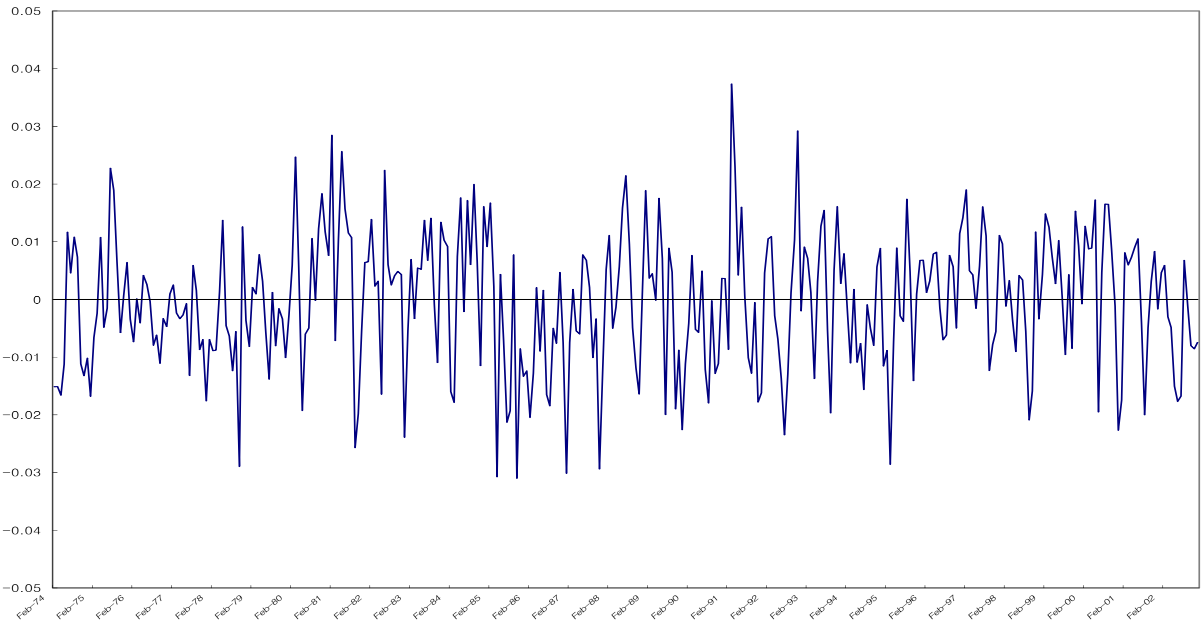

Figure 2
Deutsche Mark Returns, from Jan. 1974 through Dec. 2002.

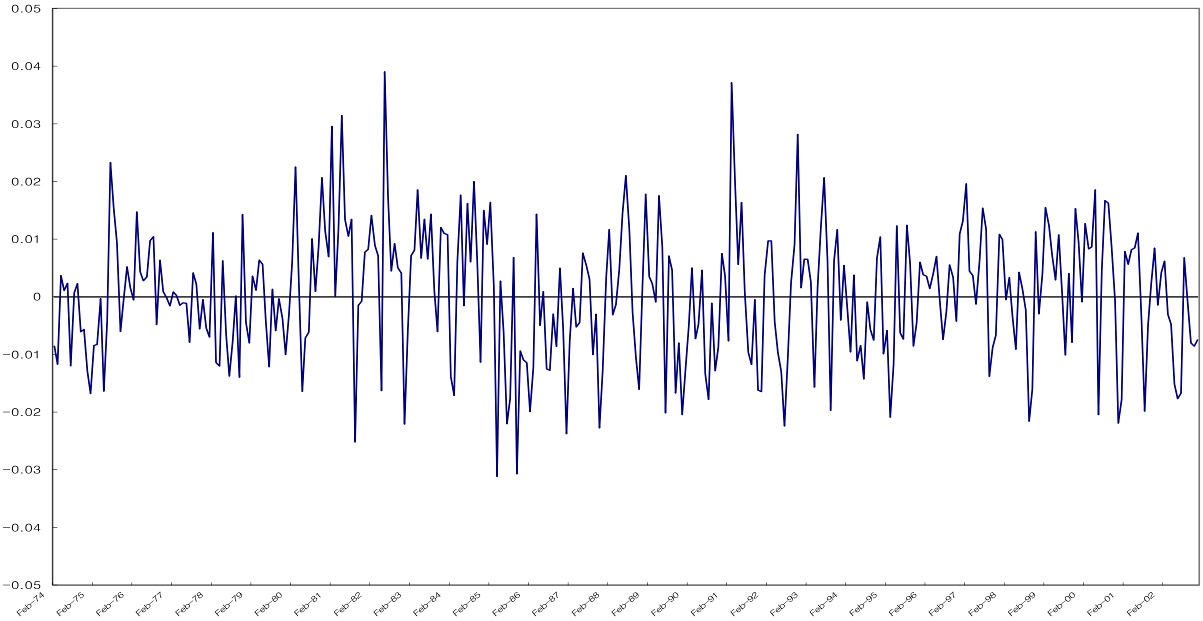

Figure 3
French Franc Returns, from Jan. 1974 through Dec. 2002.

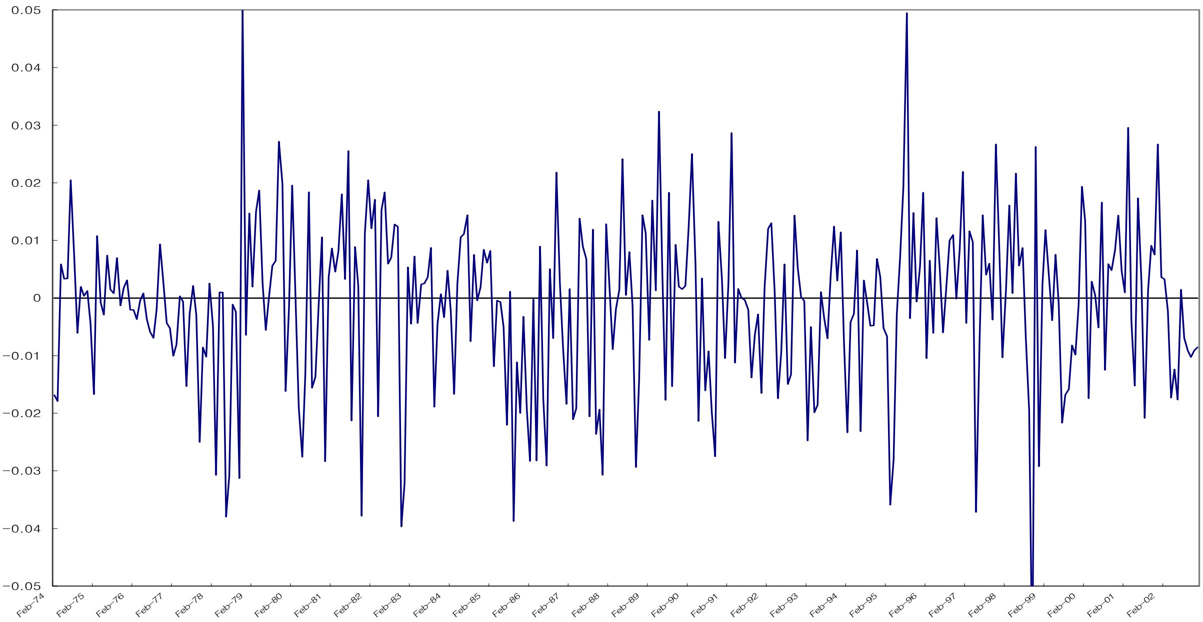

Figure 4
Japanese Yen Returns, from Jan. 1974 through Dec. 2002.

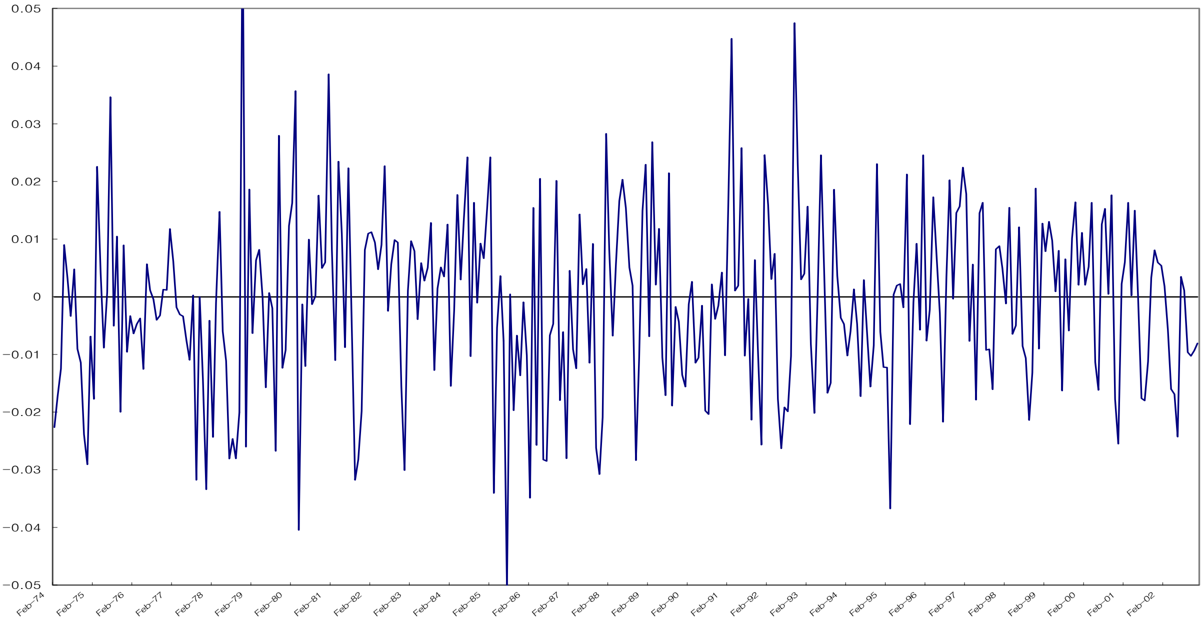

Figure 5
Swiss Franc Returns, from Jan. 1974 through Dec. 2002.

Endnotes

---

[1] Brockwell and Davis (1991) showed that when a weakly stationary process has long-term memory, its autocorrelation function exhibits hyperbolic decay. Hyperbolic decay means that for $x_t = \varphi(L)u_t$, then for any hypothetical rate r, the coefficient on L^s in $Q(L)$ is larger than r^s, for all s greater than some sufficiently large S.

[2] Note that this sample period includes the post-1999 period during which exchange rates were effectively fixed for some major European countries, as dictated by the Treaty of Maastricht. Given our comparisons with the pre-1995 sample, this should not substantively alter the interpretation our results.

[3] We tested for stationary using the Augmented Dickey-Fuller (ADF, 1981) test for the null hypothesis of a unit-root and the Kwiatkowski, Phillips, Schmidt, and Shin (1992) test for the null hypothesis of stationarity. On the basis of ADF tests, each exchange rate series in levels displays evidence of a unit root. For the return series, the null hypothesis of a unit-root was rejected and the null hypothesis of stationarity was not rejected at the 5 percent level of significance, suggesting that log-differences are sufficient to induce stationarity.

[4] The computations are performed using the Matlab toolbox *Wavekit*, of Ojanen, ("WAVEKIT: a Wavelet Toolbox for Matlab," Department of Mathematics, Rutgers University, April 1998) which can be obtained from http://www.math.rutgers.edu/~ojanen/wavekit.

[5] We also estimated the model with the Daubechies-20 wavelet, with some impact on several estimates. As indicated by Jensen (1999), this is likely the result of boundary effects associated with the larger smoothing window.