© 1993 American Statistical Association Journal of Business & Economic Statistics, January 1993, Vol. 11, No. 1

Long Memory in Foreign-Exchange Rates

Yin-Wong Cheung
Economics Board, University of California, Santa Cruz, CA 95064

Using the Geweke–Porter-Hudak test, we find evidence of long memory in exchange-rate data. This implies that the empirical evidence of unit roots in exchange rates may not be robust to long-memory alternatives. Fractionally integrated autoregressive moving average (ARFIMA) models are estimated by both the time-domain exact maximum likelihood (ML) method and the frequency-domain approximate ML method. Impulse-response functions and forecasts based on these estimated ARFIMA models are evaluated to gain insight into the long-memory characteristics of exchange rates. Some tentative explanations of the long memory found in the exchange rates are discussed.

KEY WORDS: Exchange-rate dynamics; Forecast; GPH test; Impulse-response function; Maximum likelihood estimation.

Numerous efforts have been made to understand exchange-rate dynamics since the inception of the floating-rates regime in 1973. Knowledge of the time series properties of exchange rates has significant economic implications. For instance, the empirical behavior of exchange rates provides useful information for constructing models of exchange-rate determination and can be used to evaluate the performance of exchange-rate models. The appropriate statistical tool for studying the relationship between exchange rates and other economic variables can depend on the time series properties of exchange rates. Moreover, exchange-rate dynamics are important for the determination of international trade flows, prices of tradable goods, prices of foreign-exchange futures (or options), and international-asset portfolios.

Integrated autoregressive moving average (ARIMA) models are the standard time series models used to study the intertemporal dynamics of exchange rates. Previous studies showed that exchange rates exhibit large persistence and appear to follow a martingale process. Typically, the standard tests for unit roots (e.g., the Dickey–Fuller and Phillips–Perron tests) cannot reject the null hypothesis that there is a unit root in exchange-rate data (Baillie and Bollerslev 1989; Meese and Singleton 1982). On the other hand, Booth, Kaen, and Koveos (1982), using the rescaled range technique, suggested the possibility of long memory, defined later, in exchange-rate data. As demonstrated by Aydogan and Booth (1988), Davies and Harte (1987), and Lo (1991), however, the rescaled range analysis is not a reliable tool to test for long memory because this technique is not robust to the short-range persistence and heteroscedasticity.

This article examines the time series dynamics of exchange rates in the fractionally integrated autoregressive moving average (ARFIMA) framework (Granger and Joyeux 1980; Hosking 1981), which allows for long memory in the data and is a generalization of the ARIMA model. Specifically, the differencing parameter of an ARFIMA model is not restricted to the integer domain and can assume real values. This generalization makes an ARFIMA model a parsimonious and flexible model to study long memory and short-run dynamics simultaneously. Moreover, fractional integration is a more general way to describe long-range dependence than the unit-root specification and provides an alternative perspective to examine the unit-root hypothesis. This article will provide more convincing evidence of long memory in exchange-rate changes based on a statistical procedure that is asymptotically robust to short-range dependence and has a well-defined asymptotic distribution. In addition, ARFIMA models are fitted to the data. Impulse-response functions and forecasts based on the fitted models are evaluated to gain more insight into the long-memory characteristics of exchange rates.

Section 1 introduces the concept of long memory and the ARFIMA model. Because of the complication involved in estimating ARFIMA models, we suggest testing for the presence of long memory before estimating an ARFIMA model. In this article, we use a seminonparametric method to test for long memory and two maximum likelihood (ML) procedures to estimate an ARFIMA model. These procedures are also discussed in Section 1. Section 2 presents empirical results. Section 3 offers tentative explanations of the long-memory behavior found in exchange rates, and Section 4 summarizes the article.

1. LONG-MEMORY TIME SERIES MODEL

A covariance stationary time series is said to exhibit long memory if it satisfies the following condition (McLeod and Hipel 1978, p. 492):

$$\sum_{k=-n}^{-n} |\rho(k)| \to \infty \quad \text{as} \quad n \to \infty, \tag{1}$$

where $\rho(k)$ is the autocorrelation at lag k. This infinite-sum condition suggests that correlations at long lags are not negligible. It can be shown that stationary autoregressive moving average (ARMA) processes do not possess long memory.

Long memory can be found in the class of ARFIMA processes. A time series $\mathbf{X} = \{X_1, \ldots, X_T\}$ follows an ARFIMA(p, d, q) process if

$$\Phi(B)(1 - B)^d X_t = \Theta(B)\varepsilon_t, \tag{2}$$

where $\varepsilon_t \sim \text{iid}(0, \sigma^2)$, B is the backward-shift operator, $\Phi(B) \equiv 1 - \phi_1 B - \cdots - \phi_p B^p$, $\Theta(B) \equiv 1 + \theta_1 B + \cdots + \theta_q B^q$, and $(1 - B)^d$ is the fractional differencing operator defined by $(1 - B)^d = \Sigma_{k=0}^{\infty} \Gamma(k - d)B^k/\Gamma(k + 1)\Gamma(-d)$ with $\Gamma(.)$ being the gamma function. $\mathbf{X}$ is both stationary and invertible if the roots of $\Phi(B)$ and $\Theta(B)$ are outside the unit circle and $d < |.5|$. When $d = 0$, an ARFIMA process reduces to an ARMA process.

Hosking (1981) showed that the autocorrelation, $\rho(.)$, of an ARFIMA process satisfies $\rho(k) \propto k^{2d-1}$ as $k \to \infty$. Thus the memory property of a process depends crucially on the value of d. When $d \in (0, .5)$, the autocorrelations do not have a finite sum. When $d \leq 0$, the autocorrelations have a finite sum; that is, ARFIMA processes with $d \in (0, .5)$ display long memory. Hence the existence of long memory can be conveniently determined by testing for the statistical significance of the sample differencing parameter d.

1.1 Test for Long Memory

Geweke and Porter-Hudak (hereafter abbreviated GPH) (1983) proposed a seminonparametric procedure to test for long memory. They showed that, when attention is confined to frequencies near 0, the differencing parameter d can be consistently estimated from the least squares regression

$$\ln(I(\omega_j)) = c - d \ln(4 \sin^2(\omega_j/2)) + \eta_j, \qquad j = 1, \ldots, n, \tag{3}$$

where $\omega_j = 2\pi j/T$ ($j = 1, \ldots, T - 1$), $n = g(T) << T$, and $I(\omega_j)$ is the periodogram of $\mathbf{X}$ at frequency ω_j defined by

$$I(\omega) = \frac{1}{2\pi T}\left|\sum_{t=1}^{T} e^{it\omega}(X_t - \overline{X})\right|^2, \tag{4}$$

There is evidence of long memory if the least squares estimate of $d(\hat{d}_{\text{GPH}})$ is significantly larger than 0. With a proper choice of n, the asymptotic distribution of $\hat{d}_{\text{GPH}}$ depends on neither the order of the ARMA component nor the distribution of the error term of the ARFIMA process (see also Yajima 1989). It is suggested to set $n = T^{.5}$ and to use the known variance of η_j, $\pi^2/6$, to compute the estimated variance of $\hat{d}_{\text{GPH}}$.

The GPH test can be used as a test for unit roots. Under the unit-root hypothesis, the first-differenced data follow a stationary ARMA process with $d = 0$. Hence the unit-root hypothesis can be tested by determining whether the $\hat{d}_{\text{GPH}}$ from the first-differenced data is significantly different from 0. This is very interesting for studying exchange-rate dynamics because standard unit-root tests, which have stationary ARMA models as their alternatives, usually cannot reject the presence of a unit root in exchange-rate data. By including fractional processes as the alternatives, the GPH test provides a different perspective to examine the unit-root hypothesis.

1.2 Estimating ARFIMA Models

The parameters of an ARFIMA model can be jointly estimated by either a time-domain exact ML method or a frequency-domain approximate ML method. The time-domain ML method was discussed by Li and McLeod (1986) and Sowell (1990a). The frequency-domain procedure was proposed by Fox and Taqqu (1986). The properties of these two ML estimators were also discussed by Yajima (1985, 1988).

Assuming normality, the likelihood function of the ARFIMA model is given by

$$\mathcal{L}(\mathbf{X}|.) = (2\pi)^{-T/2}|\Sigma|^{-1/2} \exp(-\mathbf{X}'\Sigma^{-1}\mathbf{X}/2), \tag{5}$$

where Σ is the $T \times T$ covariance matrix of $\mathbf{X}$ and is a function of d, θ_i's, ϕ_i's, and σ^2. Li and McLeod (1986) showed that the ML estimator, which is obtained by maximizing (5) with respect to the parameter vector $\xi = (d, \phi_i\text{'s}, \theta_i\text{'s})$ and σ^2, is consistent and asymptotically normal.

Fox and Taqqu (1986) suggested a frequency-domain method to estimate the ARFIMA model. This method amounts to minimizing

$$\sigma_T^2(\xi) = \sum_{k=1}^{T-1} I(2\pi k/T)\{f(2\pi k/T, \xi)\}^{-1} \tag{6}$$

with respect to parameter vector ξ. $I(\lambda)$ is the periodogram defined in (4), and $f(\lambda, \xi)$ is proportional to the spectral density of $\mathbf{X}$ at frequency λ and satisfies the normalization condition

$$\int_{-\pi}^{\pi} \ln\{f(\lambda, \xi)\}\, d\lambda = 0. \tag{7}$$

Given the spectral density of an ARFIMA process, (6) can be used to obtain $\hat{\xi}$. The exact relationship between $f(\lambda, \xi)$ and the spectral density is required, however, to compute the sample variance–covariance matrix of ξ and evaluate the statistical significance of the parameter estimates. It can be shown that, for ARFIMA models,

$$f(\lambda, \xi) = |1 - e^{-i\lambda}|^{-2d} \frac{|\Theta(e^{-i\lambda})|^2}{|\Phi(e^{-i\lambda})|^2} \tag{8}$$

satisfies (7) and the distribution of $\sqrt{T}(\hat{\xi} - \xi_0)$ can be approximated by

$$\sqrt{T}(\hat{\xi} = \xi_0) \to N\left(0, 2\sigma_T^2(\hat{\xi})\left[\left.\frac{\partial^2\sigma_T^2(\xi)}{\partial\xi\partial\xi'}\right|_{\hat{\xi}}\right]^{-1}\right), \tag{9}$$

where ξ_0 is the true parameter vector. The general form of (9) was given by Fox and Taqqu (1986). The frequency-domain estimation results reported in this article are based on (6), (8), and (9).

2. EVIDENCE OF LONG MEMORY IN NOMINAL EXCHANGE RATES

The time series properties of five nominal dollar spot rates—British pound (BP), Deutsche mark (DM), Swiss franc (SF), French franc (FF), and Japanese yen (JY)—are examined. End-of-week exchange rates from January 1974 to December 1989 were taken from the Chicago Mercantile Exchange Yearbooks. The FF data are only available from the fourth week of August 1974. Thursday data are used when Friday quotations are not available. The testing and estimation results are based on data from 1974 to 1987. Observations in 1988 and 1989 are reserved for out-of-sample forecast analysis.

Graphs of the first-differenced log data are presented in Figure 1. Exchange-rate changes appear to be random fluctuations around 0 and have no obvious structural break during the sample period; that is, exchange-rate data seem to be difference-stationary. Indeed, when the augmented Dickey–Fuller test and the corresponding Phillips–Perron test for a unit root were conducted on individual exchange-rate series, these tests could not reject the unit-root hypothesis at the 5% level. As noted by Diebold and Rudebusch (1991a) and Sowell (1990b), however, these standard unit-root tests have low power against fractional integration alternatives.

2.1 Test for Long Memory

The results of applying the GPH test on first-differenced log data are reported in Table 1. The unit-root null hypothesis is tested against the long-memory alternative. In addition to the $\hat{d}_{\mathrm{GPH}}$ that is based on $n = T^{.5}$, which is commonly used to test for long memory, we also report $\hat{d}_{\mathrm{GPH}}$'s based on $n = T^{.45}$ and $T^{.55}$ to check the sensitivity of our results to the choice of n. p values computed with the theoretical variance of η_j are given in parentheses below the estimates.

The results indicate that the evidence of long memory in the DM, SF, JY, and FF series is qualitatively the same across different choices of n. For the BP series, the null is marginally rejected when $n = T^{.5}$. Note that the $\hat{d}_{\mathrm{GPH}}$ corresponding to $n = T^{.55}$ is smaller than that corresponding to $n = T^{.45}$ or $T^{.50}$. This is in accord with GPH's (1983) observation that the estimation result can be "contaminated" as more periodogram ordinates are included in the regression. The d parameter estimates suggest that the memory of the DM, SF, and JY series is of comparable magnitude. The memory of the FF series is the strongest among the five series, whereas that in the BP series is the weakest.

The modified rescaled range test for long memory, suggested by Lo (1991), was also conducted on individual exchange-rate series. The result, which is available on request, is similar to that obtained from the GPH test; that is, the unit-root hypothesis is rejected in favor of the long-memory alternative.

In sum, we find evidence of long memory in exchange-rate changes though the result for the BP series is marginal. The unit-root hypothesis for the exchange-rate series is rejected in favor of the long-memory alternative because the $\hat{d}_{\mathrm{GPH}}$'s are usually significantly larger than 0. By the construction of the test, the empirical result is asymptotically robust to both the short-run dynamics and the underlying distribution of the series. This suggests that the evidence of a unit root in exchange-rate data is not robust to fractional integration alternatives. These findings warrant the use of ARFIMA models to examine exchange-rate dynamics.

2.2 Estimation of ARFIMA Models

ARFIMA(p, d, q) models with p and q less than or equal to 3 are used to capture both the long memory and the short-run dynamics of the exchange-rate data; that is, 16 different model specifications are considered for each exchange-rate series. The parameters of each model specification are estimated by the time-domain exact ML and the frequency-domain approximate ML methods. Estimates are obtained via the Davidon–Fletcher–Powell algorithm. The order of the ARFIMA model is initially determined by the Akaike information criterion (Hosking 1984; Beran 1986). Then we check for canceling factors in the AR and MA polynomials and use the Schwarz information criterion to choose the final ARFIMA specifications.

Parameter estimates and the corresponding p values obtained from the exact and approximate ML methods are presented in Table 2, a and b. Overall, the estimation results are in accordance with the GPH test results. The estimated ARFIMA models indicate that there is long memory in exchange-rate changes, and the dynamics are more complicated than implied by the random-walk hypothesis.

The frequency-domain approximate ML method was also applied to exchange-rate data filtered by two trapezoidal tapers. These two data tapers transform, respectively, 10% and 25% of the data at each end of the data series. Estimation results obtained from the tapered data are qualitatively similar to those based on nontapered data.

Note that conditional heteroscedasticity is not incorporated in these estimation methods. This may cause some loss in efficiency because it is known that exchange-rate data exhibit time-varying conditional variances. A long-memory model incorporating conditional heteroscedasticity is not easy to construct or estimate at this moment, however. Given the sample size considered here, the loss in efficiency can be small. Furthermore, results in Subsection 2.3 suggest that the dynamic behavior implied by these estimated models is consistent with observed exchange-rate movements.

2.3 Impulse-Response Functions

The impulse-response function (Campbell and Mankiw 1987; Watson 1986) of the estimated ARFIMA model

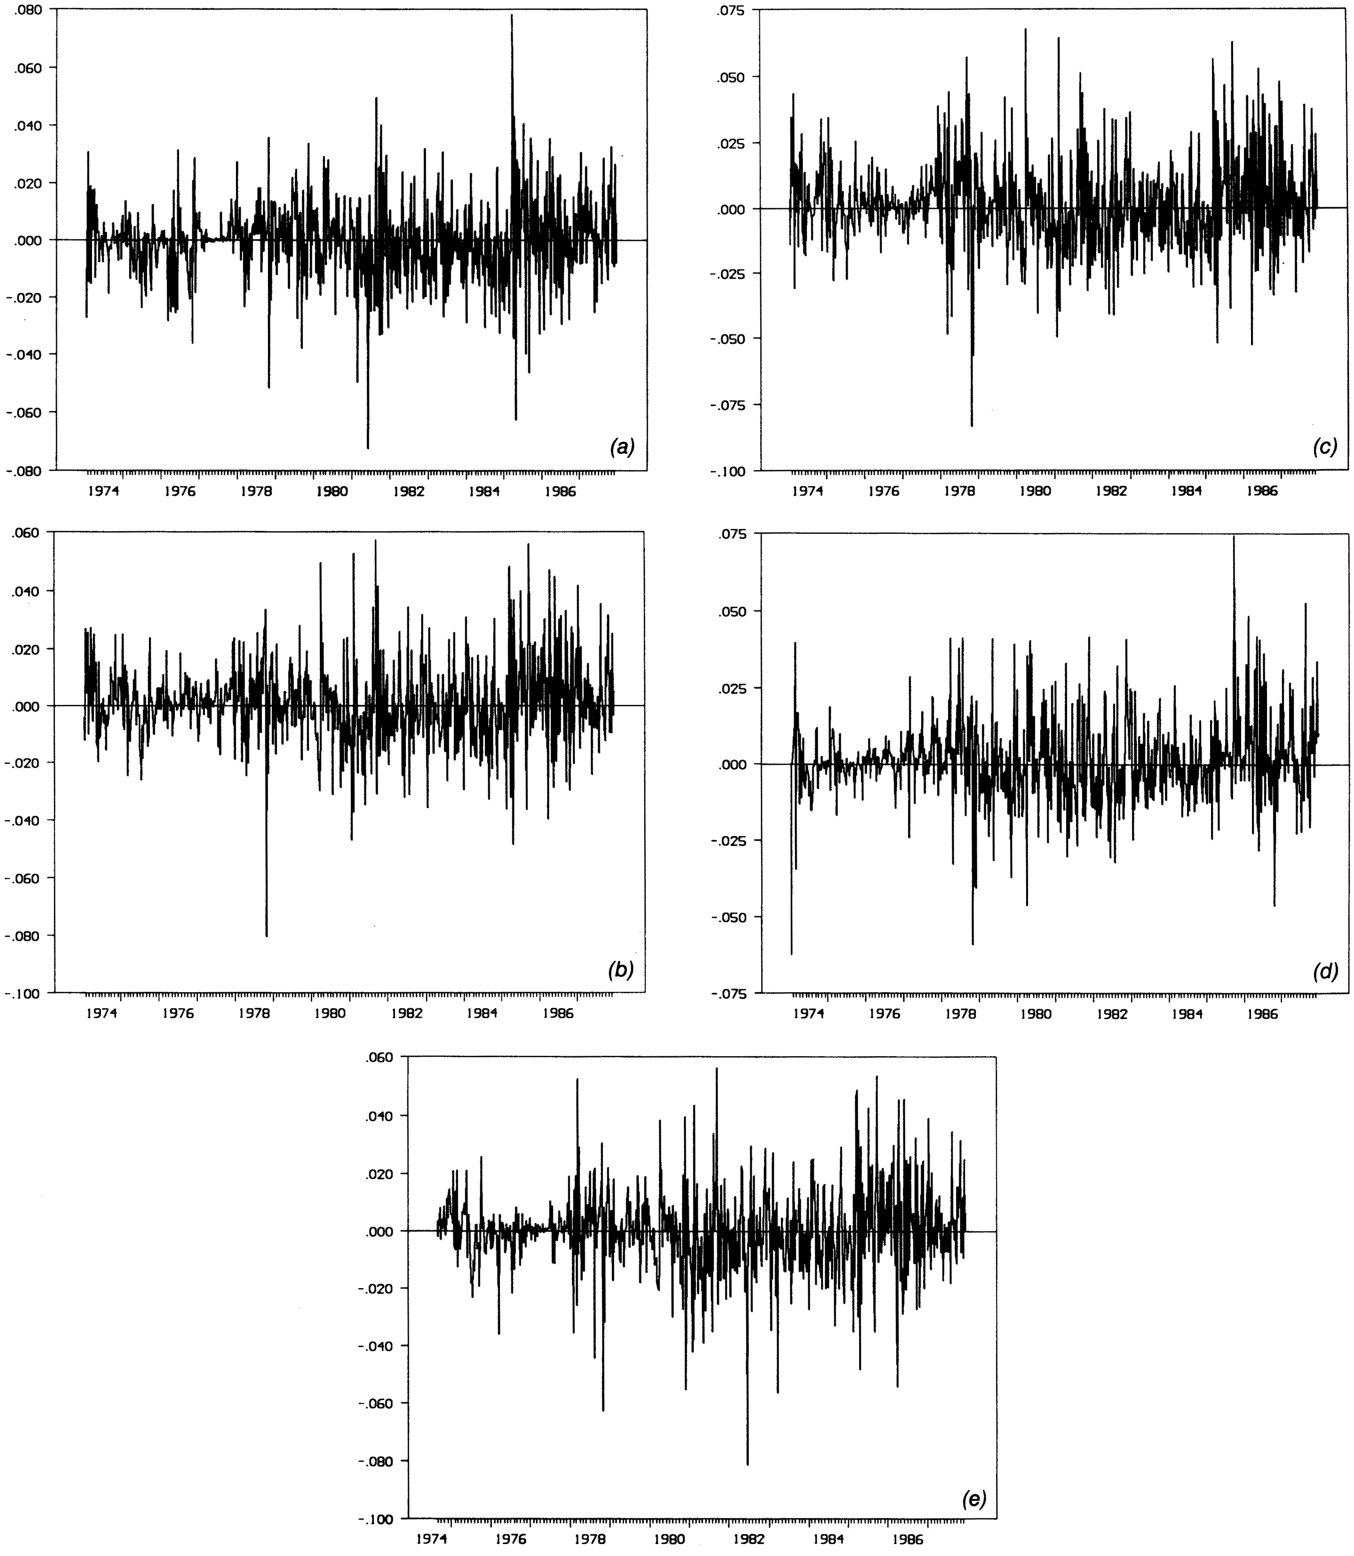

Figure 1. Weekly Exchange-Rate Change Series, 1974–1987: a, BP; b, DM; c, SF; d, JY; e, FF.

is one measure of the persistence in exchange-rate changes. Let X_t be the exchange-rate change at time t. Consider $(1 - B)X_t = A(B)\varepsilon_t$, where $A(B) \equiv (1 + a_1B + a_2B^2 + \cdots)$. The impulse-response function, $C(k)$, is given by

$$C(k) = \partial X_t/\partial \varepsilon_{t-k} = 1 + a_1 + \cdots + a_k, \qquad (10)$$

Table 1. Results of the Geweke–Porter-Hudak Test

GPH	BP	DM	SF	JY	FF
D(.45)	.1272	.3230**	.2797*	.2797*	.5046***
p value	(.249)	(.042)	(.068)	(.068)	(.004)
D(.50)	.1869	.2943**	.2870**	.2907**	.4240***
p value	(.106)	(.025)	(.028)	(.026)	(.003)
D(.55)	.1098	.2285**	.1993*	.1956*	.3084***
p value	(.186)	(.032)	(.053)	(.056)	(.007)

NOTE: D(.45), D(.50), and D(.55) give the d estimates corresponding to $n = T^{.45}$, $T^{.50}$ and $T^{.55}$. "p value" gives the p value of the no-long-memory null hypothesis ($d = 0$) against the long-memory alternative. *, ** and *** indicate significance at 10%, 5%, and 2.5%, respectively.

which measures the effect of a unit shock on the k-period-ahead exchange-rate change. For a fractional process, it can be shown that $C(\infty) = A(1) = 0$ if $d < 1$, $A(1)$ converges if $d = 1$, and $A(1)$ diverges if $d > 1$. The persistence measure $A(1)$ is known as the long-run impulse response.

Since the d estimates of the exchange-rate change series are all less than 1, $A(1) = 0$ for these series. As suggested by Campbell and Mankiw (1987) and Diebold and Rudebusch (1989), however, the impulse-response function itself contains more useful information. The sampling variability of impulse responses can be assessed by their asymptotic standard errors. For a given k, $C(k)$ is a function of the parameter vector ξ. Hence the asymptotic variance of $C(k)$ is given by $\nabla C(k)\Omega\nabla C'(k)$, where $\nabla C(k)$ is $\partial C(k)/\partial\xi$ and Ω is the variance–covariance matrix of ξ (Campbell and Mankiw 1987, appendix).

Graphs of the first 104 impulse responses and their confidence intervals derived from ARFIMA models given in Table 2 are presented in Figure 2. The estimated coefficient variance–covariance matrix was used to compute the sample variance of the impulse response. The confidence intervals are constructed from the ± 2 standard errors of $C(k)$. To obtain a proper scale to present these graphs, we omitted confidence intervals of the first five impulse responses.

Table 2. Parameter Estimates of the Selected ARFIMA Models

	d	ϕ_1	ϕ_2	θ_1
	a. Time-Domain Exact ML Estimates			
BP	.1343	−.1341	−.1086	
	(.004)	(.014)	(.009)	
DM	.0455	—	—	—
	(.017)			
SF	.0659			−.0741
	(.076)			(.111)
JY	.0616	—	.1122	—
	(.029)		(.003)	
FF	.3664	.4776	.1227	−.8657
	(.002)	(.000)	(.006)	(.000)
	b. Frequency-Domain Approximate ML Estimates			
BP	.1334	−.1335	−.1079	—
	(.004)	(.013)	(.001)	
DM	.0464			
	(.047)			
SF	.0477	−.7555	—	.7016
	(.061)	(.000)		(.000)
JY	.0550	—	.1095	
	(.042)		(.002)	
FF	.4707	.3982	.1332	−.8905
	(.003)	(.002)	(.002)	(.000)

NOTE: Coefficient estimates of the ARFIMA model given by Equation (2) in the text are presented. p values are given in parentheses.

These calculated impulse responses show that the effects of a unit shock on future exchange-rate changes are similar across five exchange-rate series. The long memory and short memory interact in such a way that the persistence, as measured by $C(k)$, is usually small and less than .1. This is also true for the FF series, which has a relatively large d estimate. The slow decay of $C(k)$ demonstrates the long-memory effect. On the other hand, the confidence intervals of these estimated impulse responses have a lower bound very close to 0 and, usually, include 0; that is, these estimated impulse responses are, at most, only marginally significant.

2.4 Out-of-Sample Forecast

This subsection compares the out-of-sample forecasts from the estimated ARFIMA models and the random-walk model, which is commonly used as a benchmark for comparing the forecasting performance of exchange-rate models (Meese and Rogoff 1983). Forecasts of the ARFIMA model were generated from the truncated AR representation of (2). Specifically, at time t, the k-period-ahead forecast is given by

$$X_{t+k|t} = \sum_{i=1}^{k-1} \psi_i X_{t+k-i|t} + \sum_{i=k}^{t+k-1} \psi_i X_{t+k-i}, \qquad (11)$$

where $(1 - \psi_1 B - \psi_2 B - \cdots) \equiv \Theta^{-1}(B)\Phi(B)(1 - B)^d$ and $X_{t+k|t}$ can be interpreted as the sample analog of the optimal predictor discussed by Peiris and Perera (1988). These forecasts are ex ante because they are generated recursively using only information actually available. The forecasting period is 1988–1989. The forecasting horizons 1, 10, 20, 30, and 40 are considered.

The mean squared forecast error (MSFE) and the mean absolute forecast error (MAFE) are used to evaluate the forecasting performance. Table 3 reports the MSFE's and MAFE's of the ARFIMA models relative to those of the random-walk model. An entry smaller (larger) than 1 means that the ARFIMA model yields a smaller (larger) MSFE or MAFE. In general, the relative MSFE and MAFE are very close to 1. The relative MSFE's and MAFE's that are less than 1 are usually found in the forecasting horizons 30 and 40. The evidence of better performance in longer horizons, however, seems not consistent across currencies. The ARFIMA models do not improve on the random walk in out-of-sample forecasts. This result corroborates those of others using different techniques and different data sets, such as Meese and Rogoff (1983), who found that structural-exchange models do not outperform the random-walk model, and Diebold and Nason (1990),

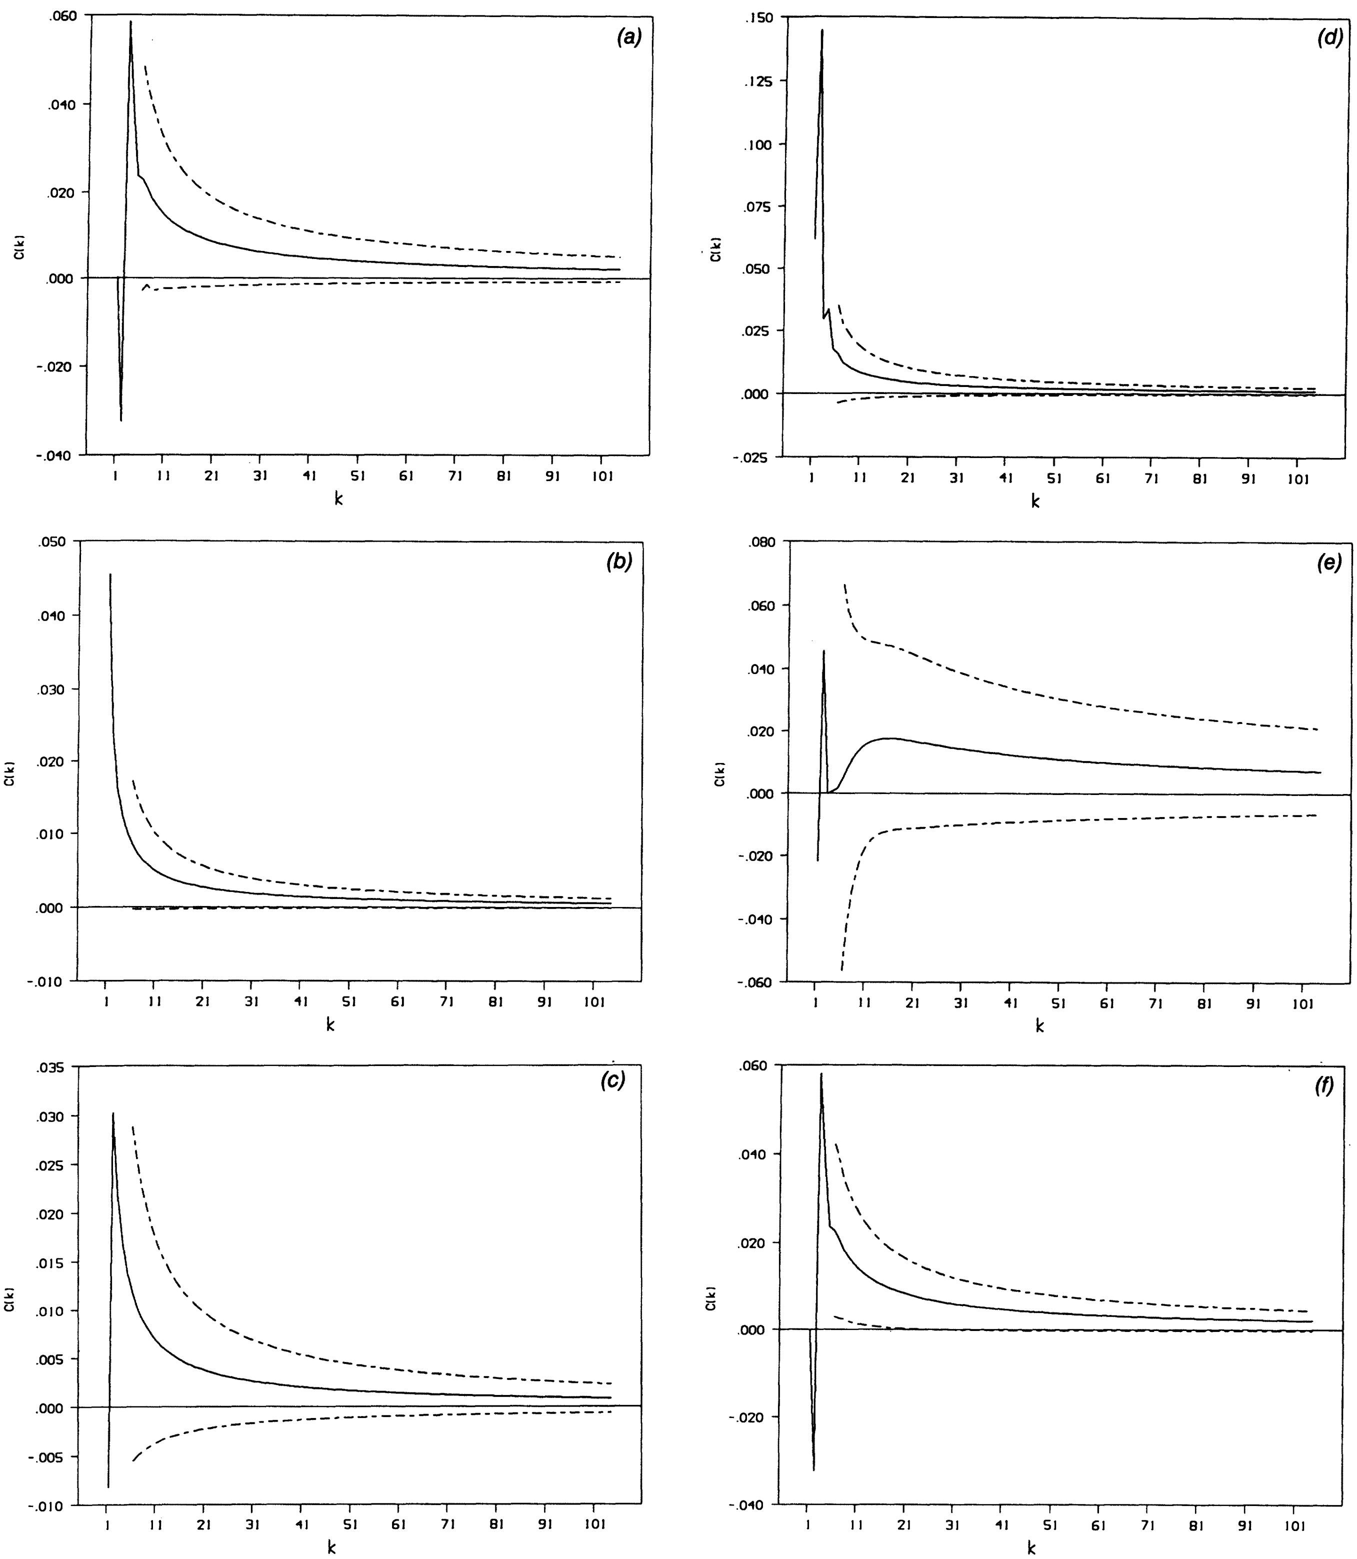

Figure 2. (continued)

who could not improve exchange-rate forecasts using nonparametric prediction techniques.

3. SOURCES OF LONG MEMORY

There are different possible sources of long memory in exchange rates. For instance, the long-memory behavior of exchange rates can be related to the dynamic properties of other economic variables. The purchasing power parity (PPP) hypothesis suggests that exchange-rate fluctuations are tied to the movements of relative national prices. Recent empirical evidence of PPP was reported by, for example, Abuaf and Jorion (1990), Cheung and Lai (in press), and Diebold, Hasted, and Rush (1991).

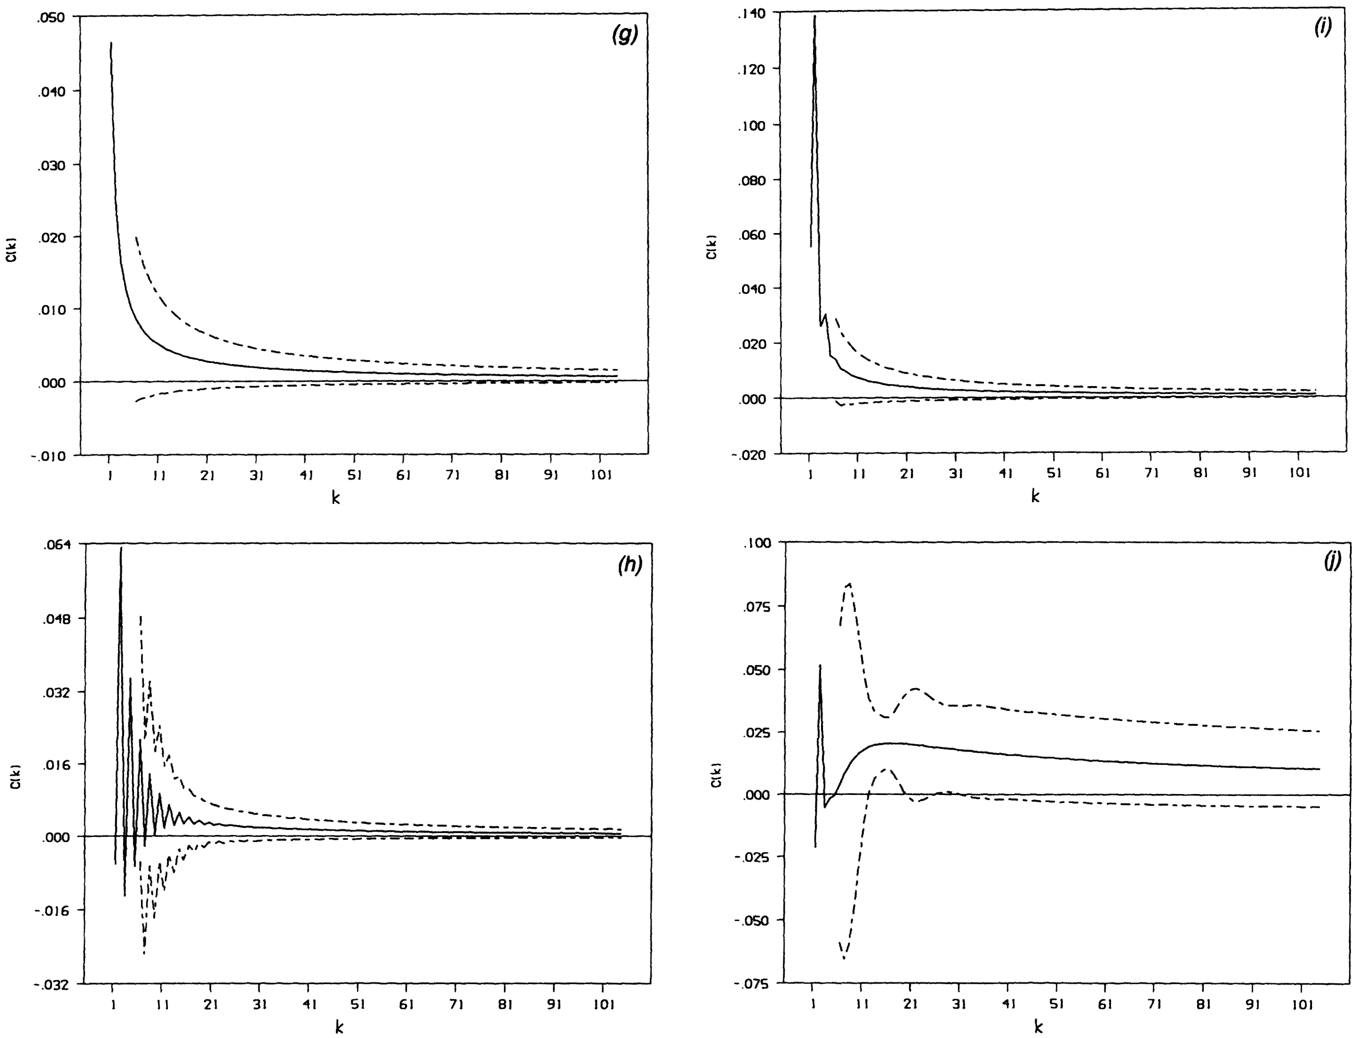

Figure 2. Impulse-Response Functions and Their Confidence Intervals. Figure 2a to 2e, presents results based on ARFIMA models reported in Table 2a and Figure 2f to j is based on ARFIMA models reported in Table 2b: a, MLBP; b, MLDM; c, MLSF; d, MLJY; e, MLFF; f, FDBP; g, FDDM; h, FDSF; i, FDJY; j, FDFF.

Table 4 reports the tests for long memory in the first differences of five relative national price series. The countries considered are the United States, the United Kingdom, West Germany, Switzerland, Japan, and France. The data used are monthly consumer price indexes taken from the International Financial Statistics data base. The relative price indexes are all against the U.S. index. The sample period covered is January 1974 to December 1987. The GPH test suggests that the five relative price series may display long memory, and the evidence is qualitatively the same across different choices of n. Note that the frequency of the price data used is different from that of the exchange rates studied. Nonetheless, the PPP hypothesis suggests that the finding of long memory in relative national price changes reinforces the evidence of long memory in exchange-rate changes.

Fundamentals may contribute to long memory found in the exchange-rate data as well. Standard exchange-rate determination models usually explain exchange-rate dynamics by movements of relative money supplies, relative outputs, and so forth. Recent empirical studies on the dynamic properties of the U.S. macroeconomic time series—such as output, consumption and money-supply data—showed that these macroeconomic series may be fractionally integrated (Diebold and Rudebusch 1989, 1991b; Haubrich 1989; Porter-Hudak 1990). When we examine the relative money supply, the relative output, and the relative interest rate of these six countries, however, we only find some evidence of long memory in the relative-money-supply series but not in the other series (results not reported here). Hence it is not clear to what extent the long memory found in the exchange-rate data can be attributed to these fundamentals. More work is required to establish the relationship of the dynamics of exchange rates and fundamentals.

4. SUMMARY

The time series properties of five major nominal exchange-rate series are examined with the general

Table 3. Out-of-Sample Forecast Analysis

	Forecast horizon	BP	DM	SF	JY	FF
Relative MSFE	1	1.0044	1.0041	1.0020	1.0510	1.0206
		1.0043	1.0043	1.0288	1.0464	1.0277
	10	1.0020	1.0028	1.0055	1.0058	1.0196
		1.0020	1.0029	1.0034	1.0051	1.0244
	20	1.0049	1.0038	1.0053	1.0050	1.0210
		1.0049	1.0031	1.0039	1.0044	1.0267
	30	.9961	.9992	.9998	1.0007	1.0045
		.9961	.9992	.9998	1.0006	1.0029
	40	1.0007	.9991	1.0002	1.0019	.9989
		1.0007	.9991	1.0001	1.0017	.9987
Relative MASFE	1	.9921	.9948	1.0021	1.0258	1.0132
		.9922	1.0047	1.0151	1.0237	1.0175
	10	1.0011	1.0021	1.0044	1.0060	1.0122
		1.0011	1.0021	1.0030	1.0055	1.0152
	20	1.0028	1.0017	1.0035	1.0056	1.0159
		1.0029	1.0018	1.0026	1.0050	1.0197
	30	1.0001	.9996	1.0004	1.0040	1.0039
		1.0001	.9996	1.0003	1.0036	1.0050
	40	1.0009	.9995	1.0004	1.0035	.9949
		1.0009	.9994	1.0003	1.0032	.9962

NOTE: Each cell reports the MSFE's (MAFE's) of an ARFIMA model relative to those of a random-walk model. The first entry in each cell is the relative MSFE (MAFE) of ARFIMA models with coefficients estimated using the time-domain exact ML method, while the second is that of models with coefficients estimated using the frequency-domain appropriate ML method. Forecasts used to generate MSFE and MASE numbers are based on parameters reported in Table 2.

ARFIMA model. The ARFIMA model, in which the differencing parameter d can assume noninteger values, provides a direct and convenient framework to study both short- and long-memory behavior. The ARFIMA model also suggests interesting long-memory alternatives to the unit-root hypothesis.

We find evidence of long memory in exchange-rate change data. The test result is asymptotically robust to both the distribution of the innovation term and the short-run dynamics in the data. This implies that the empirical evidence of unit roots in exchange-rate data is not robust to long-memory alternatives. When ARFIMA models are fitted to the data, it is found that exchange-rate dynamics are more complex than that implied by a random walk. Nonetheless, impulse-response function analysis indicates that persistence in exchange-rate changes can be difficult to detect. Furthermore, the ARFIMA model does not outperform a random walk in out-of-sample forecast.

Recent studies on the dynamic properties of macroeconomic time series indicate that the long memory found in exchange-rate data may not be an isolated incident. For instance, we found evidence of long memory in some relative national price indexes, which are closely related to exchange rates under the PPP hypothesis. More work is required, however, to establish the relationship of the dynamics of exchange rates and fundamentals.

Table 4. Tests for Long Memory in Relative Price Indexes

GPH	US/UK	US/GE	US/SW	US/JA	US/FR
$D(.45)$	.4595*	.8589***	.5928**	.9852***	.5026**
p value	(.058)	(.002)	(.021)	(.000)	(.042)
$D(.50)$	.5468**	.7971***	.5009**	.9684***	.3438*
p value	(.017)	(.001)	(.026)	(.000)	(.090)
$D(.55)$	.3035*	.3889**	.3802**	.5774***	.1206
p value	(.075)	(.033)	(.036)	(.003)	(.283)

NOTE: $D(.45)$, $D(.50)$ and $D(.55)$ give the d estimates corresponding to $n = T^{.45}$, $T^{.50}$, and $T^{.55}$. p value gives the p value of the no-long-memory null hypothesis ($d = 0$) against the long-memory alternative. *, **, and *** indicate significance at 10%, 5%, and .5%, respectively. The columns US/UK, US/GE, US/SW, US/JA, and US/FR give the statistics computed from the first differences of the consumer price indexes of the United Kingdom, Germany, Switzerland, Japan, and France relative to that of the United States.

ACKNOWLEDGMENTS

I thank Geoffrey Booth, Francis Diebold, Kevin Hassett, Kon-Sun Lai, Roberto Mariano, Marc Nerlove, Peter Pauly, Mark Rush, and Fallaw Sowell for helpful comments and suggestions. The article also benefited from the constructive comments and suggestions from an associate editor, two anonymous referees, and John Geweke. This research was supported by the GICES.

[Received June 1991. Revised July 1992.]

REFERENCES

Abuaf, N., and Jorion, P. (1990), "Purchasing Power Parity in the Long Run," *Journal of Finance*, 45, 157–174.

Aydogan, K., and Booth, G. G. (1988), "Are There Long Cycles in Common Stock Returns," *Southern Economic Journal*, 55, 141–149.

Baillie, R. T., and Bollerslev, T. (1989), "The Message in Daily Exchange Rates: A Conditional-Variance Tale," *Journal of Business & Economic Statistics*, 7, 297–305.

Beran, J. (1986), "Estimation, Testing and Prediction for Self-Similar and Related Processes," unpublished dissertation, Swiss Federal Institute of Technology.

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

Campbell, J. Y., and Mankiw, N. G. (1987), "Are Output Fluctuations Transitory?" *Quarterly Journal of Economics*, 102, 857–880.

Cheung, Y.-W., and Lai, K. S. (in press), "Long-Run Purchasing Power Parity During the Recent Float," *Journal of International Economics*, 30.

Davies, R. B., and Harte, D. S. (1987), "Tests for Hurst Effect," *Biometrika*, 95–101.

Diebold, F. X., and Nason, J. (1990), "Nonparametric Exchange Rate Prediction," *Journal of International Economics*, 28, 315–332.

Diebold, F. X., and Rudebusch, G. D. (1989), "Long Memory and Persistence in Aggregate Output," *Journal of Monetary Economics*, 24, 189–209.

——— (1991a), "On the Power of Dickey–Fuller Tests Against Fractional Alternatives," *Economics Letters*, 35, 155–160.

——— (1991b), "Is Consumption Too Smooth? Long Memory and the Deaton Paradox," *Review of Economics and Statistics*, 73, 1–17.

Diebold, F. X., Husted, S., and Rush, M. (1991), "Real Exchange Rates Under the Gold Standard," *Journal of Political Economy*, 99, 1252–1271.

Fox, R., and Taqqu, M. S. (1986), "Large-Sample Properties of Parameter Estimates for Strongly Dependent Stationary Gaussian Time Series," *The Annals of Statistics*, 14, 517–532.

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

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

Haubrich, J. G. (1989), "Consumption and Fractional Differencing: Old and New Anomalies," unpublished manuscript, Wharton School, University of Pennsylvania.

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

——— (1984), "Modeling Persistence in Hydrological Time Series Using Fractional Differencing," *Water Resources Research*, 20, 1898–1908.

Li, W. K., and McLeod, A. I. (1986), "Fractional Time Series Modelling," *Biometrika*, 73, 217–221.

Lo, A. W. (1991), "Long-Term Memory in Stock Market Prices," *Econometrica*, 59, 1279–1314.

McLeod, A. I., and Hipel, K. W. (1978), "Preservation of the Rescaled Adjusted Range 1. A Reassessment of the Hurst Phenomenon," *Water Resources Research*, 14, 491–508.

Meese, R. A., and Rogoff, K. (1983), "Empirical Exchange Rate Models of the Seventies: Do They Fit Out of Sample? " *Journal of International Economics*, 14, 3–24.

Meese, R., and Singleton, K. (1982), "On Unit Roots and the Empirical Modeling of Exchange Rates," *Journal of Finance*, 37, 1029–1035.

Peiris, M. S., and Perera, B. J. C. (1988), "On Prediction With Fractionally Differenced ARIMA Models," *Journal of Time Series Analysis*, 9, 215–220.

Porter-Hudak, S. (1990), "An Application of the Seasonal Fractionally Differenced Model to the Monetary Aggregates," *Journal of the American Statistical Association*, 85, 338–344.

Sowell, F. (1990a), "Maximum Likelihood Estimation of Stationary Univariate Fractionally Integrated Time Series Models," unpublished manuscript, Carnegie Mellon University, Dept. of Economics.

——— (1990b), "Fractional Unit Root Distribution," *Econometrica*, 50, 495–505.

Watson, M. W. (1986), "Univariate Detrending Methods With Stochastic Trends," *Journal of Monetary Economics*, 18, 49–75.

Yajima, Y. (1985), "On Estimation of Long-Memory Time Series Models," *Australian Journal of Statistics*, 27, 303–320.

——— (1988), "On Estimation of a Regression Model With Long-Memory Stationary Errors," *The Annals of Statistics*, 16, 791–807.

——— (1989), "A Central Limit Theorem of Fourier Transforms of Strongly Dependent Stationary Processes," *Journal of Time Series Analysis*, 10, 373–383.