Journal of International Economics 52 (2000) 321–330
Journal of INTERNATIONAL ECONOMICS
www.elsevier.nl/locate/econbase

On the purchasing power parity puzzle

Yin-Wong Cheung[a], Kon S. Lai[b,*]

[a]*University of California, Santa Cruz, CA 95064, USA*
[b]*Department of Economics, California State University, Los Angeles, CA 90032, USA*

Received 26 October 1998; received in revised form 18 March 1999; accepted 22 June 1999

Abstract

A puzzle concerning purchasing power parity is examined: Although the immense exchange rate volatility suggests a likely major role of nominal shocks under sticky prices, the observed half-life persistence of the real exchange rate seems excessively high to be rationalized by price stickiness. This study analyzes carefully the adjustment dynamics of real exchange rates through impulse response analysis. Half-life estimates are found to have substantial imprecision. Moreover, the dynamic response pattern suggests that the shock response is initially amplified before dissipating and that such non-monotonic dynamics can contribute to more than one-third of the observed persistence of real exchange rates.

Keywords: Real exchange rate; Half-life persistence; Shock amplification; Non-monotonic mean reversion

JEL classification: F31; F41

---

1. Introduction

The purchasing power parity (PPP) theory has had its ebbs and flows over the years. Enormous interest in the theory has emerged since the advent of flexible exchange rates in the early 1970s. The recent floating experience has not been too reassuring, however. Early studies generally fail to uncover parity reversion. With

---

*Corresponding author. Tel.: +1-323-343-2930; fax: +1-323-343-5462.
E-mail address: klai@calstatela.edu (K.S. Lai).

PII: S0022-1996(99)00056-2

the aid of improved statistical methods, some recent studies are able to unveil PPP reversion (Frankel and Rose, 1996; Wu, 1996; Papell, 1997; Cheung and Lai, 1998; Taylor and Sarno, 1998). The speed of parity reversion seems torpid, nonetheless, with reported half-life estimates mostly between 3 and 5 years (Rogoff, 1996). This half-life range implies a slow reversion rate of about 13 to 20 percent per year.

Rogoff (1996) points out the difficulty in reconciling the perplexedly high persistence of real exchange rates with their immense short-term volatility. Although slow reversion can be rationalized if real shocks are predominant, real shocks are not volatile enough over the short term to account for the vast exchange rate volatility. On the other hand, short-term exchange rate volatility can be caused by monetary shocks under sticky prices, but the estimated half-lives of PPP reversion reported in prior studies seem far too long to be explained by price stickiness. Specifically, if nominal stickiness is really responsible for short-run PPP deviations, ‘‘one would expect substantial convergence to PPP over one to two years, as wages and prices adjust to a shock’’ (Rogoff, 1996, p. 654). This poses a puzzle. No existing model seems able to consistently explain both the tremendous short-term volatility and the ‘‘excessive’’ persistence in the real exchange rate. Clarida and Gali (1994) and Rogers (1999) identify the relevance of multiple shocks in explaining the variability of real exchange rates, but their results still do not fully resolve the PPP puzzle. Instead of examining the nature of shocks, the present study explores the dynamic structure of the parity-reverting process itself and identifies its role in explaining the high persistence of the real exchange rate.

Using impulse response analysis, this study analyzes the adjustment dynamics of real exchange rates by evaluating both the sample half-life measure and its estimation accuracy. Since reporting merely point estimates does not convey the inevitable imprecision with which the adjustment speed is measured, confidence interval estimates are computed. Empirical results show that the confidence intervals for half-life estimates are generally wide, suggesting a high level of imprecision in half-life estimation. The results underline the potential deficiency in relying on point estimates of half-lives and highlight the significant uncertainty in measuring the speed of parity reversion.[1]

The impulse response analysis also shows that the shock impact tends to amplify first before it dissipates. The full impact of a shock is not felt immediately but until a few periods after the initial shock. Hence, following the shock, the real

---

[1] Instead of analyzing half-life estimates for countries individually, Cheung and Lai (1999) investigate whether systematic differences in half-lives exist across country groups using nonparametric tests of group medians, which require no assumptions on the precision of individual estimates in each group. While this recent study identifies considerable cross-sectional variability in the reversion speed, the study here examines the sampling variability of the speed estimates from individual time series themselves.

exchange rate does not revert to its long-run value monotonically. The nonmonotonic adjustment is found to contribute considerably to the persistence of the real exchange rate.

2. Measuring the PPP reversion speed

The data under study are monthly real exchange rates constructed from nominal exchange rates and consumer price indices. Specifically, the real exchange rates of four European countries – France (FR), Germany (GE), Italy (IT), and the United Kingdom (UK) – vis-à-vis the United States (US) are investigated. Taken from the International Monetary Fund's *International Financial Statistics* data CD-ROM, the data cover the sample period from April 1973 through December 1996. Following the common practice in previous studies, all the series of real exchange rates are expressed in logarithms. All the four series of real exchange rates have been found to be stationary based on an efficient unit root test (Cheung and Lai, 1999).[2] Although this supports the validity of long-run PPP, it offers little information about the speed at which deviations from PPP die out over the short or medium run. To obtain such information, explicit computation of persistence is needed and the half-life measure has often been used to quantify persistence.[3]

The computation of half-lives can be formally derived from impulse response analysis. Given that the real exchange rate is found to be stationary, its dynamics can be captured in general by an autoregressive moving average (ARMA) model as follows:

$$B(L)y_t = D(L)u_t \tag{1}$$

where L is the lag operator, all roots of $B(L)$ and $D(L)$ are stable, and u_t is the

---

[2] The French franc, German mark, Italian lira and British pound are major European currencies frequently examined in PPP studies. A question may arise concerning whether the sample selection biases our results toward finding faster reversion. Of 20 real European exchange rates examined by Cheung and Lai (1999), the unit-root hypothesis can be rejected in 8 cases, with 6 of them using an efficient test and 2 other cases using fractional analysis. Nevertheless, the half-life estimates for all the European rates are rather similar; they are mostly about 3 to 4 years, matching what we find from our data here. On the other hand, the half-life estimates for European rates are found to be generally higher when compared with those estimates for real exchange rates of other geographic regions (Africa, Asia, North America, South America and Oceania). In view of these cross-sectional findings, the use of the European data here should not bias our results in favor of finding shorter half-lives.

[3] Engel (1999) notes that the equilibrium real exchange rate may not be a constant, albeit the real rate contains a transitory part converging to PPP. Indeed, the real rate may follow a slow-moving random walk (i.e., a random walk with small innovation variance), which usual unit root tests do not rule out. It remains instructive and interesting to inquire about the reversion speed, even if the equilibrium is not constant. Given the relatively short sample period for the current float, unlike long historical data, a constant equilibrium rate may not be a bad approximation to the underlying data generating process.

random error. The persistence of the process over different horizons can be analyzed by studying the moving average representation for y_t:

$$y_t = C(L)u_t \tag{2}$$

where $C(L)=1+C_1L+C_2L^2+\dots$, obtained from $C(L)=B^{-1}(L)D(L)$. The moving average coefficients, $(C_1, C_2, \dots)$, are referred to as impulse responses. In general, C_j tracks the impact of a unit shock at time t on the level of y at time $t+j$. For a stationary process, which contains no unit root, $C_\infty=0$; the process thus has zero long-run persistence. Over time horizons much shorter than infinity, nonetheless, $C_j \neq 0$ and sizable persistence may still exist.

Instead of studying the entire sequence of C_j, $j=1, 2, \dots$, a summary measure of persistence often employed in the PPP literature is the half-life, which indicates how long it takes for the impact of a unit shock on the real exchange rate to dissipate by half. By definition, the half-life, denoted by l_h, gives that $C_j=1/2$ for $j=l_h$. The l_h value has sometimes been used to indirectly infer the underlying reversion speed. If a constant speed of adjustment is assumed, the speed can be computed from l_h as $1-\exp[\ln(1/2)/l_h]$. Such an assumption may not be proper, however, when the reversion process is characterized by non-monotonic adjustments.

The precision with which l_h can be measured hinges upon the sampling properties of the impulse response estimate. The sampling variability of C_j can be assessed by their estimated standard errors. For a given horizon, j, C_j is a nonlinear function of the ARMA parameter vector, denoted by ζ, which contains parameters in $B(L)$ and $D(L)$. Using the delta method, asymptotic standard errors can be estimated from

$$\mathrm{Var}(C_j) = \nabla C_j' \, \Omega \, \nabla C_j \tag{3}$$

where ∇C_j is $\partial C_j/\partial\zeta$ and Ω is the variance–covariance matrix of ζ (Campbell and Mankiw, 1987).

3. Empirical results

Graphs of the first 108 impulse responses – which correspond to a time span of 9 years for monthly data – and their confidence intervals are displayed in Fig. 1. The model specifications reported in Table 1 are selected based on the Akaike information criterion, and they are found to pass residual tests for serial correlation. Both 90% and 95% confidence bands are constructed using the delta method. Models of more parsimonious specifications (e.g., AR(1) models) had been tried, but they all failed to pass residual tests.

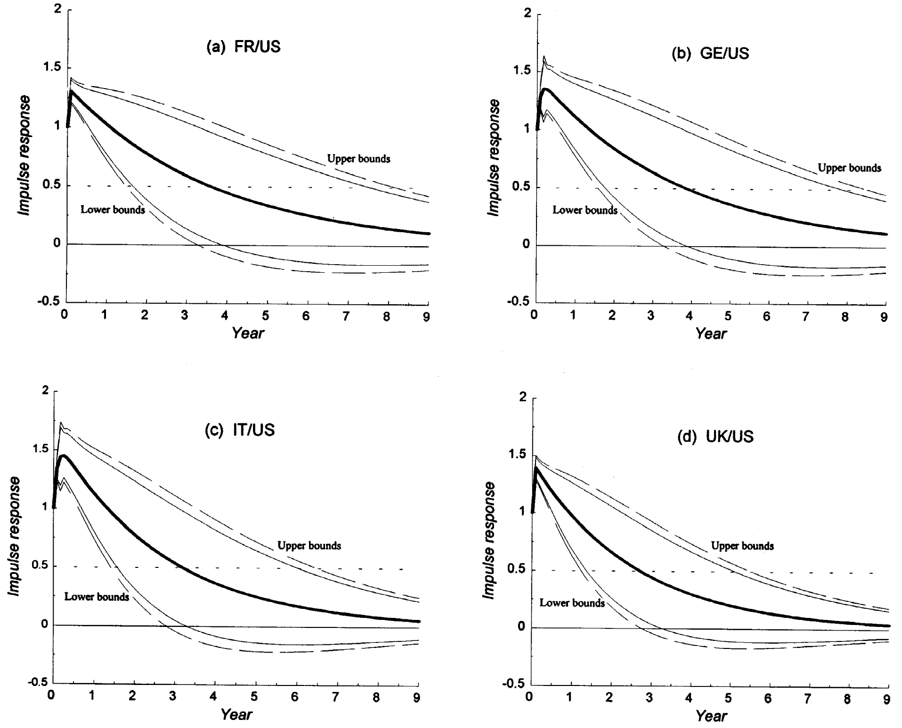

Fig. 1. Adjustment dynamics of real exchange rates. The thick solid line indicates the point estimates of impulse responses to a unit shock. The two thin solid lines give the corresponding 90% confidence band. The two broken lines give the 95% confidence band.

3.1. The uncertain half-life

To quantify explicitly the speed of parity reversion, the half-life estimates are computed, together with their 90% and 95% confidence intervals. In addition to the delta method, bootstrapping will be used to check the robustness of the confidence interval estimation in finite samples. Without resorting to asymptotic normal approximation, bootstrapping has been known to be a powerful tool for approximating the sampling distribution and variance of general statistics. In this study, parametric residual-based bootstrapping will be performed (Efron and Tibshirani, 1993). Using resampling (with replacement) techniques, the unknown distribution of the innovation term is approximated by the empirical distribution of the estimated residuals, through which the confidence regions of the estimated parameter can be computed. The confidence intervals here are obtained through bootstrapping with 10 000 replications. In contrast to the delta method, the bootstrap method can account for possible finite-sample bias and is robust to

Table 1
Specifications and parameters of estimated models[a]

Series	Lag selection		Model coefficient estimates			Ljung–Box Q-statistic		J–B	$ARCH_{10}$	$ARCH_{20}$
	p	q	b_1	b_2	d_1	Q_{10}	Q_{20}			
FR/US	1	1	0.977		0.330	7.80	14.56	22.38	6.39	12.52
			(0.129)**		(0.057)**	[0.35]	[0.63]	[0.00]	[0.78]	[0.90]
GE/US	2	0	1.287	−0.303		5.87	11.54	0.96	5.74	13.98
			(0.057)**	(0.058)**		[0.55]	[0.83]	[0.62]	[0.84]	[0.83]
IT/US	2	0	1.345	−0.364		4.70	11.69	11.95	8.90	13.79
			(0.056)**	(0.056)**		[0.70]	[0.82]	[0.00]	[0.54]	[0.84]
UK/US	1	1	0.968		0.425	4.63	17.98	11.59	19.84	25.86
			(0.015)**		(0.055)**	[0.71]	[0.28]	[0.00]	[0.03]	[0.17]

[a] The ARMA model, $B(L)y_t = D(L)u_t$ with $B(L) = 1 - b_1L - \cdots - b_pL^p$ and $D(L) = 1 + d_1L + \cdots + d_qL^q$, is fitted to the real exchange rate data. The columns beneath "p" and "q" give the AR and MA lag parameters chosen using the Akaike information criterion. The numbers in parentheses report the standard errors for the corresponding model coefficient estimates. Statistical significance is indicated by a double asterisk (**) for the 5% level. The columns beneath "Q_{10}" and "Q_{20}" present the usual Ljung–Box statistics for up to 10th and 20th order serial correlation, respectively. The column beneath "J–B" gives the Bera–Jarque statistic for a normality test of the residuals. The columns beneath "$ARCH_{10}$" and "$ARCH_{20}$" give the statistics which test for ARCH effects based on the correlation of the squared residuals up to 10th and 20th lag order, respectively. The numbers in brackets indicate the corresponding p-values of individual statistics.

non-normal innovations. The latter property is particularly desirable because diagnostic testing of the innovations in real exchange rate processes confirms the presence of serious departures from normality (significant deviations from normality can be detected in the FR, IT and UK data based on the standard Bera–Jarque test).

The half-life estimates for the various real exchange rate series are reported in Table 2. The average of all the l_h estimates is approximately 3.3 years, suggesting a reversion rate of about 19 percent per year. These half-life persistence estimates

Table 2
Half-life persistence estimates and their confidence intervals[a]

Series	l_h	Confidence intervals (the delta method)				Confidence intervals (the bootstrap method)			
		L_{95}	L_{90}	U_{90}	U_{95}	L_{95}	L_{90}	U_{90}	U_{95}
FR/US	3.60	1.46	1.64	7.55	8.20	1.20	1.40	8.14	8.72
GE/US	3.86	1.57	1.76	7.88	8.52	1.29	1.49	8.54	9.34
IT/US	3.18	1.39	1.55	6.02	6.48	1.23	1.46	6.38	7.38
UK/US	2.70	1.23	1.36	5.11	5.53	1.06	1.21	5.53	6.44

[a] The column "l_h" gives the point estimates of the half-life persistence (in years). (L_{90}, U_{90}) represents the 90% confidence interval for l_h; whereas, (L_{95}, U_{95}) indicates the 95% confidence interval for l_h.

are largely comparable to those from previous studies for dollar-based real exchange rates. For example, Abuaf and Jorion (1990) reported an average half-life of 3.3 years for eight series of real exchange rates. Lothian and Taylor (1996) estimated the half-life for the dollar-pound rate to be 4.7 years. In examining pooled data on real rates of a group of currencies, Frankel and Rose (1996) produced an estimate half-life of about 4 years. Using similar analysis, Wu (1996) found the half-life to be roughly 2.5 years, whereas Wei and Parsley (1995) obtained half-life estimates around 4.5 years. It should be noted that most of these half-life estimates from the earlier studies were obtained from long-horizon data. More recently, comparable half-lives were reported by Cheung and Lai (1999) based on the current float data. For industrial countries – which include many European countries – the median of their half-life estimates is about 3.3 years. On the other hand, this recent study also uncovered different persistence patterns between industrial countries and developing countries: The half-life estimates for developing countries tend to be much less persistent, with a median value of about 1.4 years.

Table 2 further shows that the confidence intervals for our half-life estimates are uniformly wide. In particular, the bootstrap confidence intervals are all wider than the asymptotic intervals. The bootstrap confidence intervals properly incorporate additional sampling variability due to the effects of both finite-sample correction and non-normal innovations. The statistical results confirm that the asymptotic intervals understate the actual amount of sampling variability. In general, the wide confidence intervals underscore a high level of imprecision with which half-lives are estimated. The usual emphasis on the point estimates of the half-life, therefore, fails to reveal the substantial sampling uncertainty associated with this measure of persistence. Interestingly, the lower bounds of both the asymptotic or bootstrap confidence intervals include a range of short half-lives, which can be much less than 2 years.

3.2. Non-monotonic parity reversion

All the graphs in Fig. 1 share a notable feature. The impulse responses are not a monotonic function of the adjustment horizon. They are all hump-shaped, with initial shock amplification, followed shortly by shock dissipation. The eventual shock dissipation, as reflected by the decay of C_j toward zero, confirms the existence of parity reversion. In the initial phase, however, the shock impact magnifies rather than diminishes, with the maximum C_j value occurring within a short time period after the shock. According to the sequence of C_j obtained, the amplification of the shock response generally lasts for 1 to 3 months only. Although this phase is short and reverse adjustments occur right afterward, it takes more than 1 year for the impulse response to just return back to unity. Consequently, the non-monotonicity in the initial response significantly prolongs the adjustment process.

The initial amplification of the shock response, albeit over a rather short time, has implications for interpreting the half-life measure. For a given rate of decay of PPP deviations, the half-life persistence for the real exchange rate can vary widely, depending upon the magnitude of the amplification. A large amplified response can yield a long half-life and create the appearance of slow adjustment even when PPP deviations are corrected at a relatively fast speed. The impulse response estimation results indicate the presence of shock amplification by a factor of about 1.4 on average. We calculate that the short-term shock amplification alone can account for more than a third of the observed persistence in the real exchange rate.

To illustrate further the effects of the non-monotonicity in the shock response on the persistence estimates, Table 3 reports the half-life estimates of the adjustment speed after the initial amplification, i.e., after C_j has reached a maximum value subsequent to the unit shock. As compared with the l_h estimates given in Table 2, the modified estimates reported in Table 3 are much shorter, mostly around two years long. The differences in the speed estimates confirm that the initial amplification adds substantially to the half-life persistence of the real exchange rate. Once the impact of amplification peaks, there appears to be fairly fast mean reversion. Moreover, the confidence intervals, though tighter than all those in Table 2, remain wide. The upper bounds of these confidence intervals still contain long half-lives (about four to six years), which are larger than the usual half-life estimates. On the other hand, their lower bounds can reach as low as less than one year long, which implies a decay rate of 50% per year or faster. In the UK case, for example, the bootstrap 95% confidence interval has a lower bound of 0.67 year and an upper bound of 4.3 years.

The adjustment speed (AS) at time $t=\tau$ can be directly computed as $AS_\tau = -[\mathrm{d}C_t/\mathrm{d}t]_{t=\tau}/C_\tau$, which gives the rate of decrease in the impulse response at time $t=\tau$. The real exchange rate reverts toward parity as long as $AS_\tau > 0$. For discrete-time data, AS_τ is approximated by $-(C_{\tau+1} - C_\tau)/C_\tau$. In each case, the adjustment speed starts from a negative value, reflecting the initial shock

Table 3
Adjustment speeds after the initial shock amplification[a]

Series	l_h	Confidence intervals (the delta method)				Confidence intervals (the bootstrap method)				AS
		L_{95}	L_{90}	U_{90}	U_{95}	L_{95}	L_{90}	U_{90}	U_{95}	
FR/US	2.50	1.15	1.27	5.56	6.14	0.82	0.97	5.79	6.27	27.1%
GE/US	2.60	1.29	1.42	5.11	5.58	0.85	0.98	5.74	6.01	27.5%
IT/US	1.92	1.11	1.13	3.70	4.02	0.76	0.87	3.93	4.66	36.6%
UK/US	1.77	0.89	0.97	3.50	3.86	0.67	0.77	3.65	4.30	38.5%

[a] The column "l_h" gives the point estimates of the half-life adjustment speed (in years). (L_{90}, U_{90}) represents the 90% confidence interval for l_h; whereas, (L_{95}, U_{95}) indicates the 95% confidence interval for l_h. The column "AS" gives the adjustment speed estimates (calculated as $-[\mathrm{d}C_t/\mathrm{d}t]/C_t$ from the impulse response function) in terms of an average decay rate in percentage per year.

amplification. Reverting dynamics then quickly take over and reach a steady positive speed: PPP deviations can die out at a rate of between 27% to 38% per year, much faster than what previous half-life estimates have typically implied.

4. Concluding remarks

The adjustment dynamics of real exchange rates in response to shocks to parity have been analyzed using impulse response analysis. The half-life persistence estimates are found to have a high level of imprecision. Great caution should thus be taken in making model inferences based on the point estimates alone. Moreover, the convergence toward PPP is found to be not monotonic. The non-monotonic dynamics can substantially prolong the adjustment process and augment the persistence of the real exchange rate.

It should be noted that the non-monotonic response is not consistent with the monotonic price-adjustment behavior usually assumed in Dornbusch-type sticky-price models (Dornbusch, 1976). Consider a typical price-adjustment equation (with the log of the foreign price level being normalized to be zero):

$$p_{t+1} = p_t - \theta(p_t - s_t) + E_t(s_{t+1}) - s_t, \quad 0 < \theta < 1 \tag{4}$$

where p_t is the log of the domestic price level and s_t is the log of the exchange rate. If expectations are rational, so that $s_{t+1} = E_t(s_{t+1}) + \epsilon_{t+1}$, where ϵ_{t+1} is white noise, the equation can be rewritten as: $y_{t+1} = (1-\theta)y_t + \epsilon_{t+1}$, with $y_t = s_t - p_t$ being the real exchange rate. This real exchange rate process implied by standard sticky-price models suggests a monotonic impulse response function, which is incompatible with the findings here (the authors owe this observation to Charles Engel). This raises a new question: Where does the observed non-monotonicity in the adjustment process come from? Future research to identify the source of the non-monotonicity should be of interest for a better understanding of the real exchange rate behavior.

Acknowledgements

The authors are grateful to Janice Breuer, Menzie Chinn, Casper de Vries, Clive Granger, Ulrich Muller, Hashem Pesaran and two anonymous referees for useful comments and suggestions. This paper has benefited greatly from the extensive comments provided by Charles Engel, a co-editor of this journal. Any remaining errors are ours.

References

Abuaf, N., Jorion, P., 1990. Purchasing power parity in the long run. Journal of Finance 45, 157–174.

Campbell, J.Y., Mankiw, N.G., 1987. Are output fluctuations transitory? Quarterly Journal of Economics 102, 857–880.

Cheung, Y.W., Lai, K.S., 1998. Parity reversion in real exchange rates during the post-Bretton Woods period. Journal of International Money and Finance 17, 597–614.

Cheung, Y.W., Lai, K.S., 1999. On cross-country differences in the persistence of real exchange rates. Journal of International Economics 48, forthcoming.

Clarida, R., Gali, J., 1994. Sources of real exchange rate fluctuations: how important are nominal shock? Carnegie-Rochester Conference Series on Public Policy 41, 1–56.

Dornbusch, R., 1976. Expectations and exchange rate dynamics. Journal of Political Economy 84, 1161–1176.

Efron, B., Tibshirani, R.J., 1993. An Introduction To the Bootstrap. Chapman and Hall, New York.

Engel, C., 1999. Long-run PPP may not hold after all. Journal of International Economics, forthcoming.

Frankel, J.A., Rose, A.K., 1996. A panel project on purchasing power parity: mean reversion within and between countries. Journal of International Economics 40, 209–224.

Lothian, J.R., Taylor, M.P., 1996. Real exchange rate behavior: The recent float from the perspective of the past two centuries. Journal of Political Economy 104, 488–509.

Papell, D.H., 1997. Searching for stationarity: purchasing power parity under the current float. Journal of International Economics 43, 313–332.

Rogers, J.H., 1999. Monetary shocks and real exchange rates. Journal of International Economics 49, 269–288.

Rogoff, K., 1996. The purchasing power parity puzzle. Journal of Economic Literature 34, 647–668.

Taylor, M.P., Sarno, L., 1998. The behavior of real exchange rates during the post-Bretton Woods period. Journal of International Economics 46, 28–312.

Wei, S.-J., Parsley, D.C., 1995. Purchasing power disparity during the floating rate period: Exchange rate volatility, trade barriers and other culprits. NBER Working Paper No. 5032.

Wu, Y., 1996. Are real exchange rates nonstationary? Evidence from a panel-data test. Journal of Money, Credit, and Banking 28, 54–63.