Journal of International Money and Finance 51 (2015) 285–302

Contents lists available at ScienceDirect

Journal of International Money and Finance

journal homepage: www.elsevier.com/locate/jimf

Bayesian analysis of nonlinear exchange rate dynamics and the purchasing power parity persistence puzzle[☆]

Ming Chien Lo [a,1], James Morley [b,*]

[a] Saint Cloud State University, USA
[b] School of Economics, UNSW Business School, University of New South Wales, Sydney, NSW 2052, USA

ARTICLE INFO

Article history:
Available online 23 December 2014

JEL classification:
C11
C22
F31

Keywords:
Bayesian analysis
Real exchange rate dynamics
Purchasing power parity
Nonlinear threshold models
Bayesian model averaging
Half lives

ABSTRACT

We investigate the persistence of real exchange rates using Bayesian methods. First, an algorithm for Bayesian estimation of nonlinear threshold models is developed. Unlike standard grid-based estimation, the Bayesian approach fully captures joint parameter uncertainty and uncertainty about complicated functions of the parameters, such as the half-life measure of persistence based on generalized impulse response functions. Second, model comparison is conducted via marginal likelihoods, which reflect the relative abilities of models to predict the data given prior beliefs about model parameters. This comparison is conducted for a range of linear and nonlinear models and provides a direct evaluation of the importance of nonlinear dynamics in modeling exchange rates. The marginal likelihoods also imply weights for a model-averaged measure of persistence. The empirical results for real exchange rate data from the G7 countries suggest general support for nonlinearity, with the strength of the evidence depending on which country pair is being considered. However, the model-averaged estimates of half-lives are almost

☆ We thank an anonymous referee, as well as conference and seminar participants at the 5th INFINITI conference held at Trinity College, Dublin, the 16th Annual Symposium of the Society for Nonlinear Dynamics and Econometrics at the Federal Reserve Bank of San Francisco, the 18th Midwest Econometric Group conference at the University of Kansas, and Minnesota State University at Mankato, for helpful comments. The usual disclaimers apply.

* Corresponding author. Tel.: +61 2 9385 3366; fax: +61 2 9313 6337.
E-mail addresses: mclo@stcloudstate.edu (M.C. Lo), james.morley@unsw.edu.au (J. Morley).
1 Department of Economics, Stewart Hall 380, St. Cloud State University, St. Cloud, MN 56301-4498, USA. Tel.: +1 320 308 0143; fax: +1 320 308 2228.

http://dx.doi.org/10.1016/j.jimonfin.2014.12.003
0261-5606/

always as small or smaller than for the linear models alone, suggesting that the purchasing power parity persistence puzzle is less of a puzzle than previously thought.

1. Introduction

Numerous studies, including Michael et al. (1997), Obstfeld and Taylor (1997), Sarantis (1999), Sarno et al. (2004), and Bec et al. (2010), have made use of nonlinear threshold-type autoregressive models to investigate the purchasing power parity (PPP) persistence puzzle, a notion initiated in a survey by Rogoff (1996). The motivation for using nonlinear models in this setting is that the original empirical findings used to establish the puzzle may have arisen due to model misspecification. Specifically, linear time series models restrict the degree of adjustment of real exchange rates to their PPP levels to be the same at all points of time. However, basic theory suggests that transaction costs can determine when the "law of one price" drives real exchange rates towards PPP and when it does not.[2] Hence, nonlinear models that allow for regime-switching behavior in real exchange rates may be more appropriate to study PPP. Indeed, the findings of many recent empirical studies imply that estimated PPP adjustments are faster for nonlinear models than those estimated for linear models, thus providing a potential resolution for the PPP persistence puzzle. Sarno (2003) and Taylor and Sarno (2003) provide detailed surveys of this literature.

In this paper, we adopt a Bayesian approach to investigate exchange rate nonlinearities and the PPP persistence puzzle. There are three reasons for doing this. First, standard frequentist estimation for nonlinear threshold models typically considered in the literature on exchange rates is cumbersome as it involves procedures to grid-search for the value of the parameters in nonlinear transition functions. Bayesian methods allow for joint estimation of all model parameters, as well as complicated functions of the parameters, such as the half-life measure of persistence based on generalized impulse response functions. Second, testing threshold-type nonlinearities in the frequentist setting is challenging due to the presence of nuisance parameters, with the concomitant problem that tests may be relatively uninformative in small samples due to weak power. In the Bayesian framework, model comparison via marginal likelihoods, which reflect the relative abilities of models to predict the data given prior beliefs about model parameters, is conceptually straight forward for any set of models and an inability to discriminate between models based on sample information will be evident in posterior odds ratios being close to even. Third, while frequentist inferences about exchange rate persistence can be highly sensitive to model specification, the Bayesian approach allows for model-averaged measures that address inherent uncertainty about model-specification issues such as lag order or the possible presence of nonlinear dynamics.

Our empirical findings can be summarized as follows. Based on our model comparison, there is general support for nonlinear threshold dynamics in real exchange rates for the G7 countries, although the strength of the evidence varies across country pairs. Meanwhile, our model-averaged measures of real exchange rate persistence are generally lower than for linear models alone. Thus, our analysis takes the resolution of the PPP persistence puzzle further than frequentist analysis based on nonlinear models. In the frequentist setting, the finding of lower persistence is a "knife-edge" results that depends crucially on the presence of nonlinear dynamics in real exchange rates, with tests for nonlinearity providing little support for nonlinearity across country pairs in practice. These "knife-edge" inferences are particularly worrisome given the fact that tests of nonlinearity can suffer from weak power in small samples. By contrast, our finding based on Bayesian analysis is that model-averaged measures of persistence are generally lower than those based on linear models, including in the cases where the evidence for nonlinearity is somewhat ambiguous. Specifically, we find that half-lives

[2] See, for example, Heckscher (1916), Cassel (1922), Dumas (1992) and O'Connell (1997).

for G7 real exchange rates mostly range between 2 and 3 years compared to the 3–5 years found in Rogoff (1996). This might be seen as only a partial resolution of the PPP persistence puzzle given that 2–3 year half-lives are still too long to be easily reconciled with sticky goods prices alone. However, when one considers the possibility of threshold effects, the 2–3 year unconditional half-lives become much more economically plausible as exchange rates would not be expected to adjust quickly when they are close to their PPP levels, which they often are in practice.

The remainder of this paper is organized as follows: Section 2 presents the linear and nonlinear models of the real exchange rate considered in our analysis. Section 3 discusses practical issues for Bayesian estimation for these models. Section 4 reports the empirical results for an application of these models and Bayesian methods to real exchange rate data from the G7 countries. Section 5 concludes.

2. Models

There are many different time series models of exchange rates. The main distinction between them is whether they assume linear or nonlinear dynamics. Within the realm of nonlinear models, the emphasis for exchange rates has been on models that allow for nonlinear conditional mean dynamics. However, exchange rates are asset prices, so there are also models that allow changing conditional variances to help capture fat tails in the distribution of exchange rate returns. In our analysis, we focus on the distinction between linear and nonlinear models of conditional mean dynamics. However, we also consider the effects of accounting for heteroskedasticity and fat tails on inferences about nonlinear mean dynamics and the persistence of exchange rate fluctuations.

The benchmark linear model that we consider is a finite-order autoregressive (AR) model:

$$\phi(L)(q_t - \mu) = \varepsilon_t,\ \varepsilon_t \sim i.i.d.N\left(0, \sigma^2\right), \tag{1}$$

Where q_t is the log real exchange rate, $\phi(L) = 1 - \phi_1 L - \cdots - \phi_p L^p$, and the roots of $\phi(z) = 0$ lie on or outside the unit circle. Roots outside the unit circle corresponds to the idea that PPP holds in the long run.[3] The Gaussian error assumption is driven by the need for a parametric structure in order to conduct our Bayesian analysis.[4]

In terms of nonlinear models of conditional mean dynamics for exchange rates, the existing literature has emphasized so-called "self-exciting" threshold models with discrete transitions (TAR) and smooth transitions (STAR) between different regimes for the AR dynamics (see Michael et al., 1997; Obstfeld and Taylor, 1997; Taylor et al., 2001; Sarno et al., 2004). Building on this literature and inspired by Franses and van Dijk (2000), Bec et al. (2010) develop a general multi-regime logistic STAR (MR-LSTAR) model that nests both TAR and STAR dynamics. The model, which we adopt here, starts with a Dickey-Fuller transformation of the benchmark linear AR model in (1) into an error-correction representation:

$$\Delta q_t = \beta(q_{t-1} - \mu) + u_t, \tag{1'}$$

where $\beta \equiv -\phi(1)$, $u_t \equiv \sum_{j=1}^{p-1} \phi_j^* \Delta q_{t-j}$, and $\phi_j^* \equiv -\sum_{i=j+1}^{p} \phi_i$. Nonlinear conditional mean dynamics are then allowed for by letting the error–correction coefficient β be regime-dependent as follows:

$$\Delta q_t = \sum_{r=1}^{3} F_r(q_{t-1} - \mu | \gamma, c) \beta_r (q_{t-1} - \mu) \tag{2}$$

where

[3] The strongest evidence for long-run PPP comes from the long samples of exchange rate data considered in Abuaf and Jorion (1990) and Lothian and Taylor (1996), although it should always be acknowledged that long-run PPP may not strictly hold due to the possible presence of a small random walk component (see Engel, 2000).

[4] Given that exchange rates are asset prices, a Student t distribution with a low degree of freedom for the error term might seem a more reasonable assumption. However, when we considered this alternative assumption, we found that our results were highly robust. This robustness likely reflects the fact that we consider quarterly data and a Gaussian assumption for exchange rates is somewhat more reasonable at lower frequencies than at high frequencies (i.e., accounting for fat tails would be more important for daily or weekly data). For brevity, we only report results assuming Gaussian errors.

$$F_1 = [1 + \exp(-\gamma(q_{t-1} - \mu - c)]^{-1}, \tag{3}$$

$$F_2 = [1 + \exp(\gamma(q_{t-1} - \mu + c)]^{-1}, \tag{4}$$

$$F_3 = 1 - F_1 - F_2, \tag{5}$$

with the restriction $\beta_1 = \beta_2 \equiv \beta_{out}$ and, for notational convenience, $\beta_3 \equiv \beta_{in}$. In words, the prevailing error correction coefficient at any point of time depends on the level of the lagged exchange rate relative to symmetric thresholds around the mean μ, with the width of the threshold bands determined by the threshold parameter c. The transition functions $F_r(q_{t-1} - \mu | \gamma, c)$ determine the weights put on each regime according to logistic specifications that depends on the smooth transition parameter γ, which is restricted to be positive in order to identify the regimes. Note that, as, $\gamma \rightarrow \infty$ the MR-LSTAR model approximates a band-TAR model.

Given this setup, it is straight$\gamma' \rightarrow \infty$ forward to allow other parameters to also depend on the regime, including the variance of the shocks. Thus, in order to address the possibility of heteroskedasticity, we also consider whether augmenting the models discussed above with regime-dependent variances, σ^2_{out} and σ^2_{in}, affects our inferences about exchange rate dynamics.

3. Bayesian estimation

We conduct our Bayesian estimation via a multi-block random-walk chain version of the Metropolis–Hastings (MH) algorithm. The MH algorithm is a posterior simulator in which draws are first made from an easy-to-simulate proposal distribution (e.g., a multivariate Normal distribution). Then the draws are accepted or rejected as draws from a target distribution (i.e., the posterior distribution) based on the relative densities of the draws for both the proposal and target distributions.

As with any importance-sampling algorithm, the success of the posterior simulator in providing an accurate discrete approximation of the target distribution depends on the proposal distribution. We follow a common approach in the applied literature of making our proposal a multivariate Student t distribution based on the posterior mode and the curvature of the posterior around the mode. However, some issues arise in doing so for the nonlinear MR-LSTAR model. First, just as with maximum likelihood estimation of nonlinear threshold models, there is a need for a grid search across the threshold parameter c to find the posterior mode. However, it is important to emphasize that this only applies to constructing the proposal distribution. Bayesian estimation of the threshold based on the target distribution has the benefit that it does not involve discretization of the sample space for the threshold parameter. Second, by using a grid search to estimate the threshold parameter, numerical derivatives cannot be used to evaluate the curvature of the posterior with respect to the threshold parameter. Thus, there is no guide from numerical optimization for the scale of the proposal density, even if its location can be pinned down at the posterior mode.

In our analysis, we address the problem of determining a good proposal distribution for nonlinear threshold models by considering an alternative measure of the curvature of the posterior with respect to the threshold parameter c. First, we invert the “posterior ratio” for the threshold based on a critical value of 3.84. Specifically, given diffuse priors, this is equivalent to inverting the likelihood ratio statistic for c to construct a 95% confidence interval (in a frequentist sense) under the assumption that the statistic has a standard $\chi^2(1)$ asymptotic distribution. Note that, even if the actual distribution is not $\chi^2(1)$, this approach will work because it will still capture the approximate curvature of the posterior with respect to c. In particular, parameter estimates based on draws from the target distribution (i.e., the posterior) will be robust to different assumptions about the proposal distribution as long as the proposal loosely captures the shape of the posterior. For example, if we use a critical value based on a $\chi^2(2)$ distribution instead to determine our proposal, our posterior estimates are the same as when we assume a $\chi^2(1)$ distribution. We just need to make a plausible assumption about the critical value, with

3.84 providing a convenient benchmark. Then, given an interval based on the inverted "posterior ratio", we back out an implied standard error (again, in a frequentist sense) for the threshold parameter under the assumption that the estimator has a standard asymptotic distribution. Again, our parameter estimates will be robust to rescaling of the standard error, so it is not crucial that this assumption is literally true.

To summarize, our approach for approximating the curvature of the posterior with respect to the threshold parameter proceeds as follows:

1) Construct "confidence set" for c based on inverting the posterior ratio.[5] Assuming the set is contiguous, denote the estimated 95% confidence interval as $[\hat{c}_{0.025}, \hat{c}_{0.975}]$.[6]
2) Note that, if a standard error were available and assuming asymptotic normality, another estimate of the 95% confidence interval would be $\hat{c} \pm 1.96 \times SE(\hat{c})$, meaning that $\hat{c} \pm 1.96 \times SE(\hat{c}) \approx [\hat{c}_{0.025}, \hat{c}_{0.975}]$.
3) Assuming an asymptotic equivalence of the two confidence interval estimators, construct an approximate standard error as $\hat{\sigma}_c = \frac{1}{3.92}(\hat{c}_{0.975} - \hat{c}_{0.025})$.

In terms of the smooth transition parameter γ, while it is possible to estimate it by numerical optimization, there are practical difficulties with doing so. As $\gamma \to \infty$ (i.e., as the MR-LSTAR model becomes more like a Band TAR model), γ becomes unidentified (i.e., there is no impact on the likelihood for changes in γ when it is extremely large). Bayesian analysis helps to some extent because an informative prior on γ has the implication that the posterior will change even if the likelihood does not. However, in practice, to allow for relatively diffuse priors and to aid in numerical optimization, we follow the frequentist literature (see, for example, Franses and van Dijk, 2000) and conduct a grid search for γ to obtain $\hat{\sigma}_\gamma$ for the proposal distribution. Again, it should be emphasized that the grid search is for the proposal distribution only and is only meant to loosely approximate the posterior. The draws of γ from the target distribution will be accurate even given the approximations in the proposal distribution. Meanwhile, we check the robustness of posterior moments to different assumptions for the proposal distribution for these nonlinear parameters.

Letting θ denote the vector of model parameters, the overall proposal distribution is constructed as follows:

$$\theta \sim MT\left(\mu^\theta, \Sigma^\theta, \nu^\theta\right),$$

where μ^θ is set to the previous draw for the random walk chain version of the MH algorithm and ν^θ is the degrees of freedom parameter that it set as $T-k$, where T is the sample size and k is the number of parameters. The key aspect of the proposal density is the scale matrix Σ^θ. Letting θ^L denote the "linear" parameters and θ^{NL} denote the "nonlinear" parameters (i.e., c and γ), where $\theta = (\theta^L, \theta^{NL})'$, Σ^θ is given as follows:

$$\Sigma^\theta = \kappa \begin{bmatrix} \hat{var}\left(\hat{\theta}^L\right) & 0 \\ 0 & \hat{var}\left(\hat{\theta}^{NL}\right) \end{bmatrix},$$

where κ is a "tuning" parameter for the MH algorithm, $\hat{var}(\hat{\theta}^L)$ is the variance-covariance of the "linear" parameters based on the estimated inverse expected Hessian at the posterior mode conditional on the "nonlinear" parameters and $\hat{var}(\hat{\theta}^{NL}) = (\hat{\sigma}_c^2, \hat{\sigma}_\gamma^2)' I_{2\times 2}$ is based on the indirect estimated standard deviations discussed above. Results are robust to allowing for either positive or negative correlation between the linear and nonlinear parameters. Meanwhile, we consider different parameter blocking

[5] See Hansen (1996) for his detailed discussion of the method.

[6] If the confidence set is not contiguous, we take the conservative approach of using the smallest and largest values in the set to construct a 95% confidence interval.

schemes (i.e., conditional drawing from subsets of θ) and we adjust κ to attain an acceptance rate for the MH algorithm of between 20 and 50%.

Model comparison and model weights for constructing a model-averaged measure of persistence are based on marginal likelihoods. These are proportional to the probability that a model (including priors on parameters) would have predicted the observed data. Following Chib and Jeliazkov (2001), we calculate these using the Bayes identity and the MH output. We have confirmed that marginal likelihood estimates and posterior moments are robust across multiple runs of the MH algorithm and for different starting values of the random-walk chain. For each run, we consider 20,000 draws after 10,000 burn-in draws.

4. Empirical results

4.1. Data and priors

We consider quarterly real exchange rates for eight different country pairs from the G7; these include non-euro currency exchange rates from 1974Q1 to 2014Q2 and euro currency exchange rates from 1974Q1 to 1998Q4. We calculate the real exchange rate series using nominal exchange rates and consumer price index (CPI) data from the IFS database.[7] We convert the monthly series into a quarterly frequency by taking the end-of-quarter values. When looking at long-horizon persistence properties of exchange rates, there is little benefit of considering monthly data instead of quarterly data, while there would be a cost in terms needing more complicated models to account for the fat tails and volatility clustering that is more evident in higher–frequency exchange rate data. Also, the computation, especially of marginal likelihoods, is much faster given quarterly data instead of monthly data.[8]

Five of the real exchange rate series are vis-à-vis the U.S. dollars; all are commonly examined in the literature, but only the pound–dollar exchange rate is included in Bec et al. (2010).[9] To compare with their results, we also include three series that do not involve the U.S. dollar. All real exchange rate series are converted into logarithms and re-centered. The full sample period is separated into two: 1974Q1 to 1979Q4 provides a training sample to help us with the elicitation of priors for certain parameters that depend on the scale of the data (e.g., the variance of shocks) and/or parameters for which model comparison could potentially be sensitive to what might otherwise be arbitrary assumptions (e.g., the nonlinear parameters); 1980Q1 to 2014Q2 (or 1998Q4 for the euro currency exchange rates) is used for Bayesian estimation and model comparison. We consider up to four lags for the AR specification. Because we use the error-correction and Dickey–Fuller transformation given in (1′) and (2), the AR(4) model, for example, is specified with the regressand as the first difference of the log real exchange rate and the regressors are the first lag of the log real exchange rate and three lags of the first differences.

For all of the models, the priors for the AR parameters have a truncated Normal distribution to rule out an explosive process (i.e., a draw from the proposal density can only be accepted if the roots of the characteristic equation $\phi(z) = 0$ lie on or outside the unit circle). The mean of the underlying Normal distribution for the error–correction coefficient on the lag of the log real exchange rate is set based on the OLS estimate from an AR(1) model of the log real exchange rate for the training sample (or we set it to 0 if the estimated AR(1) coefficient is greater than 1) and the standard deviation is 0.1. The mean of the underlying Normal distribution for the coefficients on the lagged differences is set to 0 and the standard deviation is 0.5. The implied mean of the prior distributions will depend on the effects of truncation, but given the relatively tight prior on the error–correction coefficient, the implied mean will be close to the OLS estimate. Meanwhile, the prior for the forecast-error variance has a Gamma

[7] For the U.K. CPI, we obtain data from FRED for the pre-1988 sample period due to lack of availability from IFS. The IFS and FRED data for U.K. CPI are very similar in the post-1988 sample.

[8] Note, however, that we have confirmed that the posterior parameter inferences are qualitatively similar (adjusting for the frequency) when considering similar models with monthly data.

[9] We exclude the German real exchange rate because the interpretation of the CPI before and after the German unification is problematic.

distribution $\sigma^2 \sim Gamma(\nu, \delta)$ where the Gamma distribution for variable x is parameterized as follows:

$$f(x|\nu,\delta) = \frac{(\delta/2)^{(\nu/2)}}{\Gamma(\nu/2)} x^{(\nu/2)-1} e^{-(\delta/2)x}.$$

For the prior on the forecast error variance, we set the rate parameter $\delta = 1$ (i.e., the Gamma distribution collapses to is a Chi-squared distribution with ν degrees of freedom) and we set the shape parameter ν to the sample variance of the forecast error in the training sample implied by the means of underlying prior distributions for the error–correction coefficient and the coefficients on the lagged differences. This implies that the prior mean for the forecast-error variance will be equal to its sample value in the training sample, but the prior is relatively uninformative prior and is common to all models, so it should not affect the model comparisons.

The elicitation of priors for the nonlinear parameters in the MR-LSTAR model is slightly more involved and requires more discussion. For the threshold parameter c, we assume a Gamma distribution and set the rate parameter $\delta = 0.25$ and the shape parameter ν to δ times the median absolute real exchange deviation (in logarithms) from the sample mean using data from the training sample period, implying the prior mean for c is equal to the median deviation from the mean real exchange rate in the training sample. For the smooth transition parameter γ, we also assume a Gamma distribution and set the shape parameter $\Delta\beta \equiv \beta_{in} - \beta_{out}$ and the rate parameter δ to $0.25 \times \nu$ times the mean difference between ranked absolute deviations from the mean real exchange rate in the training sample. This calibrates the prior on the transition function to be more precise and with more weight on a smooth transition when there is more information on the effects of small differences on real exchange rate dynamics. If the mean difference is larger, implying less information about the effects of small differences, the prior on the transition function becomes more diffuse and shifts more mass towards a relatively discrete transition function. Meanwhile, it should be noted that, despite the use of training sample information to calibrate the priors for both c and γ, the priors are relatively uninformative given low values for the rate parameter δ in practice.

The remaining parameter for the nonlinear threshold models is the change in the error correction coefficient across regimes: $\Delta\beta \equiv \beta_{in} - \beta_{out}$. For this parameter, our prior is more informative than for other parameters and is based on the transaction costs notion that the adjustment to PPP will be larger when the exchange rate is far away from its PPP level. In particular, for $\Delta\beta$, we again assume a Gamma distribution and set the rate parameter $\delta = 50$ and the shape parameter $\nu = 2$, implying the mean for the reduction in the error correction coefficient is 0.04, with a standard deviation of 0.04. To offset any impact on the implied prior for the half-life of a shock to the real exchange rate, we add 0.04 to the prior mean of the underlying Normal distribution for the error–correction coefficient β_{in} (or we set it to 0 if the adjusted value would otherwise be positive, which would correspond to an explosive process).

We justify an informative prior for the change in the error coefficient in two ways. First, we have strong theoretical reasons based on transaction costs to believe the error correction effect is larger in the outside regime when the real exchange rate is further from PPP. This is exactly the motivation for using a threshold model for real exchange rates and the ability to specify an informative prior that specifies the model according to that dynamic is a benefit of Bayesian analysis.[10] Second, even though an informative prior that places little weight on $\Delta\beta \approx 0$ might seem at first glance to push our empirical findings towards finding evidence of nonlinearity, it does not in fact do so. This is because we also consider linear models in our model comparison. Indeed, it is important for comparing linear and nonlinear models with Bayesian model comparison that there is little or no prior weight on the portion of the parameter space for the nonlinear models that corresponds to linearity. Only in this case will our true prior odds for linearity and nonlinearity be equal when considering Bayes factors (see below) to

[10] In Bayesian analysis, the prior and the model are closely related. As an example, it is possible to compare two priors given the same model specification using marginal likelihoods. In essence, the comparison is between how well two prior models predicted the sample data.

Table 1
Log marginal likelihoods and Bayes factors.

Model	Autoregressive lag order (p)			
	1	2	3	4
British Pound–U.S. Dollar				
Linear	−418.56 (0.00)	−417.55 (0.01)	−418.87 (0.00)	−419.94 (0.00)
Linear-h	−417.51 (0.01)	−416.70 (0.01)	−418.12 (0.00)	−419.11 (0.00)
MRLSTAR	−417.10 (0.01)	−415.81 (0.03)	−417.14 (0.01)	−418.08 (0.00)
MRLSTAR-h	−415.87 (0.03)	**−412.44**	−416.03 (0.03)	−417.14 (0.01)
Canadian Dollar–U.S. Dollar				
Linear	−351.03 (0.00)	−350.31 (0.00)	−350.81 (0.00)	−350.99 (0.00)
Linear-h	−352.36 (0.00)	−351.07 (0.00)	−352.39 (0.00)	−351.91 (0.00)
MRLSTAR	−349.48 (0.00)	−348.74 (0.00)	−349.28 (0.00)	**−341.88**
MRLSTAR-h	−350.68 (0.00)	−348.11 (0.00)	−347.13 (0.01)	−344.79 (0.05)
French Franc–U.S. Dollar				
Linear	−240.49 (0.00)	−239.43 (0.00)	−240.73 (0.00)	−240.13 (0.00)
Linear-h	−240.96 (0.00)	−240.03 (0.00)	−241.26 (0.00)	−240.74 (0.00)
MRLSTAR	−238.55 (0.00)	−234.00 (0.00)	−233.58 (0.00)	−238.30 (0.00)
MRLSTAR-h	−238.89 (0.00)	−236.28 (0.00)	−235.77 (0.00)	**−226.95**
Italian Lira–U.S. Dollar				
Linear	−243.35 (0.00)	−242.51 (0.00)	−243.32 (0.00)	−242.26 (0.00)
Linear-h	−243.26 (0.00)	−242.69 (0.00)	−243.36 (0.00)	−242.26 (0.00)
MRLSTAR	−239.95 (0.00)	−240.69 (0.00)	−236.25 (0.10)	**−233.99**
MRLSTAR-h	−242.08 (0.00)	−239.17 (0.01)	−241.53 (0.00)	−234.72 (0.48)
Japanese Yen–U.S. Dollar				
Linear	−437.69 (0.00)	−438.67 (0.00)	−439.69 (0.00)	−436.49 (0.01)
Linear-h	−438.06 (0.00)	−439.07 (0.00)	−440.06 (0.00)	−436.54 (0.01)
MRLSTAR	−435.65 (0.02)	−436.64 (0.01)	**−431.89**	−435.93 (0.02)
MRLSTAR-h	−436.29 (0.01)	−436.88 (0.01)	−434.13 (0.11)	−434.34 (0.09)
Canadian Dollar–British Pound				
Linear	−425.54 (0.02)	−426.16 (0.01)	−427.13 (0.00)	−428.51 (0.00)
Linear-h	−425.96 (0.01)	−426.48 (0.01)	−427.45 (0.00)	−428.75 (0.00)
MRLSTAR	−421.88 (0.62)	−424.33 (0.05)	−425.34 (0.02)	−426.65 (0.01)
MRLSTAR-h	−424.34 (0.05)	−422.75 (0.26)	−422.05 (0.52)	**−421.40**
British Pound–French Franc				
Linear	−222.23 (0.23)	−222.44 (0.19)	−223.71 (0.05)	−224.15 (0.03)
Linear-h	−223.08 (0.10)	−222.98 (0.11)	−224.27 (0.03)	−224.77 (0.02)
MRLSTAR	−220.99 (0.80)	−221.09 (0.72)	−222.39 (0.20)	−222.79 (0.13)
MRLSTAR-h	−221.24 (0.62)	**−220.76**	−222.22 (0.23)	−222.65 (0.15)
Italian Lira–French Franc				
Linear	−188.95 (0.00)	−188.69 (0.00)	−189.89 (0.00)	−190.31 (0.00)
Linear-h	−189.04 (0.00)	−188.57 (0.00)	−189.77 (0.00)	−189.91 (0.00)
MRLSTAR	−187.15 (0.00)	−183.82 (0.09)	−188.33 (0.00)	−182.65 (0.28)
MRLSTAR-h	−186.63 (0.01)	−184.11 (0.06)	−184.56 (0.04)	**−181.38**

Note: The largest log marginal likelihood value is in bold. The Bayes factor is reported in parentheses and is equal to the marginal likelihood value of a particular specification divided by the largest overall marginal likelihood value for a given country pair. Values smaller than (0.01) are reported as (0.00).

calculate posterior odds, while equal prior odds for a linear and nonlinear model would implicitly favor linearity given less informative priors on parameters in the nonlinear models related to nonlinearity.

These priors on the parameters for the linear and nonlinear models have roughly similar implications for the half-life measure of persistence of a shock to the real exchange rate. For example, assuming an autoregressive lag order of 1 and an OLS estimate of 0.85 for the AR(1) parameter in the training sample, the implied prior means and standard deviations for the half-life for the linear and nonlinear models are all about 2 years and the percentiles of the prior distributions are very similar.

4.2. Posteriors

Table 1 reports the log marginal likelihood for each model, with a corresponding Bayes factor in parentheses. Each Bayes factor is calculated as the ratio between the marginal likelihood value of a

Table 2
Linearity tests.

Model	Autoregressive lag order (p)			
	1	2	3	4
LM statistics				
British Pound–U.S. Dollar	0.8008	1.1511	0.9972	1.3307
Canadian–U.S. Dollar	0.3142	0.6407	0.2820	0.6387
French Franc–U.S. Dollar	3.3259	5.3360*	5.0672*	4.9502*
Italian Lira–U.S. Dollar	2.5983	3.4869	3.4183	2.9887
Japanese Yen–U.S. Dollar	0.6581	0.8834	0.7586	0.6552
Canadian Dollar–British Pound	0.1362	0.1927	0.0926	0.1588
British-Pound–French Franc	4.7258*	5.9124*	6.2107**	5.8198*
Italian Lira–French Franc	5.3353*	8.4937**	8.1788**	7.6529**
Heteroskedasticity robust LM statistics				
British Pound–U.S. Dollar	0.7672	1.1636	1.1029	1.3563
Canadian–U.S. Dollar	0.4344	0.9823	0.4470	0.9313
French Franc–U.S. Dollar	3.0357	4.8462*	5.1359*	5.4650*
Italian Lira–U.S. Dollar	1.5188	2.1945	2.5005	2.4644
Japanese Yen–U.S. Dollar	0.2244	0.2955	0.2581	0.2386
Canadian Dollar–British Pound	0.2212	0.3660	0.1809	0.3260
British Pound–French Franc	2.4456	3.1143	3.8321−	3.8346−
Italian Lira–French Franc	3.1241	4.3023−	4.6030−	5.1818*

Note: −, *, and ** denote level of statistical significance at 15%, 10%, and 5% respectively.

particular specification and the largest marginal likelihood value among the sixteen specifications for a given country pair. Based on these results, the MR-LSTAR models with and without heteroskedastic disturbances are well supported for all eight real exchange rates. Moreover, in every one of these cases, the second best model, using the Bayes factor as the measure, is also an MR-LSTAR specification. Linear models only have nontrivial Bayes factors for the pound-franc real exchange rate. Meanwhile, our results suggest that both nonlinear conditional mean dynamics and heteroskedasticity are often important in understanding the exchange rate data. Indeed, the MR-LSTAR-h specification has the highest Bayes factor in five out of eight cases.

It is illustrative to compare Table 1 with Table 2, which reports results for a frequentist LM-type test of linearity that was also considered in Bec et al. (2010).[11] Our sample periods are different than theirs, so the results in Table 2 are not an exact replication of their results. However, the inferences are similar. There are four series that are common in the two studies: pound-dollar, Canadian dollar-pound, pound-franc and lira-franc. Note that they fix the number of lags to 2 and adopt a slightly unconventional set of significance levels from 5% to 15%. As in their study, we find evidence of nonlinearity for the pound-franc and lira franc, but not the Canadian dollar-pound. Our results for the pound-dollar differ and this appears to be due to a difference in the sample period.

Table 2 illustrates a difficulty with frequentist hypothesis testing in this context. Based on a predetermined level of significance of even 15%, we fail to reject linearity for more than half of the series that we consider, including, for example, for the Italy/U.S. real exchange rate. Yet, we can reject linearity for the France/U.S. and the Italy/France real exchange rates. Although transitivity may not necessarily apply, this result suggests that the failure to reject may simply reflect low power of the test in a small sample setting. Unfortunately, failure to reject due to lower power is highly problematic in this setting because subsequent inferences about the persistence of shocks can be highly sensitive across country pairs depending on whether we condition on a linear or nonlinear model. By contrast,

[11] Bec et al. (2010) refer these tests as LML tests. The tests require an estimation procedure that grid searches for the maximal LM statistics over a set of γ and c in (3)–(5). The LM statistics are computed as $T(\tilde{\varepsilon}'\tilde{\varepsilon} - \hat{\varepsilon}'\hat{\varepsilon})$, where $\tilde{\varepsilon}'\tilde{\varepsilon}$ is the sum of squared residuals for the linear model and $\hat{\varepsilon}'\hat{\varepsilon}$ is the sum of squared residuals for the MR-LSTAR model. Full details of computation for the LM test statistic can be found in Appendix B of Bec et al. (2010).

Table 3
Parameter estimates for preferred nonlinear models.

	£/USD	CN$/US$	FF/US$	ITL/US$	¥/US$	CN$/£	£/FF	ITL/FF
σ^2_{out}	29.368	–	26.290	–	–	28.615	14.178	8.736
	(5.807)		(4.969)			(4.011)	(2.924)	(1.900)
σ^2_{in}	19.608	8.448	22.272	27.326	31.762	37.355	20.273	6.543
	(2.869)	(0.957)	(4.518)	(3.764)	(3.564)	(7.155)	(4.172)	(2.690)
β_{in}	0.917	0.972	0.922	0.936	0.960	0.949	0.857	0.907
$+ 1$	(0.040)	(0.018)	(0.037)	(0.037)	(0.024)	(0.031)	(0.053)	(0.040)
β_{out}	−0.039	−0.020	−0.027	−0.028	−0.021	−0.029	−0.034	−0.022
$- \beta_{in}$	(0.033)	(0.016)	(0.024)	(0.024)	(0.018)	(0.023)	(0.031)	(0.020)
c	9.777	8.850	14.176	4.923	11.163	6.624	5.020	2.379
	(5.228)	(4.446)	(4.936)	(3.046)	(6.969)	(3.405)	(2.817)	(1.936)
γ	5.011	9.691	7.344	7.906	2.969	3.494	11.199	10.601
	(4.807)	(9.628)	(7.261)	(7.616)	(2.802)	(3.712)	(10.509)	(10.519)

Note: Estimates are based on posterior means, with standard deviations in parentheses. For CN$/US$, ITL/US$ and ¥/US$, the variance σ^2_{in} prevails in both regimes.

our subsequent Bayesian inferences about persistence are more consistent across country pairs by allowing weight on both linear and nonlinear models.[12]

To provide a sense of the possible nonlinear features in the exchange rate data, we report the estimates and the empirical transition functions for the nonlinear model with the highest Bayes factor for each country pair in Table 3 and Fig. 1 respectively. The posterior means for the sum of the autoregressive coefficients within the threshold bands (i.e., $\beta_{in} + 1$) range from 0.857 to 0.972. By allowing nonlinearities, we find that the posterior means for the change in the sum of the autoregressive coefficients (i.e., $\beta_{out} - \beta_{in}$) ranges from −0.039 to −0.020, with posterior standard deviations generally about half of the prior standard deviation of about 0.04. The data are informative, but not definitive about the magnitude of the change in persistence across regimes. Meanwhile, Fig. 1 illustrates the estimated transition functions based on posterior means. The mirrored logistic function imposed by the MRLSTAR specification can converge to a discrete step function like a TAR model when γ is large enough. It can also mimic other functions, such as the exponential function. However, Fig. 1 clearly suggests that the changes in the dynamics are discrete around a threshold. In many cases, there are very few points in between 0 and 1. At the same time, it was important for a fair comparison to linear models not to prespecify the form of nonlinearity in our estimation.

Interpretation of threshold estimates is seldom easy. An intuitive but not comprehensive view is to see them as estimates for the cost of transportation in the "iceberg" form (see O'Connell and Wei, 2001). In their investigation of the "law of one price", Obstfeld and Taylor (1997) find that threshold estimates are positively related to the distance between two locations. Along these lines, our results regarding the thresholds are revealing. The European country-pairs of U.K./France and Italy/France have relatively small threshold estimates of 5.020 and 2.379, respectively. If thresholds represent the cost of arbitrage, we would expect these European countries to enjoy smaller cost of transaction amongst themselves. Not surprisingly, the Europe/U.S. pairs generally have larger threshold estimates, with the pairs of U.K./U.S. and France/U.S. having threshold estimates of 9.777 and 13.176, respectively, although Italy/U.S. has a relatively small threshold estimate of 4.923. The estimated threshold for the Japan/U.S. real exchange rate is one of the largest at 11.163. Meanwhile, the most surprising results involve the Canadian dollar. One would expect, because of geographical distances and historical trade tie, the U.S. and the Canadian economies are the closest. But the threshold estimate for the real exchange rates of the two dollars is relatively high at 8.850 and is larger than the threshold estimate for the Canadian dollar-pound real exchange rate at 6.624, which in turn is also smaller than that for the

[12] Bayesian model averaging tends to put all weight on one model when models are "sparse" in the sense of being quite different from each other. In our case, the models are similar, with the main distinction being between linear and nonlinear specifications. As a result, we find that the various models all tend to receive nontrivial weight, especially across different lag lengths, supporting the use of Bayesian model averaging as a way to combine models.

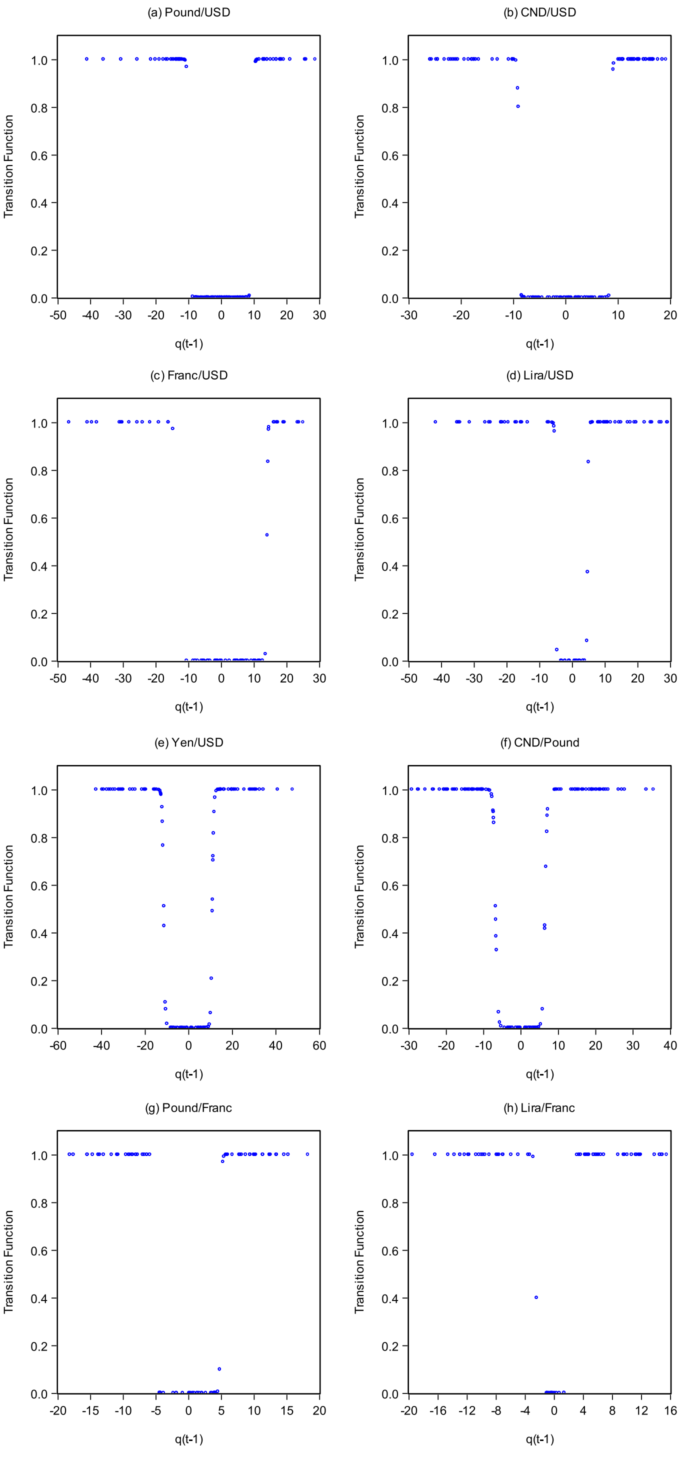

Fig. 1. Estimated transition functions for the preferred nonlinear models.

Table 4
Half-life estimates in years (posterior means and weighted averages).

Model	Autoregressive Lag order (p)				
	1	2	3	4	All Lags
British Pound–U.S. Dollar					
Linear	2.50	2.25	2.25	2.25	2.25
Linear-h	2.50	2.25	2.50	2.25	2.25
MRLSTAR	2.25	2.00	2.25	2.25	2.00
MRLSTAR-h	2.50	2.25	2.25	2.25	2.25
All Models	2.50	2.25	2.25	2.25	2.25
Canadian Dollar–U.S. Dollar					
Linear	7.25	6.50	7.00	6.50	6.75
Linear-h	7.25	6.75	7.00	6.50	6.75
MRLSTAR	5.75	5.50	5.75	5.50	5.50
MRLSTAR-h	5.75	5.50	5.50	5.25	5.25
All Models	6.00	5.50	5.50	5.50	5.50
French Franc–U.S. Dollar					
Linear	3.00	2.75	3.00	2.75	2.75
Linear-h	3.25	2.75	3.00	3.00	3.00
MRLSTAR	2.75	2.75	2.75	2.75	2.75
MRLSTAR-h	2.75	2.75	2.75	2.75	2.75
All Models	2.75	2.75	2.75	2.75	2.75
Italian Lira–U.S. Dollar					
Linear	3.50	3.25	3.50	3.25	3.25
Linear-h	3.75	3.25	3.50	3.25	3.25
MRLSTAR	3.25	3.00	3.00	3.00	3.00
MRLSTAR-h	3.00	3.00	3.00	3.00	3.00
All Models	3.25	3.00	3.00	3.00	3.00
Japanese Yen–U.S. Dollar					
Linear	4.50	4.25	4.50	4.00	4.25
Linear-h	4.25	4.25	4.50	4.00	4.00
MRLSTAR	4.00	3.75	3.75	3.75	3.75
MRLSTAR-h	4.00	3.75	3.75	3.75	3.75
All Models	4.00	3.75	3.75	3.75	3.75
Canadian Dollar–British Pound					
Linear	3.75	3.50	3.75	3.50	3.75
Linear-h	3.75	3.50	3.75	3.75	3.75
MRLSTAR	3.00	3.00	3.00	3.00	3.00
MRLSTAR-h	3.00	3.00	3.00	3.00	3.00
All Models	3.00	3.00	3.00	3.00	3.00
British Pound–French Franc					
Linear	1.50	1.50	1.50	1.50	1.50
Linear-h	1.50	1.50	1.50	1.50	1.50
MRLSTAR	1.50	1.50	1.50	1.50	1.50
MRLSTAR-h	1.50	1.50	1.50	1.50	1.50
All Models	1.50	1.50	1.50	1.50	1.50
Italian Lira–French Franc					
Linear	2.25	2.25	2.25	2.50	2.25
Linear-h	2.25	2.50	2.25	2.50	2.50
MRLSTAR	2.25	2.25	2.25	2.50	2.50
MRLSTAR-h	2.50	2.25	2.25	2.50	2.50
All Models	2.50	2.25	2.25	2.50	2.50

U.S. dollar-pound, even though both North American countries are separated from the United Kingdom by a similar distance. However, it should be noted that the standard deviations for the threshold parameters are fairly large relative to the magnitude of the posterior means. So it is difficult to draw strong conclusions about the sources of the threshold effects from correlations between their estimated values and geographical distances.

In Table 4, we report estimates expressed in years based on posterior means for the unconditional half-life of real exchange rate deviations from PPP. Given parameter values, the computation of half-

Table 5
Half-life distributions in years (posterior percentiles).

	Lag	5%	25%	Median	75%	95%
British Pound–U.S. Dollar						
Linear	1	1.25	1.75	2.00	2.50	6.00
	2	1.25	1.75	2.00	2.50	4.50
	3	1.25	1.75	2.00	2.50	5.25
	4	1.50	1.75	2.00	2.50	4.25
Linear-h	1	1.25	1.75	2.00	2.75	7.50
	2	1.25	1.75	2.00	2.50	5.00
	3	1.25	1.75	2.00	2.50	6.25
	4	1.50	1.75	2.00	2.50	5.00
MRLSTAR	1	1.25	1.75	2.00	2.50	5.00
	2	1.25	1.75	2.00	2.25	4.00
	3	1.25	1.75	2.00	2.50	4.50
	4	1.25	1.75	2.00	2.50	4.25
MRLSTAR-h	1	1.00	1.75	2.00	2.75	6.00
	2	1.25	1.75	2.00	2.50	4.50
	3	1.00	1.75	2.00	2.50	5.25
	4	1.25	1.75	2.00	2.50	4.75
Canadian Dollar–U.S. Dollar						
Linear	1	2.75	4.25	5.75	9.25	∞
	2	2.75	4.00	5.25	7.75	∞
	3	2.75	4.00	5.50	9.00	∞
	4	2.75	4.00	5.25	7.75	∞
Linear-h	1	2.75	4.00	6.00	9.75	∞
	2	2.75	4.00	5.50	8.50	∞
	3	2.75	4.00	5.50	8.50	∞
	4	2.75	4.00	5.25	7.75	∞
MRLSTAR	1	2.25	3.50	4.75	6.75	∞
	2	2.25	3.50	4.50	6.25	∞
	3	2.25	3.50	4.75	6.75	∞
	4	2.25	3.50	4.50	6.50	∞
MRLSTAR-h	1	2.25	3.50	4.75	6.75	∞
	2	2.25	3.50	4.50	6.25	∞
	3	2.50	3.50	4.50	6.25	∞
	4	2.50	3.50	4.50	6.00	14.50
French Franc–U.S. Dollar						
Linear	1	1.50	2.00	2.50	3.25	10.25
	2	1.50	2.00	2.25	3.00	7.75
	3	1.50	2.00	2.50	3.00	8.25
	4	1.75	2.25	2.50	3.00	6.50
Linear-h	1	1.50	2.00	2.50	3.25	11.25
	2	1.50	2.00	2.25	3.00	7.75
	3	1.50	2.00	2.50	3.25	8.75
	4	1.75	2.25	2.50	3.00	6.25
MRLSTAR	1	1.25	2.00	2.25	3.25	7.25
	2	1.50	2.00	2.25	3.00	6.50
	3	1.25	2.00	2.25	3.00	6.50
	4	1.50	2.25	2.50	3.00	5.50
MRLSTAR-h	1	1.25	2.00	2.25	3.00	7.00
	2	1.50	2.00	2.25	3.00	6.25
	3	1.50	2.00	2.25	3.00	6.25
	4	1.50	2.25	2.50	3.00	6.00
Italian Lira–U.S. Dollar						
Linear	1	1.50	2.00	2.75	4.00	∞
	2	1.50	2.00	2.50	3.25	10.25
	3	1.50	2.00	2.50	3.50	13.25
	4	1.75	2.25	2.75	3.25	8.50
Linear-h	1	1.50	2.00	2.75	4.00	∞
	2	1.50	2.00	2.50	3.50	10.75
	3	1.50	2.00	2.75	3.75	14.50
	4	1.75	2.25	2.50	3.25	8.75

(*continued on next page*)

Table 5 (*continued*)

	Lag	5%	25%	Median	75%	95%
MRLSTAR	1	1.50	2.00	2.50	3.50	9.25
	2	1.50	2.00	2.50	3.25	7.25
	3	1.25	2.00	2.50	3.25	8.50
	4	1.75	2.25	2.50	3.25	7.00
MRLSTAR-h	1	1.50	2.00	2.50	3.50	8.50
	2	1.50	2.00	2.50	3.25	7.75
	3	1.25	2.00	2.50	3.50	8.50
	4	1.75	2.25	2.50	3.25	6.50
Japanese Yen–U.S. Dollar						
Linear	1	2.00	2.75	3.50	5.00	∞
	2	2.00	2.75	3.25	4.50	13.50
	3	2.00	2.75	3.50	5.00	∞
	4	2.25	2.75	3.25	4.25	10.50
Linear-h	1	2.00	2.75	3.50	4.75	∞
	2	2.00	2.75	3.25	4.50	∞
	3	2.00	2.75	3.50	5.00	∞
	4	2.25	2.75	3.25	4.25	10.50
MRLSTAR	1	1.75	2.50	3.25	4.50	10.75
	2	1.75	2.50	3.25	4.25	9.75
	3	1.75	2.50	3.25	4.50	10.00
	4	2.00	2.75	3.25	4.00	8.75
MRLSTAR-h	1	1.75	2.50	3.25	4.50	11.75
	2	1.75	2.50	3.25	4.25	10.25
	3	1.75	2.50	3.25	4.25	10.75
	4	2.00	2.75	3.25	4.25	8.50
Canadian Dollar–British Pound						
Linear	1	1.75	2.25	3.00	4.00	14.75
	2	1.75	2.25	2.75	3.75	11.25
	3	1.75	2.25	2.75	4.00	∞
	4	1.50	2.25	2.75	3.75	∞
Linear-h	1	1.75	2.25	3.00	4.00	∞
	2	1.75	2.25	2.75	3.75	11.75
	3	1.75	2.25	3.00	4.00	14.25
	4	1.75	2.25	3.00	4.00	∞
MRLSTAR	1	1.50	2.00	2.75	3.50	8.00
	2	1.50	2.00	2.50	3.25	7.50
	3	1.50	2.00	2.50	3.50	8.00
	4	1.25	2.00	2.50	3.25	7.25
MRLSTAR-h	1	1.50	2.00	2.75	3.50	7.50
	2	1.50	2.00	2.50	3.25	6.75
	3	1.50	2.00	2.50	3.50	7.75
	4	1.25	2.25	2.50	3.50	7.25
British Pound–French Franc						
Linear	1	1.00	1.25	1.25	1.75	3.00
	2	1.00	1.25	1.50	1.50	2.50
	3	1.00	1.25	1.25	1.50	2.50
	4	1.00	1.25	1.50	1.75	2.50
Linear-h	1	1.00	1.25	1.50	1.75	2.75
	2	1.00	1.25	1.50	1.50	2.50
	3	1.00	1.25	1.50	1.50	2.50
	4	1.00	1.25	1.50	1.75	2.50
MRLSTAR	1	1.00	1.25	1.50	1.75	3.25
	2	1.00	1.25	1.50	1.75	2.75
	3	0.75	1.25	1.50	1.75	3.00
	4	0.75	1.25	1.50	1.75	2.75
MRLSTAR-h	1	1.00	1.25	1.50	1.75	3.00
	2	1.00	1.25	1.50	1.75	2.75
	3	1.00	1.25	1.50	1.75	2.75
	4	0.75	1.25	1.50	1.75	2.75
Italian Lira–French Franc						
Linear	1	1.25	1.75	2.00	2.50	5.75
	2	1.25	1.75	2.00	2.50	4.75

Table 5 (*continued*)

	Lag	5%	25%	Median	75%	95%
	3	1.25	1.75	2.00	2.50	4.50
	4	1.50	2.00	2.25	2.75	4.75
Linear-h	1	1.25	1.75	2.00	2.50	4.75
	2	1.50	1.75	2.00	2.50	5.25
	3	1.25	1.75	2.00	2.50	5.25
	4	1.50	2.00	2.25	2.75	4.75
MRLSTAR	1	1.25	1.75	2.00	2.50	5.25
	2	1.25	1.75	2.00	2.50	4.75
	3	1.25	1.75	2.00	2.50	5.00
	4	1.50	2.00	2.25	2.75	5.00
MRLSTAR-h	1	1.25	1.75	2.00	2.75	6.00
	2	1.25	1.75	2.00	2.50	5.25
	3	1.25	1.75	2.00	2.50	5.25
	4	1.50	2.00	2.25	2.75	4.75

Note: Simulations of generalized impulse responses allow for a maximum of 15 years. Any simulation that hit the limit is regarded as ∞.

lives for the AR models is conventional. The computation for the MR-LSTAR models requires generalized impulse response simulation as discussed in Koop et al. (1996, 2000). We randomize the initial conditions and the properties (size and sign) of the shocks. This is different to many conditional exercises as in Taylor et al. (2001) and Bec et al. (2010), but is similar to Lo (2008), which provides full details of the pros and cons of different approaches. We also use the marginal likelihood value to determine relative weights in computing different model-averaged measures of the half-lives. In Table 4, the last row labeled “All Models” and the last column labeled “All Lags” in each panel reports such measures. The former reports the weighted half-life between the linear and the nonlinear model given the same lag order. The latter reports the weighted half-life among different lag orders for the same model. The number where the “All Models” row and the “All Lags” column intersect indicate the weighted half-life for all models with all lags. These are our overall model-averaged estimates of half-lives. Because we use quarterly data, we round our estimates to the nearest 0.25 years for ease of interpretation.

Rogoff (1996) surveyed a number of studies that make use of linear models and found that half-life estimates range from 3 to 5 years for most real exchange rates between industrialized countries. From our estimates, there is an extreme case of the Canadian dollar–dollar for which the half-life estimates for the *linear* models is always more 6 years. However, for the rest of the real exchange rates, the half-life estimates for the *linear* models generally fall into Rogoff's range or a bit shorter. The results are somewhat sensitive to the lag length; in general, the larger the number of lags, the smaller the level of persistence. Using the Bayesian weighting scheme, we find that the model-averaged half-life measures for the linear models (under “All Lags” corresponding to *linear* and *linear-h*) are almost all below 4.

When we examine the results for the nonlinear models, we find even smaller half-life estimates. From the linear to the nonlinear models, the reductions are typically 0.25 years, but sometimes as much as 1.25 years. The reason for relatively shorter half-lives for nonlinear models versus linear models can be explained by the results in Table 5, which reports key percentiles from the posterior distribution for the half-lives. It turns out that the difference between the posterior *medians* for the linear and the nonlinear models is usually negligible. In a frequentist framework, Lo (2008) shows that when the MR-LSTAR model is the true data generating process, the Monte Carlo median of the half-life estimates from a linear model are not significantly different to the true unconditional half-life generated from the nonlinear model. Our findings here match this result. However, a closer examination of the other percentiles shows that the distributions for all models are skewed, resulting in our earlier finding that means are larger than medians. Importantly, the upper 95% bound for the linear models is much higher than that for the nonlinear models and hits the level of “infinity” frequently.[13]

[13] In our generalized impulse response simulations, we set a maximum of 15 years (60 quarters) horizon. When the simulated half-life hits this limit, we label it as “infinity” to capture the idea that a large fraction of a shock may never completely die out.

Table 6
Implied weights from marginal likelihood values.

Model	Autoregressive lag order (p)			
	1	2	3	4
British Pound–U.S. Dollar				
Linear	0.0019	0.0052	0.0014	0.0005
Linear-h	0.0054	0.0123	0.0030	0.0011
MRLSTAR	0.0082	0.0297	0.0078	0.0031
MRLSTAR-h	0.0280	0.8609	0.0238	0.0078
	Sum of all linear: 0.0307		Sum of all MRLSTARs: 0.9693	
Canadian Dollar–U.S. Dollar				
Linear	0.0001	0.0002	0.0001	0.0001
Linear-h	0.0000	0.0001	0.0000	0.0000
MRLSTAR	0.0005	0.0010	0.0006	0.9391
MRLSTAR-h	0.0001	0.0019	0.0049	0.0512
	Sum of all linear: 0.0007		Sum of all MRLSTARs: 0.9993	
French Franc–U.S. Dollar				
Linear	0.0000	0.0000	0.0000	0.0000
Linear-h	0.0000	0.0000	0.0000	0.0000
MRLSTAR	0.0000	0.0009	0.0013	0.0000
MRLSTAR-h	0.0000	0.0001	0.0001	0.9975
	Sum of all linear: 0.0000		Sum of all MRLSTARs: 1.0000	
Italian Lira–U.S. Dollar				
Linear	0.0001	0.0001	0.0001	0.0002
Linear-h	0.0001	0.0001	0.0001	0.0002
MRLSTAR	0.0016	0.0008	0.0652	0.6264
MRLSTAR-h	0.0002	0.0035	0.0003	0.3012
	Sum of all linear: 0.0008		Sum of all MRLSTARs: 0.9992	
Japanese Yen–U.S. Dollar				
Linear	0.0024	0.0009	0.0003	0.0078
Linear-h	0.0016	0.0006	0.0002	0.0075
MRLSTAR	(0.01)81	0.0068	0.7753	0.0136
MRLSTAR-h	0.0095	0.0053	0.0830	0.0671
	Sum of all linear: 0.0213		Sum of all MRLSTARs: 0.9787	
Canadian Dollar–British Pound				
Linear	0.0061	0.0033	0.0013	0.0003
Linear-h	0.0040	0.0024	0.0009	0.0002
MRLSTAR	0.2406	0.0207	0.0076	0.0020
MRLSTAR-h	0.0205	0.1001	0.2029	0.3870
	Sum of all linear: 0.0187		Sum of all MRLSTARs: 0.9813	
British Pound–French Franc				
Linear	0.0500	0.0405	0.0115	0.0073
Linear-h	0.0214	0.0238	0.0065	0.0040
MRLSTAR	0.1729	0.1560	0.0425	0.0285
MRLSTAR-h	0.1343	0.2175	0.0505	0.0328
	Sum of all linear: 0.1650		Sum of all MRLSTARs: 0.8350	
Italian Lira–French Franc				
Linear	0.0003	0.0004	0.0001	0.0001
Linear-h	0.0003	0.0005	0.0002	0.0001
MRLSTAR	0.0021	0.0584	0.0006	0.1889
MRLSTAR-h	0.0035	0.0436	0.0278	0.6729
	Sum of all linear: 0.0021		Sum of all MRLSTARs: 0.9979	

Note: Weights smaller than 0.0001 are reported as 0.0000.

This echoes the results in Murray and Pappell (2002) and Rossi (2005), which imply that estimation uncertainty for linear models is large. What is new here is that the MR-LSTAR models manage to not only generate smaller mean half-lives, but also less uncertainty about the range of possible half-lives.

Another new finding with our results compared to the previous literature is the overall model-averaged half-life (at the far bottom right corner for each panel in Table 4). Although we have estimates as low as 1.50 years, we also have an estimate as high as 5.50 years when Canada is involved. These model-averaged estimates are based on the weights using the marginal likelihood value for all models and lags. These weights are reported in Table 6. For certain data series, weights for a specific nonlinear model may reach above 90% (Canadian dollar–dollar and franc-dollar). Overall, nonlinear models always receive more than 80% of the weight in calculating model-averaged half-lives. The consequence of this weight on the nonlinear models is that the model-averaged half-lives are always as low or lower than in the linear case. Thus, we obtain a stronger result about the PPP persistence than is provided by the frequentist literature, which only finds less persistence when conditioning on a nonlinear model.

5. Conclusion

In this paper, we have employed Bayesian analysis to re-examine the previous empirical findings on real exchange rate persistence that were based on frequentist inferences. Our results strengthen some previous results about the importance of nonlinearities, but add important new insights about the general persistence of real exchange rates. In particular, in terms of uncertainty about half-lives, the nonlinear models yield more accurate inferences than linear models. Also, even when there is a nontrivial posterior probability for linear dynamics, there is clear evidence that the persistence of real exchange rates is lower than reported in Rogoff (1996) based on linear models, with the estimated half-life being 3 years or less for most country pairs. Thus, we confirm the frequentist results that condition on a nonlinear model and find that exchange rates are not quite as persistent as suggested by possibly misspecified linear models. Notably, however, our results imply less persistence, even when frequentist tests fail to reject linearity.

We conclude by noting that our analysis of exchange rate persistence is based on estimated models for which purchasing power parity appears to hold in the long run. It is possible, however, that there is a small random walk component in the real exchange rate (see, for example, Engel and Kim, 1999; Engel, 2000) and that explicitly modeling it would affect our inferences about the persistence of transitory deviations from the long-run equilibrium level of the real exchange rate. Incorporating nonlinear transitory dynamics in an unobserved components model that allows for stochastic permanent movements in the real exchange rate is a complicated econometric problem that we leave for future research. However, accounting for such movements should only serve to further reduce the estimated persistence of the transitory component of the real exchange rate and reinforce the empirical findings presented here.

References

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

Bec, F., Ben Salem, M., Carrasco, M., 2010. Detecting mean revision in real exchange rates from a multiple regime STAR model. Ann. Econ. Stat. 99/100, 1–33.

Cassel, G., 1922. Money and Foreign Exchange after 1914. Constable, London, U.K.

Chib, S., Jeliazkov, I., 2001. Marginal likelihood from the Metropolis-Hastings output. J. Am. Stat. Assoc. 96, 270–281.

Dumas, B., 1992. Dynamic equilibrium and the real exchange rate in spatially separated world. Rev. Financial Stud. 5, 153–180.

Engel, C., 2000. Long-run PPP may not hold after all. J. Int. Econ. 57, 243–273.

Engel, C., Kim, C.-J., 1999. The long-run U.S./U.K. real exchange rate. J. Money, Credit Bank. 31, 335–355.

Franses, P.H., van Dijk, D., 2000. Non-linear Time Series Models in Empirical Finance. Cambridge University Press, Cambridge, U.K.

Hansen, B.E., 1996. Inferences in TAR models. Stud. Nonlinear Dyn. Econ. 2 (1). Article 1.

Heckscher, E.F., 1916. Växelkursens grundval vid pappersmyntfot. Ekono-misk Tedskrift 18, 309–312.

Koop, G., Pesaran, M.H., Potter, S.M., 1996. Impulse response analysis in nonlinear multivariate models. J. Econ. 74, 119–147.

Lo, M.C., 2008. Nonlinear PPP deviations: a Monte Carlo investigation of their unconditional half-lives. Stud. Nonlinear Dyn. Econ. 12 (4). Article 5.

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

Michael, P., Nobay, A.R., Peel, D.A., 1997. Transaction costs and nonlinear adjustment in real exchange rates: an empirical investigation. J. Political Econ. 105, 862–879.

Murray, C., Pappell, D., 2002. The purchasing power parity paradigm. J. Int. Econ. 56, 1–19.

Obstfeld, M., Taylor, A.M., 1997. Nonlinear aspects of goods-market arbitrage and adjustment: Heckscher's commodity points revisited. J. Jpn. Int. Econ. 11, 441–479.

O'Connell, P.G.J., 1997. Perspectives on Purchasing Power Parity. Ph.D. dissertation. Harvard University.

O'Connell, P.G.J., Wei, S.-J., 2001. "The bigger they are, the harder they fall": retail price differences across U.S. cities. J. Int. Econ. 56, 21–53.

Potter, S.M., 2000. Nonlinear impulse response functions. J. Econ. Dyn. Control 24, 1425–1446.

Rogoff, K., 1996. The purchasing power parity puzzle. J. Econ. Literature 34, 647–668.
Rossi, B., 2005. Confidence intervals for half-life deviations from purchasing power parity. J. Bus. Econ. Stat. 23, 432–442.
Sarantis, N., 1999. Modeling non-linearities in effective real exchange rates. J. Int. Money Finance 18, 27–45.
Sarno, L., 2003. Nonlinear exchange rate models: a selective overview. Rivasta Polit. Econ. 93, 3–46.
Sarno, L., Taylor, M., Chowdhury, I., 2004. Nonlinear dynamics in deviations from the law of one Price: a Broad-based empirical study. J. Int. Money Finance 23, 1–25.
Taylor, M.P., Sarno, L., 2003. The Economics of Exchange Rates. Cambridge University Press, Cambridge, U.K.
Taylor, M.P., Peel, D.A., Sarno, L., 2001. Nonlinear mean-revision in real exchange rates: towards a solution of the purchasing power parity puzzles. Int. Econ. Rev. 42, 1015–1042.