Australian School of Business Research Paper No. 2013 ECON 05

Bayesian Analysis of Nonlinear Exchange Rate Dynamics and the Purchasing Power Parity Persistence Puzzle

Ming Chien Lo
James Morley

Bayesian Analysis of Nonlinear Exchange Rate Dynamics and the Purchasing Power Parity Persistence Puzzle*

Ming Chien Lo
Saint Cloud State University

James Morley
University of New South Wales

April 2013

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, but the strength of the evidence depends on which country pair is considered. However, the model-averaged estimates of half-lives are uniformly smaller than for the linear models alone, suggesting that the purchasing power parity persistence puzzle is less of a puzzle than previously thought.

Keywords: Bayesian Analysis; Real Exchange Rate Dynamics; Purchasing Power Parity; Nonlinear Threshold Models; Bayesian Model Averaging; Half lives
JEL Classification: C11; C22; F31

---

* **Lo:** Department of Economics, Stewart Hall 380, St. Cloud State University, St. Cloud, MN 56301-4498, 1-320-308-0143 (tel.), 1-320-308-2228 (fax), mclo@stcloudstate.edu; **Morley:** School of Economics, Australian School of Business, University of New South Wales, Sydney, NSW, 2052, +612 9385 3366 (tel.), +612 9313 6337 (fax), james.morley@unsw.edu.au. This paper has been presented 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 a seminar at Minnesota State University at Mankato. We thank conference and seminar participants for comments. The usual disclaimers apply.

1. Introduction

Numerous studies, including Michael, Nobay, and Peel (1997), Obstfeld and Taylor (1997), Sarantis (1999), Sarno, Taylor, and Chowdhury (2004), and Bec, Ben Salem, and Carrasco (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 affect when PPP is effective and when it is not.[1] 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

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

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 straightforward 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 the model and lag specification, the Bayesian approach allows for model-averaged measures that address inherent uncertainty about model-specification issues such as whether a linear or nonlinear model is more appropriate.

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 considerably across country pairs. However, our model-averaged measures of real exchange rate persistence are uniformly 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 only mixed 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 uniformly lower than those based on linear models, including in the cases where the evidence for nonlinearity is ambiguous. Specifically, we find that half-lives for G7 real exchange rates range between 2-3 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 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 stationary finite-order autoregressive (AR) model:

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

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 outside the unit circle. The stationarity assumption corresponds to the idea that PPP holds in the long run.[2] The Gaussian error assumption is driven by the need for a parametric structure in order to conduct our Bayesian analysis.[3]

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, Nobay, and Peel, 1997; Obstfeld and Taylor, 1997; Taylor, Peel, and Sarno, 2001; and Sarno, Chowdhury, and Taylor, 2004). Building on this literature and inspired by Franses and van Dijk (2000), Bec, Ben Salem, and Carrasco (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

[2] 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).

[3] 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 finding 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).

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

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

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

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

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 $\gamma \to \infty$, the MR-LSTAR model approximates a band-TAR model.

Given this setup, it is straightforward 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_{out}^2 and σ_{in}^2, 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 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 $\chi^2(1)$ assumption. 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 a parameter has a standard asymptotic distribution. Note that this calculation only applies to constructing the proposal distribution and that the posterior estimates should be robust to different assumptions for the asymptotic distribution, such as a $\chi^2(2)$, which we have verified in our empirical analysis. The important issue is that we obtain some sense of the curvature of the posterior with respect to the parameter, not to literally conduct frequentist inference. Then, we use the interval based on the inverted "posterior ratio" to 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. In particular, the approach proceeds as follows:

1) Construct "confidence set" for c based on inverting the posterior ratio.[4] Assuming the set is contiguous, denote the estimated 95% confidence interval as $[\hat{c}_{0.025}, \hat{c}_{0.975}]$.[5]

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

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

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:

$$\theta \sim MT(\mu^{\theta}, \Sigma^{\theta}, \nu^{\theta}),$$

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 $\theta = \left(\theta^{L}, \theta^{NL}\right)'$, Σ^{θ} is given as follows:

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

where κ is a "tuning" parameter for the MH algorithm, $\hat{\text{var}}(\hat{\theta}^{L})$ 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 $\hat{\text{var}}(\hat{\theta}^{NL}) = \left(\hat{\sigma}_{c}^{2}, \hat{\sigma}_{\gamma}^{2}\right)' I_{2\times 2}$ is based on the indirect estimated standard deviations discussed above. In practice, we consider different parameter blocking schemes (i.e., conditional drawing from subsets of θ) and we adjust κ to attain an acceptance rate for the MH algorithm of between 20-50%.

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 2007Q1 and euro currency exchange rates from 1974Q1 to 1998Q4. We calculate the real exchange rate series using nominal exchange rates and consumer price index data from the IFS database. 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.[6]

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

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

in Bec, Ben Salem and Carrasco (2010).[7] 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 2007Q1 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 Normal distribution. The prior means are all set to zero, except for the error-correction coefficient which is set to the OLS estimate based on an AR(1) model for the training sample. The prior standard deviations are 0.5, which are relatively uninformative, although we consider a truncation of the joint Normal distribution for the AR parameters to ensure stationarity (i.e., a draw from the proposal density can only be accepted if the roots of the characteristic equation $\phi(z) = 0$ lie outside the unit circle). The prior for the forecast error variance has a Gamma distribution $\sigma^2 \sim Gamma(\nu, \delta)$ where the Gamma distribution for variable x is parameterized as follows:

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

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

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, implying that the prior mean for the variance is equal to the sample variance in the training sample. This is a relatively uninformative prior that is common to all models and so does not affect the model comparison.

The elicitation of priors for the nonlinear parameters in the MR-LSTAR model is slightly more involved and requires some discussion. For the threshold parameter c, we assume a Gamma distribution and set the rate parameter $\delta = 0.5$ 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 in the training sample. For the smooth transition parameter γ, we also assume a Gamma distribution and set the rate parameter $\delta = 0.5$ and the shape parameter ν to $0.1 \times \delta$ times the mean non-zero absolute quarterly change in the real exchange rate over the training sample, implying the prior mean is equal to one-tenth of the mean absolute quarterly change. This calibrates the scale of the small change in the exchange rate that could produce a change in the dynamics of the exchange rate. Also, despite the use of training sample information to calibrate the means of the priors for both c and γ, the priors are still relatively uninformative given the low values for the rate parameter δ.

The remaining parameter for the nonlinear threshold models is the change in error correction coefficient across regimes: $\Delta\beta \equiv \beta_{in} - \beta_{out}$. For this parameter, our prior is somewhat 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 $\delta = 100$ and the shape parameter $\nu = 10$, implying the mean for the reduction in the error correction coefficient is 0.1, with a standard deviation of about 0.045. We justify this relatively informative prior 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.[8] Second, even though an informative prior might seem at first glance to push our empirical findings towards finding evidence of nonlinearity, it does not in fact do so.[9] 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 calculate posterior odds, while equal prior odds for a linear and

[8] 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.

[9] It should also be noted that the prior does not push the nonlinear models to imply shorter half lives for real exchange rates. This is because the prior is put on the change in persistence between regimes, not on the overall level of persistence, for which the prior is quite diffuse given the uninformative prior on the error correction coefficient.

nonlinear model would implicitly favor linearity given less informative priors on parameters in the nonlinear models related to nonlinearity.

4.2. Posteriors

Table 1 reports the logarithm of the marginal likelihood value, with the corresponding Bayes factor in the parentheses. Each Bayes factor is calculated as the ratio between the marginal likelihood value of a 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 seven out of eight series. Moreover, in every one of these cases, the second best model, using the Bayes factor as the measure, is also an MR-LSTAR specification. A linear model only has the highest Bayes factor overall for the Canadian-U.S. dollar real exchange rate. Even for that series, the MR-LSTAR(4)-h specification, where "h" denotes "heteroskedastic" disturbances, comes as a close second to the linear AR(4) specification, with a Bayes factor of 0.89. Meanwhile, our results suggest that both nonlinear conditional mean dynamics and heteroskedasticity are important in understanding the exchange rate data. Indeed, among the series in which nonlinearities are most probable, an MR-LSTAR-h specification has the highest Bayes factor in five out of seven 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, Ben Salem and Carrasco (2010).[10] Our sample periods are different than theirs, so the results in Table 2

---

[10] Bec, Ben Salem, and Carrasco (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 not an exact replication of their results. However, the inferences are almost the same. 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 all of the common series except the Canadian dollar-pound. However, Table 2 illustrates a difficulty with frequentist hypothesis testing in this context. Based on a pre-determined level of significance of even 15%, we fail to reject linearity for 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, our subsequent Bayesian inferences about persistence are much more consistent across country pairs by always giving weight to both linear and nonlinear models.[11]

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 Figure 1 respectively. The

---

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, Ben Salem and Carrasco (2010).

[11] 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, supporting the use of Bayesian model averaging as a way to combine models.

posterior means for the sum of the autoregressive coefficients within the threshold bands (i.e., $\beta_{in} + 1$) range from 0.904 to 0.982. 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 -0.093 to -0.048, with posterior standard deviations generally about half of the prior standard deviation of about 0.045. The data are informative, but not definitive about the magnitude of the change in persistence across regimes. Meanwhile, Figure 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, Figure 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: France/U.K. and Italy/France have the smallest threshold estimates 3.906 and 6.166 respectively. If thresholds represent the cost of arbitrage, we would expect European countries in the European Union enjoy smaller cost of transaction among them. Not surprisingly, all Europe/U.S. pairs with the Atlantic Ocean to separate the two continents apart have similar threshold estimates; the range is small, only from 7.015 to 8.080. The threshold for the Japan/U.S. real exchange rate is the largest among all: 14.807. The most interesting ones are series that 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 8.743, larger than any for the Europe/U.S. pairs. Also the threshold estimate for the Canadian dollar-pound real exchange rate is 10.455, which is larger than that for the U.S. dollar-pound, even though both North American countries are separated from the United Kingdom by a similar distance.

Based on draws of the parameters from their posterior distributions, we also compute and report posterior means for the unconditional half-life of real exchange rate deviations from PPP in Table 4. Given parameter values, the computation of half-lives for the AR models is conventional. The computation for the MR-LSTAR models requires generalized impulse response simulation as discussed in Koop, Pesaran, and Potter (1996) and Potter (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, Peel, and Sarno (2001) and Bec, Ben Salem, and Carrasco (2004), but is similar to Lo (2008).

We use the marginal likelihood value as relative weights to compute 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 up the decimals to the nearest 0.25 for the 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 are two extreme cases involving the Canadian dollar in which the half-life estimates for the *linear* models are above 6 years or longer. For the rest of the data series, however, the half-life estimates for the *linear* models fall right into Rogoff's range, from 3 to 4.25 years. These results are also sensitive to the lag length; in general, the larger the number of lags, the smaller the level of persistence. Using the Bayesian weighing scheme, we also find that the model-averaged half-life measures for the linear models (under "All Lags" corresponding to *linear* and *linear-h*) are 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 reduction ranges from 1 to 3 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 quantiles 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 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 quantiles 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 liner models is much higher than that for the nonlinear models and hits the level of "infinity" frequently.[12] 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 2.50 years, we also have estimate as high as 6 years when Canada is involved. These model-averaged estimates are based on the weights using the marginal likelihood value for all models and lags. The weights are reported in Table 6. For certain data series, weights for a specific nonlinear model may reach to 20% (lira-dollar) or even 25% (yen-dollar). Overall, nonlinear models receive about 70% to 80% of the weight in calculating model-averaged half-lives. The consequence of this weight on the nonlinear models is that the estimated half-lives are uniformly lower than in the linear case, even when the evidence against linear models is more ambiguous. Thus, we obtain a stronger result about the PPP persistence than provided by the frequentist literature, which only finds less persistence when conditioning on a nonlinear model.

[12] 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.

5. Conclusion

In this paper, we have employed Bayesian analysis to re-examine 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 though 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 (1995) based on linear models, with the range being between 2-3 years for most country pairs. Thus, we confirm the frequentist results that condition on a nonlinear model that one partial resolution of the purchasing power parity persistence puzzle is that exchange rates are not quite as persistent as suggested by possibly misspecified linear models. Notably, 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 the assumption that purchasing power parity holds 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, and Engel, 2000) and that allowing for 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. Of course, 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. Purchasing power parity in the long run. *Journal of 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. *Annals of Economics and Statistics* **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. *Journal of the American Statistical Association* **96**: 270-281.

Dumas, B. 1992. Dynamic equilibrium and the real exchange rate in spatially separated world. *Review of Financial Studies* **5**: 153-180.

Engel, C. 2000. Long-run PPP may not hold after all. *Journal of International Economics* **57**: 243-273.

Engel, C., Kim, C.-J. 1999. The long-run U.S./U.K. real exchange rate. *Journal of Money, Credit and Banking* **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. *Studies in Nonlinear Dynamics and Econometrics* **2**: No.1, Article 1.

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

Koop, G., Pesaran, M.H., and Potter, S.M. 1996. Impulse response analysis in nonlinear multivariate models *Journal of Econometrics* **74**: 119-147.

Lo, M.C. 2008. Nonlinear PPP deviations: a Monte Carlo investigation of their unconditional half-lives *Studies in Nonlinear Dynamics and Econometrics* **12**: No.4, Article 5.

Lothian, J.R., Taylor, M.P. Real exchange rate behavior: the recent float from the perspective of the past two centuries *Journal of Political Economics* **104**: 488-509.

Michael, P., Nobay, A.R., Peel, D.A. 1997. Transaction costs and nonlinear adjustment in real exchange rates: an empirical investigation. *Journal of Political Economy* **105**: 862-879.

Murray, C., Pappell, D. 2002. The purchasing power parity paradigm. *Journal of International Economics* **56**: 1-19.

Obstfeld, M., Taylor, A.M. 1997. Nonlinear aspects of goods-market arbitrage and adjustment: Heckscher's commodity points revisited. *Journal of the Japanese and International Economies* **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. *Journal of International Economics* **56**: 21-53.

Potter, S.M. 2000. Nonlinear impulse response functions. *Journal of Economic Dynamics and Control* **24**: 1425-1446.

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

Rossi, B. 2005. Confidence intervals for half-life deviations from purchasing power parity. *Journal of Business and Economic Statistics* **23**: 432-442.

Sarantis, N. 1999. Modeling non-linearities in effective real exchange rates. *Journal of International Money and Finance* **18**: 27-45.

Sarno, L. 2003. Nonlinear exchange rate models: a selective overview. *Rivasta di Politica Economica* **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. *Journal of International Money and 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. *International Economic Review* **42**: 1015-1042.

Table 1: Log Marginal Likelihood and Bayes Factors

British Pound-U.S. Dollar

Model	*Autoregressive Lag Order (p)*							
	1		2		3		4	
Linear	-335.48	(0.29)	-335.26	(0.36)	-336.66	(0.09)	-337.09	(0.06)
Linear-h	-335.08	(0.43)	-335.01	(0.46)	-336.44	(0.11)	-336.85	(0.07)
MRLSTAR	-335.26	(0.36)	-334.67	(0.64)	-336.03	(0.16)	-336.35	(0.12)
MRLSTAR-h	-334.50	(0.77)	-334.23	(1.00)	-335.62	(0.25)	-336.05	(0.16)

Canadian Dollar-U.S. Dollar

Model	*Autoregressive Lag Order (p)*							
	1		2		3		4	
Linear	-255.36	(0.30)	-255.84	(0.18)	-257.29	(0.04)	-254.15	(1.00)
Linear-h	-256.16	(0.13)	-256.67	(0.08)	-258.20	(0.02)	-254.95	(0.45)
MRLSTAR	-255.94	(0.17)	-256.27	(0.12)	-257.76	(0.03)	-254.27	(0.89)
MRLSTAR-h	-256.64	(0.08)	-256.98	(0.06)	-258.55	(0.01)	-254.90	(0.47)

French Franc-U.S. Dollar

Model	*Autoregressive Lag Order (p)*							
	1		2		3		4	
Linear	-241.54	(0.11)	-240.45	(0.32)	-241.74	(0.09)	-241.18	(0.16)
Linear-h	-241.91	(0.08)	-241.06	(0.18)	-242.31	(0.05)	-241.72	(0.09)
MRLSTAR	-240.56	(0.29)	-239.33	(1.00)	-240.65	(0.27)	-239.95	(0.54)
MRLSTAR-h	-240.63	(0.27)	-239.52	(0.83)	-240.56	(0.29)	-240.02	(0.50)

Italian Lira-U.S. Dollar

Model	*Autoregressive Lag Order (p)*							
	1		2		3		4	
Linear	-244.28	(0.09)	-243.48	(0.20)	-244.22	(0.10)	-243.25	(0.26)
Linear-h	-244.21	(0.10)	-243.63	(0.18)	-244.29	(0.09)	-243.13	(0.29)
MRLSTAR	-243.60	(0.18)	-242.50	(0.54)	-243.47	(0.21)	-241.92	(0.97)
MRLSTAR-h	-243.67	(0.17)	-242.71	(0.44)	-243.56	(0.19)	-241.89	(1.00)

Note: The Bayes factor is reported in parentheses; it is equal to the marginal likelihood value for each model and lag specification divided by the largest marginal value in the group.

Continued from the last page

Japanese Yen-U.S. Dollar

Model	*Autoregressive Lag Order (p)*							
	1		2		3		4	
Linear	-346.35	(0.18)	-346.74	(0.12)	-347.88	(0.04)	-346.02	(0.26)
Linear-h	-346.45	(0.17)	-346.92	(0.10)	-347.97	(0.04)	-345.72	(0.34)
MRLSTAR	-345.95	(0.27)	-346.35	(0.18)	-347.28	(0.07)	-345.14	(0.61)
MRLSTAR-h	-345.72	(0.34)	-345.85	(0.30)	-347.11	(0.09)	-344.65	(1.00)

Canadian Dollar-British Pound

Model	*Autoregressive Lag Order (p)*							
	1		2		3		4	
Linear	-337.16	(0.05)	-335.70	(0.21)	-336.47	(0.10)	-337.61	(0.03)
Linear-h	-337.27	(0.04)	-335.73	(0.21)	-336.49	(0.10)	-337.68	(0.03)
MRLSTAR	-336.29	(0.12)	-334.46	(0.74)	-335.60	(0.23)	-336.56	(0.09)
MRLSTAR-h	-336.07	(0.15)	-334.15	(1.00)	-335.04	(0.41)	-336.18	(0.13)

British Pound-French Franc

Model	*Autoregressive Lag Order (p)*							
	1		2		3		4	
Linear	-223.31	(0.09)	-223.31	(0.09)	-224.59	(0.02)	-224.90	(0.02)
Linear-h	-223.09	(0.11)	-222.52	(0.19)	-223.82	(0.05)	-224.47	(0.03)
MRLSTAR	-222.44	(0.21)	-222.26	(0.25)	-223.45	(0.08)	-223.72	(0.06)
MRLSTAR-h	-221.68	(0.44)	-220.87	(1.00)	-222.24	(0.25)	-222.75	(0.15)

Italian Lira-French Franc

Model	*Autoregressive Lag Order (p)*							
	1		2		3		4	
Linear	-187.17	(0.17)	-186.98	(0.21)	-188.23	(0.06)	-188.74	(0.04)
Linear-h	-187.33	(0.15)	-186.91	(0.22)	-188.02	(0.07)	-188.07	(0.07)
MRLSTAR	-186.77	(0.26)	-186.57	(0.32)	-187.71	(0.10)	-188.09	(0.07)
MRLSTAR-h	-186.42	(0.36)	-185.41	(1.00)	-186.80	(0.25)	-187.12	(0.18)

Table 2: Linearity Tests

LM Statistics

Model	*Autoregressive Lag Order (p)*			
	1	2	3	4
British Pound-U.S. Dollar	2.7959	4.4016 -	3.9650 -	5.4480*
Canadian-U.S. Dollar	0.3140	0.6371	0.4627	0.9430
French Franc-U.S. Dollar	3.3300	5.3418*	5.0728*	4.9556*
Italian Iira-U.S. Dollar	2.5983	3.4869	3.4183	2.9888
Japanese Yen-U.S. Dollar	1.3232	1.9722	1.8125	1.7733
Canadian Dollar-British Pound	0.8961	1.4968	1.1165	1.1604
British-Pound-French Franc	3.7321	4.6134*	4.8375*	4.8730*
Italian Lira-French Franc	5.3371*	8.4954**	8.1807**	7.6548**

Heteroskedasticity Robust LM Statistics

Model	*Autoregressive Lag Order (p)*			
	1	2	3	4
British Pound-U.S. Dollar	3.1279	5.0530*	4.7306*	5.8211*
Canadian-U.S. Dollar	0.3114	0.6576	0.4522	0.9246
French Franc-U.S. Dollar	3.0394	4.8514*	5.1414*	5.4710*
Italian Iira-U.S. Dollar	1.5188	2.1944	2.5005	2.4644
Japanese Yen-U.S. Dollar	0.4528	0.6649	0.6536	0.7372
Canadian Dollar-British Pound	1.2295	2.1482	1.7331	1.8727
British-Pound-French Franc	2.6618	3.3910 -	4.1377 -	4.4314 -
Italian Lira-French Franc	3.1254	4.3026 -	4.6046**	5.1828**

Note: -, *, and ** denote significance at 15%, 10%, and 5%, respectively.

Table 3: Posterior Means for the Best Nonlinear Models

	£/USD	CN$/US$	FF/US$	ITL/US$	¥/US$	CN$/£	£/FF	ITL/FF
σ_{out}^2	28.160	--	--	28.702	33.309	27.224	12.891	9.057
	(4.783)			(4.508)	(5.377)	(4.835)	(2.264)	(2.807)
σ_{in}^2	21.766	5.505	25.936	19.673	26.401	30.480	22.954	6.320
	(3.997)	(0.739)	(3.519)	(5.246)	(4.470)	(5.880)	(4.729)	(2.432)
β_{in}+1	0.961	0.982	0.959	0.959	0.972	0.967	0.904	0.961
	(0.029)	(0.014)	(0.030)	(0.031)	(0.021)	(0.025)	(0.053)	(0.029)
β_{out}-β_{in}	-0.071	-0.048	-0.070	-0.071	-0.060	-0.067	-0.093	-0.075
	(0.028)	(0.018)	(0.027)	(0.029)	(0.022)	(0.026)	(0.039)	(0.030)
c	8.743	8.080	7.576	7.015	14.807	10.455	3.906	6.166
	(4.706)	(5.020)	(5.281)	(6.252)	(8.038)	(5.032)	(1.953)	(3.357)
γ	5.459	10.204	7.846	8.850	3.313	3.537	12.303	10.753
	(5.400)	(10.276)	(7.695)	(8.894)	(3.350)	(3.645)	(11.687)	(10.124)

Note: Standard deviations in parentheses. For CN$/US$ and FF/US$, the variance σ_{in}^2 prevails in both regimes.

Table 4: Posterior Half-Life (Means and Weighted Averages)

British Pound-U.S. Dollar

Model	*Autoregressive Lag Order (p)*				
	1	2	3	4	All Lags
Linear	3.75	3.25	3.50	3.25	3.50
Linear-h	4.00	3.25	3.50	3.25	3.50
MRLSTAR	2.50	2.50	2.50	2.50	2.50
MRLSTAR-h	2.75	2.50	2.50	2.50	2.50
All Models	3.25	2.75	2.75	2.75	3.00

Canadian Dollar-U.S. Dollar

Model	*Autoregressive Lag Order (p)*				
	1	2	3	4	All Lags
Linear	8.00	7.75	7.50	6.75	7.00
Linear-h	8.00	7.75	7.75	6.75	7.25
MRLSTAR	5.00	4.75	4.75	4.50	4.50
MRLSTAR-h	5.00	4.75	4.75	4.75	4.75
All Models	7.00	6.50	6.50	5.75	6.00

French Franc-U.S. Dollar

Model	*Autoregressive Lag Order (p)*				
	1	2	3	4	All Lags
Linear	3.25	3.00	3.25	3.00	3.00
Linear-h	3.50	3.25	3.25	3.00	3.25
MRLSTAR	2.25	2.25	2.25	2.50	2.25
MRLSTAR-h	2.25	2.25	2.25	2.50	2.25
All Models	2.50	2.50	2.50	2.50	2.50

Italian Lira-U.S. Dollar

Model	*Autoregressive Lag Order (p)*				
	1	2	3	4	All Lags
Linear	4.25	3.50	3.75	3.50	3.50
Linear-h	4.00	3.50	4.00	3.50	3.75
MRLSTAR	2.75	2.75	2.50	2.75	2.75
MRLSTAR-h	2.50	2.50	2.25	2.50	2.50
All Models	3.00	2.75	3.00	2.75	3.00

Continued from the last page.

Japanese Yen-U.S. Dollar

Model	*Autoregressive Lag Order (p)*				
	1	2	3	4	All Lags
Linear	3.75	3.25	3.50	3.25	3.50
Linear-h	4.00	3.25	3.50	3.25	3.50
MRLSTAR	2.50	2.50	2.50	2.50	2.50
MRLSTAR-h	2.75	2.50	2.50	2.50	2.50
All Models	3.25	2.75	2.75	2.75	3.00

Canadian Dollar-British Pound

Model	*Autoregressive Lag Order (p)*				
	1	2	3	4	All Lags
Linear	8.00	7.75	7.50	6.75	7.00
Linear-h	8.00	7.75	7.75	6.75	7.25
MRLSTAR	5.00	4.75	4.75	4.50	4.50
MRLSTAR-h	5.00	4.75	4.75	4.75	4.75
All Models	7.00	6.50	6.50	5.75	6.00

British Pound-French Franc

Model	*Autoregressive Lag Order (p)*				
	1	2	3	4	All Lags
Linear	3.25	3.00	3.25	3.00	3.00
Linear-h	3.50	3.25	3.25	3.00	3.25
MRLSTAR	2.25	2.25	2.25	2.50	2.25
MRLSTAR-h	2.25	2.25	2.25	2.50	2.25
All Models	2.50	2.50	2.50	2.50	2.50

Italian Lira-French Franc

Model	*Autoregressive Lag Order (p)*				
	1	2	3	4	All Lags
Linear	4.25	3.50	3.75	3.50	3.50
Linear-h	4.00	3.50	4.00	3.50	3.75
MRLSTAR	2.75	2.75	2.50	2.75	2.75
MRLSTAR-h	2.50	2.50	2.25	2.50	2.50
All Models	3.00	2.75	3.00	2.75	3.00

Table 5: Posterior Half-Life (Medians, Quartiles and 95% Confidence Intervals)

British Pound-U.S. Dollar

	Lag	5%	25%	Median	75%	95%
Linear	1	1.50	2.25	2.75	4.00	∞
	2	1.50	2.00	2.75	3.50	12.00
	3	1.50	2.25	2.75	3.75	∞
	4	1.75	2.25	2.50	3.25	10.25
Linear-h	1	1.50	2.25	3.00	4.25	∞
	2	1.50	2.00	2.75	3.50	12.25
	3	1.50	2.00	2.75	3.75	∞
	4	1.50	2.25	2.50	3.50	10.75
MRLSTAR	1	1.00	1.75	2.25	3.00	6.00
	2	1.00	1.75	2.25	2.75	5.25
	3	1.00	1.75	2.25	2.75	5.00
	4	1.25	2.00	2.25	2.75	5.00
MRLSTAR-h	1	1.00	1.75	2.25	3.00	6.50
	2	1.00	1.75	2.25	2.75	5.50
	3	1.00	1.75	2.25	2.75	5.50
	4	1.00	2.00	2.25	2.75	5.25

Canadian Dollar-U.S. Dollar

	Lag	5%	25%	Median	75%	95%
Linear	1	2.75	4.50	6.50	11.50	∞
	2	3.00	4.50	6.25	10.25	∞
	3	2.75	4.25	6.25	10.00	∞
	4	3.00	4.25	5.50	7.75	∞
Linear-h	1	3.00	4.75	6.75	11.50	∞
	2	2.75	4.50	6.25	10.00	∞
	3	2.75	4.50	6.25	10.00	∞
	4	3.00	4.25	5.50	8.00	∞
MRLSTAR	1	1.25	3.00	4.00	6.00	∞
	2	1.25	3.00	4.00	5.75	14.75
	3	1.25	3.00	4.00	5.50	14.25
	4	1.50	3.25	4.00	5.00	11.00
MRLSTAR-h	1	1.25	3.00	4.00	5.75	∞
	2	1.25	3.00	4.00	5.50	14.25
	3	1.25	3.00	4.00	5.50	14.00
	4	1.50	3.25	4.00	5.25	12.00

Note: Our simulations of generalized impulse responses allow for a maximum of 15 years (or 60 quarters). Any simulation that hit the limit is regarded as ∞.

Continued from the last page.

	French Franc-U.S. Dollar					
	Lag	5%	25%	Median	75%	95%
Linear	1	1.50	2.00	2.50	3.50	12.00
	2	1.50	2.00	2.50	3.25	10.00
	3	1.50	2.00	2.50	3.25	10.25
	4	1.75	2.25	2.50	3.25	7.75
Linear-h	1	1.50	2.00	2.50	3.75	∞
	2	1.50	2.00	2.50	3.25	10.75
	3	1.50	2.00	2.50	3.25	10.50
	4	1.75	2.25	2.50	3.25	7.75
MRLSTAR	1	1.00	1.75	2.00	2.75	5.50
	2	1.25	1.75	2.00	2.50	4.75
	3	1.00	1.75	2.00	2.75	4.75
	4	1.25	2.00	2.25	2.75	4.50
MRLSTAR-h	1	1.25	1.75	2.00	2.50	4.50
	2	1.25	1.75	2.00	2.50	4.50
	3	1.00	1.75	2.25	2.75	4.75
	4	1.25	2.00	2.25	2.75	4.25

	Italian Lira-U.S. Dollar					
	Lag	5%	25%	Median	75%	95%
Linear	1	1.50	2.25	3.00	4.50	∞
	2	1.50	2.00	2.75	3.75	13.25
	3	1.50	2.25	2.75	4.00	∞
	4	1.75	2.25	2.75	3.75	12.00
Linear-h	1	1.50	2.25	2.75	4.25	∞
	2	1.50	2.25	2.75	3.75	∞
	3	1.50	2.25	3.00	4.25	∞
	4	1.75	2.25	2.75	3.75	12.75
MRLSTAR	1	1.00	1.75	2.25	3.00	7.00
	2	1.00	1.75	2.25	3.00	7.00
	3	1.00	1.75	2.25	3.00	7.00
	4	1.00	2.00	2.50	3.00	6.00
MRLSTAR-h	1	1.00	1.75	2.25	3.00	6.25
	2	1.00	1.75	2.25	2.75	6.25
	3	1.00	1.75	2.25	2.75	5.00
	4	1.25	2.00	2.25	3.00	5.00

Continued from the last page.

Japanese Yen-U.S. Dollar

	Lag	5%	25%	Median	75%	95%
Linear	1	2.00	3.00	4.00	6.25	∞
	2	2.00	3.00	3.75	5.50	∞
	3	2.00	3.00	4.00	6.00	∞
	4	2.25	3.00	3.75	5.00	∞
Linear-h	1	2.00	3.00	4.00	6.25	∞
	2	2.00	3.00	3.75	5.50	∞
	3	2.00	3.00	4.00	6.00	∞
	4	2.25	3.00	3.75	5.00	∞
MRLSTAR	1	1.25	2.25	3.00	4.00	9.25
	2	1.25	2.25	2.75	3.75	8.00
	3	1.00	2.25	2.75	3.75	8.25
	4	1.50	2.50	3.00	3.75	7.25
MRLSTAR-h	1	1.00	2.25	3.00	4.25	10.00
	2	1.25	2.25	3.00	4.00	8.75
	3	1.00	2.25	3.00	4.00	8.75
	4	1.25	2.50	3.00	4.00	8.25

Canadian Dollar-British Pound

	Lag	5%	25%	Median	75%	95%
Linear	1	1.75	2.50	3.50	5.25	∞
	2	1.75	2.50	3.00	4.50	∞
	3	1.75	2.50	3.50	5.00	∞
	4	1.50	2.50	3.25	5.25	∞
Linear-h	1	1.75	2.50	3.50	5.50	∞
	2	1.75	2.50	3.00	4.25	∞
	3	1.75	2.50	3.25	5.00	∞
	4	1.25	2.50	3.25	5.00	∞
MRLSTAR	1	1.00	2.00	2.50	3.25	6.50
	2	1.25	2.00	2.50	3.00	6.00
	3	1.00	2.00	2.50	3.00	6.00
	4	1.00	1.75	2.50	3.00	6.00
MRLSTAR-h	1	1.00	1.75	2.50	3.25	6.50
	2	1.25	2.00	2.25	3.00	5.50
	3	1.00	2.00	2.50	3.25	6.00
	4	1.00	1.75	2.50	3.00	6.00

Continued from the last page.

British Pound-French Franc

	Lag	5%	25%	Median	75%	95%
Linear	1	1.00	1.25	1.50	2.00	5.00
	2	1.00	1.25	1.50	1.75	3.50
	3	1.00	1.25	1.50	1.75	3.75
	4	1.00	1.50	1.50	1.75	3.50
Linear-h	1	1.00	1.25	1.50	1.75	3.75
	2	1.00	1.25	1.50	1.75	2.75
	3	1.00	1.25	1.50	1.75	3.25
	4	1.00	1.25	1.50	1.75	2.75
MRLSTAR	1	0.75	1.25	1.50	1.75	3.25
	2	0.75	1.25	1.50	1.75	3.00
	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	0.75	1.25	1.50	1.75	2.75
	2	1.00	1.25	1.25	1.50	2.50
	3	0.75	1.25	1.25	1.50	2.50
	4	0.75	1.25	1.50	1.75	2.50

Italian Lira-French Franc

	Lag	5%	25%	Median	75%	95%
Linear	1	1.50	2.00	2.75	3.75	∞
	2	1.50	2.25	2.75	3.75	14.00
	3	1.50	2.00	2.75	3.75	14.25
	4	1.75	2.25	2.75	3.75	11.75
Linear-h	1	1.50	2.00	2.75	4.00	∞
	2	1.50	2.25	2.75	4.00	∞
	3	1.50	2.00	2.75	3.75	∞
	4	1.50	2.25	2.75	3.75	11.50
MRLSTAR	1	1.00	1.75	2.25	3.00	6.50
	2	1.25	1.75	2.25	2.75	5.50
	3	1.00	1.75	2.25	2.75	5.50
	4	1.25	2.00	2.25	2.75	4.75
MRLSTAR-h	1	1.00	1.75	2.25	3.00	6.75
	2	1.00	1.75	2.25	3.00	6.50
	3	1.00	1.75	2.25	3.00	5.75
	4	1.00	2.00	2.50	3.00	5.75

Table 6: Weights from Marginal Likelihood Values

British Pound-U.S. Dollar

Model	*Autoregressive Lag Order (p)*			
	1	2	3	4
Linear	0.0539	0.0671	0.0166	0.0107
Linear-h	0.0804	0.0862	0.0206	0.0137
MRLSTAR	0.0668	0.1208	0.0310	0.0226
MRLSTAR-h	0.1438	0.1880	0.0470	0.0306
	Sum of all linear:	0.3494	Sum of all MRLSTARs:	0.6506

Canadian Dollar-U.S. Dollar

Model	*Autoregressive Lag Order (p)*			
	1	2	3	4
Linear	0.0737	0.0457	0.0107	0.2481
Linear-h	0.0331	0.0200	0.0043	0.1110
MRLSTAR	0.0414	0.0298	0.0067	0.2202
MRLSTAR-h	0.0205	0.0147	0.0030	0.1171
	Sum of all linear:	0.5465	Sum of all MRLSTARs:	0.4535

French Franc-U.S. Dollar

Model	*Autoregressive Lag Order (p)*			
	1	2	3	4
Linear	0.0216	0.0640	0.0176	0.0309
Linear-h	0.0150	0.0350	0.0100	0.0181
MRLSTAR	0.0577	0.1974	0.0526	0.1062
MRLSTAR-h	0.0539	0.1636	0.0578	0.0985
	Sum of all linear:	0.2123	Sum of all MRLSTARs:	0.7877

Italian Lira-U.S. Dollar

Model	*Autoregressive Lag Order (p)*			
	1	2	3	4
Linear	0.0184	0.0407	0.0195	0.0511
Linear-h	0.0197	0.0352	0.0181	0.0579
MRLSTAR	0.0361	0.1085	0.0414	0.1944
MRLSTAR-h	0.0336	0.0878	0.0377	0.1997
	Sum of all linear:	0.2607	Sum of all MRLSTARs:	0.7393

Continued from the last page.

Japanese Yen-U.S. Dollar

Model	*Autoregressive Lag Order (p)*			
	1	2	3	4
Linear	0.0444	0.0301	0.0097	0.0620
Linear-h	0.0402	0.0253	0.0088	0.0831
MRLSTAR	0.0662	0.0444	0.0175	0.1487
MRLSTAR-h	0.0834	0.0731	0.0208	0.2423
	Sum of all linear:	0.3036	Sum of all MRLSTARs:	0.6964

Canadian Dollar-British Pound

Model	*Autoregressive Lag Order (p)*			
	1	2	3	4
Linear	0.0136	0.0582	0.0272	0.0087
Linear-h	0.0122	0.0567	0.0266	0.0081
MRLSTAR	0.0324	0.2026	0.0644	0.0248
MRLSTAR-h	0.0404	0.2746	0.1134	0.0362
	Sum of all linear:	0.2111	Sum of all MRLSTARs:	0.7889

British Pound-French Franc

Model	*Autoregressive Lag Order (p)*			
	1	2	3	4
Linear	0.0287	0.0287	0.0080	0.0058
Linear-h	0.0358	0.0633	0.0172	0.0090
MRLSTAR	0.0684	0.0820	0.0248	0.0190
MRLSTAR-h	0.1460	0.3290	0.0837	0.0504
	Sum of all linear:	0.1966	Sum of all MRLSTARs:	0.8034

Italian Lira-French Franc

Model	*Autoregressive Lag Order (p)*			
	1	2	3	4
Linear	0.0488	0.0593	0.0170	0.0101
Linear-h	0.0415	0.0637	0.0209	0.0199
MRLSTAR	0.0730	0.0894	0.0285	0.0195
MRLSTAR-h	0.1034	0.2835	0.0704	0.0511
	Sum of all linear:	0.2812	Sum of all MRLSTARs:	0.7188

Figure 1: Empirical Transition Functions for the Best Nonlinear Models