

Kilian, Lutz; Zha, Tao

Working Paper

Quantifying the half-life of deviations from PPP: The role of economic priors

Working Paper, No. 99-21

Provided in Cooperation with:

Federal Reserve Bank of Atlanta

Suggested Citation: Kilian, Lutz; Zha, Tao (1999) : Quantifying the half-life of deviations from PPP: The role of economic priors, Working Paper, No. 99-21, Federal Reserve Bank of Atlanta, Atlanta, GA

This Version is available at:

<https://hdl.handle.net/10419/100887>

Standard-Nutzungsbedingungen:

Die Dokumente auf EconStor dürfen zu eigenen wissenschaftlichen Zwecken und zum Privatgebrauch gespeichert und kopiert werden.

Sie dürfen die Dokumente nicht für öffentliche oder kommerzielle Zwecke vervielfältigen, öffentlich ausstellen, öffentlich zugänglich machen, vertreiben oder anderweitig nutzen.

Sofern die Verfasser die Dokumente unter Open-Content-Lizenzen (insbesondere CC-Lizenzen) zur Verfügung gestellt haben sollten, gelten abweichend von diesen Nutzungsbedingungen die in der dort genannten Lizenz gewährten Nutzungsrechte.

Terms of use:

Documents in EconStor may be saved and copied for your personal and scholarly purposes.

You are not to copy documents for public or commercial purposes, to exhibit the documents publicly, to make them publicly available on the internet, or to distribute or otherwise use the documents in public.

If the documents have been made available under an Open Content Licence (especially Creative Commons Licences), you may exercise further usage rights as specified in the indicated licence.

FEDERAL
RESERVE
BANK
of ATLANTA

**Quantifying the Half-Life of Deviations from PPP:
The Role of Economic Priors**

Lutz Kilian and Tao Zha

Working Paper 99-21
December 1999

Working Paper Series

Quantifying the Half-Life of Deviations from PPP: The Role of Economic Priors

Lutz Kilian and Tao Zha

Federal Reserve Bank of Atlanta
Working Paper 99-21
December 1999

Abstract: The half-life of deviations from purchasing power parity (PPP) plays a central role in the ongoing debate about the ability of macroeconomic models to account for the time series behavior of the real exchange rate. The main contribution of this paper is a general framework in which alternative priors for the half-life of deviations from PPP can be examined. We show how to incorporate formally the prior views of economists about the half-life. In our empirical analysis we provide two examples of such priors. One example is a consensus prior consistent with widely held views among economists with a professional interest in the PPP debate. The other example is a relatively diffuse prior designed to capture a large degree of uncertainty about the half-life. Our methodology allows us to make explicit probability statements about the half-life and to assess the likelihood that the half-life exceeds a given number of years, without taking a stand on whether or not the data have a unit root. We find only very limited support for the common view in the PPP literature that the half-life is between three and five years.

JEL classification: F41, C11

Key words: purchasing power parity, half-life, mean reversion, persistence, likelihood

The authors gratefully acknowledge helpful comments from Bob Barsky, Kathryn Dominguez, Charlie Evans, John Geweke, Gordon Hanson, Jan Kmenta, Chris Sims, Linda Tesar, Harald Uhlig, Dan Waggoner, Chuck Whiteman, and seminar participants at the 1999 NBER Summer Institute in Boston. They also thank Dave Backus, Dick Baillie, Bob Barsky, Menzie Chinn, Alan Deardorff, Kathryn Dominguez, Sebastian Edwards, Charles Engel, Jürgen von Hagen, Gordon Hanson, Miles Kimball, Ken Kletzer, Bob Lippens, Nelson Mark, Mike Melvin, David Papell, Ken Rogoff, Andrew Rose, Matthew Shapiro, Chris Sims, Linda Tesar, and Jaume Ventura for participating in the survey. The views expressed here are the authors' and not necessarily those of the Federal Reserve Bank of Atlanta or the Federal Reserve System. Any remaining errors are the authors' responsibility.

Please address questions regarding content to Lutz Kilian, Department of Economics, University of Michigan, Ann Arbor, Michigan 48109-1220, 734/764-2320, 734/764-2769 (fax), lkilian@umich.edu; or Tao Zha, Research Department, Federal Reserve Bank of Atlanta, 104 Marietta Street, NW, Atlanta, Georgia 30303-2713, 404/521-8353, 404/521-8956 (fax), tz.zha@atl.frb.org.

To receive notification about new papers, please use the publications order form on this Web site, or contact the Public Affairs Department, Federal Reserve Bank of Atlanta, 104 Marietta Street, NW, Atlanta, Georgia 30303-2713, 404/521-8020.

Quantifying the Half-Life of Deviations from PPP: The Role of Economic Priors

1. Introduction

The half-life of deviations from purchasing power parity (PPP) is a commonly-used measure of the degree of mean reversion in real exchange rates.¹ The half-life may be obtained from the impulse response function of the univariate time series representation of the real exchange rate. It is defined as the number of years that it takes for deviations from PPP to subside permanently below 0.5 in response to a unit shock in the level of the series. This particular notion of the degree of mean reversion in real exchange rates plays an important role in the ongoing debate about the ability of macroeconomic models to account for the time series behavior of the real exchange rate. The following quote by Rogoff (1996, p. 664, emphasis added) conveys the essence of the debate:

“It would seem hard to explain the short-term volatility [of real exchange rates] without a dominant role for shocks to money and financial markets. But given that such shocks should be largely neutral in the medium run, it is hard to see how this explanation is consistent with a *half-life for PPP deviations of three to five years.*”

Despite the prominent role of the half-life in this debate, there does not exist a methodology for assessing the probability that the value of the half-life is contained in a given range. This paper provides a general Bayesian framework for this purpose. Our methodology allows us to derive the entire probability distribution of the half-life. Specifically, we may ask

¹ See Frankel (1986), Frankel and Froot (1987), Abuaf and Jorion (1990), Diebold, Husted and Rush (1991), Wei and Parsley (1995), Froot and Rogoff (1995), Parsley and Wei (1996), Lothian and Taylor (1996), Wu (1996), Rogoff (1996), Papell (1997), Taylor and Peel (1998), Caner and Kilian (1999), Cheung and Lai (1999), Murray and Papell (1999).

what the probability is that the half-life exceeds a given number of years. Such questions arise naturally in economic discussions of long-run PPP, because different theoretical models have different implications for the half-life. For example, it is often argued that models without real rigidities can account for a half life of at most one or two years (see Cheung and Lai 1999). In contrast, models with real rigidities may account for much longer half-lives. In addition, researchers may be interested in the likelihood that the half-life is contained in a range of, say, three to five years, as suggested by Rogoff (1996).

Our methodology differs considerably from the existing literature on the subject of long-run PPP. First, we are concerned with the probability distribution of the half-life rather than the question of a unit root in real exchange rates. Indeed, our analysis does not depend on whether the process has a unit root or not. Second, our analysis allows economists to incorporate formally into the analysis their prior views about the half-life. Finally, our analysis also differs from recent work by Murray and Papell (1999) who provide classical confidence intervals for the half-life. As shown by Sims and Uhlig (1991) and Sims and Zha (1999), classical confidence intervals are not in general suited for making probability statements about model parameters.

We illustrate the practical feasibility of our approach in the empirical section of the paper. We consider two examples of prior distributions. The first example is a prior designed to be consistent with widely held views among economists with a professional interest in the PPP debate. The other example is a relatively diffuse prior designed to capture a large degree of uncertainty about the half-life.

The remainder of the paper is organized as follows: In section 2, we outline the advantages of the half-life in assessing the evidence of long-run PPP. In section 3, we describe the general framework of our analysis. Section 4 contains a discussion of the prior specifications

chosen for the empirical analysis. The empirical findings are presented in section 5 and the concluding remarks in Section 6.

2. Why the Half-Life is of Economic Interest

It is well known that the existence of long-run PPP is inconsistent with unit roots in the real exchange rate process. As a result, much of the attention of the profession has been focused on the question of whether the unit root hypothesis can be rejected or not. While there is increasing evidence against the unit root hypothesis, it has proved difficult to unambiguously reject the unit root null hypothesis for the recent floating rate period. In response, the profession has embarked on a quest for ever more powerful unit root tests in an attempt to resolve the PPP debate, including multivariate tests (see Edison et al. 1997; Taylor and Sarno 1998), panel data tests (see Wu and Wu 1998; Koedijk et al. 1998; Papell and Theodoridis 1998) and asymptotically efficient tests (see Cheung and Lai 1998).

We view attempts to resolve the PPP debate by means of unit root tests as misguided for two reasons. First, to economists long-run PPP means more than the absence of a unit root. It means a sufficient degree of mean reversion in real exchange rates for the predictions of theoretical models based on the PPP assumption to provide an adequate description of reality at the horizons of interest. A rejection of the unit root null hypothesis is consistent with any stationary process, including processes with a root very close to unity. Conversely, processes with small unit root components may nevertheless be strongly mean-reverting over the horizon of interest to economists. Thus, tests of the unit root null hypothesis by construction are unable to provide guidance to economists as to whether the PPP assumption should be abandoned in economic modeling or not. Even if we were able to reject the unit root hypothesis for all real

exchange rates, we would have learned little about the validity of the long-run PPP assumption in economic modeling.²

Second, what matters from an economic point of view is the *degree* of mean reversion in real exchange rates over the horizons of economic interest. Tests of the null hypothesis of a unit root in real exchange rates by construction are not suited for examining the speed of mean reversion. It may seem that the speed of mean reversion is adequately captured by the magnitude of the estimated root of the process. However, in higher-order processes the largest root may have little relation to the degree of mean reversion over horizons of economic interest.

This paper marks a fundamental departure from the preoccupation with the dominant root of the autoregressive representation of real exchange rates. We focus on the half-life of deviations from PPP because it provides a measure of the degree of mean reversion. This change in focus allows us to bypass the many difficulties in interpreting unit root tests and to address directly various questions of interest to international economists.

As noted in the introduction, the half-life of the real exchange rate process is defined as the number of years that it takes for deviations from PPP to subside permanently below 0.5 in response to a unit shock in the level of the series. Let f denote the sampling frequency of the data ($f = 1$ for years; $f = 4$ for quarters; $f = 12$ for months, etc.). Let $\phi(i)$ denote the response of the real exchange rate to a unit shock i periods ago. In practice, we proceed as follows: First, find the largest i in the range $i = 1, \dots, 40f$ for which $|\phi(i)| = 0.5$. Denote that i by h .³ Second, verify that $|\phi(j)| < 0.5$ for all $j > h$ for at least another forty years. This condition effectively

² Similar problems arise in tests of the $I(0)$ null hypothesis against $I(1)$ alternatives (e.g., Culver and Papell, 1997).

³ We focus on a horizon of 40 years because this horizon is a conservative upper limit on the horizons of interest to macroeconomists.

rules out unstable or explosive oscillatory pattern about 0.5.⁴ If h satisfies this second condition, we say that h is the half-life. If not, we say that the half-life is not reached within forty years. Note that in general there is no direct relationship between the root of an autoregressive process and the half-life. In fact, an AR(p) process may have an arbitrarily short half-life despite the presence of a unit root component.

3. A Bayesian Framework for Evaluating the Probability Distribution of the Half-Life

We first introduce some notation. Assume that the real exchange rate may be represented as an AR(p) process

$$y_t = c + \rho_1 y_{t-1} + \dots + \rho_p y_{t-p} + u_t, \quad t = 1, \dots, T, \quad (1)$$

with $u_t \sim N(0, \sigma^2)$. Let

$$b_{(p+1) \times 1} = \begin{bmatrix} \rho_1 \\ \vdots \\ \rho_p \\ c \end{bmatrix}, \quad \gamma = \frac{1}{\sigma^2}, \quad x_{(p+1) \times 1} = \begin{bmatrix} y_{t-1} \\ \vdots \\ y_{t-p} \\ 1 \end{bmatrix}, \quad X_{Tx \times (p+1)} = \begin{bmatrix} x_0' \\ \vdots \\ x_{T-1}' \end{bmatrix}, \text{ and } Y_{Tx \times 1} = \begin{bmatrix} y_1 \\ \vdots \\ y_T \end{bmatrix}. \quad (2)$$

Then the likelihood function can be written as

$$L(y_1, \dots, y_T | x_0, b, \gamma) \propto \gamma^T \exp \left\{ \frac{\gamma^2}{2} (b' X' X) b - 2b' (X' Y) + Y' Y \right\}. \quad (3)$$

The priors on b and γ take the form

$$p(\gamma) = \varphi(0, \sigma_\gamma^2), \text{ and } p(b) = \sum_{i=1}^n p_i \varphi(m_{b,i}, V_{b,i}), \quad (4)$$

where $p_1 + \dots + p_n = 1$ and $\varphi(x, y)$ is the (properly scaled) normal density function with mean x

⁴ The choice of 40 years in the second step also is conservative. Our empirical results are robust to further increases in the time horizon.

and covariance γ . $p(\gamma)$ is a diffuse normal prior distribution intended to capture our ignorance about the magnitude of the innovations. $p(b)$ is a mixed normal distribution. By appropriately choosing the number of mixed normals and their parameters we are able to approximate a wide variety of distributional shapes for the half-life. This flexibility in turn will allow us to incorporate the prior views of economists into the empirical analysis.

Our prior specification has the additional advantage that we do not dogmatically rule out the non-stationary region of the parameter space. Rather we assign little probability mass to this region a priori. This fact is important because, as we will show, some economists attach positive probability weight to the nonstationary region.

In related work on the AR(1) model Schotman and van Dijk (1991, p. 208) stress the importance of specifying a weakly informative prior for the process mean that smoothly blends into an uninformative prior as $\rho_1 \rightarrow 1$. The concern is that an AR(1) process that is close to a random walk will tend to exhibit trending behavior, unless c is close to zero. We address this concern by using a dummy observation prior, as suggested by Sims and Zha (1998).

Specifically, we add one dummy observation of the form

$$\tau \bar{y}_0 = b, \begin{bmatrix} \tau \bar{y}_0 \\ \vdots \\ \tau \bar{y}_0 \\ \tau \end{bmatrix}, \quad (5)$$

where \bar{y}_0 is the average of the initial values and τ is a hyperparameter. Expression (5) can be reduced to $(1 - \rho_1 - \dots - \rho_p) \bar{y}_0 = c$, which implies that a unit root exists if and only if $c = 0$.

Thus, when the data imply a near-unit root, the constant term must be small.

Combining (3) and (4) and letting $\lambda = \gamma^2$, one can derive the following posterior:

$$p(b, \lambda | Y, x_0) = p(\lambda | Y, x_0) p(b | \lambda, Y, x_0),$$

$$p(\lambda | Y, x_0) = \sum_{i=1}^n p_i(\lambda | Y, x_0), \quad (6)$$

$$p(b | Y, x_0, \lambda) = \sum_{i=1}^n \pi_i(\lambda | Y, x_0) \phi(\tilde{m}_b^i, \tilde{V}_b^i), \quad (7)$$

where,

$$p_i(\lambda | Y, x_0) = (\lambda)^{(T-1)/2} |\tilde{V}_{b,i}|^{1/2} C_i \exp \left\{ -\frac{\lambda}{2} \left(\frac{1}{\sigma_\gamma^2} + Y'Y \right) + \frac{1}{2} \left(V_{b,i}^{-1} m_{b,i} + \lambda (X'Y) \right)' \tilde{m}_{b,i} \right\}, \quad (8)$$

$$\pi_i(\lambda | Y, x_0) = \frac{p_i(\lambda | Y, x_0)}{p(\lambda | Y, x_0)}, \quad (9)$$

$$\tilde{V}_{b,i} = \left(V_{b,i}^{-1} + \lambda (X'X) \right)^{-1}, \text{ and } \tilde{m}_{b,i} = \tilde{V}_{b,i} \left(V_{b,i}^{-1} m_b^i + \lambda X'Y \right). \quad (10)$$

C_i is a constant term chosen such that the right hand side of (8) integrates to 1 over $(0, \infty)$. The term C_i ensures that $p_i(\lambda | Y, x_0)$ is a properly scaled density function. This term can be easily computed by standard numeral integration methods because $p_i(\lambda | Y, x_0)$ is one-dimensional and its shape is quite smooth.

The form of the posterior distribution in expressions (6) and (7) facilitates the simulation of the posterior of the half-life. We first generate draws from the marginal posterior distribution of λ in (6). Then we generate draws from the posterior distribution (7) of b conditional on λ . Finally, we construct the impulse response function associated with each draw for b and read off the half-life. The only complication is that the marginal posterior distribution (6) of λ is non-standard and cannot be simulated directly. We solve this problem by using a Metropolis-Hastings algorithm specifically designed for this problem. Consider the density function

$$J(\lambda|Y, x_0) = \sum_{i=1}^n p_i J_i(\lambda|Y, x_0), \quad (11)$$

where

$$J_i(\lambda|Y, x_0) = \frac{\hat{\Omega}_i^{T/2}}{2^{T/2} \Gamma(T/2)} \lambda^{T/2-1} \exp\left\{-\frac{\hat{\Omega}_i}{2} \lambda\right\}, \quad (12)$$

$$\hat{\Omega}_i = (Y - X\hat{B}_i)'(Y - X\hat{B}_i),$$

$$\hat{B}_i = \tilde{V}_{b,i} \left(V_{b,i}^{-1} m_{b,i} + \hat{\lambda}_i X'Y \right),$$

$$\hat{\lambda}_i = \arg \max p_i(\lambda|Y, x_0).$$

Note that $\Gamma(x)$ in (12) is the standard gamma function, i.e.,

$$\Gamma(x) = \int_0^{\infty} e^{-t} t^{x-1} dt.$$

The density (11) may be obtained as a weighted sum of draws generated from (12) for each $i = 1, \dots, n$. To obtain draws from (12) first draw a vector $z = (z_1, \dots, z_{T-1})'$ identically and independently from a normal distribution with mean zero and variance $1/\hat{\Omega}_i$. Then form $\lambda_i^* = z'z$. The random draw λ_i^* so obtained is from the distribution (12). The resulting density function (11) is used as the jumping kernel for our Metropolis-Hastings procedure. Specifically, our procedure involves four steps:

Metropolis-Hastings Algorithm. Initialize an arbitrary value λ^0 in R^+ (this value is usually set equal to the maximum point $\hat{\lambda}$). For $n = 1, \dots, N_1 + N_2$,

- (a) generate λ^* from $J(\lambda|Y, x_0)$ and u from the uniform distribution $U(0,1)$;
- (b) compute

$$q = \min \left\{ \frac{p(\lambda^* | Y, x_0) / J(\lambda^{(n-1)} | Y, x_0)}{p(\lambda^{(n-1)} | Y, x_0) / J(\lambda^* | Y, x_0)}, 1 \right\};$$

(c) if $u \leq q$, set $\lambda^{(n)} = \lambda^*$; if not, set $\lambda^{(n)} = \lambda^{(n-1)}$;

(d) collect the simulated sequence $\{\lambda^{(1)}, \dots, \lambda^{(N_1+N_2)}\}$ but retain only the last N_2 values of the sequence.

Geweke (1995) shows that this algorithm generates a sequence of random draws whose distribution converges to the target distribution $p(\lambda | Y, x_0)$.⁵ The remaining steps of the simulation of the posterior distribution of the half-life are straightforward. Given a set of posterior draws for λ , the posterior distribution (7) of b conditional on λ is mixed normal and can be easily simulated. Finally, for each draw for b , we calculate the associated impulse response function and compute h as described in section 2.

4. Construction of the Priors

Our Bayesian analysis differs from the existing literature in that it allows us to incorporate formally into the analysis prior views of economists about the half-life of deviations from PPP. A natural starting point in formulating an economically informed prior is the existing PPP literature. It is well known that theoretical models with intertemporal smoothing of consumption goods (see Rogoff 1992) or cross-country wealth redistribution effects (see Obstfeld and Rogoff 1995) imply highly persistent deviations from PPP. Similarly, terms-of-trade shocks may have long-lasting effects by affecting the structure of the economy. On the

⁵ We use the evaluation method of Waggoner and Zha (1999) to examine the convergence properties of this Metropolis-Hastings Algorithm for our data. We find that for 1.4 draws of the parameter λ in the Metropolis-Hastings procedure correspond to at least 1 effective draw.

other hand, monetary models of the exchange rate with only nominal rigidities imply much shorter half-lives. For a review of this literature see Rogoff (1996).

Despite the diversity of theoretical models, the prominent view in the recent PPP literature is that 3-5 years are likely values for the half-life of shocks to the real exchange rate under the recent float. For example, Abuaf and Jorion (1990, p. 173) suggest a half-life of 3-5 years for the post-Bretton Woods era. Rogoff (1996, pp. 657-658) conjectures that deviations from PPP dampen out at the rate of about 15 percent per year. Cheung and Lai (1999) present confidence intervals with lower bounds corresponding to a half-life of less than 1.5 years. Wei and Parsley (1995) arrive at half-life estimates of about 4.25 years for non-EMS countries using panel data models (and 4.75 for EMS countries). Wu (1996) and Papell (1997) find half-lives of 2.5 years using panel data methods. Froot and Rogoff (1995), p. 1645) consider a half-life of 3-5 years plausible. Murray and Papell (1999) note substantial sampling uncertainty in half-life estimates based on univariate models, but provide strong evidence based on panel data consistent with the claim of a half-life of 3-5 years. While the information provided by these sources is necessarily incomplete, it suggests some essential features of a prior density for the half-life. Clearly, for economists a half-life of 1 year or of 100 years is less reasonable a priori than a half-life of 3, 4 or 5 years. Hence, we know that our prior for h must be informative in a way that gives less weight to extremely persistent processes and to processes with low persistence. Moreover, the half-life is bounded from below at zero and unbounded from above, and the prior distribution is likely to be asymmetric with a long right tail. Finally, the mode of the prior density corresponds to a half-life of about 4 years.

An alternative approach is to elicit the prior distribution directly from economists in the field of international economics and macroeconomics with a professional interest in the PPP

debate. We conducted such a survey in July and August of 1999. Table 1 shows the average prior density based on the responses of 22 economists with a professional interest in the PPP question. The participants were asked to specify probability weights for ranges of half-lives in response to a reduced form innovation (or forecast error) in the univariate time series representation of the real exchange rate. The average survey responses display the key features suggested by the literature review, notably a single peak at about 4 years and a long right tail of the prior probability density with little probability mass allocated to half-lives in excess of 20 years.

The survey beliefs, though reflecting in part the time-series patterns of observed real exchange rates and being necessarily correlated, differ from what is implied by the likelihood of our parametric model (see Table 2). For example, for the U.K. (Canada), the likelihood implies a probability of 52.0% (0.0%) that the half-life does not exceed 2 years compared with 14.1% in the survey prior. Even after averaging the probabilities implied by the likelihood across countries, substantial differences remain. For example, the average probability that the half-life does not exceed 4 years is 75.2 % in the data compared with 49.6% in the survey.

In the empirical analysis, we will postulate a prior density that is intended to represent beliefs about the half-life widely held among economists with a professional interest in the PPP debate. We view the survey data in Table 1 as an approximation to these beliefs.⁶ The general framework proposed in section 3 allows us to approximate prior half-life densities like the survey

⁶ Our prior is not meant to capture a single, isolated belief but rather provides a convenient tool for reporting features of the data based on widely held beliefs (see Sims 1999). This approach has been used extensively in the Bayesian VAR literature. See, for example, Litterman (1986) and Sims and Zha (1998).

density in Table 1 by choosing a suitable prior specification for the slope parameters of the autoregressive model. We will refer to the resulting prior as the consensus prior.⁷

Although this consensus prior provides a natural benchmark for our statistical analysis of the half-life of responses to real exchange rate disturbances, it is merely an illustration and other prior views could be adopted just as easily. For example, one may wish to take a more agnostic view about the value of the half-life. Large subjective uncertainty about the half-life may be represented in the form of a diffuse half-life prior. We show that, while it is impossible to specify a completely flat prior for the half-life, it is possible to construct a quite diffuse prior. Contrary to what one might expect, we will show that diffuse priors on the half-life have a tendency to dominate the posterior at the expense of the likelihood.

4.1. The Consensus Prior for the Half-Life

The relationship between the informative half-life prior and conventional priors on the slope parameters is most transparent in the AR(1) case. In that case, $\rho_2 = \dots = \rho_p = 0$ and there is a one-to-one mapping between ρ_1 and the half-life h (measured in years) for any given frequency f . As noted earlier, the half-life can be expressed as $h = i/f$ where i is the solution to $|\rho^i| = 0.5$ (see Caner and Kilian 1999). Thus, under the (plausible) additional assumption that $\rho_1 > 0$, given a prior distribution for h , the implied prior distribution for ρ_1 will be uniquely determined by the one-to-one mapping: $\rho_1 = 0.5^{(1/hf)}$.

It is immediately clear that a half-life prior such as the survey prior in Table 1 implies a highly informative prior distribution on ρ_1 . This observation suggests that it may not be

⁷ Of course, this prior reflects no explicit agreement among the participants in the survey. In fact, it does not coincide exactly with any one of the individual responses in the survey, nor does it allow for possible differences across countries. Nevertheless, its key features are close to the views of most economists in the survey. This makes it a convenient example to illustrate the methodology proposed in this paper in a realistic setting.

desirable from an economic point of view to postulate a diffuse prior on ρ_1 . For example, a flat prior on the interval $(-1,1)$ for ρ_1 effectively concentrates high probability mass on implausibly short half-lives. This point is demonstrated in Figure 1. Figure 1 displays the half-life prior density implied by the prior $\rho_1 \sim U(-1,1)$. The shape of the implied half-life density is highly asymmetric with a mode close to zero. These features clearly contradict the views expressed in the survey.⁸ Figure 1 illustrates the important point that conventional notions of what constitutes an “informative” or “uninformative” prior do not apply in analyzing the half-life. A prior can be both informative and uninformative, depending on the dimension of interest.

This discussion illustrates the difficulties, even in the simplest models, of constructing a prior on the parameter vector that is consistent with the survey prior on the half-life. In the AR(p) case, these problems are further compounded by the fact that the one-to-one relationship between ρ_1 and h breaks down. There is no compelling reason in general for ρ_1 to be bounded by unity, even for a stationary process, and economics provides no guidance as to the joint prior distribution of the autoregressive slope parameters.

This problem may be solved within the framework discussed in section 3 by postulating a mixed Gaussian prior on ρ_1 and a tight Gaussian prior on ρ_2, \dots, ρ_p . A tight prior on higher-order lags reflects the belief that the impulse response of the real exchange rate is close to the response that would occur if the true process were an AR(1) process. Specifically, we postulate that for each of the i normal distributions $\varphi(m_{b,i}, V_{b,i})$ the higher-order lag coefficients all have prior mean zero and identical variances. For each i , the variances of the higher-order lag coefficients are deterministically related across lags by a decay function $d(t)$, as postulated by

⁸ Similarly, the standard approach of specifying a weakly informative (diffuse) Gaussian prior centered on zero (or one) implies a half-life prior that is inconsistent with views widely held among economists.

Litterman (1986). Thus, the $\varphi(m_{b,i}, V_{b,i})$ distributions differ only by the parameterization of the first element:⁹

$$m_{b,i} = \begin{bmatrix} m_{\rho_{1,i}} \\ 0 \\ \vdots \\ 0 \\ 0 \end{bmatrix}_{p \times 1}, \text{ and } V_{b,i} = \text{diag} \begin{bmatrix} \sigma_{\rho_{1,i}}^2 \\ d(2)\sigma_{\rho_2}^2 \\ \vdots \\ d(p)\sigma_{\rho_2}^2 \\ \sigma_c^2 \end{bmatrix}.$$

We propose to select the parameters of this prior subject to the constraint that the implied prior for the half-life in the AR(p) model does not differ substantially from the half-life prior provided by the economist. In the empirical section, we set the prior means of γ and c to zero and the corresponding standard deviations to 50 and 4, respectively. This prior is diffuse enough for the results to be insensitive to further increases in the standard deviations. We follow the convention in the PPP literature of using $p = 12$ for monthly data. For illustrative purposes, the values of σ_{ρ_2} and of $p_i, m_{\rho_{1,i}}, \sigma_{\rho_{1,i}}, i = 1, \dots, n$, are selected such that the implied half-life prior for the AR(p) model resembles the survey prior on the half-life. We specify four mixed normals with $p_1 = 0.68, p_2 = 0.1, p_3 = 0.17, p_4 = 0.05, m_{\rho_{1,1}} = 0.9904, m_{\rho_{1,2}} = 0.9809, m_{\rho_{1,3}} = 0.97, m_{\rho_{1,4}} = 0.9, \sigma_{\rho_{1,1}} = 0.005, \sigma_{\rho_{1,2}} = 0.001, \sigma_{\rho_{1,3}} = 0.002, \sigma_{\rho_{1,4}} = 0.003$ and $\sigma_{\rho_2} = 0.0015$. Figure 2 shows that this prior specification provides a reasonable approximation to the survey prior over the range of 0-40 years. Moreover, the probability of half-lives in excess of 40 years is 4.8 percent under our specification which is roughly consistent with the entry in Table 1.¹⁰ The cdf of the consensus prior is given in Table 3. While this prior does not match the survey data in all

⁹ For quarterly data it is common to postulate that $d(t) = 1/t$ (see Litterman 1986). For our monthly data, we postulate that $d(t)$ declines exponentially such that the decay rate in the first month matches that in the first quarter and the decay rate in the 12th month matches that in the 4th quarter (see Zha (1998), Robertson and Tallman (1999) for details).

regards, it is reasonably close for our purposes. An even better fit could be obtained by increasing the number of mixed normals.

4.2. The Diffuse Half-Life Prior

Although the survey-based prior seems representative for the views of many economists, some applied researchers may prefer a more diffuse prior. Note that a completely flat prior on the half life is ill-defined in our model. To see this point, consider the AR(1) case. When $p(h) \propto \text{constant}$ for $h > 0$, by the Jacobian transformation,

$$p(\rho) \propto \frac{|\log(a)|}{\rho \log^2 \rho}. \quad (14)$$

It is clear that expression (14) can never be proportional to a density function because it does not integrate to a constant over the parameter space. We can, however, modify the prior density (7) such that the prior density of h , although not completely flat, will be quite diffuse. This diffuse specification may be achieved using the same type of smoothness prior already adopted in section 4.1. Our leading example of a diffuse prior is based on seven mixed normals with $p_1 = 0.59$, $p_2 = 0.263$, $p_3 = 0.007$, $p_4 = 0.07$, $p_5 = 0.04$, $p_6 = 0.02$, $p_7 = 0.01$, $m_{\rho_1,1} = 0.9994$, $m_{\rho_1,2} = 0.998$, $m_{\rho_1,3} = 0.9952$, $m_{\rho_1,4} = 0.9904$, $m_{\rho_1,5} = 0.9809$, $m_{\rho_1,6} = 0.97$, $m_{\rho_1,7} = 0.5$, $\sigma_{\rho_1,1} = 0.002$, $\sigma_{\rho_1,2} = 0.002$, $\sigma_{\rho_1,3} = 0.001$, $\sigma_{\rho_1,4} = 0.002$, $\sigma_{\rho_1,5} = 0.001$, $\sigma_{\rho_1,6} = 0.002$, $\sigma_{\rho_1,7} = 0.15$ and $\sigma_{\rho_2} = 0.0015$. The scale of this prior density in the range 0-40 years depends on the probability mass assigned to half-lives larger than 40 years. The latter probability is about 47% under our specification. This probability seems unreasonably large on a priori grounds. We therefore

¹⁰ The proposed procedure can be carried out on standard PCs. For 20,000 draws, it takes about 2 ½ hours on a 450 Mhz computer to simulate the posterior distribution of the half-life for a given country.

slightly modify the procedure of section 4.1 by adding one more step designed to rescale the diffuse prior such that $\Pr(h > 40)$ is 4.8 %, consistent with the consensus prior in Table 3.

Specifically, we partition the distribution of h into two regions. In region R_0 , the value of h is less than 40 years; in region R_1 , h is equal to or greater than 40 years. Let

$$\Theta_0 = h^{-1}(R_0) \text{ and } \Theta_1 = h^{-1}(R_1).$$

The prior probabilities associated with these two regions under $p(\theta)$ are q and $1 - q$, respectively. Now consider an alternative set of prior probabilities, denoted by q^* and $1 - q^*$. This new prior may reflect the belief that $\Pr(h \in \Theta_1)$ is smaller than implied by the original prior. The new prior $p^*(\theta)$ is related to the old prior by

$$p^*(\theta) = \begin{cases} s_0 p(\theta) & \theta \in \Theta_0 \\ s_1 p(\theta) & \theta \in \Theta_1 \end{cases},$$

where $s_0 = q^*/q$ and $s_1 = (1 - q^*)/(1 - q)$. In our example, $s_0 = 95.2/53.3$ and $s_1 = 4.8/46.7$.

The appropriately re-weighted diffuse prior density is shown in Figure 3. The corresponding cdf is shown in Table 3.

The posterior associated with this re-weighted prior density may be simulated using a weighting method described in Tierney (1994) and Geweke (1997). As shown in Geweke (1997), given the empirical distribution of θ (and functions of θ) under $p(\theta)$, the corresponding posterior under $p^*(\theta)$ can be computed without much additional effort.

Following Tierney (1994) and Geweke (1997) define the probability weight

$$\varpi^*(\theta) = \frac{p^*(\theta|x_0, Y)}{p(\theta|x_0, Y)} = \frac{p^*(\theta)L(\theta|x_0, Y)}{p(\theta)L(\theta|x_0, Y)} = \frac{p^*(\theta)}{p(\theta)}.$$

where $\varpi^*(\theta)$ is bounded. Then the matrix

$$\begin{bmatrix} \theta^{(1)} & w^*(\theta^{(1)}) & h(\theta^{(1)}) \\ \vdots & \vdots & \vdots \\ \theta^{(N)} & w^*(\theta^{(N)}) & h(\theta^{(N)}) \end{bmatrix}, \quad (14)$$

stores the empirical distribution function of θ under the prior $p^*(\theta)$ where

$$w^*(\theta^{(i)}) = \frac{\varpi^*(\theta^{(i)}) w(\theta^{(i)})}{\sum_{i=1}^N \varpi^*(\theta^{(i)}) w(\theta^{(i)})}.$$

and $w(\theta^{(i)})$ is the corresponding weight under the original prior.

This reweighting procedure allows us to rescale the prior in an efficient manner because there are only two regions, and the weights (s_0 or s_1) within each region are the same. It may seem that we could have used a similar technique in section 3 instead of deriving analytically the posterior under the mixed normal prior. Specifically, one could simulate the posterior under the flat prior ($p(\theta) \propto \text{constant}$) and then reweight the draws by w^* for any given proper target density $p^*(\theta)$. In our case, $p^*(\theta)$ is mixed normal. It turns out that this target prior is quite different from the posterior under the flat prior. As a consequence, all but a few elements of the weight vector w^* tend to be zero, and the reweighting procedure becomes very inefficient. This problem is a familiar one in any importance-sampling procedure. We avoid this computational inefficiency by deriving a convenient analytic form of the posterior density function under a mixed normal prior from which we can generate draws directly without importance sampling.

5. Empirical Analysis

The real exchange rate data are constructed from the IMF's *International Financial Statistics* data base on CD-ROM. They are based on the end-of-period nominal U.S. dollar spot exchange rates and the U.S. and foreign consumer price indices. The data set comprises 292

monthly observations for 17 countries (Austria, Belgium, Canada, Denmark, Finland, France, Germany, Greece, Italy, Japan, The Netherlands, Norway, Portugal, Spain, Sweden, Switzerland, and the United Kingdom). The sample period is 1973.1-1997.4.

5.1. Results for the Consensus Prior

The consensus prior specification in Table 3 is designed to be consistent with common beliefs of economists as expressed in the literature and in the survey data. The resulting posterior median as well as interval estimates of the half-life for each of the 17 countries in our sample are shown in Table 4(a).¹¹ The median value of the half-life ranges from 3.7 to 4.7 years, depending on the country. The corresponding 68% (90%) error bands for most countries range from as low as 2.0 (0.7) years to 10.0 (24.5) years in one case. With the notable exception of Canada, our results are remarkably consistent across countries.

A comparison of the prior distribution in Table 3 and the posterior distribution in Tables 4(a) and 5(a) shows that the posterior probabilities are not just driven by the prior. There is evidence of a shift in probability mass from the tails of the distribution towards the center. Moreover, the median half-life rises from 3.7 years in the consensus prior to an average of 4.0 years in the posterior. Even more striking is the drop in the upper bound of the 68% (90%) error bands from 10.2 years to 7.4 years (from 37.4 years to 13.8 years). The corresponding lower bounds rise from 1.7 (0.7) years to 2.1 (1.3) years.

There are three questions of interest that can be addressed based on the information in Table 5(a). First, we are able to examine the likelihood that Rogoff's conjecture is correct that the half-life lies between 3 and 5 years. We find that only about 33 % of the probability mass is concentrated in this region. The average 68 % error band in Table 4(a) suggests a lower bound

of about 2 years and an upper bound close to 7 years, indicating more uncertainty about the half-life than is recognized in the PPP literature.

This finding does not mean, however, that the notion of long-run PPP must be abandoned. Although the mean reversion in real exchange rates is slow, there is strong evidence that mean reversion is taking place. Table 5(a) shows that the probability that the half-life exceeds 10 years is rarely greater than 10 percent, and the probability that the half-life exceeds 40 years is only about 2 %. On average, with about 50% (76%, 91%) probability the half-life is reached within 4 (6, 10) years. Put differently, the half-life is about three times as likely to be shorter than 6 years than to exceed 6 years. Thus, to the extent that there are permanent shocks, they seem to play a limited role at the horizons of economic interest.

Second, we can examine how the evidence of mean reversion differs across countries. The strongest evidence in favor of short half-lives is obtained for the United Kingdom ($\Pr(h \leq 4) = 56\%$), Switzerland ($\Pr(h \leq 4) = 54\%$) and Norway ($\Pr(h \leq 4) = 53\%$). The weakest evidence is obtained for Canada ($\Pr(h \leq 4) = 40\%$), Japan ($\Pr(h \leq 4) = 46\%$) and Portugal ($\Pr(h \leq 4) = 46\%$). This ranking casts doubt on explanations based on relative productivity growth, relative income growth or differences in fiscal policy (see Rogoff 1996). The countries with the shortest (longest) half-lives appear to be sufficiently diverse to rule out such simple hypotheses.

Japan is a particularly interesting case. Japan has long been singled out for the apparent slow mean reversion of its deviations from PPP. Indeed, Table 5(a) shows evidence of below average probabilities especially for half-lives of at most 4 and at most 5 years. However, the probability that the half-life for Japan exceeds 10 years is not much different from the probability

¹¹ We report the median as opposed to the mode, since we do not have information on the shape of the posterior density for half-lives larger than 40 years.

for other countries (9.5% compared with 8.9% for the average country). This evidence suggests that Japan is not as different from other countries, as it commonly believed.

Finally, we provide some tentative conclusions about the type of model needed to explain real exchange rate data. Table 5(a) shows that there is little evidence in favor of half-lives that are commonly considered short enough to be explained by monetary models with nominal rigidities alone. For example, the probability that the half-life does not exceed 2 years is between 10 and 15 percent, and there is only about 5% probability mass on half-lives of one year or less. In contrast, the half-life probabilities support the view that real rigidities play an important role in determining the real exchange rate. Further research will be needed to develop more fully the implications of particular mechanisms (such as wealth redistribution or consumption smoothing) for the duration of deviations from PPP. Our methodology enables researchers to compare directly the implications of theoretical models for the degree of mean reversion in real exchange rates to the data and to find out which models are most appropriate.

5.2. Results for the Diffuse Half-Life Prior

Some economists have no strong prior views about the half-life of deviations from PPP. In that case, a prior more diffuse than the consensus prior is called for. The analysis of diffuse priors is interesting for two reasons. First, it may seem that a diffuse half-life prior should have little effect on the posterior of the half-life. However, that conclusion would be erroneous. The fact that a half-life prior is diffuse does not imply that it will not heavily influence the posterior probability distribution of the half-life. In fact, a diffuse prior actually may have greater influence on the posterior than does the consensus prior. We illustrate this point with an example.

Consider the diffuse prior of section 4.2 (see Table 3). Compared with the consensus prior, this prior spreads out the probability mass more evenly, with slowly declining weights for half-lives in excess of three years and sharply declining weights for half-lives smaller than 2 years (see Figure 3). How does the corresponding posterior distribution for the half-life in Table 5(b) compare with the likelihood in Table 2? Even a cursory glance shows that the posterior distribution in Table 5(b) assigns substantially more probability to very long half-lives than the likelihood in Table 2. For example, the posterior probability of the half-life exceeding 6 years is 82.2% under the diffuse prior. This figure is much closer to the 80% postulated in the diffuse prior distribution in Table 3 than to the 11.6% based on the likelihood in Table 2. Upon reflection, this outcome is not surprising. The diffuse prior in Table 3 puts such large cumulative probability mass on very large half-lives that the prior effectively dominates the posterior distribution at the expense of the likelihood.

The empirical finding that a diffuse prior on h may be highly influential for the posterior is another illustration of the more general point that in a parametric model a prior density that is diffuse in one dimension need not be diffuse in other dimensions. This fact is most transparent in the AR(1) model with slope coefficient $\rho_1 > 0$. In that case, a diffuse prior on h translates into a highly informative prior on ρ_1 with most probability mass located very close to the unit root. In contrast, the consensus prior on h actually would be much less informative in the ρ_1 dimension. Thus, as noted in section 4, conventional notions of what constitutes an “informative” or “uninformative” prior do not apply in analyzing the half-life. A prior is always both informative and uninformative, depending on the dimension of interest.

The second reason for analyzing the diffuse prior in Table 3 is to obtain a better appreciation for the sensitivity of the results to the choice of prior. Consider the differences

between Tables 5(a) and 5(b). The median of the posterior distribution increases sharply from 4.0 under the consensus prior to 12.0 years under the diffuse prior (close to the prior median of 13.1). The corresponding 68 % error bands increase from 2.1-7.4 years under the consensus prior to 5.6-24.4 years under the diffuse prior. Finally, the likelihood that the half-life is contained in the range of 3-5 years shrinks from 33% to 7%.

This example shows that the posterior probabilities can be highly sensitive to the choice of prior. This by itself is not a surprising finding. More important is the observation that this dramatic change in posterior probabilities occurred for a prior that cannot be ruled out as economically implausible. This example shows that, if an economist has prior views that differ from the consensus prior suggested in section 4.1, it is essential that we be able to examine how this difference affects the posterior. Our methodology is designed to facilitate this type of sensitivity analysis and to allow economists to impose the priors with which they feel comfortable.

6. Conclusion

The literature on open economy macroeconomics concludes that different classes of theoretical models have very different implications for the persistence of deviations from PPP. Models with purely nominal rigidities, for example, are unlikely to generate much persistence in deviations from PPP. In contrast, models with persistent differences in productivity growth across countries, intertemporal consumption smoothing or cross-country wealth redistribution effects are consistent with much slower mean reversion in real exchange rates. Economists are faced with the task of deciding which of these models, if any, are consistent with the data. It is important, therefore, to be able to assess the likelihood that the degree of mean reversion in the real exchange rate data is consistent with that implied by a given class of theoretical models.

This paper proposed a general methodology capable of addressing this question. Our approach departs from the existing literature on long-run PPP in that we shift the discussion away from unit root test issues and toward a methodology for assessing the degree of mean reversion in real exchange rates at horizons of economic interest. We followed the convention in international finance of parameterizing the degree of mean reversion in terms of the half-life of deviations from PPP. The framework we proposed is designed to make explicit probability statements about the half-life. These probability statements allow economists to assess the likelihood that the half-life of deviations from PPP exceeds a given number of years without taking a stand on whether the data have a unit root or not.

The proposed methodology allows alternative priors for the half-life to be examined. In our empirical analysis, we illustrated this point by examining two priors. One prior captured widely held views about the half-life expressed in the literature and in a survey of economists with a professional interest in the PPP debate. Based on this benchmark prior, we found compelling evidence of mean reversion over horizons of economic interest. For example, the half-life is about three times as likely to be shorter than 6 years than to exceed 6 years. There is only a 33% probability, however, that the half-life is contained in the range of 3-5 years, as conjectured by Rogoff (1996).

Not all economists are likely to share the consensus view. In fact, some economists will be rather uncertain about the prior distribution of the half-life, in which case a more diffuse prior is called for. The second prior considered in our empirical analysis was of this type. We showed that diffuse half-life priors, contrary to what one might expect, tend to exert a strong influence on the posterior distribution of the half-life. This second example also demonstrated that, if an economist has prior views that differ from the consensus views about the half-life, it is important

that one be able to examine how this difference affects the posterior distribution. Our framework is designed to facilitate this type of sensitivity analysis.

References

- Abuaf, N., and P. Jorion (1990), "Purchasing Power Parity in the Long-Run," *Journal of Finance*, 45, 157-174.
- Caner, M., and L. Kilian (1999), "Size Distortions of Tests of the Null Hypothesis of Stationarity: Evidence and Implications for the PPP Debate," Department of Economics Working Paper No. 99-05, University of Michigan.
- Cheung, Y.-W., and K. S. Lai (1998), "Parity Reversion in Real Exchange Rates During the Post-Bretton Woods Period," *Journal of International Money and Finance*, 17, 597-614.
- Cheung, Y.-W., and K. S. Lai (1999), "On the Purchasing Power Parity Puzzle," manuscript, UC Santa Cruz.
- Culver, S.E., and D.H. Papell (1997), "Long-Run Purchasing Power Parity with Short-Run Data: Evidence with a Null Hypothesis of Stationarity," forthcoming: *Journal of International Money and Finance*.
- Diebold, F.X., S. Husted, and M. Rush (1991), "Real Exchange Rates under the Gold Standard," *Journal of Political Economy*, 99, 1252-1271.
- Edison, H., J. Gagnon, and W. Melick (1997), "Understanding the Empirical Literature on Purchasing Power Parity: The Post-Bretton Woods Era," *Journal of International Money and Finance*, 16, 1-17.
- Frankel, J. (1986), "International Capital Mobility and Crowding Out in the U.S. Economy: Imperfect Integration of Financial Markets or of Goods Markets?" in R. Hafer (ed.): *How Open is the U.S. Economy?*, Lexington Books.
- Frankel, J.A., and K.A. Froot (1987), "Using Survey Data to Test Standard Propositions Regarding Exchange Rate Expectations," *American Economic Review*, 77, 133-153.
- Froot, K.A., and K. Rogoff (1995), "Perspectives on PPP and Long-Run Real Exchange Rates," in: G. Grossman and K. Rogoff (eds.), *Handbook of International Economics*, vol. III, Elsevier Science.
- Froot, K.A., and K. Rogoff (1991), "The EMS, the EMU, and the Transition to a Common Currency," in: Fisher, S., and O. Blanchard (eds.), *NBER Macroeconomics Annual*, 269-317.
- Geweke, J. (1995), "Monte Carlo Simulation and Numerical Integration," in H. Amman, D. Kendrick and J. Rust (eds.), *Handbook of Computational Economics*, Amsterdam: North-Holland.

- Geweke, J. (1997), "Posterior Simulators in Econometrics," in David M. Kreps and Kenneth F. Wallis (eds.): *Advances in Economics and Econometrics: Theory and Applications*, Econometric Society Monographs Seventh World Congress Vol. III, 128-165, Cambridge University Press.
- Koedijk, K.G., P.C. Schotman, and M.A. Van Dijk (1998), "The Re-Emergence of PPP in the 1990s," *Journal of International Money and Finance*, 17, 51-61.
- Litterman, R.B. (1986), "Forecasting with Bayesian Vector Autoregressions - Five Years of Experience," *Journal of Business and Economic Statistics*, 4, 25-38.
- Lothian, J.R., and M.P. Taylor (1996), "Real Exchange Rate Behavior: The Recent Float from the Perspective of the Past Two Centuries," *Journal of Political Economy*, 104, 488-510.
- Lothian, J.R., and M.P. Taylor (1997), "Real Exchange Rate Behavior: The Problem of Power and Sample Size," *Journal of International Money and Finance*, 16, 945-954.
- Murray, C.J., and D.H. Papell (1999), "The Purchasing Power Parity Persistence Paradigm," manuscript, Department of Economics, University of Houston.
- Obstfeld, M., and K. Rogoff (1995), "Exchange Rate Dynamics Redux," *Journal of Political Economy*, 104, 488-509.
- Papell, D. (1997), "Searching for Stationarity: Purchasing Power Parity Under the Current Float," *Journal of International Economics*, 43, 313-332.
- Papell, D., and Theodoridis (1998), "Increasing Evidence of Purchasing Power Parity Over the Current Float," *Journal of International Money and Finance*, 17, 41-50.
- Parsley, D.C., and Wei, S.-J. (1996), "Convergence to the Law of One Price without Trade Barriers or Currency Fluctuations," *Quarterly Journal of Economics*, 111, 1211-1236.
- Robertson, J.C., and E.W. Tallman (1999), "Vector Autoregressions: Forecasting and Reality," *Federal Reserve Bank of Atlanta Economic Review*, First Quarter, 4-18.
- Rogoff, K. (1996), "The Purchasing Power Parity Puzzle," *Journal of Economic Literature*, 34, 647-668.
- Rogoff, K. (1992), "Traded Goods Consumption Smoothing and the Random Walk Behavior of the Real Exchange Rate," *Bank of Japan Monetary and Economic Studies*, 10, 1-29.
- Schotman, P., and H.K. van Dijk (1991), "A Bayesian Analysis of the Unit Root in Real Exchange Rates," *Journal of Econometrics*, 49, 195-238.
- Sims, C.A. (1999), "Using A Likelihood Perspective To Sharpen Economic Discourse: Three Examples," manuscript, Department of Economics, Princeton University.

- Sims, C.A., and H. Uhlig (1991), "Understanding Unit Rooters: A Helicopter Tour," *Econometrica*, 67, 1113-1155.
- Sims, C.A., and T. Zha (1998), "Bayesian Methods for Dynamic Multivariate Models," *International Economic Review*, 39, 949-986.
- Sims, C.A., and T. Zha (1999), "Error Bands for Impulse Responses," *Econometrica*, 67, 1113-1155.
- Taylor, M.P., and D. Peel (1998), "Nonlinear Mean Reversion in Real Exchange Rates: Toward a Solution to the Purchasing Power Parity Puzzles," manuscript, Oxford University.
- Taylor, M.P. and L. Sarno (1998), "The Behavior of Real Exchange Rates During the Post-Bretton Woods Period," *Journal of International Economics*, 46, 281-312.
- Tierney, L. (1994), "Markov chains for exploring posterior distributions," *Annals of Statistics* 22, 1701-62.
- Waggoner, D.F., and T. Zha (1999), "Does Normalization Matter for Inference," mimeo, Federal Reserve Bank of Atlanta.
- Wei, S.-J., and D.C. Parsley (1995), "Purchasing Power Disparity During the Floating Rate Period: Exchange Rate Volatility, Trade Barriers, and Other Culprits," *NBER Working Paper No. 5032*.
- Wu, S., and J.L. Wu (1998), "Purchasing Power Parity Under the Current Float: New Evidence from Panel Data Unit Root Tests," mimeo, Department of Economics, Ohio State University.
- Wu, Y. (1996), "Are Real Exchange Rates Nonstationary? Evidence from a Panel-Data Test," *Journal of Money, Credit and Banking*, 28, 54-63.
- Zha, T. (1998), "A Dynamic Multivariate Model for Use in Formulating Policy," *Federal Reserve Bank of Atlanta Economic Review*, First Quarter, 16-29.

Table 1. – Survey Prior Probabilities for Half-Life under the Recent Float

| Half-Life | PDF | Half-Life | CDF |
|------------------|------|-----------------|-------|
| $0 < h \leq 1$ | 4.6 | $h \leq 1$ | 4.6 |
| $1 < h \leq 2$ | 9.5 | $h \leq 2$ | 14.1 |
| $2 < h \leq 3$ | 17.3 | $h \leq 3$ | 31.4 |
| $3 < h \leq 4$ | 18.2 | $h \leq 4$ | 49.6 |
| $4 < h \leq 5$ | 14.4 | $h \leq 5$ | 64.0 |
| $5 < h \leq 6$ | 11.8 | $h \leq 6$ | 75.8 |
| $6 < h \leq 10$ | 8.1 | $h \leq 10$ | 83.9 |
| $10 < h \leq 20$ | 7.1 | $h \leq 20$ | 91.0 |
| $20 < h \leq 40$ | 3.1 | $h \leq 40$ | 94.1 |
| $h > 40$ | 5.9 | $h \leq \infty$ | 100.0 |

NOTES: Average probabilities based on a survey of 22 economists with a professional interest in the PPP question. The survey was conducted by the authors in July and August of 1999.

Table 2. – Probabilities for Half-Life Implied by Likelihood

| Country | $h \leq 1$ | $h \leq 2$ | $h \leq 3$ | $h \leq 4$ | $h \leq 5$ | $h \leq 6$ | $h \leq 10$ | $h \leq 20$ | $h \leq 40$ | $h > 40$ |
|-------------|------------|------------|------------|------------|------------|------------|-------------|-------------|-------------|----------|
| Austria | 0.0 | 10.4 | 58.2 | 80.2 | 87.7 | 91.5 | 95.9 | 97.7 | 98.3 | 1.7 |
| Belgium | 0.0 | 3.1 | 47.6 | 71.6 | 83.5 | 88.3 | 95.2 | 97.6 | 98.4 | 1.6 |
| Canada | 0.0 | 0.0 | 7.4 | 34.4 | 56.7 | 68.5 | 87.2 | 94.0 | 96.0 | 4.0 |
| Denmark | 0.0 | 8.2 | 56.3 | 79.5 | 88.1 | 91.8 | 96.3 | 98.0 | 98.5 | 1.5 |
| Finland | 0.0 | 21.5 | 75.5 | 89.1 | 94.0 | 96.3 | 98.4 | 99.2 | 99.4 | 0.6 |
| France | 0.1 | 23.0 | 68.4 | 83.5 | 89.4 | 92.5 | 96.3 | 97.9 | 98.4 | 1.6 |
| Germany | 0.0 | 10.4 | 59.1 | 81.2 | 88.3 | 91.8 | 96.0 | 97.8 | 98.3 | 1.7 |
| Greece | 0.2 | 20.1 | 52.4 | 67.9 | 77.1 | 81.2 | 88.5 | 92.4 | 93.9 | 6.1 |
| Italy | 0.0 | 21.5 | 73.0 | 86.7 | 92.1 | 94.2 | 97.3 | 98.4 | 98.8 | 1.2 |
| Japan | 0.0 | 0.6 | 22.4 | 51.0 | 66.4 | 76.0 | 89.1 | 94.5 | 96.3 | 3.7 |
| Netherlands | 0.0 | 18.1 | 67.3 | 84.1 | 90.4 | 93.3 | 96.8 | 98.2 | 98.7 | 1.3 |
| Norway | 0.7 | 49.0 | 80.1 | 88.8 | 92.8 | 94.4 | 96.8 | 97.9 | 98.4 | 1.6 |
| Portugal | 0.0 | 3.4 | 30.7 | 54.3 | 67.7 | 75.0 | 85.9 | 91.1 | 93.1 | 6.9 |
| Spain | 0.0 | 5.8 | 46.6 | 71.6 | 82.3 | 87.5 | 94.0 | 96.5 | 97.6 | 2.4 |
| Sweden | 0.0 | 8.0 | 48.8 | 72.0 | 82.7 | 87.3 | 93.7 | 96.4 | 97.3 | 2.7 |
| Switzerland | 0.0 | 28.8 | 78.5 | 89.1 | 93.6 | 95.8 | 98.1 | 99.0 | 99.3 | 0.7 |
| U.K. | 0.5 | 52.0 | 87.6 | 93.9 | 96.3 | 97.4 | 98.7 | 99.2 | 99.4 | 0.6 |
| Mean | 0.1 | 16.7 | 56.5 | 75.2 | 84.1 | 88.4 | 94.4 | 96.8 | 97.7 | 2.3 |

NOTES: Monthly IFS data.

Table 3. – Prior Probability Distributions Used in the Empirical Analysis

| Prior | $h \leq 1$ | $h \leq 2$ | $h \leq 3$ | $h \leq 4$ | $h \leq 5$ | $h \leq 6$ | $h \leq 10$ | $h \leq 20$ | $h \leq 40$ | $h > 40$ |
|-----------|------------|------------|------------|------------|------------|------------|-------------|-------------|-------------|----------|
| Consensus | 5.2 | 14.5 | 28.9 | 44.3 | 56.6 | 66.1 | 83.1 | 92.2 | 95.2 | 4.8 |
| Diffuse | 1.8 | 3.9 | 9.1 | 13.4 | 16.3 | 20.2 | 37.1 | 71.0 | 95.2 | 4.8 |

NOTES: Based on the specifications discussed in section 4.

Table 4. – Posterior Median for Half-Life and Error Bands

| Country | (a) Under Consensus Half-Life Prior | | | (b) Under Diffuse Half-Life Prior | | |
|-------------|-------------------------------------|---------------------|---------------------|-----------------------------------|---------------------|---------------------|
| | Median | 68 % Error Bands | 90 % Error Bands | Median | 68 % Error Bands | 90 % Error Bands |
| Austria | 3.9 | [2.1, 7.2] | [1.5, 13.2] | 11.9 | [5.5, 24.2] | [2.7, 35.7] |
| Belgium | 4.1 | [2.1, 7.7] | [0.7, 14.7] | 12.1 | [5.7, 24.7] | [2.8, 35.8] |
| Canada | 4.7 | [2.2, 10.0] | [0.7, 24.5] | 13.0 | [5.8, 26.1] | [2.9, 37.2] |
| Denmark | 4.0 | [2.1, 7.4] | [1.6, 14.3] | 12.0 | [5.7, 24.3] | [2.7, 35.7] |
| Finland | 3.8 | [2.0, 7.0] | [1.6, 12.9] | 11.9 | [5.6, 24.2] | [2.6, 35.7] |
| France | 3.9 | [2.1, 7.2] | [1.5, 13.8] | 12.1 | [5.5, 24.2] | [2.7, 35.6] |
| Germany | 4.0 | [2.1, 7.2] | [1.6, 13.8] | 12.0 | [5.6, 24.6] | [2.7, 35.8] |
| Greece | 4.0 | [2.1, 7.4] | [1.5, 14.2] | 12.3 | [5.5, 25.2] | [2.7, 36.5] |
| Italy | 3.9 | [2.0, 7.1] | [0.7, 12.9] | 11.9 | [5.6, 23.9] | [2.7, 35.4] |
| Japan | 4.2 | [2.2, 7.8] | [0.7, 14.2] | 11.7 | [5.8, 23.6] | [2.9, 34.9] |
| Netherlands | 3.9 | [2.1, 7.2] | [1.6, 13.7] | 12.1 | [5.6, 24.6] | [2.7, 35.7] |
| Norway | 3.7 | [2.0, 6.9] | [1.6, 12.6] | 12.0 | [5.5, 24.4] | [2.6, 36.0] |
| Portugal | 4.2 | [2.1, 8.0] | [1.6, 16.5] | 12.4 | [5.6, 25.2] | [2.7, 36.5] |
| Spain | 4.1 | [2.1, 7.6] | [1.6, 14.3] | 11.9 | [5.6, 24.1] | [2.7, 35.4] |
| Sweden | 4.1 | [2.1, 7.6] | [1.5, 14.8] | 12.1 | [5.6, 24.6] | [2.7, 35.8] |
| Switzerland | 3.7 | [2.0, 6.7] | [0.7, 11.7] | 11.7 | [5.6, 23.7] | [2.6, 35.1] |
| U.K. | 3.7 | [2.0, 6.6] | [0.7, 11.7] | 11.7 | [5.4, 23.7] | [2.6, 35.3] |
| Mean | 4.0 | [2.1, 7.4] | [1.3, 13.8] | 12.0 | [5.6, 24.4] | [2.7, 35.8] |

NOTES: Monthly IFS data.

Table 5. - Posterior Probabilities for Half-Life**(a) Under Consensus Half-Life Prior**


| Country | $h \leq 1$ | $h \leq 2$ | $h \leq 3$ | $h \leq 4$ | $h \leq 5$ | $h \leq 6$ | $h \leq 10$ | $h \leq 20$ | $h \leq 40$ | $h > 40$ |
|-------------|------------|------------|------------|------------|------------|------------|-------------|-------------|-------------|----------|
| Austria | 5.0 | 12.9 | 32.3 | 50.4 | 65.7 | 77.0 | 91.8 | 97.2 | 98.6 | 1.4 |
| Belgium | 5.1 | 11.9 | 30.7 | 47.7 | 62.8 | 74.4 | 90.7 | 96.6 | 98.2 | 1.8 |
| Canada | 5.3 | 9.5 | 27.4 | 40.2 | 52.8 | 64.0 | 84.1 | 93.8 | 96.6 | 3.4 |
| Denmark | 4.9 | 12.6 | 31.6 | 49.6 | 65.0 | 76.2 | 91.0 | 96.8 | 98.2 | 1.8 |
| Finland | 4.8 | 13.8 | 33.8 | 52.3 | 67.6 | 78.2 | 92.2 | 97.4 | 98.7 | 1.3 |
| France | 4.9 | 13.3 | 32.7 | 50.9 | 66.6 | 77.3 | 91.6 | 97.0 | 98.3 | 1.7 |
| Germany | 4.9 | 13.1 | 32.2 | 50.0 | 65.3 | 76.4 | 91.5 | 97.0 | 98.5 | 1.5 |
| Greece | 5.0 | 12.7 | 31.9 | 49.6 | 64.6 | 75.9 | 90.9 | 96.8 | 98.3 | 1.7 |
| Italy | 5.1 | 13.5 | 33.0 | 51.5 | 67.0 | 77.7 | 92.3 | 97.2 | 98.5 | 1.5 |
| Japan | 5.4 | 10.4 | 29.6 | 45.7 | 61.4 | 73.2 | 90.5 | 97.1 | 98.6 | 1.4 |
| Netherlands | 4.9 | 13.0 | 32.5 | 50.9 | 66.4 | 77.2 | 91.7 | 96.9 | 98.2 | 1.8 |
| Norway | 4.8 | 14.1 | 34.3 | 53.4 | 68.7 | 79.0 | 92.6 | 97.4 | 98.6 | 1.4 |
| Portugal | 4.9 | 11.5 | 29.7 | 46.1 | 60.9 | 72.6 | 89.3 | 95.9 | 97.7 | 2.2 |
| Spain | 5.0 | 12.0 | 31.4 | 48.2 | 63.7 | 74.9 | 90.8 | 96.8 | 98.3 | 1.7 |
| Sweden | 5.0 | 12.6 | 31.2 | 48.2 | 63.2 | 74.4 | 90.5 | 96.6 | 98.1 | 1.9 |
| Switzerland | 5.1 | 13.9 | 34.5 | 54.3 | 69.9 | 80.3 | 93.4 | 97.8 | 98.9 | 1.1 |
| U.K. | 5.1 | 15.2 | 35.7 | 55.7 | 71.3 | 80.9 | 93.3 | 97.8 | 98.9 | 1.1 |
| Mean | 5.0 | 12.7 | 32.0 | 49.7 | 64.9 | 75.9 | 91.1 | 96.8 | 98.3 | 1.7 |

(b) Under Diffuse Half-Life Prior

| Country | $h \leq 1$ | $h \leq 2$ | $h \leq 3$ | $h \leq 4$ | $h \leq 5$ | $h \leq 6$ | $h \leq 10$ | $h \leq 20$ | $h \leq 40$ | $h > 40$ |
|-------------|------------|------------|------------|------------|------------|------------|-------------|-------------|-------------|----------|
| Austria | 0.0 | 1.7 | 6.8 | 10.7 | 13.8 | 18.3 | 40.1 | 76.9 | 97.5 | 2.4 |
| Belgium | 0.0 | 1.3 | 6.3 | 9.9 | 13.2 | 17.4 | 39.1 | 76.2 | 97.5 | 2.5 |
| Canada | 0.0 | 1.2 | 5.6 | 10.3 | 13.2 | 16.8 | 35.2 | 72.8 | 96.6 | 3.4 |
| Denmark | 0.0 | 1.6 | 6.6 | 10.2 | 13.2 | 17.4 | 39.5 | 76.3 | 97.5 | 2.5 |
| Finland | 0.0 | 2.0 | 6.6 | 10.2 | 13.8 | 17.6 | 39.5 | 76.5 | 97.5 | 2.5 |
| France | 0.0 | 2.1 | 7.2 | 11.0 | 14.0 | 18.3 | 39.9 | 76.3 | 97.5 | 2.5 |
| Germany | 0.0 | 1.8 | 7.1 | 10.5 | 13.8 | 17.9 | 39.4 | 76.5 | 97.5 | 2.5 |
| Greece | 0.0 | 1.9 | 6.5 | 11.1 | 14.5 | 18.1 | 38.4 | 74.7 | 97.0 | 3.0 |
| Italy | 0.0 | 1.8 | 6.4 | 10.6 | 13.9 | 17.7 | 39.7 | 76.9 | 97.6 | 2.4 |
| Japan | 0.0 | 1.2 | 5.6 | 9.6 | 12.8 | 17.0 | 40.5 | 77.8 | 98.0 | 2.0 |
| Netherlands | 0.0 | 1.9 | 6.6 | 10.7 | 13.7 | 17.4 | 38.8 | 75.9 | 97.4 | 2.6 |
| Norway | 0.0 | 2.4 | 6.9 | 10.7 | 14.3 | 18.2 | 39.5 | 76.1 | 97.3 | 2.7 |
| Portugal | 0.0 | 1.4 | 6.8 | 10.9 | 13.7 | 17.4 | 37.5 | 74.8 | 97.1 | 2.9 |
| Spain | 0.0 | 1.4 | 6.5 | 10.2 | 13.3 | 17.9 | 39.3 | 76.7 | 97.6 | 2.4 |
| Sweden | 0.0 | 1.5 | 6.6 | 10.2 | 13.6 | 18.0 | 39.5 | 76.1 | 97.4 | 2.6 |
| Switzerland | 0.0 | 2.1 | 6.6 | 10.1 | 13.8 | 18.1 | 41.0 | 78.2 | 97.7 | 2.3 |
| U.K. | 0.0 | 2.6 | 7.0 | 11.0 | 14.2 | 18.4 | 40.8 | 77.4 | 97.6 | 2.4 |
| Mean | 0.0 | 1.8 | 6.6 | 10.5 | 13.7 | 17.8 | 39.3 | 76.2 | 97.4 | 2.6 |


NOTES: Monthly IFS data.

Figure 1. - Prior P.D.F. for Half-Life Implied by $\rho_1 \sim U(-1,1)$




NOTES: Example based on AR(1) model.

Figure 2. – Consensus Prior for the Half-Life



NOTES: Parametric approximation based on AR(p) model prior described in section 4.

Figure 3. – Diffuse Half-Life Prior



NOTES: Parametric approximation based on AR(p) model prior described in section 4.2. The diffuse prior shown has been re-weighted as discussed in section 4.2.