

QUANTIFYING THE UNCERTAINTY ABOUT THE HALF-LIFE OF DEVIATIONS FROM PPP

LUTZ KILIAN^{a*} AND TAO ZHA^b

^a *Department of Economics, University of Michigan, Ann Arbor, MI 48109-1220, USA and Centre for Economic Policy Research, London, UK*

^b *Research Department, Federal Reserve Bank of Atlanta, Atlanta, GA 30303, USA*

SUMMARY

We propose a Bayesian framework in which the uncertainty about the half-life of deviations from purchasing power parity can be quantified. Based on the responses to a survey study, we propose a prior probability distribution for the half-life under the recent float intended to capture widely held views among economists. We derive the posterior probability distribution of the half-life under this consensus prior and confirm the presence of substantial uncertainty about the half-life. We provide for the first time a comprehensive formal evaluation of several nonnested hypotheses of economic interest, including Rogoff's (1996) claim that the half-life is contained in the range of 3 to 5 years. We find that no hypothesis receives strong support from the data. Copyright © 2002 John Wiley & Sons, Ltd.

1. INTRODUCTION

In recent years, the half-life of deviations from purchasing power parity (PPP) has become a commonly used measure of the degree of mean reversion in real exchange rates.¹ Understanding just how uncertain half-life estimates are is important for economists. For example, Rogoff (1996, p. 664, emphasis added) stresses that:

...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*.

Rogoff's claim about the value of half-life is based on an informal reading of the empirical literature on the half-life under the recent float. In this paper, we present a Bayesian framework in which the uncertainty about the half-life can be quantified and alternative hypotheses about the half-life may be formally evaluated.

The literature on open economy macroeconomics concludes that different classes of theoretical models have different implications for the persistence of deviations from PPP. Models with purely nominal rigidities, for example, are unlikely to generate much persistence in deviations

* Correspondence to: Lutz Kilian, Department of Economics, University of Michigan, Ann Arbor, MI 48109-1220, USA. E-mail: lkilian@umich.edu

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 (2000), Cheung and Lai (2000a,b), Murray and Papell (2000a,b), Ng (2000).

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. It is not surprising therefore that many economists are uncertain about what degree of mean reversion to expect in the real exchange rate. We set out to capture this prior uncertainty in the form of a prior distribution for the half-life. The starting point of our analysis is a survey study conducted by the authors in the summer of 1999, in which economists with an interest in the PPP debate were questioned about their subjective prior probability distribution for the half-life under the recent float. Based on the survey responses, we propose a prior probability distribution for the half-life intended to capture widely held views among economists. This *consensus prior* is used as a benchmark in our analysis and formally compared to a number of alternative views about the half-life. The use of economically informed priors in studying the uncertainty about the half-life has not been attempted before.

Our Bayesian framework allows us to incorporate formally such prior information. We first derive the posterior probability distribution of the half-life under the consensus prior. Our results confirm the presence of substantial uncertainty about the half-life of deviations from PPP, similar to the conclusion of a recent study by Murray and Papell (2000) based on classical confidence intervals for the half-life. On average, for the 17 countries analysed, the 68% (90%) posterior error bands cover values between about 2.4 (1.9) and 7.5 (15.1) years.

Our second finding is that the posterior is heavily influenced by the consensus prior. This fact makes it important to study the extent of the empirical support for the consensus prior relative to other plausible views about the half-life. Through the lens of our econometric model, we provide for the first time a comprehensive formal evaluation of several nonnested hypotheses of economic interest, including Rogoff's (1996) claim that the half-life is contained in the range of three to five years. These hypotheses are parameterized as priors about the half-life. Such nonnested comparisons have not been carried out in the existing PPP literature. We show that the data do not strongly favour Rogoff's claim nor any other of a number of hypotheses of economic interest including the consensus prior. These empirical results are new and would not have been readily apparent without formal analysis. They suggest that PPP studies of one country and one variable at a time are unlikely to shed light on the questions of economic interest, and they motivate the investigation of structural models of the real exchange rate that involve multiple equations, a task to which Bayesian methods seem eminently suited.

The remainder of the paper is organized as follows. In Section 2, we motivate the use of the half-life concept in assessing the evidence of long-run PPP. In Section 3, we describe the survey data underlying the consensus prior. In Section 4, we discuss the general econometric framework of our analysis. Section 5 contains the empirical results. Conclusions are presented 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).

This literature, however, ignores the fact that 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. 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. 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. It may seem that the speed of mean reversion is adequately captured by the magnitude of the estimated root of the process (or alternatively by the cumulative impulse response function). However, in higher-order processes these measures may have little relation to the degree of mean reversion over horizons of economic interest.

In this paper, we therefore pursue an alternative approach that has recently gained popularity in the empirical literature (e.g. Cheung and Lai, 2000a,b; Murray and Papell, 2000a,b). Instead of relying on unit root tests, we parameterize the degree of mean reversion in terms of the half-life of deviations from PPP. Although the half-life is neither a unique measure of mean reversion nor necessarily an optimal one, the half-life concept is familiar to many economists and already plays an important role in the PPP debate (see footnote 1 for examples). The half-life of the real exchange rate process is commonly defined as the number of years required for the impulse response to a unit shock to dissipate by one half (see e.g. Diebold *et al.*, 1991, p. 1267; Rogoff, 1996, p. 656; Murray and Papell, 2000a, pp. 3,7; Cheung and Lai, 2000b, p. 324).

Let y_t denote the log real exchange rate. Define $\phi_i = \partial y_{t+i} / \partial u_t$ for $i = 0, 1, 2, \dots$, where $u_t = y_t - E_{t-1}(y_t)$ and $E_{t-1}(y_t)$ denotes the linear projection of y_t on its past. Then ϕ_i is the impulse response of y_{t+i} to a one-unit change in u_t at date t . In general it is impossible to compute impulse responses for all horizons. In practice, we lump together all half-lives greater than 40 years, because from an economic point of view there is little difference between these half-life values. Specifically, we proceed as follows. First, find the largest i in the range $i = 1, \dots, 40f$ for which $|\phi_{i-1}| \geq 0.5$ and $|\phi_i| < 0.5$, where f denotes the sampling frequency of the data ($f = 1$ for years; $f = 4$ for quarters; $f = 12$ for months, etc.). If there is no i that satisfies this condition, we say that the half-life is not reached within forty years (denoted by $h > 40$ years). If the condition is satisfied, denote the largest i by k and verify that $|\phi_j| < 0.5$ for all j in the range $k + 1, \dots, k + 40f$. This second condition provides a practical way of ruling out unstable or explosive patterns for the horizon of economic interest. If the second condition is met, we say that the half-life is reached in $h = k/f$ years. If not, we say that the half-life is not reached within forty years (denoted by $h > 40$ years).

3. THE CONSENSUS PRIOR

Our 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. An important practical question is that of how to elicit such priors. In July and August of 1999 we conducted a survey of economists working in the fields of international economics and macroeconomics. The survey participants were asked by email to specify probability weights for given ranges of

half-lives that correspond to their subjective prior beliefs about the half-life under the recent float. The half-life was defined in terms of the response to an innovation in the univariate time series representation of the real exchange rate. Respondents were assured that their individual responses would be treated confidentially and that only average survey responses would be made public.²

Table I shows the prior density based on the average responses of 22 economists with a professional interest in the PPP question. The average survey response displays a single peak at about 4 years and a long right tail with little probability mass allocated to half-lives in excess of 20 years. Although the survey data in Table I do not reflect an explicit consensus, we view these data as representative for widely held beliefs about the half-life in the economics profession.³

As with all such empirical studies, an important concern is to what extent the survey prior in Table I simply reflects the results of the existing empirical literature on long-run PPP as opposed to introspection. Although our survey explicitly requested subjective prior probabilities about the half-life, we cannot rule out that some respondents were influenced by the post-Bretton Woods data. There is evidence, however, that this concern is not as serious as it may seem.

Most empirical studies of the half-life arrive at estimates in the range of three to five years. This is what prompted Rogoff's quote in the Introduction. Suppose for the sake of argument that a survey respondent had relied mainly on Rogoff's well-publicized view that the half-life is expected to lie between three and five years. It seems reasonable to interpret Rogoff as giving an implicit one-standard error (or possibly two-standard error) band for the half-life. If that were the case we would have expected our survey respondents (some of whom are leading contributors to the empirical PPP literature) to place at least 68% (and possibly as much as 95%) of the probability mass into the three-to-five year range. This is far from the actual numbers. In fact, only one survey respondent allocated more than 45% probability to this region, and the mean response was 33%. This comparison suggests that the prior in Table I, although perhaps not purely subjective, was not unduly influenced by the data.

Similarly, suppose that the survey respondents had access to the classical confidence intervals reported by Murray and Papell (2000a). The 95% confidence intervals in Murray and Papell tend to range from as low as one quarter year to infinity. These intervals are so wide that they convey virtually no information about the economically interesting regions of the half-life. There is no

Table I. Survey prior probabilities for half-life under the recent float

	$h \leq 1$	$h \leq 2$	$h \leq 3$	$h \leq 4$	$h \leq 5$	$h \leq 6$	$h \leq 10$	$h \leq 20$	$h \leq 40$	$h > 40$
Per cent	4.6	14.1	31.4	49.6	64.0	75.8	83.9	91.0	94.1	5.9

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 1999.

² We made an effort to approach a broad pool of economists working in macroeconomics and international economics. We emailed our questionnaire to close to forty economists. More than half agreed to provide their prior probability distributions in a table format similar to Table I. The full list of participants is given in the Acknowledgements to this paper. There were no noticeable biases in the response pattern to the survey, except that the response rate was naturally higher at Michigan where one of the authors is based. We also studied individual responses for potential biases, but could not detect any systematic patterns. Finally, we deleted one response at a time and recalculated the average prior to verify that the average result is robust to outliers.

³ 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, 2000). This approach has been used extensively in the Bayesian VAR literature. See, for example, Litterman (1986) and Sims and Zha (1998).

Table II. 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
UK	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

Note: Monthly IFS data.

obvious way of mapping this information into a probability distribution for the half-life of the kind shown in Table I.

Finally, it is straightforward to show that the survey beliefs about the half-life differ from the probability distribution that would be obtained by inspecting the likelihood of monthly post-Bretton Woods data. We fit an AR(12) model under a flat (improper) prior on the parameters of the autoregressive model. In that case, the posterior distribution will be proportionate to the likelihood. Table II shows the implied probability distribution for the half-life. We find that for the UK (Canada), for example, the likelihood implies a posterior 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 posterior probability that the half-life does not exceed 4 years is 75.2% in Table II compared with 49.6% in the survey. Again there is no evidence that the prior in Table I was obtained from inspecting the post-Bretton Woods data.

4. A BAYESIAN FRAMEWORK FOR EVALUATING THE PROBABILITY DISTRIBUTION OF THE HALF-LIFE

Define the real exchange rate $y_t = s_t + p_t^* - p_t$, where s_t denotes the end-of-period nominal US dollar spot exchange rate and p_t^* and p_t are the foreign and US consumer price indices. All variables are in logs.⁴ We first introduce some notation. Assume that the real exchange rate may

⁴ The real exchange rate data are constructed from the IMF's *International Financial Statistics*. 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.

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

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

where b is the vector of all regression coefficients and λ denotes the precision of the innovations. Then the likelihood function conditional on the initial observations is

$$L(Y|b, \lambda) \propto \lambda^{T/2} \exp \left\{ -\frac{\lambda}{2} (b'(X'X)b - 2b'(X'Y) + Y'Y) \right\}.^5 \quad (3)$$

The priors on λ and b take the form

$$p(\lambda) = \sqrt{\frac{\beta}{\pi}} \lambda^{-\frac{1}{2}} e^{-\beta\lambda} \quad \text{and} \quad p(b) = \sum_{i=1}^n p_i \varphi(m_{b,i}, V_{b,i}) \quad (4)$$

where $p_1 + \dots + p_n = 1$ and $\varphi(x, y)$ is the (properly scaled) normal density function with mean x and covariance y . $p(\lambda)$ is a conventional Gamma prior density for the precision λ . The smaller β , the more diffuse is $p(\lambda)$. $p(b)$ is a mixed normal distribution. Note that unlike in a normal mixture model, in which the probability weights are treated as parameters to be estimated, the weights $p_i, i = 1, \dots, n$, in equation (4) are chosen in advance to approximate a given half-life prior. By appropriately choosing the number of mixed normals and their parameters we are able to approximate a variety of distributional shapes for the half-life. This flexibility in turn allows us to incorporate the prior views of economists into the empirical analysis.

In related work on the AR(1) model, Schotman and van Dijk (1991, p. 208) stress the importance of specifying a prior that forces c to approach zero as $\rho_1 \rightarrow 1$. In our AR(p) model, we address this issue by incorporating the dummy observation prior of Sims and Zha (1998). Specifically, we add one dummy observation to the data set such that

$$\tau \bar{y}_0 = c\tau + \rho_1 \tau \bar{y}_0 + \dots + \rho_p \tau \bar{y}_0 + u \quad (5)$$

where \bar{y}_0 is the mean of the initial values and τ is a hyperparameter. As $\tau \rightarrow \infty$, this dummy observation implies that a unit root exists if and only if $c = 0$. We adopt a value of τ that is standard in the Bayesian VAR literature (see Sims and Zha, 1998 for further discussion).

⁵ We follow the convention of taking the initial observation x_0 as given in the autoregressive analysis.

The derivation of the posterior distribution follows from an application of the marginal likelihood identity of Chib (1995). Combining equation (3) and (4), we obtain:

$$p(\lambda, b|Y) = \sum_{i=1}^n w_i p_i(\lambda|Y) p_i(b|Y, \lambda) \tag{6}$$

where

$$w_i = \frac{p_i L_i(Y)}{\sum_{i=1}^n p_i L_i(Y)}$$

$$p_i(\lambda|Y) = \lambda^{(T-1)/2} |\tilde{V}_{b,i}|^{1/2} C_i \exp \left\{ -\frac{\lambda}{2} (2\beta + Y'Y) + \frac{1}{2} (V_{b,i}^{-1} m_{b,i} + \lambda(X'Y))' \tilde{m}_{b,i} \right\} \tag{7}$$

$$p_i(b|Y, \lambda) = \varphi(\tilde{m}_{b,i}, \tilde{V}_{b,i}) \tag{8}$$

$$\tilde{V}_{b,i} = (V_{b,i}^{-1} + \lambda(X'X))^{-1} \text{ and } \tilde{m}_{b,i} = \tilde{V}_{b,i} (V_{b,i}^{-1} m_{b,i} + \lambda X'Y)$$

The term $L_i(Y)$ is the marginal likelihood under the i th component of the prior

$$p_i(\lambda, b) = p(\lambda) \varphi(m_{b,i}, V_{b,i}) \tag{9}$$

The constant term C_i in equation (7) is chosen such that the right-hand side of equation (7) integrates to 1 over $(0, \infty)$, as required for a properly scaled density function. This term can be easily computed by standard numeral integration methods because $p_i(\lambda|Y)$ is one-dimensional and its shape is quite smooth. Given the value of C_i , it is straightforward to calculate $L_i(Y)$ for any value of λ and b without further simulations:

$$L_i(Y) = \frac{L(Y|\lambda, b) p_i(\lambda, b)}{p_i(\lambda|Y) p_i(b|Y, \lambda)} \tag{10}$$

The relationship between a given half-life prior and the prior on the slope parameters is most transparent in the AR(1) case. If $\rho_2 = \dots = \rho_p = 0$, according to our earlier definition, the half-life can be expressed as $h = i/f$ where i is the solution to $|\rho_1^{i-1}| \geq 0.5$ and $|\rho_1^i| < 0.5$ and for $0 < \rho_1 < 1$ there is a one-to-one mapping between the half-life h (measured in years) and ρ_1 for any given frequency f (see Froot and Rogoff, 1995, p. 1656; Abuaf and Jorion, 1990, p. 159; Caner and Kilian, 2000, p. 13; Murray and Papell, 2000b pp. 2–3). In the higher-order autoregressive model, however, the one-to-one relationship between ρ_1 and h breaks down. Moreover, 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.⁶

We address this problem by postulating a mixed Gaussian prior on ρ_1 and a relatively 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 on average to the response that would occur if the true

⁶Note that in general there is no direct relationship between the root of an autoregressive process and the half-life. In fact, an autoregressive process of higher order may have an arbitrarily short half-life despite the presence of a unit root component.

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 Litterman (1986). Thus, the $\varphi(m_{b,i}, V_{b,i})$ distributions differ only by the parameterization of the first element:⁷

$$m_{b,i} = \begin{bmatrix} m_{\rho_1,i} \\ 0 \\ \vdots \\ 0 \\ 0 \end{bmatrix}_{p \times 1} \quad \text{and} \quad V_{b,i} = \text{diag} \begin{bmatrix} \sigma_{\rho_1,i}^2 \\ d(2)\sigma_{\rho_2}^2 \\ \vdots \\ d(p)\sigma_{\rho_2}^2 \\ \sigma_c^2 \end{bmatrix}$$

We propose to select the parameters of this prior by trial and error 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 mean for the intercept to zero and the corresponding standard deviation to 4. We follow the convention in the PPP literature of using $p = 12$ for monthly data. 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. Specifically, we employ 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 $\sigma_{\rho_2} = 0.0015$. We set $\beta = 0.0004$, resulting in a very diffuse prior for the precision λ . Figure 1 shows that such a 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% under our specification which is roughly consistent with the entry in Table I. In the remainder of the paper, we will refer to this parametric approximation as the *consensus prior*.⁸

The form of the posterior distribution in expressions (6)–(8) facilitates the simulation of the posterior of the half-life. For each component i of the prior, $i = 1, \dots, n$, we first generate draws from the marginal posterior distribution (7) of λ and then generate draws from the posterior distribution (8) of b conditional on λ . According to (6), the proportion of draws from the posterior under the i th component of the prior is w_i .⁹ Finally, we construct the impulse response function associated with each draw for $b = [\rho_1 \dots \rho_p c]'$. The impulse response coefficients may be obtained recursively from $\phi_i = \sum_{j=1}^i \phi_{i-j} \rho_j, i = 1, 2, \dots$, starting with $\phi_0 = 1$ and setting $\rho_j = 0$ for $j > p$. The half-life may be read off the impulse response function as discussed in Section 2.

The only complication is that the marginal posterior distribution (7) 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 following proposal density for the posterior

⁷ 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 months matches that in the first quarter and the decay rate in the twelfth month matches that in the fourth quarter (see Zha, 1998 and Robertson and Tallman, 1999 for details).

⁸ 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 and it provides a useful benchmark.

⁹ This is equivalent to simulating the same number of draws from each posterior and then weighting them by w_i .




Figure 1. Consensus prior for the half-life: dotted line: survey prior, solid line: consensus prior. *Note:* The consensus prior is a parametric approximation to the survey prior based on the AR(12) model prior described in section 4

under the *i*th component of the prior:

$$J_i(\lambda|Y) = \frac{\hat{\Omega}_i^{T/2}}{2^{T/2}\Gamma(T/2)} \lambda^{T/2-1} \exp\left\{\frac{-\hat{\Omega}_i}{2} \lambda\right\} \tag{11}$$

where

$$\begin{aligned} \hat{\Omega}_i &= (Y - X\hat{B}_i)'(Y - X\hat{B}_i) \\ \hat{B}_i &= \tilde{V}_{b,i}(V_{b,i}^{-1}m_{b,i} + \hat{\lambda}_i X'Y) \end{aligned}$$

The term $\Gamma(T/2)$ in (11) is the standard gamma function evaluated at $T/2$. To obtain draws from (11) 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 (11). Our Metropolis–Hastings procedure involves the following four steps for each $i = 1, \dots, n$:

Metropolis–Hastings Algorithm Initialize an arbitrary value λ^0 in R^+ (this value is usually set equal to the maximum point $\hat{\lambda}$). For $m = 1, \dots, M_1 + M_2$,

- (a) Generate λ^* from $J_i(\lambda|Y)$ and u from the uniform distribution $U(0,1)$;
 (b) Compute

$$q = \min \left\{ \frac{p_i(\lambda^*|Y)/J_i(\lambda^{(m-1)}|Y)}{p_i(\lambda^{(m-1)}|Y)/J_i(\lambda^*|Y)}, 1 \right\}$$

- (c) If $u \leq q$, set $\lambda^{(m)} = \lambda^*$; if not, set $\lambda^{(m)} = \lambda^{(m-1)}$;
 (d) Collect the simulated sequence $\{\lambda^{(1)}, \dots, \lambda^{(M_1+M_2)}\}$ but retain only the last M_2 values of the sequence.

The convergence to the target distribution $p(\lambda|Y)$ of the distribution of these draws follows from arguments in Geweke (1995) and Chib and Greenberg (1995).¹⁰

5. EMPIRICAL RESULTS

This section will first address the question of how much posterior uncertainty there is about the half-life, given common prior beliefs among economists as summarized by the consensus prior. We will then evaluate the empirical support for the consensus prior relative to plausible alternative hypotheses about the half-life.

5.1. How Much Uncertainty Is There About the Half-life?

Table III shows the posterior median of the half-life as well as regions of high posterior probability for each of the 17 countries in our sample.¹¹ The median value of the half-life ranges from 3.1 to 5.7 years, depending on the country. With the notable exception of Canada, our results are remarkably consistent across countries. These median results may at first sight look supportive of Rogoff's claim that the half-life is between three and five years, but—even excluding Canada—the corresponding 68% (90%) error bands range from as low as 1.9 (1.7) years to as high as 8.7 (16.5) years. These intervals imply substantial uncertainty about the half-life under the recent float. Our results confirm the qualitative conclusions of a recent study by Murray and Papell (2000a) based on classical confidence intervals.

5.2. How Much Empirical Support Is There for the Consensus Prior?

It is useful to compare the posterior distribution to the consensus prior distribution displayed in Figure 2. Although the median half-life is virtually unchanged, there is some evidence of a shift in probability mass from the tails of the distribution towards the center. The upper bound of the 68% (90%) error bands drops from 10.2 (37.4) years in the prior to an average of 7.5 (15.1) years in the posterior. The corresponding lower bounds rise from 1.7 (0.7) years in the prior to an average of 2.4 (1.9) years in the posterior. Nevertheless, the overall shape of the posterior is

¹⁰ 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 1.4 draws of the parameter λ in the Metropolis–Hastings procedure correspond to at least 1 effective draw.

¹¹ We simulate 20,000 draws from the posterior distribution of the half-life using the method proposed in Section 4. The procedure takes about 2 1/2 hours on a 450 Mhz computer for a given country. 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.

Table III. Posterior median for half-life and error bands under the consensus prior

Country	Median	68% error bands	90% error bands
Austria	4.0	[2.4, 7.3]	[1.8, 13.7]
Belgium	4.4	[2.8, 8.1]	[2.0, 16.5]
Canada	5.7	[3.6, 11.8]	[2.9, 35.1]
Denmark	4.1	[2.5, 7.6]	[1.9, 15.2]
Finland	3.6	[2.1, 6.7]	[1.8, 11.9]
France	3.7	[2.1, 6.9]	[1.8, 12.7]
Germany	3.9	[2.3, 7.2]	[1.8, 14.2]
Greece	3.8	[2.1, 7.3]	[1.8, 14.1]
Italy	3.8	[2.1, 6.9]	[1.8, 13.2]
Japan	4.8	[3.2, 8.7]	[2.6, 16.1]
Netherlands	3.8	[2.1, 7.0]	[1.8, 13.7]
Norway	3.2	[1.9, 6.2]	[1.7, 10.7]
Portugal	4.6	[2.9, 8.5]	[2.0, 17.1]
Spain	4.3	[2.7, 7.9]	[2.0, 14.8]
Sweden	4.3	[2.6, 8.1]	[1.9, 16.5]
Switzerland	3.4	[2.0, 6.1]	[1.7, 11.2]
UK	3.1	[1.9, 5.7]	[1.7, 10.1]
Mean	4.0	[2.4, 7.5]	[1.9, 15.1]

Note: Monthly IFS data.

heavily influenced by the consensus prior. The similarity of the posterior and the consensus prior is cause for concern because it suggests that the data have little influence on the posterior. It is therefore important to study formally the extent of the empirical support for the consensus prior relative to other plausible hypotheses of economic interest. Our methodology is well suited to this purpose.

Specifically, we construct Bayes factors to evaluate pairs of nonnested hypotheses of economic interest. Suppose we have two models A and B , which differ only in terms of their priors. Furthermore, suppose that we have already obtained the posterior output of model A under the consensus prior and now consider an alternative model B under a prior that is a truncated version of the consensus prior, assigning probability 1 to a given range of the parameter region (say half-lives between 3 and 5 years). Then the Bayes factor in favour of model B over model A is

$$\frac{p(Y|B)}{p(Y|A)} = E \left[\frac{p(\lambda, b|B)}{p(\lambda, bA)} \middle| Y, A \right] \quad (12)$$

(see Geweke, 1999).¹² Table IV shows the Bayes factors for each pair of hypotheses and country. Bayes factors in excess of 20:1 for a given hypothesis are typically considered strong evidence in favor of this hypothesis against the alternative. We begin by comparing the empirical support for the consensus prior to that for Rogoff's claim that $3 < h \leq 5$. The first column in Table IV shows that the data are inconclusive. All Bayes factors are between 1.16 and 1.56.

Rogoff's claim of course does not exhaust the set of economically interesting alternative hypotheses. For example, we may be interested in formulating hypotheses about the half-life that

¹² The normalizing constant that makes $p(\lambda, b|B)$ a proper pdf is calculated directly from the simulated draws that produce the empirical pdf of the consensus prior in Figure 2. Because $p(\lambda, b|B)$ is a truncated version of $p(\lambda, b|A)$, the reweighting function $p(\lambda, b|B)/p(\lambda, b|A)$ is well behaved for the reasons discussed in the Technical Appendix.




Figure 2. Half-life prior densities: dotted line: consensus prior, solid line: diffuse half-life prior. *Note:* The diffuse prior shown has been reweighted as discussed in the Technical Appendix

are of economic interest because they capture implications of certain classes of models. We know that in the absence of any rigidities mean reversion is likely to take place very quickly. Thus, some source of friction is needed to explain the apparent slow mean reversion in real exchange rates. For example, 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 other hand, monetary models of the exchange rate with only nominal rigidities imply much shorter half-lives.¹³

In practice, we proceed by categorizing open economy models into four groups, starting with models implying relatively short half-lives. For example, standard monetary models of the exchange rate rely on sticky nominal prices and/or nominal wages to explain the persistence of real exchange rates. It has been argued that models with such purely nominal rigidities are consistent with half-lives of up to two years (see Cheung and Lai, 2000a, Murray and Papell 2000a). Thus, we refer to the hypothesis $0 < h \leq 2$ as 'no real rigidities'. To explain larger half-lives, some degree of real rigidity in the economy is required. We divide the remaining parameter region somewhat arbitrarily into the hypotheses $2 < h \leq 6$ ('some real rigidities'), $6 < h \leq 40$ ('strong

¹³ For a review of this literature see Rogoff (1996).

Table IV. Bayes factors for selected hypotheses

Country	Rogoff	Real rigidities				Diffuse
		None	Some	Strong	Severe	
		3–5 years	0–2 years	2–6 years	6–40 years	
Relative to consensus half-life prior						
Austria	1.43	0.69	1.33	0.79	0.32	0.95
Belgium	1.48	0.39	1.30	0.93	0.48	1.15
Canada	1.21	0.06	1.03	1.47	0.96	1.90
Denmark	1.45	0.59	1.32	0.83	0.38	1.01
Finland	1.32	1.10	1.32	0.65	0.27	0.84
France	1.37	0.88	1.33	0.70	0.30	0.91
Germany	1.41	0.71	1.33	0.77	0.37	0.98
Greece	1.35	0.84	1.30	0.78	0.34	0.91
Italy	1.41	0.84	1.34	0.70	0.36	0.94
Japan	1.56	0.11	1.28	1.09	0.37	1.30
Netherlands	1.37	0.88	1.33	0.72	0.32	0.92
Norway	1.21	1.46	1.29	0.55	0.23	0.74
Portugal	1.44	0.36	1.27	1.00	0.44	1.20
Spain	1.48	0.38	1.33	0.91	0.40	1.15
Sweden	1.46	0.48	1.30	0.91	0.42	1.17
Switzerland	1.33	1.24	1.34	0.55	0.23	0.74
UK	1.16	1.73	1.27	0.48	0.15	0.66
Mean	1.38	0.75	1.29	0.81	0.37	1.03

Note: Monthly IFS data.

real rigidities') and $40 < h \leq \infty$ ('several real rigidities'. Together these hypotheses are likely to capture the main hypotheses of economic interest.

Table IV shows that regardless of which of these four hypotheses we pit against the consensus prior, the evidence is inconclusive. There is systematic evidence against the hypothesis $40 < h \leq \infty$ for all countries, but even against such an extreme hypothesis, the Bayes factor never reaches a factor of more than about 7. There is weak evidence in favour of some real rigidities for all countries, but the Bayes factors never exceed 1.34. For the other hypotheses the Bayes factors range from 0.06 to 1.73 (in favour of no real rigidities) and from 0.48 to 1.47 (in favour of strong real rigidities). Despite the apparent high persistence of deviations from PPP, with the exception of Canada, we cannot formally rule out any of the hypotheses, not even ranges commonly associated with standard monetary models of the exchange rate.

Although the consensus prior provides a natural benchmark for our statistical analysis, some economists may prefer a half-life prior that is more diffuse in the sense of spreading out the probability mass for the half-life more evenly than the consensus prior. We propose one such 'diffuse' half-life prior in Figure 2.¹⁴ Because this prior assigns substantially more cumulative probability mass to half-lives greater than six years than the consensus prior, one would hope that the data would be able to discriminate between these two hypotheses. The last column of Table IV, however, shows that for no country the Bayes factors favours either prior by more than a factor of 2.¹⁵

¹⁴ For details of the construction of this diffuse half-life prior the reader is referred to the Technical Appendix.

¹⁵ The marginal likelihood under each prior is computed as described in Section 2. The ratio of these two marginal likelihood functions gives the Bayes factor.

5.3. Further Sensitivity Analysis

The consensus prior is of special interest because it reflects the expressed prior beliefs of economists. It is not clear, however, whether our finding that the data do not favour any one prior in particular will also extend to comparisons not involving the consensus prior. We therefore complete the analysis by comparing a number of hypotheses not including the consensus prior.

These Bayes factors may be computed as a by-product of our earlier analysis. As in expression (12), let model A denote the model under the consensus prior and consider an alternative model B under a prior that assigns probability 1 to a given range of the parameter region. Given two Bayes factors in favour of model B over model A and in favour of model \bar{B} over model A , where \bar{B} denotes the complement of B , it is straightforward to compute the Bayes factor for model B over model \bar{B} by taking the ratio of the two original Bayes factors. Table V shows that the data do not favour any of the hypotheses, with few exceptions. For example, for the UK the Bayes factor in favour of the hypothesis $0 < h \leq 40$ is about 7; for Japan and for Canada, the Bayes factors against the hypothesis $0 < h \leq 2$ are 20 and 10, respectively. The evidence for the other countries, however, is mixed and inconclusive.

5.4. Interpretation

These results are illustrative of a general problem. Undoubtedly, one could refine the hypotheses to reflect the implications of specific classes of theoretical economic models for the degree of mean reversion in real exchange rates. Tables IV and V, however, make it clear that the data are unlikely to be strongly supportive of any one hypothesis, at least as long as we analyse one

Table V. Bayes factors for selected hypotheses not involving the consensus prior

Country	Rogoff	Real rigidities			
		None	Some	Strong	Severe
		3–5 years	0–2 years	2–6 years	6–40 years
		Relative to complement			
Austria	1.71	0.65	2.03	0.73	0.31
Belgium	1.82	0.36	1.92	0.90	0.47
Canada	1.32	0.05	1.06	1.82	0.96
Denmark	1.75	0.55	2.01	0.78	0.37
Finland	1.50	1.12	1.99	0.56	0.26
France	1.60	0.86	2.05	0.62	0.29
Germany	1.67	0.67	2.04	0.71	0.36
Greece	1.56	0.82	1.90	0.71	0.33
Italy	1.67	0.82	2.10	0.63	0.35
Japan	1.99	0.10	1.83	1.14	0.36
Netherlands	1.59	0.86	2.06	0.64	0.31
Norway	1.31	1.58	1.89	0.46	0.22
Portugal	1.73	0.33	1.79	1.00	0.43
Spain	1.82	0.34	2.03	0.87	0.39
Sweden	1.76	0.44	1.89	0.87	0.41
Switzerland	1.52	1.30	2.10	0.46	0.22
UK	1.24	1.98	1.80	0.39	0.14
Mean	1.62	0.75	1.91	0.78	0.36

Note: Monthly IFS data.

country and one variable at a time. The contribution of this paper has been to establish formally the point that there is only very limited information in the post-Bretton Woods data about the hypotheses of economic interest. The fact that the posterior for the half-life in general will be heavily influenced by the prior distribution suggests that there is no reason to expect empirical work to settle disagreements about which range of the half-life is most likely, unless we depart from the modelling framework pursued in this paper.

One possible direction for future research would be to move away from univariate linear reduced-form models of the real exchange rate and to focus instead on the responses to specific shocks in a fully structural model involving several macroeconomic variables (see Ng, 2000 for a recent example). Economic theory suggests that movements in the real exchange rate are due to a variety of underlying economic forces. Different driving forces (such as monetary shocks and productivity shocks) are likely to have different effects on deviations from PPP. Clearly, building such structural models is a daunting research agenda because it inevitably involves controversial identifying assumptions. Another possible response would be to pool data for several countries and to focus on bilateral or possibly multilateral relationships between countries. Given the large number of parameters such models will contain, Bayesian methods based on hierarchical priors seem well suited to this purpose.

6. CONCLUSIONS

This paper has analysed the uncertainty about the half-life of deviations from PPP under the recent float. In contrast to recent related work in the classical paradigm by Murray and Papell (2000a) we brought to bear additional prior information about the half-life based on survey information collected from economists with an interest in the PPP debate. We approximated widely held views among the survey respondents by a consensus prior for the half-life. We derived the posterior distribution of the half-life under this consensus prior. Our results confirmed the view that the value of the half-life under the recent float is highly uncertain, a conclusion consistent with that reached by Murray and Papell (2000a) using very different methods.

We also found that the posterior distribution, although not identical to the prior, was greatly influenced by the consensus prior. This fact motivated us to study the extent of the empirical support for the consensus prior relative to other plausible hypotheses of economic interest. We provided for the first time a comprehensive formal evaluation of several nonnested hypotheses of economic interest, including Rogoff's claim that the half-life is contained in a range between three and five years. We also considered a battery of hypotheses corresponding to half-life values in ranges that might be expected in the presence of no, some, strong or severe real rigidities in the economy. Finally, we compared the consensus prior to a more diffuse half-life prior that assigned substantially more probability mass to longer half-lives. With rare exceptions, we found that the data were not conclusive.

Our results are useful in two ways. First, they suggest that researchers interested in the half-life of deviations from PPP will look in vain for conclusive empirical evidence from studies of one country and one variable at a time. Second, our results motivate the analysis of structural models of the real exchange rate that involve multiple countries. Bayesian methods seem well suited to dealing with such highly parameterized models. Our approach in this paper represents a first step in that direction.

TECHNICAL APPENDIX: REWEIGHTING SCHEME FOR 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. We can, however, postulate a prior density for h that, although not completely flat, will be quite diffuse. This diffuse specification may be achieved using the same type of prior already adopted in Section 4. 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 slightly modify the procedure of Section 4 by rescaling the diffuse prior such that $\Pr(h > 40)$ is 4.8%, consistent with the consensus prior.

Let $\theta = (\lambda, b)$ denote the parameter vector of the model. We partition the distribution of h into two regions. In region R_0 , the value of h is less than or equal to 40 years; in region R_1 , h is 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 $\Pr(h \in R_1)$ 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 reweighted diffuse prior density is shown in Figure 2 together with the consensus prior. The posterior associated with this re-weighted prior density may be simulated using the weighting method described by Tierney (1994) and Geweke (1997). As shown in Geweke (1997), given the posterior distribution of θ (and functions of θ) under the prior $p(\theta)$, the corresponding posterior under the prior $p^*(\theta)$ can be computed without much additional effort. Following Tierney (1994) and Geweke (1997) define the probability weight

$$\bar{w}^*(\theta) = \frac{p^*(\theta|Y)}{p(\theta|Y)} = \frac{p^*(\theta)L(\theta|Y)}{p(\theta)L(\theta|Y)} = \frac{p^*(\theta)}{p(\theta)}$$

where $\bar{w}^*(\theta)$ is bounded. Then the matrix

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

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

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

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 4 instead of deriving analytically the posterior under the mixed normal prior. Specifically, one could have simulated the posterior under the flat prior ($p(\theta) \propto \text{constant}$) and then reweighted the draws by w^* for any given proper target density $p^*(\theta)$. In our case, however, $p^*(\theta)$ 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^* will 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.

ACKNOWLEDGEMENTS

We acknowledge helpful comments from Bob Barsky, Kathryn Dominguez, Charlie Evans, Gordon Hanson, Jan Kmenta, Chris Sims, Linda Tesar, Harald Uhlig, Dan Waggoner, Chuck Whiteman, and especially from John Geweke and three anonymous referees. We are indebted to John Geweke for discovering an error in an earlier draft of the paper. We also thank seminar participants at the Chicago Fed, the University of Michigan, the 1999 NBER Summer Institute and the 2000 European Summer Symposium in Macroeconomics organized by the CEPR. We are grateful to 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 in this paper do not necessarily reflect the views of the Federal Reserve Bank of Atlanta or the Federal Reserve System.

REFERENCES

- Abuaf N, Jorion P. 1990. Purchasing power parity in the long-run. *Journal of Finance* **45**: 157–174.
- Caner M, Kilian L. 2000. Size distortions of tests of the null hypothesis of stationarity: evidence and implications for the PPP debate. *Journal of International Money and Finance* (forthcoming).
- Cheung Y-W, Lai KS. 1998. Parity reversion in real exchange rates during the post-Bretton Woods period. *Journal of International Money and Finance* **17**: 597–614.
- Cheung Y-W, Lai KS. 2000a. On cross-country differences in the persistence of real exchange rates. *Journal of International Economics* **50**: 375–397.
- Cheung Y-W, Lai KS. 2000b. On the purchasing power parity puzzle. *Journal of International Economics* **52**: 321–330.

- Chib S. 1995. Marginal likelihood from the Gibbs output. *Journal of the American Statistical Association* **90**: 1313–1321.
- Chib S, Greenberg E. 1995. Understanding the Metropolis–Hastings algorithm. *American Statistician* **49**: 327–335.
- Diebold FX, Husted S, Rush M. 1991. Real exchange rates under the gold standard. *Journal of Political Economy* **99**: 1252–1271.
- Edison H, Gagnon J, Melick W. 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 *How Open is the U.S. Economy?* Hafer R (ed.). Lexington Books: Lexington.
- Frankel JA, Froot KA. 1987. Using survey data to test standard propositions regarding exchange rate expectations. *American Economic Review* **77**: 133–153.
- Froot KA, Rogoff K. 1995. Perspectives on PPP and long-run real exchange rates. In *Handbook of International Economics* (Vol. III), Grossman G, Rogoff K (eds). Elsevier Science: Amsterdam: 1647–1688.
- Geweke J. 1995. Monte Carlo simulation and numerical integration. In *Handbook of Computational Economics*, Amman H, Kendrick D, Rust J (eds). North-Holland: Amsterdam.
- Geweke J. 1997. Posterior simulators in econometrics. In *Advances in Economics and Econometrics: Theory and Applications*, Kreps DM, Wallis KF (eds). Econometric Society Monographs Seventh World Congress Vol. III, Cambridge University Press: New York; 128–165.
- Geweke J. 1999. Simulation methods for model criticism and robustness analysis. Forthcoming in *Bayesian Statistics 6*, Berger JO, Bernardo JM, David AP, Smith AFM (eds). Oxford University Press: Oxford.
- Koedijk KG, Schotman PC, Van Dijk MA. 1998. The re-emergence of PPP in the 1990s. *Journal of International Money and Finance* **17**: 51–61.
- Litterman RB. 1986. Forecasting with Bayesian vector autoregressions—five years of experience. *Journal of Business and Economic Statistics* **4**: 25–38.
- Lothian JR, Taylor MP. 1996. Real exchange rate behavior: the recent float from the perspective of the past two centuries. *Journal of Political Economy* **104**: 488–510.
- Murray CJ, Papell DH. 2000a. The purchasing power parity persistence paradigm. *Journal of International Economics* (forthcoming).
- Murray CJ, Papell DH. 2000b. The purchasing power parity persistence puzzle is worse than you think: a note on long run real exchange rate behavior, manuscript, Department of Economics, University of Houston.
- Ng S 2000. Can sticky prices account for the variations and persistence in real exchange rates, manuscript, Department of Economics, Boston College.
- Obstfeld M, Rogoff K. 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, Theodoridis H. 1998. Increasing evidence of purchasing power parity over the current float. *Journal of International Money and Finance* **17**: 41–50.
- Parsley DC, 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 JC, Tallman EW. 1999. Vector autoregressions: forecasting and reality. *Federal Reserve Bank of Atlanta Economic Review* First Quarter: 4–18.
- 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.
- Rogoff K. 1996. The purchasing power parity puzzle. *Journal of Economic Literature* **34**: 647–668.
- Schotman P, van Dijk HK. 1991. A Bayesian analysis of the unit root in real exchange rates. *Journal of Econometrics* **49**: 195–238.
- Sims CA. 2000. Using a likelihood perspective to sharpen economic discourse: three examples. *Journal of Econometrics* **95**: 443–462.
- Sims CA, Zha T. 1998. Bayesian methods for dynamic multivariate models. *International Economic Review* **39**: 949–986.
- Taylor MP, Peel D. 1998. Nonlinear mean reversion in real exchange rates: toward a solution to the purchasing power parity puzzles, manuscript, Oxford University.

- Taylor MP, Sarno L. 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 DF, Zha T. 1999. Does normalization matter for inference?, manuscript, Federal Reserve Bank of Atlanta.
- Wei S-J, Parsley DC. 1995. Purchasing power disparity during the floating rate period: exchange rate volatility, trade barriers, and other culprits. NBER Working Paper No. 5032.
- Wu S, Wu JL. 1998. Purchasing power parity under the current float: new evidence from panel data unit root tests. *Journal of Money, Credit and Banking* (forthcoming).
- 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.