

July 2004

The Power of the "Objective" Bayesian Unit-Root Test

Francis W. Ahking
University of Connecticut

Follow this and additional works at: http://digitalcommons.uconn.edu/econ_wpapers

Recommended Citation

Ahking, Francis W., "The Power of the "Objective" Bayesian Unit-Root Test" (2004). *Economics Working Papers*. 200414.
http://digitalcommons.uconn.edu/econ_wpapers/200414



University of
Connecticut

Department of Economics Working Paper Series

The Power of the "Objective" Bayesian Unit-Root Test

Francis W. Ahking
University of Connecticut

Working Paper 2004-14

July 2004

341 Mansfield Road, Unit 1063
Storrs, CT 06269-1063
Phone: (860) 486-3022
Fax: (860) 486-4463
<http://www.econ.uconn.edu/>

Abstract

Some researchers, for example, Koop (1992), and Sims (1988), advocated for Bayesian alternatives to unit-root testing over the classical approach using the augmented Dickey-Fuller test (ADF). This paper studies the power of what Koop (1992) has called the "Objective" Bayesian approach to unit-root testing. Koop's "objective" Bayesian test is interesting in light of the call by Phillips (1991a, 1991b) for more objective Bayesian analysis of time series. We apply the "objective" Bayesian unit-root test to a study of long-run purchasing power parity (PPP) in the post-Bretton Woods era and also Monte Carlo simulations. Overall, our results suggest that the "objective" Bayesian test is biased in favor of trend-stationarity. We conclude that, at least for the "objective" Bayesian test, it is not better than the classical ADF approach in unit-root tests, and because of its bias, the "objective" priors suggested by Koop is not appropriate.

Journal of Economic Literature Classification: C11, C22, F31

I wish to thank Professor Stephen M. Miller, participants at the Southern Economic Association Annual Meetings, and Department of Economics, University of Connecticut, brown-bag seminar for comments on an earlier draft of this paper. Remaining errors are my sole responsibilities.

The Power of the “Objective” Bayesian Unit-Root Test

1. Introduction

The question of how best to characterize the growth component of macroeconomic time series has been hotly debated since the publication of the study by Nelson and Plosser (1982). The most widely held view before 1982 was that most macroeconomic time series could be characterized as trend-stationary (TS) series. Nelson and Plosser’s study called into question that widely held view by demonstrating that thirteen out of fourteen U.S. macroeconomic time series in their study contained a unit root, and thus were difference-stationary (DS) series. Although the debate is far from settled,¹ a related debate is concerned with the statistical techniques to use in testing between TS and DS series. Several economists, e.g., DeJong and Whiteman (1991), Koop (1992), and in particular, Sims (1988), and Sims and Uhlig (1991), have advocated forcefully for Bayesian alternatives over the more traditional classical approach such as the ADF tests (Dickey and Fuller, 1981) in unit-root testing. These economists cited several advantages of the Bayesian approach over the classical approach. For example, it is well known that ADF tests have low power against plausible alternatives, especially against trend-stationary alternative. The Bayesian approach, on the other hand, would reveal that both the unit root and the trend-stationary hypotheses would receive similar posterior probabilities. Thus, the Bayesian approach provides a more reasonable summary of sample information than the classical approach. Another problem with the classical unit-root tests is the discontinuity of the classical asymptotic theory when there is a unit root [see Sims (1988) especially for a good discussion of this point]. The Bayesian approach, on the other hand, since it is based on the likelihood function, does not have the same discontinuity problem. Finally, Koop (1994) has also pointed out that, in the classical approach, where small sample critical values are used in practice, could frequently differ substantially from the asymptotic critical values. The Bayesian approach, on the other hand, since it is conditional on the observed sample, provides exact small sample results. Indeed, an entire issue of the *Journal of Applied Econometrics* (JAE, 1991) was devoted to a discussion of these and other issues.

Despite the apparent advantages of the Bayesian approach over the classical approach in unit-root testing, only a relatively small number of studies have appeared using the Bayesian approach. The reasons may be that the Bayesian approach requires a likelihood function and the use of prior information. The perception is that other than for the simplest cases, the Bayesian approach is computationally burdensome, and the use of prior information is by

far the most controversial. Phillips (1991a, 1991b) identified the need for priors as the biggest obstacle to Bayesian analysis and argued for more objective Bayesian analysis in time series.

In this paper, we study what Koop (1992) has called the “objective” Bayesian test. Koop’s (1992) Bayesian unit-root test is based on the work of Zellner and Siow (1980). There are several features of this approach that make it attractive for economists looking for an alternative to the classical unit-root tests. First, it is computationally simple. Second, it requires informative priors, thus avoiding the problems associated with using non-informative priors, which are frequently improper. Third, it does not require significant subjective prior information, only that all competing hypotheses have equal prior probability, leading Koop (1992) to call this an “objective” Bayesian test.

Koop’s (1992) Bayesian test appears to address one of the biggest criticism of Bayesian econometrics and is very much in the spirit of Phillips’s (1991a, 1991b) call for more objective analysis in Bayesian time series. However, the usefulness of Koop’s (1992) Bayesian unit-root test, and in particular, whether or not it provides a better alternative to classical unit root tests, and how appropriate are his objective priors are questions that have not been examined carefully. We seek to answer these two related questions objectively by providing a Monte Carlo study of the power of Koop’s (1992) “objective” Bayesian unit-root test. We start our empirical analysis by applying the “objective” unit-root test to a study of long-run purchasing power parity (PPP) to see how this test would perform in actual empirical application. This is also motivated by the fact that frequently, the usefulness of an empirical test is determined, to a large extent, by how well it performs in actual application. Next, we perform a number of simulations to assess the power of the test and thus provide information on both the usefulness of the test and the appropriateness of the “objective” priors.

As we will argue below (see footnote 3), however, even-though we assign equal prior probability to all the competing hypotheses, the unit-root hypothesis actually receives the highest prior probability. We provide a brief discussion of Koop’s (1992) “objective” Bayesian unit-root test in the next section. We present our empirical results on long-run PPP in section 3. We study the properties of the “objective” Bayesian unit-root test using Monte Carlo simulations in section 4. Finally, our summary and conclusions are in section 5.

2. Koop’s “objective” Bayesian Approach

The modeling objective of the Bayesian approach is not to reject or fail to reject a hypothesis based on a pre-determined level of significance, but to determine how probable a hypothesis is relative to other competing

hypotheses. There are several ways of comparing hypotheses using Bayesian methods. The most common method is to calculate posterior odds ratios for various competing hypotheses based on prior and sample information. This gives the researcher the odds in favor of one hypothesis relative to other competing hypotheses. The objective in Koop's "objective" Bayesian approach is to find the linear model that would best describe the time series,

$\underline{q} = (q_1, q_2, \dots, q_T)'$. We consider three hypotheses:

$$H_1 : \quad q_t = \alpha_0 + \sum_{i=1}^k \alpha_i q_{t-i} + \alpha_{k+1} t + \xi_{1t}, \quad (1)$$

$$H_2 : \quad q_t = \alpha_0 + \sum_{i=1}^k \alpha_i q_{t-i} + \xi_{2t}, \quad (\alpha_{k+1} = 0) \quad (2)$$

$$H_3 : \quad \Delta q_t = \alpha_0 - \left(\sum_{i=2}^k \alpha_i \right) \Delta q_{t-1} - \left(\sum_{i=3}^k \alpha_i \right) \Delta q_{t-2} - \dots - \alpha_k \Delta q_{t-(k-1)} + \xi_{3t}, \quad (3)$$

$$\left(\sum_{i=1}^k \alpha_i = 1, \alpha_{k+1} = 0 \right)$$

where t = a linear deterministic time trend; and ξ_{jt} , $j = 1, 2, 3$ is a serially uncorrelated error process with zero-mean and constant variance.

Hypothesis 1 (H_1) is the null model. It hypothesizes a trend-stationary auto-regressive process of order k , i.e., a trend-stationary AR(k) process. H_2 hypothesizes a stationary AR(k) process, while H_3 hypothesizes an AR(k) process with a unit root. Note that both H_2 and H_3 are special cases of H_1 with linear restrictions (given in parentheses next to the respective equations), imposed on the null model. The trend-stationary hypothesis is included because it is the leading alternative to unit-root non-stationarity in macroeconomics time series. The stationary alternative is also included to see how well the "objective" Bayesian test can distinguish between non-stationary series and stationary series with a high degree of persistence, as is frequently encountered in macroeconomic time series.

We compare the three hypotheses, based on both prior and sample information, by calculating the posterior odds ratios:

$$K_{1j} = \frac{P(H_1)P(H_1|\underline{q})}{P(H_j)P(H_j|\underline{q})}, j = 2, 3, \quad (4)$$

where $P(H_1)/P(H_j)$ is called the prior odds ratio, and $P(H_i|\underline{q}) = \int P(\theta_i|H_i)L(\theta_i|q, H_i)d\theta_i, i = 1, 2, 3,$

is the posterior probability that $H_i, i = 1, 2, 3,$ were true given the sample data $\underline{q}, P(\theta_i|H_i)$ and

$L(\theta_i|\underline{y}, H_i)$ are the prior density and likelihood function, respectively, under each hypothesis, and $\underline{\theta}$ is a vector of parameters. Thus, the posterior odds ratio gives the ratio of the probabilities of the two hypotheses holding given the sample data.² On the assumption that all three hypotheses have equal prior probability,³ i.e.,

$P(H_1) = P(H_2) = P(H_3),$ equation (4) reduces to

$$K_{1j} = \frac{P(H_1|\underline{q})}{P(H_j|\underline{q})}, j = 2, 3. \quad (5)$$

Following Koop (1992), we calculate the posterior odds ratio for testing a set of exact linear restrictions with a formula suggested by Zellner and Siow (1980). The Zellner-Siow posterior odds ratio is calculated approximately as:

$$K_{1j} \cong \frac{(\pi^{0.5} / \Gamma[(r+1)/2])(v/2)^{r/2}}{(1+rF/v)^{(v-1)/2}}, j = 2, 3, \quad (6)$$

where $\Gamma[\cdot]$ = the Gamma function, $v = T - n,$ T = the total number of observations, n = the number of regressors in the null model, r = the number of linear restrictions tested, and F = the usual F-statistics for testing the set of linear restrictions. After we obtained the posterior odds ratios, we can then calculate the posterior probability for each of the three hypotheses. We apply the ‘‘objective’’ Bayesian test to a study of long-run PPP in the next section.

3. Testing for Long-Run PPP

There is a vast literature on empirical long-run PPP, but we are aware of only three studies using a Bayesian approach [Schotman and van Dijk (1991), Whitt (1992), and Ahking (1997)]. Long-run PPP requires that the real exchange rate is a stationary time series. Empirical evidence has been mixed and not very robust, however, especially using data from the post-Bretton Woods era [see the recent survey by Rogoff (1996), and Sarno and Taylor (2002)].

We define the real exchange rate in natural logarithm form as:

$$q_t = e_t + P_t^* - P_t,$$

where q_t is the natural logarithm of the real exchange rate, e_t is the natural logarithm of the nominal exchange rate, defined as the domestic currency price of one unit of foreign currency, and $P_t (P_t^*)$ is the natural logarithm of an index of the domestic (foreign) price level. A test for long-run PPP is a test of whether or not q_t is a stationary time series. The source of our data is SourceOECD, and consist of the OECD G-7 countries - the U.S., the U.K., Canada, Germany, Italy, France, and Japan. Our data consist of monthly observations from April 1973 to February 1999 for the G-7 countries, and are not seasonally adjusted. The sample ends when several of the G-7 countries became part of the Euro zone. Nominal exchange rates are bilateral and are monthly averages. In all cases, we use the consumer price index as our measure of the average price level. Thus, we have a total of twenty-one real exchange rates in our sample.⁴

We start our empirical tests by first presenting our test for unit-root using the ADF test. The ADF unit-root test results will provide a comparison to our Bayesian approach to unit-root testing. The ADF regression actually estimated is:

$$\Delta q_t = \beta_0 + \phi q_{t-1} + \sum_{i=1}^l \beta_i \Delta q_{t-i} + \beta_{l+1} t + v_t, \quad (7)$$

where t = a linear deterministic time trend, and v_t is a serially uncorrelated error process with zero mean and constant variance. We have included a linear deterministic time trend since several researchers have found that the stochastic processes of some of the real exchange rates in our study could not be adequately modeled without the inclusion of a linear deterministic time trend [see Cheung and Lai (1998), and Koedijk, Schotman, and Van Dijk (1998) for recent examples]. Furthermore, since the time series of real exchange rates typically show evidence of a trend, the trend-stationary hypothesis offers an alternative to the unit-root hypothesis as the source of the trend in the data. The linear deterministic time trend in the real exchange rate is generally interpreted as representing systematic differences in productivity growth between tradable and non-tradable goods in the two countries [see also Cheung and Lai (1998) for a recent discussion]. Thus, the presence of a linear deterministic time trend in the real exchange rate time series is generally not interpreted as a violation of long-run PPP.⁵ The lag length for the lagged first-differences is determined by using a general-to-specific method recommended by Ng and Perron (1995) and Perron

(1997). We start by estimating Equation (7) with a pre-determined maximum lag length of 12 and sequentially drop the last included lag if it is not statistically significant at the 10% significance level. If, however, the lag length determined is the same as the maximum lag length, we start over with a maximum lag length of 14.⁶

4. Empirical Results

Table 1 reports the empirical results using the ADF test. Column 2 reports the lag length selected for each real exchange rate, columns 3 and 4 report the t-statistic for the hypothesis $H_0 : \phi = 0$ without and with a trend, respectively. We use the 5% and the 10% critical values from Fuller (1976) and the lag-adjusted critical values for exact sample size from Cheung and Lai (1995). The results are as expected and consistent with earlier studies using the ADF tests. The evidence against long-run PPP is rather strong without the inclusion of a time trend and slightly better when a time trend is included.

Table 2 presents the results of the “objective” Bayesian tests. In order to focus on the three main hypotheses, we use a non-Bayesian method to determine the lag length.⁷ Since it is not clear how best to determine the lag length for an equation such as equation (1), we keep it simple by starting with a lag length of one, and then two, and so on, and stopping when we obtain a white-noise error process.⁸ We also experimented with using a uniform lag of four for all the real exchange rates, the results are consistent with those reported in Table 2. What we find striking is that, of the three hypotheses, the trend-stationary hypothesis receives the highest posterior probability in all cases except for the Japanese yen/German mark real exchange rate. In that case, the stationary hypothesis receives the highest posterior probability. The French franc/German mark real exchange rate also deserves mention because it is the only case where the trend-stationary and the stationary hypotheses receive approximately the same posterior probabilities. Interestingly, these are the same two real exchange rates that Cheung and Lai (1998) have found to be well characterized by stationary or trend-stationary processes using the conventional ADF tests. The unit root hypothesis, on the other hand, receives no significant posterior probability. In sum, the Bayesian results strongly support the hypothesis that the real exchange rate are trend-stationary AR processes. Or, put differently, the restrictions on the null model are not strongly supported by the data. In the next section, we will investigate the reliability of the “objective” Bayesian results by performing a number of simulations.

5. The power of the “objective” Bayesian unit-root test

In this section, we perform a number of Monte Carlo experiments to study the power of the “objective” Bayesian unit-root test. This will allow us to assess the reliability of our results in Section 4, investigate the usefulness the “objective” Bayesian approach in empirical application, and whether or not it is better than the classical ADF approach in unit-root tests. The power of the “objective” Bayesian unit-root test in these experiments is defined as the percentage of time that the hypothesis corresponding to the data generating model (DGM) receives the highest posterior probability in repeated samples. For all our experiments, we use a DGM of the following form:

$$y_t = \delta + \rho y_{t-1} + \lambda t + u_t, \quad u_t \sim iid N(0, \sigma_u^2),$$

where $\underline{\theta} = (\delta, \rho, \lambda)$ is a vector of coefficients. For the first experiment, we use

$\underline{\theta} = (0.059317, 0.95 \leq \rho \leq 0.99, -0.0000201)$, $u_t \sim iid N(0, 0.000564)$, and ρ varies from 0.95 to 0.99, in increment of 0.01. Thus, the data generating process is a trend-stationary AR(1) model. The constant, the time-trend coefficient, and the variance of the error process are averages from estimating the null model (Equation 1) with data from the twenty-one real exchange rate series. Thus, they represent economically plausible parameters. The values of ρ that we choose are also based on what we consider to be plausible alternatives for monthly series. For example, Sims (1988) argued that for monthly data, it is reasonable to concentrate the prior odds on the interval (0.98, 1) as opposed to, say (0.5, 1). Second, there appears to be a consensus among economists that deviations from long-run PPP have a half-life of about three to five years [see e.g., Abuaf and Jorion (1990), Rogoff (1996)].⁹ According to Caner and Kilian (2001), this corresponds to a ρ value of 0.98 and 0.99, respectively, for a half-life of three and five years, respectively, using monthly data. In sum, we believe that the parameter values that we choose for our Monte Carlo experiment are economically plausible.

For all our simulations, we first generate $T + 100$ observations, where T is the actual sample size used in our simulations, and discard the first 100 observations to avoid the initialization problem. All experiments are replicated 5,000 times for each ρ value. We start our simulation by letting $T = 311$, which corresponds to the sample size of our real exchange rates. The results are reported in the upper panel of Table 3. The results reveal a surprising pattern. For $0.95 < \rho < 0.98$, the trend-stationary hypothesis receives the highest posterior probability the largest proportion of the time. The power of the test increases from $\rho = 0.95$ and peaks at

$\rho = 0.97$. At $\rho = 0.99$, however, the unit-root hypothesis receives the highest posterior probability the largest proportion of the time, suggesting an extremely low power of the “objective” Bayesian test at this value of ρ .

We perform a second experiment using a DGM model with $\underline{\theta} = (0.059317, 0.95 \leq \rho \leq 1.00, 0.0)$, i.e., an AR(1) model, and u_t has the same properties as the first experiment. In this experiment, however, we vary ρ from 0.95 to 1.00, in increment of 0.01. Thus, the DGM is a stationary AR(1) model and a random-walk with drift model. The results are reported in the lower panel of Table 3. Interestingly, for $0.95 < \rho < 0.99$, the results mirror the results reported in the upper panel of Table 3. This, perhaps, is not surprising since the effect of the time trend is extremely small, thus making the two DGMs almost the same. At $\rho = 1.00$, the trend-stationary hypothesis receives the highest posterior probability the largest percentage of the time, followed by the stationary hypothesis, while the unit-root hypothesis receives the highest posterior probability less than 1% of the time.

The results in Table 3 are not very encouraging to the “objective” Bayesian unit-root test. First, for $0.95 \leq \rho \leq 0.98$, it cannot distinguish between a highly persistent trend-stationary model from a highly persistent stationary AR model. There appears to be a bias in favor of the trend-stationary model. At $\rho = 0.99$ the unit-root hypothesis is favored regardless of the DGM used. When the DGM is a random-walk with drift model, the “objective” Bayesian unit-root is biased in favor of the trend-stationary model. Thus, just as the classical ADF test is criticized frequently for its bias in favor of finding a unit-root, it appears that the “objective” Bayesian unit-root test can also be criticized for its bias in favoring trend-stationarity. Given the Monte Carlo results, it is impossible to draw any conclusions regarding long-run PPP since the underlying DGMs of the real exchange rates are unknown. Furthermore, the simulation results suggest that the “objective” Bayesian unit-root test does not provide a better statistical approach than the classical ADF test in unit-root testing.

The above simulations were done using parameter values suggested by the real exchange rate sample. To gain further insight into the “objective” Bayesian test, we performed additional simulations to determine if the “objective” Bayesian unit-root test may be sensitive to sample size, the size of the constant, or the size of the trend coefficient. In particular, we use sample sizes of $T = 60, 240, 720$, which are approximately the sample sizes of post war annual, quarterly, and monthly data, respectively. For the constant term, we use $\delta = 1, 0, \text{ and } -1$, and for the trend-stationary model, we use $\lambda = 0.05, \text{ and } -0.05$. We perform simulations for all possible combinations of

T , δ , and λ , using both the AR(1) and the trend-stationary DGMs, and using $\sigma_u^2 = 1$ in all cases. For both the trend-stationary and the AR(1) DGMs, the results are not very sensitive to the constant and the trend coefficient that we use but are sensitive to the sample size. Rather than reporting the large number of results, we will instead report in Tables 4 and 5 only the representative results. Table 4 shows the results for the trend-stationary AR(1) DGM with $\underline{\theta} = (1.00, 0.95 \leq \rho \leq 0.99, 0.05)$. For $T = 60$, the “objective” Bayesian unit-root test cannot easily distinguish between the trend stationary and the stationary hypotheses. Its power, however, increases gradually as ρ approach 0.99. For the two larger samples, the power of the “objective” test drops off significantly, for $0.95 \leq \rho \leq 0.98$, the stationary hypothesis receives the largest posterior probability the largest percentage of the time. For $\rho = 0.99$, the “objective” test favors the unit-root hypothesis almost 100 percent of the time, a result which is similar to that of the top panel of Table 3. Turning now to Table 5 where we use an AR(1) DGM with $\underline{\theta} = (1.00, 0.95 \leq \rho \leq 1.00, 0.00)$, we see that the power of the “objective” Bayesian test is extremely low. For $T = 60$ and $T = 240$, there is a clear bias in favor of the trend stationary hypothesis, and also for $T = 720$ and $0.98 \leq \rho \leq 1.00$. The unit root hypothesis is favored when $0.95 \leq \rho \leq 0.97$ for $T = 720$.

Finally, we are also interested in knowing whether or not the power of the “objective” Bayesian test would change if the DGM is less persistent. We repeated our simulations for the three sample sizes with the AR(1) DGM, and $\underline{\theta} = (1.00, \rho = 0.70, 0.75, 0.80, 0.85, 0.90)$. The results are shown in Table 6. Once again, the power of the “objective” Bayesian test is extremely low, and interestingly, for the two larger sample sizes, the unit-root hypothesis is favored most of the time.

6. Summary and Conclusions

Researchers generally agree that the Bayesian approach offers an alternative and useful way to the classical approach in empirical modeling. In unit-root testing, Sims (1988), Sims and Uhlig (1991), and Koop (1992, 1994) have advocated the Bayesian approach over the classical ADF tests. Phillips (1991a, 1991b) identified one of the criticisms of the Bayesian approach is the need for subjective priors, and called for more objective Bayesian analysis. Koop (1992) offered one approach which requires less subjective input from the researcher, is computationally simple, and thus offers an attractive alternative to the classical approach. In this paper, we study the power of Koop’s “objective” unit-root test. In particular, we are interested in whether or not it provides a better

alternative to the classical ADF unit-root test, and whether or not the use of “objective” priors are appropriate. We first test for long-run PPP using both the classical ADF and the “objective” Bayesian tests. The results with the classical ADF suggest little support for long-run PPP. Our “objective” Bayesian results, on the other hand, provide a stark contrast to the ADF results. In all cases, the hypothesis of a unit root does not receive significant posterior probability. Rather, sample information appear to strongly support the hypothesis of trend-stationarity for all cases except the Japanese yen/German mark real exchange rate where the sample information suggest a stationary time series. The French franc/German mark real exchange rate is the other case where the stationary hypothesis receives significant posterior probability.

Next, we study the power of the “objective” Bayesian test with Monte Carlo simulations. The results are not very encouraging. In particular, using economically plausible parameters for monthly data of real exchange rates for our DGMs, we find that the “objective” Bayesian test cannot distinguish between a trend-stationary AR model from a stationary AR model when the time trend effect is relatively small, and the time series is highly persistence. The bias is in favor of finding a trend-stationary model. When $\rho = 0.99$, the “objective” Bayesian test is biased in favor of a unit-root. On the other hand, when the DMS is a random-walk with drift model, the test is biased in favor of the trend-stationary hypothesis. Additional simulation results suggest that the “objective” Bayesian test is sensitive to sample sizes, and is biased in favor of the trend-stationary hypothesis regardless of the DGM used. For relatively moderate ρ values, the power of the “objective” Bayesian test actually decreases with sample size.

There is no evidence to suggest that the “objective” Bayesian test is better than the classical ADF tests in unit-root tests. Moreover, the “objective” priors do not appear to be appropriate since they tend to produce results that are biased in favor of the trend-stationary hypothesis. Thus, unfortunately, while there is a need for more objective analysis of Bayesian time series, Koop’s “objective” Bayesian test does not appear to move us closer to that goal.

References

- Abuaf, Niso, Jorion, Philippe, 1990. Purchasing power parity in the long run. *The Journal of Finance* **45**, 157-174.
- Ahking, Francis W., 1997. Testing long-run purchasing power parity with a Bayesian unit root approach: the experience of Canada in the 1950s. *Applied Economics* **29**, 813-819.
- Caner, Mehmet, Kilan, Lutz, 2001. Size distortions of tests of the null hypothesis of stationarity: evidence and implications for the PPP debate. *Journal of International Money and Finance* **20**, 639-657.
- Cheung, Yin-Wong, Lai, Kon S., 1995. Lag order and critical values of the augmented Dickey-Fuller test. *Journal of Business and Economic Statistics* **13**, 277-280.
- Cheung, Yin-Wong, Lai, Kon S., 1998. Parity reversion in real exchange rates during the post-Bretton Woods period. *Journal of International Money and Finance* **17**, 597-614.
- Culver, Sara, Papell, David H., 1999. Long-run purchasing power parity with short-run data: evidence with a null hypothesis of stationarity. *Journal of International Money and Finance* **18**, 751-768.
- DeJong, David N., Whiteman, Charles H., 1991. The case for trend-stationarity is stronger than we thought. *Journal of Applied Econometrics* **6**, 413-421.
- Dickey, David A., Fuller, Wayne A., 1981. Likelihood ratio statistics for autoregressive time series with a unit root. *Econometrica* **49**, 1057-1072.
- Fuller, Wayne A., 1976. *Introduction to Statistical Time Series*. John Wiley and Sons, New York.
- Journal of Applied Econometrics. October-December, 1991.
- Koedijk, Kees G., Schotman, Peter C., Van Dijk, Mathijs A., 1998. The re-emergence of PPP in the 1990s. *Journal of International Money and Finance* **17**, 51-61.
- Koop, Gary, 1992. 'Objective' Bayesian unit root tests. *Journal of Applied Econometrics* **7**, 65-82.
- Koop, Gary, 1994. Recent progress in applied Bayesian econometrics. *Journal of Economic Survey* **8**, 1-34.
- Murray, Christian J., Nelson, Charles R., 2000. The uncertain trend in U.S. GDP. *Journal of Monetary Economics* **46**, 79-95.
- Murray, Christian J., Papell, David H. 2002. The purchasing power parity persistence paradigm. *Journal of International Economics* **56**, 1-19.
- Ng, Serena, Perron, Pierre, 1995. Unit root tests in ARMA models with data-dependent methods for the selection of the truncation lag. *Journal of the American Statistical Association* **90**, 268-281.
- Nelson, Charles, Plosser, Charles, 1982. Trends and random walks in macroeconomics time series: Some evidence and implications. *Journal of Monetary Economics* **10**, 139-162.
- Papell, David H., 1997. Searching for stationarity: purchasing power parity under the current float. *Journal of International Economics* **43**, 313-332.
- Perron, Pierre, 1997. Further evidence on breaking trend functions in macroeconomic variables. *Journal of Econometrics* **80**, 355-385.

- Phillips, P. C. B., 1991a. To criticize the critics: an objective analysis of stochastic trends. *Journal of Applied Econometrics* **6**, 333-364.
- Phillips, P. C. B., 1991a. Bayesian unit roots: de rebus prioribus simpler est disputandum. *Journal of Applied Econometrics* **6**, 333-364.
- Rogoff, Kenneth, 1996. The purchasing power parity puzzle. *Journal of Economic Literature* **34**, 647-668.
- Sarno, Lucio, Taylor, Mark, P., 2002. Purchasing power parity and the real exchange rate. *IMF Staff Papers* **49**, 65-105.
- Schotman, Peter, Van Dijk, Herman K., 1991. A Bayesian analysis of the unit root in real exchange rates. *Journal of Econometrics* **49**, 195-238.
- Sims, Christopher A., 1988. Bayesian skepticism on unit root econometrics. *Journal of Economic Dynamics and Control* **12**, 463-474.
- Sims, Christopher A., Uhlig, Harald, 1991. Understanding unit rooters: A helicopter tour. *Econometrica* **59**, 1591-1599.
- Whitt, Joseph A., 1992. The long-run behavior of the real exchange rate. *Journal of Money, Credit and Banking* **24**, 72-82.
- Zellner, Arnold, Siow, A., 1980. Posterior odds for selected regressions. In: Bernardo, J. M., DeGroot, M. H., Lindley, D. V., Smith, A. F. M. (eds), *Bayesian Statistics*, University Press, Valencia, 583-603.

Footnotes

¹ For a recent contribution to the debate, see Murray and Nelson (2000).

² Koop also considered what he called the g -prior approach. This approach, however, requires the subjective input by the researcher on the g values, and thus is less “objective” than the approach considered here.

³ One could argue that even-though the three competing hypotheses have equal prior probability, the unit-root hypothesis is the most favored, and the trend-stationary hypothesis is the least favored. For example, the prior probability for the unit-root hypothesis, i.e., $\rho = 1$ is 33.33 percent, while the stationary interval ($0 < \rho < 1$) also receives 33.33 percent prior probability, distributed uniformly in that interval, resulting in each point in that interval receiving extremely low prior probability. For the trend-stationary hypothesis, however, the 33.33 percent prior probability, in addition to being distributed on the stationary interval, must also be distributed on the infinite interval of values that the time-trend coefficient may take.

⁴ The bilateral nominal exchange rates available have the U.S. dollar as the base currency, i.e., foreign currency per U.S. dollar. Non-U.S. dollar based exchange rates are computed as cross-rates. This assumes cross-rate equality except for transaction costs. This is probably a valid assumption for the G-7 countries. Alternatively, as long as the measurement error is a stationary process, our tests for unit-root will not be affected.

⁵ Other researchers, Papell (1997), and Culver and Papell (1999), for example, argued that the presence of a linear deterministic time trend is inconsistent with long-run PPP, however.

⁶ We have also tried other methods of lag-length determination, e.g., Akaike’s Information Criteria. While the lag lengths determined tend to be much shorter, they do not produce significantly different results than those reported in Table 1, however.

⁷ We use a non-Bayesian method, since the number of hypotheses increases rather rapidly if the same Bayesian approach is used to determine the lag lengths. This reduces the number of alternative hypotheses we need to consider and allow us to allocate the entire prior probability to the three hypotheses in question which are what we are most interested in. To check the robustness of our Bayesian results, we have also used a uniform four lags for all the real exchange rate series. The results are not very different from those reported in the paper.

⁸ Our later simulation results also suggest that the results in Table 2 are unlikely the result of the choice of the lag-length.

⁹ A recent paper by Murray and Papell (2002) has shown that the half-life estimates are extremely unreliable, however.

Table 1
Univariate ADF Test Results

Real Exchange Rate	l	$t(\phi)$	
		without trend	with trend
Canada/U.S.	12	-1.373	-2.025
U.K./U.S.	11	-2.594**	-2.725
Germany/U.S.	10	-2.331	-2.472
Italy/U.S.	10	-2.507	-2.545
Japan/U.S.	11	-1.932	-2.731
France/U.S.	3	-2.249	-2.244
U.K./Canada	11	-2.014	-2.719
Germany/Canada	10	-2.642**	-2.660
Italy/Canada	1 ⁺	-2.057	-3.298**
Japan/Canada	11	-1.779	-3.416**
France/Canada	5	-2.329	-2.663
Germany/U.K.	12	-2.196	-2.793
Italy/U.K.	11	-2.356	-2.351
Japan/U.K.	10	-2.035	-2.595
France/U.K.	1	-2.273	-2.630
Italy/Germany	2	-1.654	-1.917
Japan/Germany	8	-1.668	-3.598*
France/Germany	9	-3.066**	-4.499*
Japan/Italy	3	-2.433	-3.101
France/Italy	3	-1.966	-1.991
France/Japan	8	-1.983	-3.791*

Note: *,** Denote the rejection of the null hypothesis at the 5% and the 10% significance levels, respectively.

⁺ The lag length is 10 for the Canadian dollar/lira real exchange rate when estimated without a linear time trend.

Table 2
Posterior Probabilities

Real Exchange Rate	k	Trend Stationary	Stationary	Unit Root
Canada/U.S.	11	0.83	0.14	0.03
U.K./U.S.	2	0.84	0.06	0.10
Germany/U.S.	2	0.91	0.06	0.03
Italy/U.S.	2	0.91	0.05	0.04
Japan/U.S.	2	0.85	0.11	0.04
France/U.S.	2	0.93	0.04	0.02
U.K./Canada	2	0.74	0.20	0.06
Germany/Canada	2	0.92	0.04	0.03
Italy/Canada	2	0.53	0.29	0.17
Japan/Canada	2	0.70	0.25	0.05
France/Canada	2	0.89	0.07	0.04
Germany/U.K.	2	0.80	0.15	0.05
Italy/U.K.	2	0.87	0.04	0.09
Japan/U.K.	3	0.89	0.07	0.04
France/U.K.	2	0.83	0.09	0.08
Italy/Germany	2	0.83	0.14	0.04
Japan/Germany	2	0.21	0.65	0.15
France/Germany	2	0.39	0.36	0.25
Japan/Italy	2	0.63	0.18	0.18
France/Italy	2	0.91	0.07	0.02
France/Japan	3	0.65	0.24	0.11

Note: The posterior probabilities may not sum to one because of rounding.

Table 3
The Power of the “Objective” Bayesian Unit-Root Test

(a) Data Generating Model: $y_t = 0.059317 + \rho y_{t-1} - 0.0000201t + u_t$,

Proportion of times that a hypothesis receives the largest posterior probability:

ρ	Trend Stationary	Stationary	Unit Root
0.99	3.64	0.02	96.34
0.98	68.18	16.36	15.46
0.97	73.00	14.80	12.20
0.96	66.76	10.30	22.94
0.95	55.44	6.16	38.40

(b) Data Generating Model: $y_t = 0.059317 + \rho y_{t-1} + u_t$,

Proportion of times that a hypothesis receives the largest posterior probability:

ρ	Trend Stationary	Stationary	Unit Root
1.00	67.22	32.32	0.46
0.99	5.08	0.10	94.82
0.98	75.26	3.84	20.90
0.97	79.76	2.16	18.08
0.96	69.60	1.36	29.04
0.95	56.34	0.52	43.13

Table 4
 Data Generating Model: $q_t = \delta + \rho y_{t-1} + \lambda t + u_t$, $u_t = (0, 1)$

$T = 60$			
Proportion of times that a hypothesis receives the largest posterior probability:			
ρ	Trend Stationary	Stationary	Unit Root
0.99	51.00	33.40	15.60
0.98	49.52	48.68	1.80
0.97	48.06	50.90	1.04
0.96	45.40	53.52	1.08
0.95	43.98	55.10	0.92

$T = 240$			
Proportion of times that a hypothesis receives the largest posterior probability:			
ρ	Trend Stationary	Stationary	Unit Root
0.99	0.10	0.00	99.90
0.98	49.02	50.72	0.26
0.97	38.74	61.24	0.02
0.96	26.26	73.74	0.00
0.95	16.10	83.90	0.00

$T = 720$			
Proportion of times that a hypothesis receives the largest posterior probability:			
ρ	Trend Stationary	Stationary	Unit Root
0.99	0.00	0.00	100.00
0.98	13.64	86.36	0.00
0.97	2.04	97.96	0.00
0.96	0.06	99.94	0.00
0.95	0.00	100.00	0.00

Table 5
 Data Generating Model: $q_t = \delta + \rho y_{t-1} + u_t$, $u_t \sim (0, 1)$

T = 60			
Proportion of times that a hypothesis receives the largest posterior probability:			
ρ	Trend Stationary	Stationary	Unit Root
1.00	51.00	46.66	2.34
0.99	59.94	32.44	7.62
0.98	69.16	20.52	10.32
0.97	70.88	17.22	11.90
0.96	70.62	15.30	14.08
0.95	71.42	13.66	14.92

T = 240			
Proportion of times that a hypothesis receives the largest posterior probability:			
ρ	Trend Stationary	Stationary	Unit Root
1.00	64.80	34.58	0.62
0.99	75.02	11.24	13.74
0.98	82.92	7.22	9.86
0.97	81.78	4.30	13.92
0.96	76.80	3.58	19.62
0.95	68.22	2.08	29.70

T = 720			
Proportion of times that a hypothesis receives the largest posterior probability:			
ρ	Trend Stationary	Stationary	Unit Root
1.00	74.92	24.86	0.22
0.99	63.38	1.16	35.46
0.98	71.60	1.34	27.02
0.97	43.58	4.20	52.22
0.96	15.00	0.00	85.00
0.95	2.60	0.00	97.40

Table 6
 Data Generating Model: $q_t = \delta + \rho y_{t-1} + u_t$, $u_t \sim (0, 1)$

T = 60			
Proportion of times that a hypothesis receives the largest posterior probability:			
ρ	Trend Stationary	Stationary	Unit Root
0.90	66.38	6.94	26.68
0.85	55.70	3.24	41.06
0.80	40.68	2.22	57.10
0.75	24.28	2.20	73.52
0.70	12.92	0.04	87.04

T = 240			
Proportion of times that a hypothesis receives the largest posterior probability:			
ρ	Trend Stationary	Stationary	Unit Root
0.90	14.18	0.04	85.78
0.85	0.46	0.00	99.54
0.80	0.00	0.00	100.00
0.75	0.00	0.00	100.00
0.70	0.00	0.00	100.00

T = 720			
Proportion of times that a hypothesis receives the largest posterior probability:			
ρ	Trend Stationary	Stationary	Unit Root
0.90	0.00	0.00	100.00
0.85	0.00	0.00	100.00
0.80	0.00	0.00	100.00
0.75	0.00	0.00	100.00
0.70	0.00	0.00	100.00