PPP Strikes Out:
The effect of common factor shocks on the real exchange rate

Nelson C. Mark *
The University of Notre Dame and NBER

Donggyu Sul
The University of Auckland

July 19, 2008

Abstract

Recent panel-data studies on PPP have measured the half-life of convergence to range from less than a year to about three years. These are estimates of convergence speeds in response only to idiosyncratic shocks to a country's real exchange rate. A large proportion of real exchange rate variation–that due to variation in a persistent common factor–has been neglected by these studies. Once the influence of the common factor is taken into account, the evidence weighs heavily against PPP. This paper reports evidence that the common factor accounts for 79 percent of the variation in the real exchange rate and that it is unit root nonstationary. Among the 19 relative prices that make up the overall real exchange rate, only the prices of 7 goods show evidence of long-run convergence to the law of one price.

Keywords: Purchasing power parity, Common factor, Law-of-one price, Convergence, Cross sectional dependence.

JEL Classification: F3

*Mark is grateful for financial support from the National Science Foundation. Some of this work was done while Mark was visiting at the City University of Hong Kong.

Introduction

In the latest salvo fired in the back and forth on PPP convergence, Imbs et al. (2005) (hereafter IMRR) argued that the unreasonably long half-life of convergence to PPP can be explained by bias induced by aggregating sectoral or individual goods prices that differed in their own convergence rates towards the law-of-one price. After adjusting for aggregation bias, they obtain an estimate of the half-life to convergence of less than one year, which many economists view as reasonable. We argue in this paper that the IMRR half-life estimate is seriously understated, not because of their accounting for aggregation bias.[1] Our point is that their half-life calculation does not take into account the dynamics of a quantitatively important and highly persistent component of the real exchange rate. We note at the outset that our critique is not leveled specifically at IMRR. Every other panel-data study on PPP that explicitly addresses the cross-sectional correlation of the observations (including earlier work of out own) that we are aware of produces a similar understatement of the half-life.[2] The different panel-data PPP studies in the literature each consider slightly different empirical specifications so to focus on the core issues, and because we use their data, we take the IMRR paper as a reference point.

The issue that we address is this. Suppose that the researcher specifies the log of the i-th real exchange rate in a panel data set follows an autoregressive process. The innovation will be correlated across i. This is obvious if for no other reason than all of the deviations from the law-of-one price in the panel are defined relative to the same base (numeraire) country. Any shocks that originate in the base country will affect all of the deviations. If this cross-sectional correlation is not accounted for, the estimates of the autoregressive coefficients will be biased.

At this point, the precise way that one controls for the cross-sectional correlation is not critical so let us suppose that the error term is given by a three-component model that contains a country fixed effect, a common factor, and a purely idiosyncratic factor. Let us also suppose that one controls for the common factor by including a set of time dummies. The researcher then obtains estimates of the autoregressive coefficients which (s)he uses to compute the half-life either directly or through impulse response analysis. The result is the half-life to convergence following a shock only to the idiosyncratic component. The half-life obtained in this way would be a satisfactory measure if the cross-sectionally correlated factor were serially uncorrelated, as is typically assumed in the error-components representation, or possibly if the variation in the real exchange rate due to

[1] Chen and Engel (2005) have argued against the quantitative importance of IMRR's correction for aggregation bias in computing the half-life.

[2] Choi, Mark, and Sul (2006). Other papers that are subject to our critique include Fleissig and Strauss (2000), Frankel and Rose (1996), Koedijk, Schotman, and Van Dijk (1998), O'Connell (1996), Murray and Papell (2005), Papell (2006) and Crucini and Shintani (2006). Cecchetti, Mark and Sonora (2002), on the other hand, computed the half life of real exchange rate deviations from the cross-sectional avearage.

variation in the common factor is small. Unfortunately, neither is the case. Using IMRR's data, we find that the common factor is highly serially correlated–so much so that it may plausibly be unit-root nonstationary, and that the variation in the common factor accounts for about 79 percent of the variation in the deviation from the law-of-one price. Calculating the half-life of adjustment while ignoring the effects of the common factor is not exactly the same as, but in some ways analogous to using only the autoregressive coefficients from an ARMA model to infer the adjustment dynamics of a process while ignoring the moving average part.

Given the importance of the common factor, we revisit the question of whether PPP holds at all in the long run. Panel unit-root tests in this context hold no advantage over time-series unit-root tests if there is a single common factor that drives the nonstationarity of law-of-one price deviations because ultimately what one is testing for is whether the common factor (a univariate time-series process) has a unit root. Instead, we cast the issue in a different context. We employ a concept of long-run convergence to the law-of-one price in which the cross-sectional dispersion of law-of-one price deviations decreases over time and apply a statistical test for whether such convergence is present in the data. An underlying mechanism that would predict long-run (asymptotic) convergence to the law-of-one price is a process where globalization and dismantling of trade barriers increases commodity market integration over time.

Once the effects of the common factor are taken into consideration, the implications for PPP convergence are not good. For 12 of the 19 goods in the data set, we find no evidence for long-run convergence to the law-of-one price. Since these prices are aggregated into the calculation of a country's overall price level, our final conclusion from the IMRR data must be that the true half-life is infinite and that PPP fails in the long run. We caution the reader that these conclusions apply only to the IMRR data that we have analyzed and that favorable evidence of PPP may yet be found in other data sets perhaps with more refined and disaggregated categories of goods.

The empirical evidence that we report on the importance of the common factor in driving real exchange rate dynamics, its persistence, and understatement of the half-life to convergence one obtains by ignoring the common factor is relatively straightforward. In terms of the proposed statistical tests of long-run convergence, we believe that our convergence concept is a useful and appropriate way to characterize the law-of-one price in the long run. However, we do recognize that this may not be universally shared amongst readers. Whether one is sympathetic to equating the convergence concept with the law-of-one price or not, this analysis does establish that for the majority of the goods in the data set there has been a tendency for the cross-sectional dispersion of their prices measured in U.S. dollars to increase over time.

The remainder of this short paper is structured as follows. The next section discusses how existing panel-data studies have neglected the influence of the common factor. In Section 2, we obtain estimates of the common factor in the real exchange rate. We examine its time-series behavior and measure its importance in explaining real exchange rate dynamics. Section 3 discusses a test of growth convergence developed by Phillips and Sul (2007b) and how we adapt this methodology to test the hypothesis that the law-of-one price holds in the long run. Section 4 concludes. Technical details of many of the arguments made in the text are consigned to the appendix.

1 Half-life understatement from panel data studies

We analyze the same data used by IMRR, which we obtained from their website. The data consists of prices from sectors $i = 1, \ldots, I$, which we'll refer to as goods, and countries indexed $c = 1, \ldots, C$. We employ data on prices for $I = 19$ goods and $C = 10$ countries plus the U.S., which serves as the numeraire country and is indexed as $c = 0$. The log relative price of good i and country c is $q_{ict} = \ln(S_{ct}P_{ict}/P_{0ct})$, where S_{ct} is the nominal U.S. dollar price of a unit of country c's money and P_{ict} is the country c currency price of good i in country c. The IMRR data set is described more fully in the appendix.

To fix ideas, consider IMRR's specification of the dynamics governing country $c's$ relative price of good i,

$$q_{ict} = \gamma_{ic} + \sum_{j=1}^{p} \rho_{icj} q_{ict-j} + \sum_{h=0}^{H} \lambda_{ich} \bar{q}_{t-h} + \epsilon_{ict}. \tag{1}$$

γ_{ic} is a good-and-country fixed effect, ϵ_{ict} is the purely idiosyncratic shock, and the cross-sectional average of the panel $\bar{q}_t = \frac{1}{CI}\sum_{i=1}^{I}\sum_{c=1}^{C} q_{ict}$, is their measure of a factor that is common to all of the relative prices in the panel, which is responsible for the cross-sectional correlation. By including lags of $\bar{q}_t$, the common factor is acknowledged to be serially correlated. Shocks to the common factor are allowed to have a differential impact on the relative price common factor by allowing λ_{ich} to vary across i and c. For each relative price, they subtract current and lagged cross sectional averages of the relative prices to soak up variation caused by the unobserved common factor.

The concept of PPP convergence, however, applies to the real exchange rate between the U.S. and country c,

$$Q_{ct} = \sum_{i=1}^{I} \omega_{ic} q_{ict}, \quad \text{where} \quad \sum_{i=1}^{I} \omega_{ic} = 1,$$

and not the adjustment of the relative price of good i. The link between the dynamics of relative price ic in (1) and the real exchange rate is made by employing Peseran's mean group common correlated elements (MG-CCE) estimator of the autoregressive coefficients for the real exchange rate Q_{ct}. It estimates the autoregressive coefficients for Q_{ct}, assuming identical coefficients across countries c, and controls for cross-sectional correlation of the error terms. The MG-CCE estimator for the k-th autoregressive coefficient ρ_k is

$$\hat{\rho}_k = \frac{1}{IC} \sum_{i=1}^{I} \sum_{c=1}^{C} \hat{\rho}_{ick}$$

The researcher would typically compute the half-life of convergence to PPP using these estimates of the autoregressive coefficients in Q_{ct}. In the case of an AR(1), then the half life estimate would be $-\ln(2)/\ln(\hat{\rho}_1)$. For the general AR(K) case, the half life can be estimated by Monte Carlo simulation. IMRR employ the simulation methodology and find that the half life for PPP convergence is less than one year.

The calculations just described, trace the impulse response of the real exchange rate following shocks to the idiosyncratic components of the relative prices. However, if the common factor accounts for a large proportion of the variation in the q_{ict}, then an important component of relative price dynamics has been ignored. To obtain a full picture of the half-life, one must also analyze the time for half of a shock to both the idiosyncratic effect and the common factor to dissipate.

2 How does the common factor behave and how important is it?

In this section, we assess the importance of common factor variations in driving the dynamics of the price of good i across countries c. We are interested in seeing how persistent is the common factor, and how important it is in driving the dynamics of deviations from the law-of-one price and therefore the real exchange rate.

Let us cast the dynamics of the relative prices in a slightly more general form. We'll assume that q_{ict} is given by the two-component model,[3]

$$q_{ict} = e_{ict} + \lambda_{ic} F_t \tag{2}$$

[3] This is a representation that has been used by Stock and Watson (2002), Bai and Ng (2002, 2004), and Bai (2003).

where e_{ict} is orthogonal to F_t and each component is a p-th order autoregression,

$$e_{ict} = \gamma_{ic} + \sum_{j=1}^{p} \rho_{icj} e_{ict-j} + \epsilon_{ict} \tag{3}$$

$$F_t = \sum_{j=1}^{p} \phi_j F_{t-j} + v_t \tag{4}$$

$\epsilon_{ict} \overset{iid}{\sim} \left(0, \sigma^2_{\epsilon ic}\right)$, $v_t \overset{iid}{\sim} \left(0, \sigma^2_v\right)$. For convenience, we assume that both the idiosyncractic and common components are at most AR(p) processes.

e_{ict} captures purely idiosyncratic dynamics of the relative price. An ϵ_{ict} shock propagates through the autocorrelation structure of (3). While IMRR and other panel data PPP research has analyzed the implied dynamics of the real exchange rate in response to idiosyncratic shocks, they have neglected the implied dynamics following a shock to the potentially serially correlated common factor F_t. As we mentioned above, IMRR is not alone in this omission.

In work such as Murray and Papell (2005) or Choi, Mark, and Sul (2006), an explicit error components model for the error term is not written down but the authors do recognize that the errors are correlated across individual countries and apply a generalized least squares (GLS) correction in estimation. However, the estimation assumes that the errors are serially uncorrelated which a check on the residuals would in all likelihood reveal that they are not (serially uncorrelated). The half-lives estimated by these papers are understated because they employ only the estimated autocorrelation coefficients while ignoring the serial correlation in the residuals. An analogy is if one were to compute the correlegram of an autoregressive moving average (ARMA) model using only the autoregressive coefficients and ignores the moving average component.

The two-component model (2)-(4) also has the observationally equivalent representation

$$q_{ict} = \gamma_{ic} + \sum_{j=1}^{p} \rho_{icj} q_{ict-j} + \lambda_{ic} F_t - \sum_{j=1}^{p} \lambda_{ic} \rho_{icj} F_{t-j} + \epsilon_{ict}, \tag{5}$$

which is equivalent to the IMRR specification in (2) if the common factor is measured by the cross-sectional mean $F_t = \bar{q}_t$.

We present four alternative measurements of F_t. The first measure follows IMRR in using the cross-sectional average of the relative prices in the panel. This measure is increasingly accurate as the cross-sectional dimension of the panel increases. Taking the grand average of the q_{ict} in the context of the model laid out in (2) and (3) gives

$$\bar{q}_t = \bar{\lambda}_1 F_t + O_p\left(\frac{1}{\sqrt{IC}}\right) \tag{6}$$

where $\bar{\lambda}_1 = (1/(IC))\sum_i \sum_c \lambda_{ic1}$. Our second measure is the first principle component estimate of the factor which we will denote by $\hat{F}_t$. The third measure is the cross-sectional average of the log nominal exchange rate, $\bar{s}_t = (1/C)\sum_{c=1}^{C} \ln(S_{ct})$ and the fourth is the cross-sectional average across countries of their real exchange rates, $\bar{q}_t' = \left(1/C\sum_{c=1}^{C} \ln(S_{ct}P_{ct}/P_{0t})\right)$. These alternative measures of F_t are shown in Figure 1.

Only three series are distinguishable in the figure. The principle component estimator $\hat{F}_t$ is nearly identical to the cross-sectional average used by IMRR and one of the lines lies on top of the other. It is obvious that all of these measures are highly correlated with each other and all are highly serially correlated. They look like plots of unit-root nonstationary processes.

To delve further into the persistence of these series, we fit by least-squares, a twelfth order autoregression to the IMRR measure,

$$\bar{q}_t = a + \sum_{k=1}^{12} \rho_k \bar{q}_{t-k} + \epsilon_t. \tag{7}$$

The sum of the point estimates is $\sum_{k=1}^{12} \hat{\rho}_k = 0.9814$ which implies a half–life of the $\bar{q}_t$ dynamics of 37 months. But as is well known, the least-squares estimator of the autoregressive coefficients are biased downwards in small samples so 37 months probably understates the truth. When we do a Jackknife bias correction, the Jackknifed sum of the autoregressive coefficients is 1.028 which leads to the conclusion that $\bar{q}_t$ is unit-root nonstationary.[4] The implication then is that individual relative price inherit this unit-root nonstationarity as must the real exchange rates.

So it appears that the common factor is persistent and may be a unit-root process. The next question is how big is the unit root in the real exchange rate? To assess its relative importance in explaining the variation of the real exchange rate over time, let us write the two-component model of eq.(2) in terms of deviations from their time-series averages. Subtracting the time-series average eliminates the fixed effects γ_{ic} which gives,

$$\left(q_{ict} - \frac{1}{T}\sum_{t=1}^{T} q_{ict}\right) = \left(e_{ict} - \frac{1}{T}\sum_{t=1}^{T} e_{ict}\right) + \lambda_{ic}\left(F_t - \frac{1}{T}\sum_{t=1}^{T} F_t\right). \tag{8}$$

Since F_t is orthogonal to e_{it}, the short-run variance of the relative price of the $i-$th good q_i can

[4] Details about the Jackknife used can be found in the appendix.

simply be decomposed as

$$\mathrm{Var}\left(q_i\right)=\mathrm{Var}\left(e_i\right)+\mathrm{Var}\left(\lambda_i F_t\right).$$

Table 1 reports this variance decomposition for the log deviation from the law-of-one price (price of $i-$th good in c relative to the U.S.). These calculations use the cross-sectional average $\bar{q}_t$ to measure F_t and regression estimates of the factor loading coefficients λ_{ic}.[5]

For goods such as fruits, communications, tobacco, the fraction of relative price variation due to common factor shocks ranges from 0.23 to 0.42, which is relatively small. For the other 15 goods, variation in the common component accounts for upwards of 75 percent of the relative price variation. The average proportion across all 18 goods is 79 percent. Extant measures of the half-life of PPP deviations drawn from panel data studies have evidently ignored the most important source of real exchange rate dynamics–and a source that appears to be unit-root nonstationary.

3 A log-t test of law-of-one price convergence

The evidence from the last section suggests a re-examination of PPP in panel data is in order. Given the factor structure of this data set, panel unit root tests will have no advantages over standard time-series unit root tests if the F_t component has a unit root because the nonstationary behavior in all of the relative prices are being driven by the same univariate time series. If the e_{ict} are all stationary, the data form a cointegrated panel and a panel unit root test can only succeed in informing us about a unit root to the extent that it either accepts or rejects a unit root in the single time series F_t.

We will therefore not revisit this issue with conventional unit root tests. Instead, we'll undertake an examination of relative convergence to the law-of-one price using the log-t test suggested by Phillips and Sul (2007b). The concept of relative convergence in the present context refers to the idea that the cross-sectional dispersion of the U.S. dollar price of a good should show a tendency to decrease over time. It is called the log-t test because we regress a measure of the cross-sectional dispersion of prices on $\ln(t)$. The technical details of the test are described in the appendix. Here, we will simply give an intuitive account of the procedure.

Test of convergence to law-of-one price. We will be working with the U.S. dollar prices of the different goods in every country of the sample so in this section, we have $C = 11$ countries which

[5] See the appendix. The sample counterparts are $\mathrm{Var}(q_i)=\mathrm{E}\left(\frac{1}{CT}\sum_{c=1}^{C}\sum_{t=1}^{T}\left(q_{ict}-\frac{1}{T}\sum_{t=1}^{T}q_{ict}\right)^2\right)$, $\mathrm{Var}(e_i)$ $=\mathrm{E}\left(\frac{1}{CT}\sum_{c=1}^{C}\sum_{t=1}^{T}\left(e_{ict}-\frac{1}{T}\sum_{t=1}^{T}e_{ict}\right)^2\right)$, and $\mathrm{Var}(\lambda_i F)=\mathrm{E}\left(\frac{1}{CT}\sum_{c=1}^{C}\sum_{t=1}^{T}\lambda_{ic}^2\left(F_t-\frac{1}{T}\sum_{t=1}^{T}F_t\right)^2\right)$

are indexed by $c = 1, ..., 11$. We will set the U.S. as country 1. Let p_{ict} be the U.S. dollar price of good i in country c at time t,

$$p_{ict} = S_{ct} P_{ict},$$

where for the U.S., $S_{1t} = 1$, and assume that the price of i in country c is generated according to the two-component model

$$\ln p_{ict} = \delta_{ict} F_{it} + e_{ict}, \tag{9}$$

where F_{it} is a common component for good i and e_{ict} is a stationary component. We note that there is a subtle difference between the model in (9), which applies to absolute prices and eq.(2) which applies to the relative price $q_{ict} = \ln(p_{ict}) - \ln(p_{i1t})$. The data will exhibit relative convergence to the law-of-one price if $\delta_{ict} \to \delta_i$. This says that the in the long run, for all c and r, $\ln(p_{ict})$ and $\ln(p_{irt})$ have a common trend which can be stochastic or deterministic. If the common trend $\delta_i F_{it}$ is a stochastic trend, then $\ln(p_{ict})$ and $\ln(p_{irt})$ is cointegrated in the long run with cointegrating vector $(1, -1)$. This convergence concept does not require $\delta_{ict} = \delta_{irt}$ in any finite sample, but instead looks for evidence that

$$\delta_{ict} \to \delta_i \quad \text{as } t \to \infty,$$

for countries $c = 1, ...C$. To set up such a test of asymptotic convergence to the law-of-one price, we let the time-varying factor loadings have the parametric form

$$\delta_{ict} = \delta_i + \psi_{ict}, \tag{10}$$

where

$$\psi_{ict} \overset{iid}{\sim} \left(0, \frac{(t)^{-2\alpha_i}}{\ln(t)} \sigma^2_{\psi ic}\right). \tag{11}$$

The data will exhibit asymptotic convergence to the law-of-one price if $\alpha_i \geq 0$. This is because $\text{Var}(\psi_{ict}) \to 0$ as $t \to \infty$ for any positive value of α_i which in turn implies $\delta_{ict} \to \delta_i$ as $t \to \infty$.[6] If $\alpha_i < 0$, it will be the case that $\text{Var}(\psi_{ict}) \to \infty$ as $t \to \infty$. If the variance diverges then there can be no relative convergence to the law-of-one price. The parameter α_i can be viewed as convergence rate to the 'asymptotic' law-of-one price since it regulates the speed at which the

[6] Absolute convergence is said to occur if $\ln(\pi_{ict}) - \ln(\pi_{irt}) \to 0$, and relative convergence is said to occur if $\ln(\pi_{ict}) / \ln(\pi_{irt}) \to 1$. What's the difference between the two concepts? Suppose the e_{ict} component is white noise which we'll ignore and $\delta_{ict} = 1$ and $\delta_{irt} = \left(1 - 1/\sqrt{t}\right)$. The prices exhibit relative convergence since $\ln(\pi_{ict}) / \ln(\pi_{irt}) = \left(1/\left(1 - 1/\sqrt{t}\right)\right) \to 1$. If the common factor is a time trend, $F_t = t$,then there is no absolute convergence because $\ln(\pi_{ict}) - \ln(\pi_{irt}) = F_t/\sqrt{t} = \sqrt{t}$ which diverges. Our primary interest is in testing for relative convergence.

variance $\mathrm{Var}(\psi_{ict})$ converges.

To connect α_i to the data, let

$$v_{ict} = \frac{p_{ict}}{\frac{1}{C}\sum_{k=1}^{C} p_{ikt}},$$

and

$$V_{it} = \frac{1}{C}\sum_{c=1}^{C}(v_{ict} - 1)^2 .$$

v_{ict} has the flavor of the price of good i in country c normalized by (controlling for) the common factor F_{it} and V_{it} has the dimensionality of the cross-sectional variance of p_{ict}. Under the null hypothesis of convergence to the law-of-one price, the cross-sectional variance V_{it} will be decreasing as t increases. Let us normalize this quasi-variance to be 1 when $t = 1$. In the appendix, it is shown that the relation between V_{it} and ψ_{ict} is

$$\ln(V_{i1}/V_{it}) - 2(\ln(\ln(t))) = a_i + 2\alpha_i \ln(t) + u_{it}, \tag{12}$$

Thus to test the composite null hypothesis that there is relative convergence to the law-of-one price for i,

$$\mathcal{H}_0 : \delta_{ict} \to \delta_i \quad \text{for all } c \in [1, 11],$$

against the alternative that there is no relative convergence in at least one instance,

$$\mathcal{H}_A : \delta_{ict} \nrightarrow \delta_i \quad \text{for some } c.$$

one regresses $\ln(V_{i1}/V_{it}) - 2(\ln(\ln(t)))$ on $\ln(t)$. One then rejects the null hypothesis of long-run convergence to the law-of-one price if the slope coefficient is significantly negative.

Before we run the test, let us observe the behavior of the cross-sectional variances V_{it}. Phillips and Sul (2007b) show that elimination of the cyclical components of the data improves the finite sample power and size of the test. We follow that recommendation here and use the Hodrick–Prescott (1997) trend of p_{ict} to construct V_{it}. Another data-related issue that we mention is that the price data are indices rather than price levels. If the last year of the sample is chosen as the base year, then the p_{ict} will all appear to converge to a single point. Instead, we use the first year as the base year so that the prices will diverge from this point. Figure 2 gives us a view of the V_{it} constructed in this way. They seem to stabilize around 1985 or 1986.

We do not want the test to be affected by initial conditions created by the base year so we discard the first half of the sample. So from the original 175 time-series observations, we work with the 87 observations from October 1988 to December 1995. From this subsample, we discard an additional fraction r of the sample.[7]

We run the test using a range of $r \in (0.2, 0.32)$. For $r = 0.2$, the sample begins at 1990.02 whereas for $r = 0.32$ the sample begins at 1990.12. Table 2 shows the t-statistics for the slope coefficient from estimating eq.(12). Let us examine the first column of the table. Here, the sample begins in 1990.02. For 12 of the 19 goods (the first 12 listed in the table) the evidence weighs against law-of-one price convergence. Interestingly, these 12 commodity categories are composed predominately of traded goods. The strongest evidence against price convergence shows up for clothing where the t-statistic is -364.4. Nontraded goods prices, on the other hand, such as for public transportation and hotels show evidence of law-of-one price convergence. The prices for leisure goods, books, alcohol, dairy and fuel also show statistically significant evidence of long-run law–of-one price convergence. As we view the columns of the table moving towards the right, it is seen that the general conclusion about which goods prices show evidence of long-run convergence and which do not is robust to the variation of the sample starting point. Roughly speaking, relative convergence of to the law-of-one price is found for only one-third of the goods in the data set. Due to the nonconvergent behavior of the remaining two-thirds, it probably cannot be said that PPP holds in the long run. We will return to the issue of overall PPP convergence below.

We can observe the nonconvergent behavior in the log prices of one of the divergent commodities, say for clothing (the v_{ict} for i =clothing) in Figure 3. Recall that the p_{ict} are HP trend values of the prices. This gives the v_{ict} a very smooth appearance. As can be seen, the cross-sectional average of the dollar price of clothing widens over time. To contrast the behavior of these prices with those of fuel (which converges), we plot the v_{ict} for fuel in Figure 4. As can be seen, the cross-sectional dispersion of these prices declines over time. We split up the cross-sectional variances (the V_{it}) among the convergent and nonconvergent goods and plot these in Figures 5 and 6.

Test for convergence to PPP. The test results in Table 2 say that there is no evidence for the prices of 2/3 of the goods in the IMRR data set to converge to the law-of-one price. This then suggests that PPP doesn't hold in the long run. We can apply the convergence test to the overall price levels across countries measured in U.S. dollars. Here, we are looking for evidence that the cross-sectional dispersion in national price levels is decreasing over time. In regressing the PPP counterpart of $\ln\left(\frac{V_1}{V_t}\right) - 2\ln(\ln t)$ on $\ln(t)$, we obtain a point estimate of the slope of -1.42 with a t-ratio of -81.4 so the data show no evidence of asymptotic convergence to PPP.

[7] Phillips and Sul (2007b) also suggest setting r between 0.2 and 0.3.

4 Conclusion

Estimation of real exchange rate dynamics in panel-data studies on PPP must deal with cross-sectionally correlated error. In doing so, measurements of the speed of convergence have been distracted by techniques to control for cross-sectionally correlated error terms and has ignored the time-dependent nature in a common factor component in the dynamics of the real exchange rate.

We argue that the common factor, which gives rise to the cross-sectional correlation, is a quantitatively large component of law-of-one price deviations and that it is a persistent and possibly unit-root nonstationary process. We find that ignoring the dynamics of the common factor, researchers have understated the half-life to PPP convergence and once the common factor dynamics are taken into account in evaluation of real exchange rate adjustment, the evidence does not support long-run PPP. An analysis of the evolution of the cross-sectional dispersion of law-of-one price deviations shows no tendency for this dispersion to decline over time for the majority of the goods in the IMRR data set. The purchasing power parity puzzle as discussed by Rogoff (1996) still stands and may even be more of a puzzle than previously thought.

Appendix

The IMRR data

The IMRR data set has prices for 19 sectors: fuel, alcohol, clothing, footwear, bread, books, public transportation, leisure, vehicles, communication, fruits, meat, dairy, tobacco, furniture, domestic appliances, rent, sound, and hotels. Country coverage includes the U.S. (US), Belgium (BE), Germany (DE), Denmark (DK), Spain (ES), Italy (IT), France (FR), Greece (GR), the Netherlands (NL), Portugal (PT), and the U.K. (UK).

The observations begin in 1960, but due to missing values, IMRR report that only 180 observations are usable. To form a balanced panel, we were able to use 175 observations.

We exclude Finland because their data becomes available only in 1985:01. Except for furniture and domestic appliance, all data are available from 1981:01 to 1995:12. Furniture and domestic appliance data start from 1981:01 but ends to 1994:09 for 11 countries. IMRR performed their analysis on an unbalanced panel. We work with a balanced panel. To make balanced panels, we are able to use the data from 1981:01 to 1994:09. Rent and sound data are not available for Portugal, and tobacco data are not available for Italy.

Proofs of statements made in the text

We adopt the assumptions made in Bai and Ng (2002, 2004). The key assumptions are that the idiosyncratic errors e_{ict} are cross sectionally independent (or weakly dependent) and the common factors are independent of the idiosyncratic errors. We refer the reader to Assumptions A through E in Bai and Ng (2004) in the case of a nonstationary common factor and Bai and Ng (2002) for a stationary common factor.

Derivation of eq. (6). We are assuming that the relative price of good i q_{ict} has one common factor. Taking cross sectional average of eq. (3) gives,

$$\frac{1}{IC}\sum_{i=1}^{I}\sum_{c=1}^{C} q_{ict} = \frac{1}{IC}\sum_{i=1}^{I}\sum_{c=1}^{C} e_{ict} + \frac{1}{IC}\sum_{i=1}^{I}\sum_{c=1}^{C} \lambda_{ic} F_t$$

If e_{ict} is cross sectionally independent, then the law of large numbers implies that the first term on the right hand side is bounded by $O_p\left(N^{-1/2}\right)$ where $N = I \times C$. This remains true if e_{ict} is weakly correlated with e_{jct} in the Bai and Ng (2002) sense. The result is,

$$\bar{q}_t = \frac{1}{IC}\sum_{i=1}^{I}\sum_{c=1}^{C} q_{ict} = \bar{\lambda}_1 F_t + O_p\left(\frac{1}{\sqrt{IC}}\right),$$

which is eq. (6).

Bias-adjustment with the Jackknife. It is well known that the least squares estimator for autoregressive coefficients, $\hat{\rho}_k$, are biased downward due to estimation of the regression constant. In general, the order of the bias is given by

$$\mathrm{E}\left(\hat{\rho}_k - \rho_k\right) = O\left(T^{-1}\right) = \frac{a}{T} + O\left(T^{-2}\right),$$

where a is the asymptotic expansion coefficient. Many bias correction methods are available but a few of them have asymptotic justification both under stationary and nonstationary environment. Jackknifing is one of the methods. See Quenouille (1956) for detailed explanation. Rewrite eq. (1) as

$$\bar{q}_t = a + \rho\bar{q}_{t-1} + \sum_{k=1}^{11}\delta_k \Delta\bar{q}_{t-k} + \epsilon_t,$$

where $\rho = \sum_k^{12} \rho_k$. The Jackknife estimate is obtained first by splitting the sample into halves Let $\hat{\rho}_1$ be the first OLS estimator from the first subsample and $\hat{\rho}_2$ be the estimator from the second subsample. Then

$$\mathrm{E}\left(\hat{\rho} - \rho\right) = \frac{a}{T} + O\left(T^{-2}\right),\ \mathrm{E}\left(\hat{\rho}_1 - \rho\right) = \frac{a}{T_1} + O\left(T_1^{-2}\right),\ \ \mathrm{E}\left(\hat{\rho}_2 - \rho\right) = \frac{a}{T_2} + O\left(T_2^{-2}\right)$$

where T_1 and T_2 are the sample sizes for the first and second sample. When $T_1 = T_2 = T/2$, the Jackknife estimator is

$$\hat{\rho}_{\mathrm{J}} = 2\hat{\rho} - \frac{\hat{\rho}_1 + \hat{\rho}_2}{2}$$

The Jackknife achieves a first-order reduction in bias. The bias of $\hat{\rho}_{\mathrm{J}}$ is

$$\begin{aligned}
\mathrm{E}\left(\hat{\rho}_{\mathrm{J}} - \rho\right) &= \mathrm{E}\left(2\hat{\rho} - \frac{\hat{\rho}_1 + \hat{\rho}_2}{2} - \rho\right) = \mathrm{E}\left(2\hat{\rho} - 2\rho - \left[\frac{\hat{\rho}_1 + \hat{\rho}_2}{2} - \rho\right]\right) \\
&= \frac{2a}{T} + O\left(T^{-2}\right) - \frac{1}{2}\left(\frac{a}{T_1} + \frac{a}{T_2}\right) + O\left(T^{-2}\right) \\
&= \frac{2a}{T} + O\left(T^{-2}\right) - \frac{1}{2}\left(\frac{2a}{T} + \frac{2a}{T}\right) + O\left(T^{-2}\right) = O\left(T^{-2}\right)
\end{aligned}$$

Variance Decomposition. As we saw from Figure 1, the principle component estimate for F_t is almost identical to the cross sectional average of q_{ict}.Since it won't make any difference which one we use, we employed the cross sectional average, $\bar{q}_t$, as an approximated factor to estimate the factor loadings.

The factor loading coefficients are estimated by running the following regression for each c and i

$$q_{ict} = a_{ic} + \lambda_{ic}\bar{q}_t + e_{ict}, \tag{13}$$

then the sample contemporaneous variance for the idiosyncratic components for each i is calculated by

$$\widehat{Var}\left(e_i\right) = \frac{1}{C}\sum_{c=1}^{C}\frac{1}{T}\sum_{t=1}^{T}\hat{e}_{ict}^2.$$

Let $q_{i_t} = \frac{1}{C}\sum_{c=1}^{C} q_{ict}$. The sample contemporaneous variance for q_{it} is

$$\widehat{Var}\left(q_{it}\right) = \frac{1}{C}\sum_{c=1}^{C}\frac{1}{T}\sum_{t=1}^{T}\left(q_{ict} - \frac{1}{T}\sum_{t=1}^{T} q_{ict}\right)^2$$

From the two measures, the variance for the common factor components are calculated by

$$\widehat{Var}\left(\lambda_i F_t\right) = \widehat{Var}\left(q_i\right) - \widehat{Var}\left(e_i\right)$$

There is no cross product term since $\hat{e}_{ict}$ is orthogonal to $\bar{q}_t$. If $\bar{q}_t$ is unit-root nonstationary and e_{ict} is stationary, the data form a cointegrated panel. That is, q_{ict} is cointegrated with q_{jst} for all $i \neq j$ and $c \neq s$ with cointegrating vector of $\left(\lambda_{ic}, -\lambda_{js}\right)$. Moreover, by assumption in the common factor literature, $\bar{q}_t$ is not correlated with e_{ict}. Hence the LS regression in (13) becomes a cointegrating regression. Naturally the LS estimator $\hat{\lambda}_{ic}$ is superconsistent.

Further note that the sample variance of q_{it} increases as T increases if $\bar{q}_t$ is unit-root nonstationary. Phillips and Sul (2007a) provide the exact variance formula for this case – nonstationary common factor under cross section dependence.

Time Varying Factor Loading Representation. Let p_{ict} be the U.S. dollar price of good i in country c at time t,

$$p_{ict} = S_{ct}P_{ict},$$

where $S_{1t} = 1$, and assume that p_{ict} is generated according to the two-component model

$$\ln(p_{ict}) = \delta_{ic} F_t + e_{ict},$$

Then we can rewrite as

$$\ln(p_{ict}) = \left(\delta_{ic} + \frac{e_{ict}}{F_t}\right) F_t = \delta_{ict} F_t$$

where the time varying factor loading coefficient implies the time varying economic distance between $\ln(p_{ict})$ and the common factor F_t.

Alternatively when $\ln(p_{ict})$ contains two common factors with time varying factor loadings, for example

$$\ln(p_{ict}) = \delta_{1,ict} F_{1t} + \delta_{2,ict} F_{2t} + e_{ict}$$

it can be rewritten as

$$\ln(p_{ict}) = \left(\delta_{1,ict} + \delta_{2,ict}\frac{F_{2t}}{F_{1t}} + \frac{e_{ict}}{F_{1t}}\right) F_{1t} = \delta_{ict} F_t$$

where F_{1t} becomes a dominant factor, F_t. Hence regardless of parametric structures of $\ln(p_{ict})$, we can rewrite as a single time varying common factor model.

In practice, however, it useful and realistic to include a small transitory component such as a cyclical or seasonal component e_{ict} where

$$\ln(p_{ict}) = \delta_{ict} F_t + e_{ict},$$

Phillips and Sul (2007) show that the asymptotic properties of their test are invariant to the filtering method.

Convergence Test. Let us represent the time-varying factor loading in eq. (10) as

$$\begin{aligned} \delta_{ict} &= \delta_i + \sigma_{ict}\xi_{ict} && (14) \\ &= \delta_i + \frac{\sigma_{ic}\xi_{ict}}{\ln(t)\, t^{\alpha}} = \delta_i + \frac{\psi_{ict}}{\ln(t)\, t^{\alpha}}, && (15) \end{aligned}$$

for some $\sigma_{ic} > 0$, $t \geq 1$. For good i,the null hypothesis H_0 that the effect of the common trend has a homogenous effect implies $\delta_{ic} = \delta_i$ for all c and $\delta_i \neq 0$. Then,

$$v_{ict} - 1 = \frac{1}{\ln(t)\, t^{\alpha}} \frac{\psi_{ict} - \psi_{it}}{\delta_i + \frac{1}{\ln(t)t^{\alpha}}\psi_{it}},$$

$$(v_{ict}-1)^2 = \frac{(\psi_{ict}-\psi_{it})^2}{\psi_{it}^2+\ln(t)^2 t^{2\alpha}\delta_i^2+2\delta_i \ln(t)\, t^{\alpha}\psi_{it}},$$

and

$$V_{it} = \frac{1}{C}\sum_{c=1}^{C}(v_{ict}-1)^2 = \frac{\frac{1}{C}\sum_{c=1}^{C}(\psi_{ict}-\psi_{it})^2}{\psi_{it}^2+\ln(t)^2 t^{2\alpha}\delta_i^2+2\delta_i \ln(t)\, t^{\alpha}\psi_{it}}. \tag{16}$$

As $C, T \to \infty$ under Phillips and Sul's regularity conditions, we have

$$V_{it} = \left(\frac{1}{\ln(t)^2 t^{2\alpha}}\right)\frac{\sigma^2_{\psi_{it}}/\delta_i^2}{1+\ln(t)^{-2} t^{-2\alpha}\psi_{it}^2/\delta_i^2+2\ln(t)^{-1} t^{-\alpha}\psi_{it}/\delta_i}, \tag{17}$$

and

$$V_{i1} = \frac{\sigma^2_{\psi_{i1}}}{\psi_{i1}^2+\ln(1)^2\,\delta_i^2+2\delta_i \ln(1)\,\psi_{i1}},$$

which is independent of α_i. . Taking logs yields

$$\ln\frac{V_{i1}}{V_{it}} = \ln V_{i1} - \ln V_{it}, \tag{18}$$

and, using (**??**), we have

$$\ln V_{it} = -2\ln(\ln(t)) - 2\alpha_i \ln t + \ln\left\{\frac{v^2_{\psi_i C}}{\delta_i^2}\right\} + \epsilon_{it}, \tag{19}$$

Finally we have

$$\ln\frac{V_{i1}}{V_{it}} - 2\ln\ln(t) = a_i + 2\alpha_i \ln t + \epsilon_{it}$$

Under the alternative, $\delta_{ic} \neq \delta_i$, we have

$$\ln V_{it} = 2\ln\frac{\sigma_i}{\bar{\delta}} + \epsilon_{it},$$

However, note that $\ln t$ is negatively correlated with ϵ_t so that the regression coefficient of $\ln t$ will be negative. A more careful and detailed derivation is presented in Phillips and Sul (2007b).

References

[1] Bai, Jushan. 2003. "Inferential Theory for Factor Models of Large Dimensions," *Econometrica*, 71,135–171.

[2] Bai, Jushan and Serena Ng, 2002. "Determining the Number of Factors in Approximate Factor Models," *Econometrica*, 70, 191–221.

[3] Bai, Jushan and Serena Ng, 2004. "A PANIC Attack on Unit Roots and Cointegration," *Econometrica*, 72, 1127–1177.

[4] Cecchetti, Stephen G., Nelson C. Mark, and Robert Sonora, 2002. "Price Level Convergence Among United States Cities: Lessons for the European Central Bank," *International Economic Review,* November, 43, 1081-1099.

[5] Chen, Shiu-Sheng and Charles Engel, 2005. 'Does "Aggregation Bias" Explain the PPP Puzzle?' *Pacific Economic Review* 10, February, 49-72.

[6] Choi, Chi-Young, Nelson C. Mark and Donggyu Sul, 2006. "Unbiased Estimation of the Half-Life to PPP Convergence in Panel Data," *Journal of Money, Credit, and Banking*, June, 38, 921–938.

[7] Fleissig, Adrian R. and Jack Strauss, 2000. "Panel Unit Root Tests of Purchasing Power Parity for Price Indices," *Journal of International Money and Finance,* 19, 489-506.

[8] Frankel, Jeffery A. and Andrew K. Rose, 1996. "A Panel Project on Purchasing Power Parity: Mean Reversion within and between Countries." *Journal of International Economics* 40, 209–224.

[9] Hodrick, Robert J. and Edward C. Prescott, 1997. "Postwar U.S. Business Cycles: An Empirical Investigation," *Journal of Money, Credit, and Banking,* 29, 1–16.

[10] Imbs, Jean, Haroon Mumtaz, Morten O. Ravn, and Hélène Rey, 2005. "PPP Strikes Back: Aggregation and the Real Exchange Rate," *Quarterly Journal of Economics* 120, 1–43.

[11] Koedijk, Kees G., Peter C. Schotman, and Mathijs A. Van Dijk, 1998. "The Re-Emergence of PPP in the 1990s," *Journal of International Money and Finance*, 17,51–61.

[12] Murray, Christian and David Papell, 2005. "Do Panels Help Solve the Purchasing Power Parity Puzzle?" *Journal of Business and Economic Statistics*, October, 410–415.

[13] O'Connell, Paul G.J. 1998. "The Overvaluation of Purchasing Power Parity," *Journal of International Economics*, 44, February, 1-19.

[14] Papell, David, 2006. "The Panel Purchasing Power Parity Puzzle," *Journal of Money, Credit, and Banking,* March, 447-467

[15] Phillips, Peter C.B. and Donggyu Sul, 2007a. "Bias in Dynamic Panel Estimation with Fixed Effects, Incidental Trends and Cross Section Dependence," *Journal of Econometrics*, 138, 162-188.

[16] Phillips, Peter C.B. and Donggyu Sul, 2007b. "Transition Modelling and Econometric Convergence Tests," *Econometrica,* Vol 75, 1771-1855

[17] Quenoille, M. H., 1956, "Notes on Bias in Estimation," *Biometrika*, Vol 43, 353-360.

[18] Rogoff, Kenneth, 1996. "The Purchasing Power Parity Puzzle," *Journal of Economic Literature*, 34, 647-668.

[19] Stock, James and Mark Watson, 2002 ."Forecasting Using Principal Components from a Large Number of Predictors," *Journal of the American Statistical Association*, 97 1167–1179.

Table 1: Decomposition of $\mathrm{Var}(q_i)$

i	Total	Common	Ratio
Fruits	2.816	0.651	0.231
Comm	3.182	1.175	0.369
Tobacco	2.706	1.165	0.430
Sound	4.251	3.189	0.750
Hotel	2.825	2.252	0.797
Pub.Trans.	3.309	2.741	0.828
Dairy	3.083	2.555	0.829
Fuel	2.926	2.433	0.832
Rents	4.255	3.609	0.848
Meat	2.278	1.932	0.848
Clothing	5.308	4.508	0.849
Footwear	4.968	4.221	0.850
Vehicles	3.743	3.182	0.850
Bread	3.242	2.803	0.864
Leisure	3.101	2.684	0.866
Book	4.123	3.581	0.868
Alcohol	4.841	4.246	0.877
Dom.Appl.	4.160	3.743	0.900
Furniture	5.120	4.631	0.904
Average	3.697	2.911	0.787

Table 2: $t-$ratios for slope coefficients in
$\ln\left(V_{i1}/V_{it}\right) - 2\left(\ln\left(\ln\left(t\right)\right)\right) = a_i + 2\alpha_i \ln\left(t\right) + u_{it}$

	First Observation in the regression										
	90.2	90.3	90.4	90.5	90.6	90.7	90.8	90.9	90.10	90.11	90.12
Clothing	-364	-355	-331	-307	-287	-267	-243	-215	-185	-157	-131
Vehicles	-160	-153	-148	-139	-120	-96.1	-73.5	-56.8	-45.9	-39.3	-35.3
Fruits	-45.5	-46.0	-45.9	-44.5	-41.5	-37.7	-33.2	-28.7	-24.4	-20.6	-17.4
Comm.	-25.7	-25.6	-25.2	-24.5	-23.8	-23.0	-22.1	-21.3	-20.6	-19.9	-19.2
Dom.Appl.	-25.4	-24.4	-23.5	-22.6	-21.9	-21.2	-20.5	-20.0	-19.4	-19.0	-18.6
Footwear	-17.5	-18.0	-18.5	-19.1	-19.7	-20.3	-21.0	-21.7	-22.5	-23.4	-24.2
Meat	-11.9	-12.4	-13.1	-13.8	-14.6	-15.6	-16.6	-17.7	-19.0	-20.5	-22.1
Sound	-9.53	-10.1	-10.7	-11.5	-12.5	-13.6	-14.9	-16.5	-18.3	-20.3	-22.5
Furniture	-13.75	-13.2	-12.7	-12.4	-12.1	-12.0	-11.9	-11.9	-12.0	-12.3	-12.7
Bread	-12.76	-11.6	-10.7	-9.92	-9.36	-8.71	-8.26	-7.88	-7.57	-7.33	-7.15
Tobacco	-14.02	-12.2	-10.7	-9.50	-8.47	-7.61	-6.90	-6.31	-5.83	-5.43	-5.11
Rent	0.49	0.14	-0.31	-0.90	-1.67	-2.62	-3.74	-4.97	-6.14	-6.99	-7.54
Pub.Trans.	1.63	1.46	1.28	1.09	0.91	0.72	0.53	0.34	0.14	-0.06	-0.27
Leisure	**3.32**	**3.16**	**3.01**	**2.88**	**2.75**	**2.64**	**2.53**	**2.42**	**2.33**	**2.23**	**2.15**
Hotel	1.33	1.61	1.90	**2.19**	**2.49**	**2.81**	**3.15**	**3.52**	**3.94**	**4.42**	**4.97**
Book	**8.38**	**7.21**	**6.36**	**5.73**	**5.25**	**4.88**	**4.58**	**4.33**	**4.12**	**3.94**	**3.79**
Alcohol	1.35	**2.97**	**5.61**	**7.80**	**8.86**	**8.45**	**6.90**	**5.39**	**4.38**	**3.75**	**3.35**
Dairy	**17.3**	**16.6**	**15.9**	**15.3**	**14.7**	**14.1**	**13.5**	**13.0**	**12.5**	**12.0**	**11.5**
Fuel	**13.9**	**14.2**	**14.5**	**14.8**	**15.1**	**15.3**	**15.5**	**15.8**	**16.1**	**16.4**	**16.7**

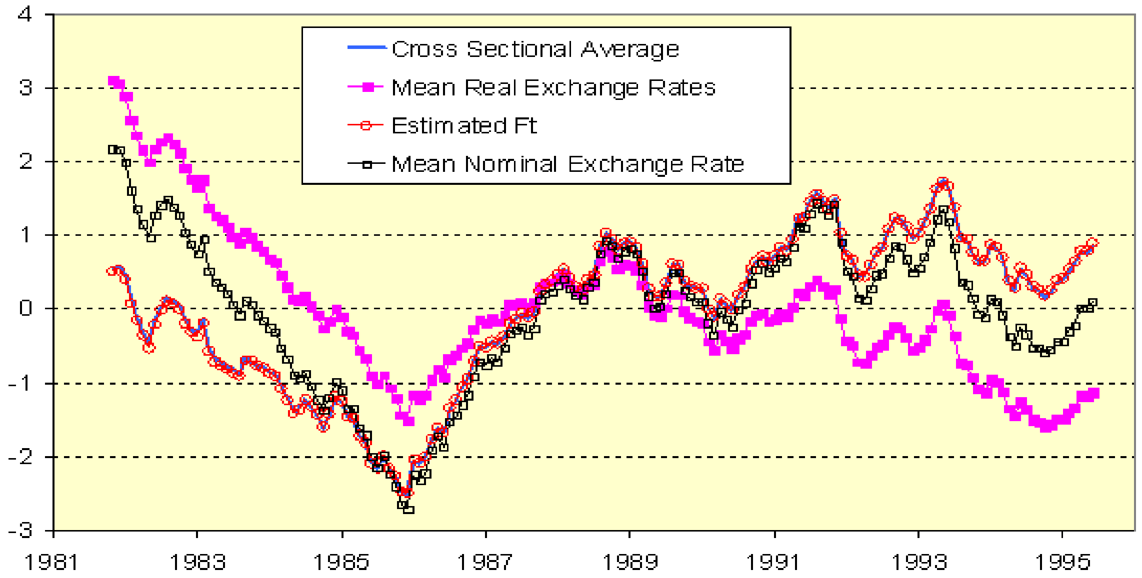

Figure 1: Alternative measurements of the common factor F_t.

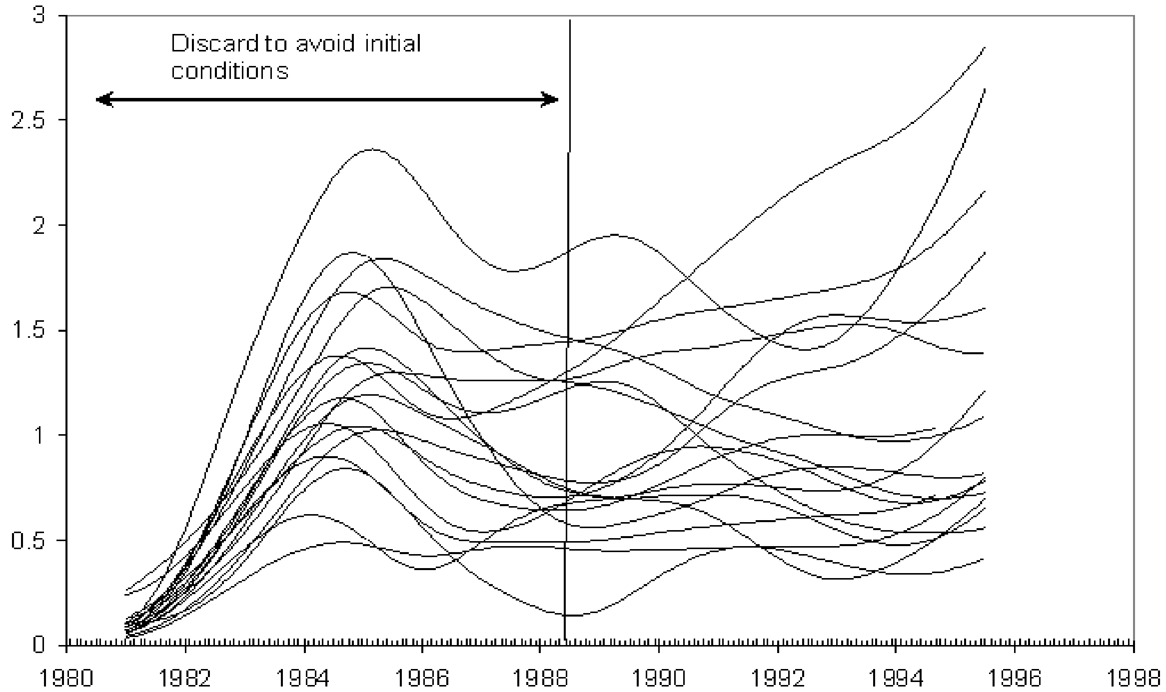

Figure 2: V_{it} for $i = 1, ..., 19$ across 12 countries. First half of sample discarded to control for initial conditions. $V_{it} = V_{it} = \frac{1}{C}\sum_{c=1}^{C}(v_{ict} - 1)^2$, where $v_{ict} = p_{ict}/\left(\frac{1}{C}\sum_{k=1}^{C} p_{ikt}\right)$, p_{ict} is the Hodrick–Prescott trend of the U.S. dollar price of good i in country c at time t. The observations do not start from the exact same point in 1981.05 due to variations in the cyclical part of the prices which have been subtracted off.

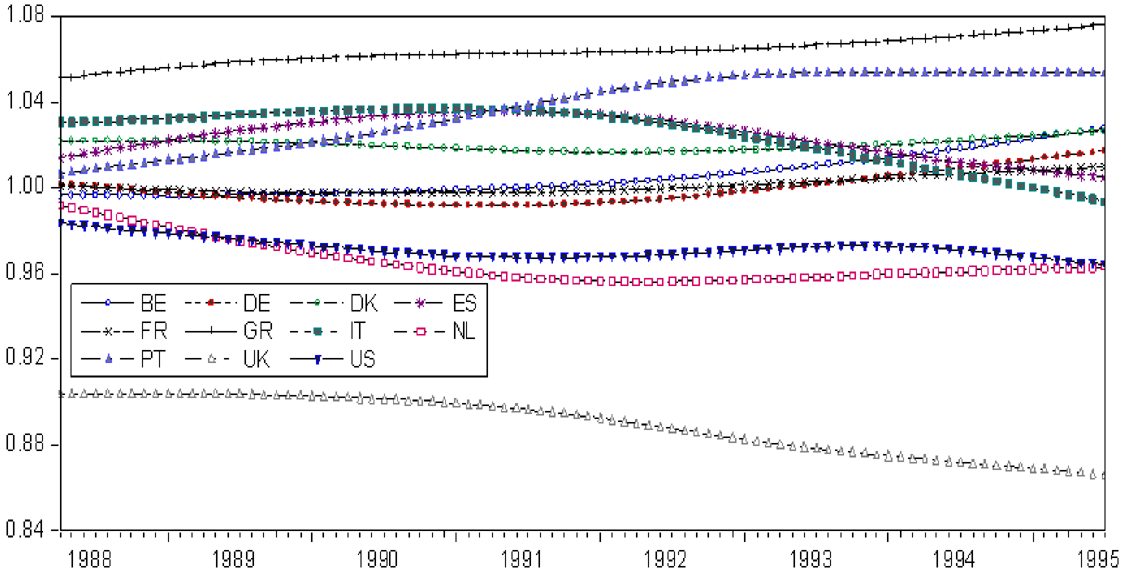

Figure 3: v_{ict} for the U.S. dollar price of i =clothing for each of 11 countries

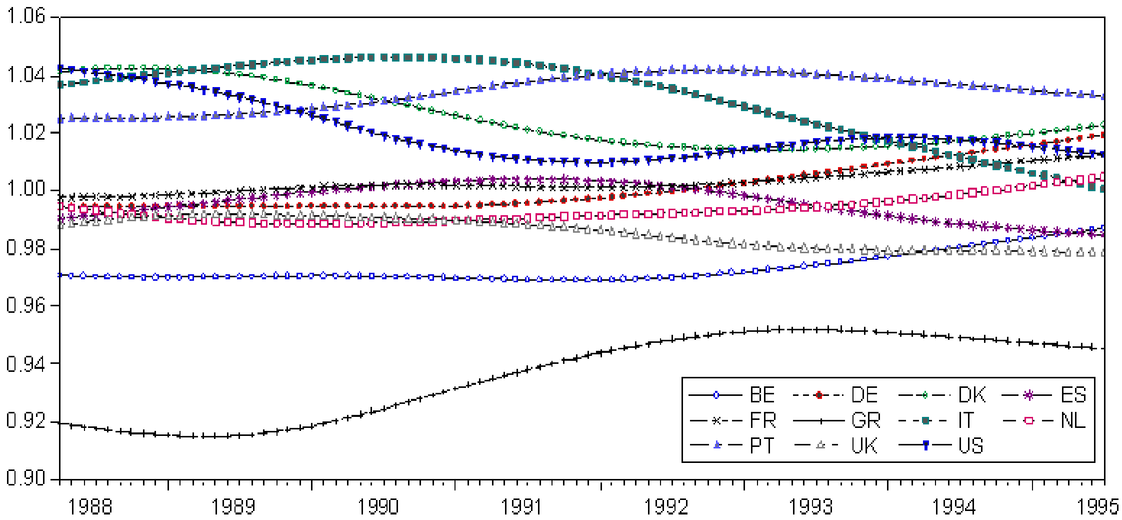

Figure 4: v_{ict} for the U.S. dollar price of i =fuel for each of 11 countries

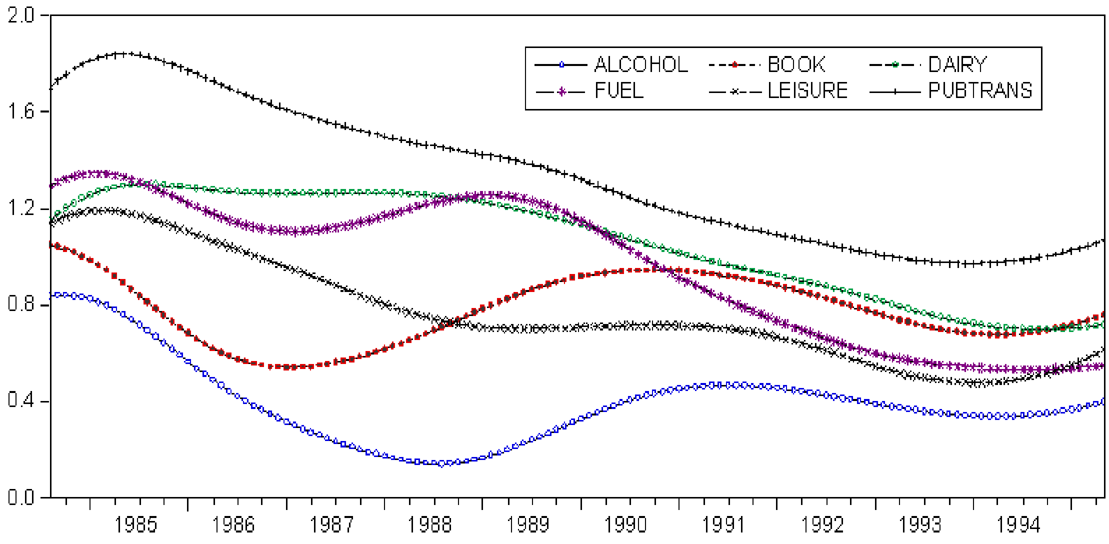

Figure 5: Cross-sectional variance of prices for convergent group of goods

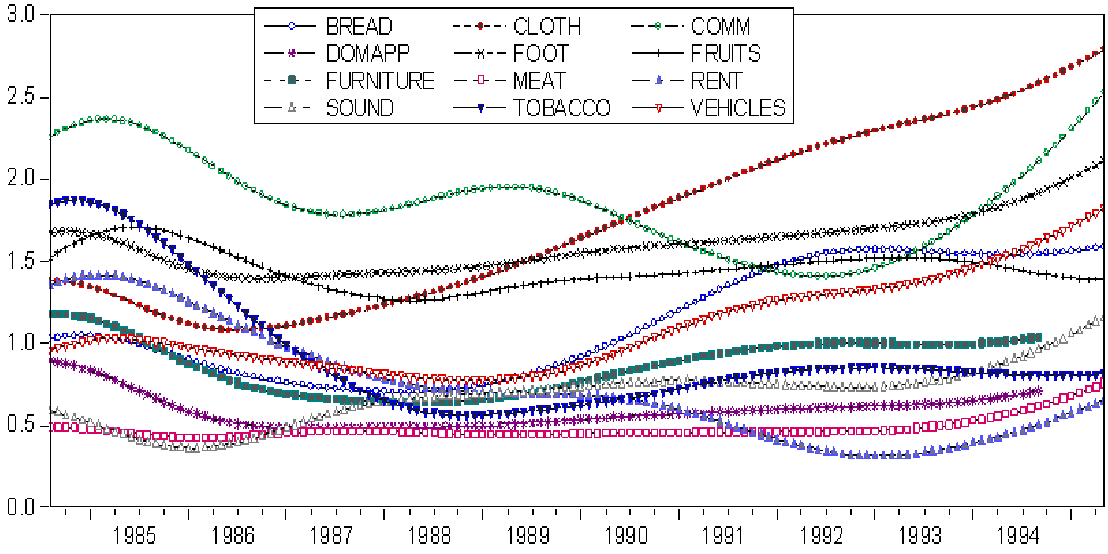

Figure 6: Cross-sectional variance of prices of divergent group of goods