Persistence in Law-of-One-Price Deviations: Evidence from Micro-data

Mario J. Crucini and Mototsugu Shintani*

Original version: December 2001
Latest version: July 2002

Abstract

We study the dynamics of good-by-good real exchange rates using a micro-panel of 270 goods prices across 90 international cities and 13 cities within the U.S., annually from 1990 to 2000. The picture of relative price adjustment that emerges from our analysis is that price adjustment is very rapid both across cities within countries and across cities of the world. In terms of persistence of Law-of-One-Price deviations, national borders appear not to matter. What national borders do matter for are the magnitudes of the long run deviations from the Law-of-One-Price. Across U.S. cities the deviations are economically small while across international cities the deviations are economically large.

1. Introduction

We study the time series and cross-sectional behavior of 270 goods prices across 90 international and 13 U.S. cities annually from 1990 to 2000. Our goal is to use this novel data source to shed light on the Purchasing Power Parity Puzzle, which Kenneth Rogoff (1996) describes as follows.

> "How can one reconcile the enormous short-term variability of real exchange rates with the extremely slow rate at which shocks appear to damp out? Most explanations for short-term exchange rate volatility point to financial factors such as changes in portfolio preferences, short-term asset price bubbles, and monetary shocks (see, for example, Maurice Obsfteld and Rogoff (1998, 2000). Such shocks can have substantial effects on the real economy in the presence of sticky nominal wages and prices. Consensus estimates for the rate at which PPP deviations

*Department of Economics, Vanderbilt University. Crucini gratefully acknowledges the financial support of the National Science Foundation (SES-0136979). We thank David Parsley and seminar participants at University of Tokyo, 2002 Midwest Macroeconomics Conference, and 2002 North American Summer Meeting of Econometric Society for helpful comments.

damp, however, suggest a half-life of three to five years, seemingly far too long to be explained by nominal rigidities. It is not difficult to rationalize slow adjustment if real shocks – shocks to tastes and technology – are predominant. But existing models based on real shocks cannot account for the short-term exchange rate volatility." Rogoff (JEL, June 1996, pp. 647-48.)

The main thesis of this paper is that the half-lives estimated in the existing empirical literature confound two distinct properties of relative prices across locations. At any point in time the price differential of a particular good across locations must be a combination of a transitory deviation and a permanent price difference. An example of a transitory deviation in prices across locations is one arising from a temporary disturbance that shifts supply or demand in either location from its steady state level. Examples of sources of permanent price differences at the retail level may include transportation costs, differences in sales taxes across locations or differences in the costs of non-traded inputs across locations. All of these factors work to prevent arbitrage from driving relative prices toward the Law-of-One-Price prediction in the long run. Because data limitations have presented formidable barriers to investigations of the behavior of absolute prices – as opposed to price indices – these two properties of prices have been impossible to distinguish in much of the existing literature. In contrast, our study uses absolute price data at the level of individual retail goods and services enabling us to estimate the long-run levels of price differences and the rate at which prices converge back to these levels once a disturbance hits a particular market.

We examine the dynamics of relative prices under two common views about their long run means. We describe these alternative views as *absolute* and *conditional price convergence*. *Absolute convergence* implies that in the long run the prices of all goods, across all locations, obey the Law-of-One-Price. *Conditional convergence* implies that in the long run prices converge to a stationary distribution, with each point in that distribution representing the price of a good in a particular city relative to the average price of that good across all cities. Implicit in these definitions is the assumption that relative prices across locations are stationary regardless of the levels of their long-run means.

Employing these alternative assumptions about the long-run turns out to have dramatic implications for the estimated speed of price adjustment. We find:

- *absolute price convergence* is slow (under the null that it occurs at all) both within and across countries; half-lives average 5.4 years across international cities and 3.3 years across U.S. cities.

- *conditional price convergence* is rapid, with half-lives averaging between 9-10 months across international cities and between 9-13 months across U.S. cities.

- the unit root null is rejected for virtually all goods in the international data and is rejected for almost all goods within the United States.

- the null of absolute price convergence is rejected for most international micro-prices, but not rejected for about half of the goods for intranational relative prices (relative prices across locations within the United States).

Thus, the primary feature of relative price behavior that distinguishes locations within countries from those we observe internationally is not the persistence of the stochastic fluctuations of relative prices around their long run means, but the magnitude of the long run deviations themselves. The relative price of a typical good between New York City and Los Angeles adjusts about as quickly to a disturbance as the relative price of a typical good between New York and Madrid; the difference is that the long run differential is close to zero for most goods in the former case, but very substantial for most goods in the latter case.

Our work is mainly focused on assessing estimates of half-lives of international price differences, which Rogoff (1996) places at between 3 and 5 years after surveying the vast literature on the topic. Less closely related, but also relevant, is an emerging literature that studies the cross-sectional and time series properties of deviations from the Law-of-One-Price at the retail level. Many of these studies utilize individual goods or subsets of the consumption basket. Examples include Cumby (1996) who studies Big Mac hamburgers, Froot, Kim and Rogoff (1995) who study wheat, butter and charcoal, Ghosh and Wolf (1994) who study the Economist Magazine, Haskel and Wolf (1998) who study IKEA furniture, and Lutz (2001) who studies dispersion in automobile prices across European countries.[1] One robust finding in these papers is that Law-of-One-Price deviations often exceed what is reasonable to attribute to transportation costs and the deviations differ significantly across goods even within fairly narrow commodity groupings.

One limitation of these data sources for the purposes of our work is they are limited in either the number of goods available for examination or in the number of locations the price survey is conducted. Having a large number of goods is important to avoid having the results depend on price dynamics in a particular sector. For example, while Cumy (1996) has found price convergence in Big Mac hamburgers to be quite rapid, it is difficult to know if these generalizes to other goods in the consumption basket. Studies with a limited set of locations make it difficult to obtain reliable inferences because researchers are unable to exploit asymptotic properties of the estimators in the cross-sectional dimension (the time samples are almost uniformly short).

The intranational part of our analysis is most closely related to work by Parsley and Wei (1996) and Cecchetti, Mark and Sonora (1999). Parsley and Wei utilize the ACCRA price survey and study the dynamics of 51 retail prices across 48 U.S. cities quarterly from 1975 to 1994. They find less rapid convergence rates than we do ranging from 45 months quarters for services to 12 to 15 months for goods in terms of the half-lives. Importantly, though, despite using different goods, a different frequency and a somewhat different methodology (bilateral relative prices as opposed to normalizing to the U.S. average price), the half-lives are much lower than observed internationally, consistent with our findings. Cecchetti, Mark and Sonora utilize an almost century long panel of U.S. CPI data for major cities and find half-lives of real exchange rates on the order of 8 to 10 years. While these half-lives are implausible from the point of view of sluggish nominal price adjustment, they do indicate that the half-lives of 3.3 years for the U.S. under absolute price convergence are not at the high end of estimates existing in the literature. A third study which uses a similar

[1] See also, Crownover, Pippenger and Steigerwald (1996), Isard (1977), Giovannini (1988), Engel (1993), Engel and Rogers (1996), Rogers and Jenkins (1995).

methodology and data is Rogers (2001); he focuses on European countries and finds quite rapid and dramatic price convergence from 1990 to 2000.

2. Theory

2.1. A Model of Retail Price Determination

Debreu (1959) includes the location as one of the dimensions along which a commodity-space is defined. What most economists have in mind when thinking of location in terms of prices is the transportation cost of moving the good from one location to another. Thus traded prices are – in the absence of tariffs and other official barriers to trade – typically assumed to satisfy the Law-of-One-Price up to a transportation cost. Because this study utilizes retail prices, there is another aspect of location that needs to be considered. An important difference between a bottle of beer consumed at a home and a bottle of beer consumed at a restaurant is that the "rental cost" of providing the space for consumption is counted in the housing part of the consumer price index in the former case and in the price of a bottle of beer away from home in the latter case. One would expect that the deviation from the Law-of-One-Price to be bound by transportation costs and taxes in the case of home consumption but not in the case of consumption away from home (unless rental costs happen to be equal across locations). Moreover, the rental costs often embody something about the location of consumption that make the beer at home and away from home different "goods" in Debreu's terminology. In this section we develop a simple model of price determination that recognizes some of these features of retail markets.

Formally, we begin with the cost minimization problem of the firm. Here a firm is a retail outlet that is assumed to sell a single good in a single location. Managers of retail outlets take the prices of inputs as given. Inputs include both non-traded and traded goods and/or services. Non-traded inputs are locally provided while traded inputs are either exported or imported.

The cost function for the individual retail firm is the solution to the following minimization problem solved at each date:

$$\min_{\left\{N_{it}^j, \mathbf{X}_{it}^j\right\}} C_{it}^j = W_{it}N_{it}^j + \sum\nolimits_{k=1}^{K(j)} P_{it}^k X_{it}^k \tag{2.1}$$

$$\text{s.t. } f^j(N_{it}^j, \mathbf{X}_{it}^j) \geq Y_t \equiv (N_{it})^{\alpha_j} (\prod_{k=1}^{K(j)} (X_{it}^j)^{\theta_k^j})^{(1-\alpha_j)} \tag{2.2}$$

where C_{it}^j is the cost of producing good j in city i at time t. w_{it} is the wage in city i and P_{it}^k is the price of traded intermediate input k in location i at time t. It is natural to think of retailers paying for non-traded inputs beyond labor services (e.g., retail space, utilities, advertising, among others), but the wage cost will suffice for expositional purposes.

The distinction between non-traded and traded inputs serves two purposes. First, as implicit in the indexing we have adopted two standard assumptions. The first is that factor mobility is much higher across sectors within a location than across locations – w_{it} is location-specific, not good-specific. The second assumption is that retailers in all locations produce good j using the same production technology – α_j and θ_k^j are good-specific, not location-specific.

Under constant returns to scale, the cost function will take the form: $Y_t \cdot C(\mathbf{P}_t, 1)$ where Y_t is the desired real output level and $C(\mathbf{P}_t, 1)$ is the unit cost function, the object of interest. We allow the retailer to set a – possibly time-varying – markup over marginal cost, $b_{it}^j \geq 1$, so the per unit retail price of good j faced by a consumer located in country i, P_{it}^j, is:

$$P_{it}^j = b_{it}^j (w_{it})^{\alpha_j} (\prod_{k=1}^{K(j)} (P_{it}^k X_{it}^k)^{\theta_k^j})^{(1-\alpha_j)} \tag{2.3}$$

Taking logarithms of (2.3) and subtracting an analogous expression for the geometric average price of the same good across all locations (we use q_{it}^j to denote the relative price of good j in location i relative to the geometric average price of the same good across all locations: $\ln(e_{it} P_{it}^j / P_t^j)$ gives us an expression for the Law-of-One-Price deviation of good j in location i relative to its geometric mean across all locations (all locations means all international cities or all U.S. cities in what follows)

$$q_{it}^j = \mathbf{b}_{it}^j + \alpha_j \mathbf{w}_{it}^j + (1 - \alpha_j) \sum_{k=1}^{K(j)} \theta_k^j q_{it}^k. \tag{2.4}$$

where $\mathbf{b}_{it}^j \equiv \ln(b_{it}^j / b_{it})$ and $\mathbf{w}_{it}^j \equiv \ln(w_{it}^j / w_{it})$.

2.2. Absolute and Conditional Price Convergence

The data needed to implement the retail model is well beyond that currently available. We have the retail prices that appear on the left-hand-side of (2.3) but we lack much of what is described on the right-hand-side of that equation. Our approach is to cast the major implications of the model for international and intranational retail price behavior across locations and time based on very mild restrictions governing the evolution of the variables that are predicted to determine the retail price.

Toward this end, we formally define conditional and absolute international price convergence using the notation and structure of the model. In what follows we assume the markup, relative wage and relative prices are traded goods are all stationary stochastic processes. We denote the steady-state values of relative markups and prices as: $\mathbf{b}_i^j$, $\mathbf{w}_i$,$\{q_i^1, \cdots,.q_i^K\}$; these objects are the theoretical population means of the time series observations or equivalently the level each economic variable are assumed to reach asymptotically if current and all future disturbances were set to zero.

Remark 1. *The stationary random variables* $\left\{\mathbf{b}_{it}^j, \mathbf{w}_{it}, q_{it}^1, \cdots q_{it}^K\right\}$ *and time-invariant parameters* $\{\alpha_j, \theta_1, \cdots \theta_k\}$ *constitute a sufficient information set to determine the retail price* q_{it}^j *at any point in time. This follows directly from equation (2.4).*

Definition 1. *Absolute price convergence. International prices are said to obey absolute convergence if, in the absence of disturbances that alter markups and input prices over time, the distribution of real exchange rates* $F(q_{it}^j)$ *converges asymptotically to a degenerate distribution at the point* $q_i^j = 0 \ \forall i, j$.

Remark 2. *A set of sufficient conditions for absolute convergence in the non-stochastic steady-state are the following:* $\mathbf{b}_i^j = 0$, $\forall i, j$ $\mathbf{w}_i = 0$ $\forall i$, $q_{it}^k = 0$ $\forall i, k \in K(j)$. *The first constraint is that the markup of price over marginal cost for each good be the same across locations. The second constraint is wage convergence across locations. The last constraint is the convergence of traded good's prices to the Law-of-One-Price across all locations.*

Remark 3. *The conditions for absolute price convergence may be relaxed by considering the following cases. Case 1: the total cost of intermediate inputs is the same across locations:* $\sum_{k=1}^{K(j)} \theta_k q_i^k = 0$. *The convergence of this sum might be plausible considering that* $q_i^1 = \tau_i^1$, *if intermediate input 1 is imported and* $q_i^2 = -\tau_i^2$, *if intermediate input 2 is exported. Averaging over a sufficiently large set of traded intermediate inputs would drive this number toward zero. Case 2: If relative markups and relative wages were negatively correlated in the cross-section, the first and second term in equation (2.4) will tend to average to zero. The negative correlation would mean that low wage countries we subjected to higher markups in product markets than higher wage countries.*

Definition 2. *Conditional price convergence. International prices are said to obey conditional convergence if, in the absence of disturbances that alter markups and input prices over time, the distribution of real exchange rates* $F(q_{it}^j)$ *converges asymptotically to the non-degenerate distribution* $G(q_i^j)$ *where the distribution* $G(q_i^j)$ *is determined by the steady-state distribution of markups, wages and trade costs across countries. The mapping between the geographic distribution of input prices and retail prices is determined, element-by-element, from*

$$q_i^j = \mathbf{b}_i^j + \alpha_j \mathbf{w}_i + (1 - \alpha_j) \sum_{k=1}^{K(j)} \theta_k^j q_i^k. \tag{2.5}$$

Much of the empirical literature on PPP and the Law-of-One-Price focuses on the trade costs as embodied in the last term of (2.5). Using a more extensive cross-section of goods than we study here, Crucini, Telmer and Zachariadis (2001) document show that – consistent with the formulation above – market structure, tradeability and the share of non-traded inputs into production are quantitatively important features of geographic price dispersion across European cities. Similar evidence for the U.S. via–a-vis Brazil may be found in Burstein, Neves, and Rebelo (2001). Our notion of conditional convergence is meant to capture these features of the price distribution.

Our empirical analysis will also contrast the behavior of relative prices across cities within the United States. The theory outlined above is applicable to both intranational and international contexts, but the quantitative implications are likely to be quite different. For example, we would expect wage dispersion to be lower across regions within a country compared to that observed internationally. The relative importance of transportation costs across and within countries is not so obvious – overland transportation costs tend to be higher than overseas transportation costs for many goods. Another feature that distinguishes intranational and international markets is that the nominal exchange rate is irrevocably fixed in the former case but often fluctuates dramatically in the latter case. Unless nominal prices adjust instantaneously to changes in the nominal exchange rate we should expect micro-real exchange rates to change when the nominal exchange rate changes. Moreover, if wages adjust more slowly to currency fluctuations than do traded goods prices we would expect retail prices with larger wage costs to exhibit a higher degree of comovement with the nominal exchange rate.

2.3. Implications

Since the retail price is a log-linear combination of the markup, wage rate and traded intermediate good's prices, the dynamics of the retail price is a function of the dynamic properties of these variables. In our empirical work we restrict ourselves to first-order autoregressive (AR(1)) processes:

$$\begin{array}{ll} q_{it}^j = \rho^j q_{it-1}^j + \eta_i^j + v_{it}^j & \text{conditional convergence} \\ q_{it}^j = \rho^j q_{it-1}^j + v_{it}^j & \text{absolute convergence.} \end{array}$$

The individual effect, η_i^j, is specific to the location and the good and non-zero values of these parameters is what distinguishes conditional from absolute convergence. In terms of the parameters of our retail model the individual effect and long run mean of a good-specific real exchange rate are related in the following way:

$$q_i^j = \frac{\eta_i^j}{1-\rho^j} = \mathbf{b}_i^j + \alpha_j \mathbf{w}_i + (1-\alpha_j) \sum_{k=1}^{K(j)} \theta_k^j q_i^k.$$

In words: the long-run price of good j in location i relative to the mean (across countries or cities within a country) is equal to the relative markup plus production-shared weighted averages of the relative wage and traded intermediate prices.

The theory has a number of testable implications and parameters of interest. Among the testable restrictions are the stationarity of good-by-good real exchange rates and the absence or the presence of the individual effects. Stationarity is perhaps the weakest restriction theoretically, requiring bounds on absolute differences in markups, wages and traded good prices. It is natural to think of traded good prices as bounded by transportation costs and tariffs while the entry and exit of firms would bound the magnitude of markups across locations. Stationarity of the relative wage may be more problematic, at least in an international context since the costs of migration are formidable. However, for countries will similar factor endowments and technology the wage levels might be expected to be quite similar. Across cities within a country we would expect all of the determinants of retail prices to be less geographically variable, with perhaps the greatest difference being in wages where factor mobility reduces the disparities substantially.

With both the theoretical and statistical model as a background it is productive to step back and think about the implications of our model for geographic price dispersion. At any point in time the differences in prices across locations will reflect the combined influence of long run differences in relative prices in the steady-state, the persistence of movements in prices along the transition path, and the magnitude of the disturbances themselves. The principle advantage of having time series and cross-sectional data on *absolute* prices is that we are able to incorporate all three of these properties in our estimation technique. Thus we will be able to provide an estimate of how much of the dispersion we see in prices at a point in time is due to stochastic elements (such as business cycle variation and nominal exchange rate variation) and how much is the result of different long run prices across locations.

3. Data

3.1. Measuring Law-of-One-Price Deviations

The price data is collected by the *Economist Intelligence Unit* in its annual international price survey. The survey spans 301 individual retail goods and services across 122 cities located in 78 countries. The greater number of cities than countries reflects the fact that in some countries price data is gathered for more than one city. We conduct our international analysis using one city from each country (though which city we choose varies somewhat across goods) while for our intranational analysis, we use cities within the continental U.S. We chose the U.S. for the simple reason that it has by far the largest number of cities surveyed at 13, compared to the next largest number of cities surveyed which equals 5 in Australia, China and Germany. The sample period begins in 1990 and ends in 2000.

The raw data are in domestic currency units so we begin by converting all of the prices from domestic currency into U.S. dollars at the average exchange rate prevailing for the year the price observation is recorded. Next, we compute the deviation of the price of each good from its average across all locations to place all prices on a common scale. Moreover, this normalization avoids problems inherent in choosing an arbitrary numeraire location.

For example, suppose the U.S. dollar price of good j in city i is P_{it}^j. Let the number of available cities for good j be N, then the relative price (in logs) q_{it}^j is the price of good j in U.S. dollars in city i relative to global mean of that good, $\overline{p}_t^j$, at time t. Namely,

$$q_{it}^j = p_{it}^j - \overline{p}_t^j = \ln P_{it}^j - \frac{1}{N}\sum_{i=1}^{N} \ln P_{it}^j$$

which is exactly the same definition of a real exchange rate for good j in location i at period t, applied in the theoretical section.

Since the raw data contain a number of missing observations and we want to work with balanced panels, we select goods and locations in the following way. First, when the exchange rate for a particular city is not available the city is removed from all panels. Second, for each good, cities that contain missing observations are removed.

In our estimation we pool our data across locations and time but estimate a separate regression for each good. The number of cities we use ranges from 23 to 69 and averages 58 in the international context; it ranges from 10 to 13 and averages 13 in the intranational context. In selecting the city to use in our international analysis we rely on the first city alphabetically from each country and add cities when necessary to achieve a balanced panel for a particular good.

Table 1 presents the cities used in our study ordered by region and then alphabetically by country. The regions are Africa, Asia, Central America, Europe, North America, Oceania and South America. The 13 U.S. cities used in the analysis are also reported in Table 1.

The number of goods for which a particular city is used in the estimation is noted in parentheses. Note that in most cases the city coming first in the alphabet is used for almost all of the analysis. For example, Germany is used as a location in all of the 270 (goods) panels that we estimate, but due to missing observations we use Berlin as the German city for 265 of those goods and Dusseldorf for the remaining 5 goods.

Table 2 presents the breakdown of our sample of goods by our own classification in terms of major consumption categories. The span of goods and services appears broadly representative of items found in a typical Consumer Price Index (the U.S. CPI for example) with the number of price observations by major consumption category varying considerably from a high of 112 (food and beverages) to a low of 3 (domestic help and salaries).

3.2. Price Dispersion

Figure 1 plots the distribution of Law-of-One-Price deviations (q_{it}^{j}) pooling all goods and locations. The figure presents lines for the distribution of prices within the United States and across international cities for 3 of the 10 years available (1990, 1995 and 2000). The densities are centered at zero, by construction and estimated using a Gaussian kernel. If the Law-of-One-Price held for all goods and across all locations, the densities would be degenerate at zero. Obviously there are significant deviations from the Law-of-One-Price even though many goods come close (e.g. within 10% of the prediction). Particularly striking are the differences in the international and intranational distributions. The dispersion of prices across cities in the U.S. is significant but much lower than what we observe internationally; in 1990 the respective standard deviations are 0.281 and 0.685. The standard deviation is quite stable over time as is evident in comparing the densities across the five year intervals.

The price dispersion in Figure 1 could be entirely due to differences in price levels across cities, common to all goods. For example, the cities in the richest countries might be located in the right-hand tail of the distribution while the cities in the poorer countries might be located in the left-hand tail of the distribution and all the price deviations are due to these city-effects. To avoid presenting two dimensions to price dispersion in a one-dimensional figure we average the good-by-good real exchange rates to obtain something like a PPP level by city. Figure 2 plots the distribution of $\overline{q}_{it} = \frac{1}{M}\sum_{j=1}^{M} q_{it}^{j}$ within and across countries for 1990, 1995 and 2000. If prices were <u>on average</u> equal across cities, these densities would be degenerate at zero. Obviously this is not the case; it is also clear that price levels vary much more across countries than they do across U.S. cities.

Alternatively, the dispersion in prices could be entirely due to differences in Law-of-One-Price deviations across goods. For example, traded goods may come close to satisfying the Law-of-One-Price prediction and therefore tend to cluster around the point 0 whereas non-traded goods have larger deviations and are to be found in the tails of the distribution. Figure 3 shows that Law-of-One-Price deviations are larger for non-traded goods than for traded goods in both international and intranational distribution. Comparing Figures 2 and 3, city effects appear more important than the traded/non-traded classification.

One explanation for the differences we observe in international and intranational price dispersion is that nominal exchange rate fluctuations move real exchange rates across time and since the U.S. is a common currency area, dispersion in price levels across cities (and goods) is lower. Another explanation is that arbitrage costs prevent the Law-of-One-Price from holding across locations and these costs are higher internationally than intranationally. One of the main goals of the paper is to examine the cross-sectional time series behavior of micro-real exchange rates to see if these alternative views can be supported.

In particular we hope to determine the relative importance of permanent deviations from the Law-of-One-Price and stochastic fluctuations around those long run levels for the shape of the price densities both within and across countries. At one extreme, real exchange rates may be constant over time but the Law-of-One-Price is grossly violated; thus the long run and period-by-period densities are one in the same. At another extreme, the long run density is degenerate, satisfying the Law-of-One-Price restriction for all goods, locations and time periods. In this case, price distribution is entirely driven by temporary fluctuations around a degenerate long run distribution.

4. The Empirical Results

4.1. Unit Root Tests

We begin our formal statistical work by conducting unit root tests for each good in our cross-section since the economic interpretation of the sample means are not very useful if real exchange rates are not stationary. With at most 11 time series observations for the real exchange rate of a given city, conventional panel unit root tests are not applicable. Instead, we pool all cities for each good and employ a panel unit root test with fixed time dimension as developed by Richard D.F. Harris and Elias Tzavalis (1999). They considered the following three different models:

Model with no constant:	$q_{it} = \rho q_{it-1} + v_{it}$
Model with constant:	$q_{it} = \mu_i + \rho q_{t-1} + v_{it}$
Model with constant and trend:	$q_{it} = \mu_i + \beta_i t + \rho q_{t-1} + v_{it}$

where q_{it} is the real exchange rate for a particular good in location i at date t. We suppress the good index since this will be the basic unit of study in what follows. Corresponding unit root test statistics for the null hypothesis of $\rho = 1$ are $C_1^{-1/2}\sqrt{N}(\widehat{\rho}_{POOL} - 1)$, $C_2^{-1/2}\sqrt{N}(\widehat{\rho}_{LSDV} - 1 - B_2)$ and $C_3^{-1/2}\sqrt{N}(\widehat{\rho}_{LSDVT} - 1 - B_3)$, respectively, where $\widehat{\rho}_{POOL}$ is the least squares pooled estimator, $\widehat{\rho}_{LSDV}$ is the least squares dummy variable (LSDV) estimator, $\widehat{\rho}_{LSDVT}$ is the least squares dummy variable with trend estimator.[2] As N grows, each statistic has been shown to follow a standard normal distribution under the null hypothesis of a unit root. In the context of our analysis, the first test can be viewed as a test of no convergence against the alternative of absolute convergence while the second and third tests can be viewed as a test of no convergence against the alternative of conditional convergence.

Table 3 summarizes the results of unit root tests. In the international data we are able to reject the null of a unit root is almost every case. Rejection rates range from 91% when we include a constant and time trend and test at the 1% level of significance to 100% when we include a constant but no time trend. The intranational evidence is not as overwhelming in the rejection rates. However, fewer rejections for intranational data compared to international evidence can be simply explained by the lack of power due to small sample size. In addition, even in the case with a constant and time trend (which we expect to have lowest power), we reject a unit root for more than one half of the goods

[2] $C_1 = 2\left\{T(T-1)\right\}^{-1}$, $B_2 = -3(T+1)^{-1}, C_2 = \left\{3(17T^2 - 20T + 17)\right\}\left\{5(T-1)(T+1)^3\right\}^{-1}$, $B_3 = -15\left\{2(T+1)\right\}^{-1}, C_3 = \left\{15(193T^2 - 728T + 1147)\right\}\left\{112(T+2)^3(T-2)\right\}^{-1}$.

in our sample (55%). Overall, these results are more favorable to the proposition of long run price convergence than much of the existing literature, though not so different when compared to results that utilize panels. The fact that we are using micro-data on goods that are comparable across locations as well as the fact that we normalize to the cross-sectional mean are likely to be additional important factors that distinguish our results from those in existing work. This evidence is reassuring from a theoretical perspective since stationarity is a weak property for international prices to satisfy, a much stronger condition is that the Law-of-One-Price holds exactly in every time period, an issue to which we now turn.

4.2. Mean Real Exchange Rates

Table 4 reports the mean real exchange rate (average across goods) for beginning and ending points of our sample period, 1990 and 2000. Given our normalization the means are zero when we present statistics across all international cities or all intranational (U.S.) cities. Our focus is on the conditional means by region, income level and the type of good.

At the regional level, we find a tendency for PPP deviations to fall between 1990 and 2000. Price levels, as we define them, range from 21% below the world average in Africa to 9% above the world average in North America in 2000. The variation is somewhat greater in 1990 with the Central/South America region having an average price level 33% below the world average and Europe being 21% above the average.

One limitation of using regional aggregates such as these is that within most regions there are wide income disparities; which in light of the pioneering work of Irving B. Kravis and Robert E. Lipsey (1983) suggests that we may be averaging out a great deal of the cross-country variation. The lower panel of the table breaks the regions into high and low-to-middle income regions to address this issue.[3] The variation in PPP is now obvious. For example, while Asia as a whole is close to the world mean, the high income countries within that region have, on average, price levels 31% above the world mean while the low-to-middle income countries in Asia have price levels averaging 26% below the world mean. Europe, the only other region with countries falling into both income groupings has a price level for the high income region 41% above the world mean while the low income region has a price level 27% below the world mean. Thus, conditioning on income levels is a key element in international price dispersion, at least in terms of the average good.[4]

Focusing on the relative price levels disaggregated by income level, the changes in relative prices over the decade of the 1990's we do not see strong evidence of price level convergence

[3]The classification of countries into these income groups is based on GNP per capita in Table 1. Basic Indicators in the 1992 World Development Report (the GNP numbers are as of 1990).

[4]In contrast to what we find, the price level differences across rich and poor countries found using the International Price Comparison project data are on the order of a factor of 3 (rich country price levels are about three times the price level in poor countries). Without delving into the details of how goods are selected into the ICP construction, it is difficult to do anything more than speculate about the potential sources of the differences. The two most obvious to us are: i) the fact that the EIU data focus on urban as opposed to rural price data and ii) the goods selected into the EIU survey are designed to give international corporations a basis for compensating their employees in the major cities in which they work around the globe. The basket that emerges is not representative of what the average person in the country of origin is likely to purchase.

or divergence. Using the means that condition on income levels, 3 groups are diverging from the mean while 5 are converging toward the mean; in terms of the country-level data the divergent group involves 25 countries while the convergent group involves 65 countries which is a bit more favorable to the price level convergence hypothesis.

The last four columns of Table 4 look at the same measures but using only traded or non-traded goods in the computations. We would expect non-traded goods to exhibit larger deviations from the Law-of-One-Price than traded goods. We find strong evidence of this: 28 of the 32 measures have larger absolute values for non-traded goods than traded goods. One would not want to interpret the quantitative importance of the deviations using these numbers because the price difference tend to average out across goods. However, the ranking would hold up so long as this tendency was no more severe in traded goods than non-traded goods. To get a clearer sense of dispersion we compute standard deviations of prices across goods and locations within each of the bins of Table 4.

Table 5 reports these findings and includes the U.S. for comparison purposes (the U.S. was not included in Table 4 for the simple reason that the means are zero, by construction). Looking at the geographic dispersion in micro-real exchange rates pooling across all goods we see that the differences are not particularly great across regions or when comparing 1990 to 2000. We do see some tendency for price dispersion to be lower across high income countries than across low income countries. For example, the standard deviation of real exchange rates averages 0.452 for high-income European countries compared to 0.617 for low- or middle-income European countries (in 2000).

More obvious and perhaps more interesting are the differences when we condition on tradeability. Pooling all international regions and goods the standard deviation of real exchange rates is 0.626 for traded goods and 0.855 for non-traded goods (in 1990). Thus we see larger deviations from the Law-of-One-Price for non-traded goods as we would expect. Two numbers that stand out in Table 5 relate to the high income European group where dispersion in traded goods drops from 0.367 in 1990 to 0.107 in 2000. Presumably this has to do with policies promoting European Union economic integration.

Traded good price dispersion is lower than non-traded good price dispersion in both in the international and intranational data as one would expect, but the dispersion of prices in the U.S. is strictly less than that observed internationally even controlling for the type of good under consideration. The dispersion prices across locations (within the U.S.) for non-traded goods is 0.344 in 1990 significantly lower even than the dispersion of traded goods prices internationally which has a dispersion of 0.626.

The ratio of variance is between 5 and 6 when moving from intranational price dispersion to international price dispersion. In other words, if we take the variance of traded goods prices in the year 2000 for both the U.S. and average internationally, the geographic variance of prices within the U.S. was about 5.5% while that observed internationally was 27.8%. To place these numbers in perspective, if each city specialized in the production of a particular good and trade took place across all locations, these variances would be consistent with transportation costs of 5.5% and 27.8%, respectively. For non-traded goods the corresponding numbers would be 10% and 60%, which seems to indicate that we would have to go beyond transport costs and recognize the role of non-traded inputs and other factors that contribute to price dispersion to account for the international data.

Recall that our retail pricing model predicts that greater wage dispersion across locations would give rise to greater retail price dispersion across locations, but with the effect magnified for retail goods which used larger shares of non-traded inputs (such as labor). The broad features of the dispersion measures in Table 5 are consistent with this prediction.

4.3. Real Exchange Rate Dynamics

Based on the unit root tests, we found the evidence against nonstationarity of Law-of-One-Price deviations. This implies the possibility of either absolute or conditional convergence of the price deviations. To compare the two classes of price convergence, we consider the estimation of AR(1) models with and without good-specific city effects. All the models are estimated for each good separately. If the Law-of-One-Price holds in the long-run for each good then absolute convergence is said to prevail and we consider the following specification:

$$q_{it} = \mu + \rho q_{i,t-1} + v_{it}. \tag{4.1}$$

To estimate this model, we employ the least squares pooled estimator ($\widehat{\rho}_{POOL}$) where it is understood that pooling is across locations which are indexed by i and not goods (Goods are indexed by j therefore we again drop the superscript for goods).

Alternatively, long run LOP may fail to hold for each good, then conditional convergence is said to occur and the following AR(1) model with unobserved individual-specific effects is appropriate:

$$q_{it} = \rho q_{i,t-1} + \eta_i + v_{it} \tag{4.2}$$

for $i = 1, ..., N$ and $t = 2, ..., T$, and $|\rho| < 1$. We assume that individual effect η_i and the time-varying error v_{it} are independently distributed across i and

$$E(\eta_i) = 0, \quad E(v_{it}) = 0, \quad \text{and} \quad E(v_{it}\eta_i) = 0 \text{ for } i = 1, ..., N \text{ and } t = 2, ..., T.$$

We allow heterosckedasticity of v_{it}, but exclude the possibility of serial correlation

$$E(v_{it}v_{is}) = 0 \quad \text{for } i = 1, ..., N \text{ and for all } t \neq s.$$

For the initial conditions q_{i1}, we assume

$$E(q_{i1}v_{it}) = 0 \quad \text{for } i = 1, ..., N \text{ and } t = 2, ..., T.$$

One way of estimating this model is to use the LSDV estimator $\widehat{\rho}_{LSDV}$ which is known to be consistent with large T. However, with the short time span of our data, the presence of bias in this estimator should be taken into consideration when interpreting the results. To avoid this inconsistency issue, it is common practice to use the GMM estimator ($\widehat{\rho}_{GMM}$) based on the first difference transformation,

$$q_{it} - q_{i,t-1} = \rho\left(q_{i,t-1} - q_{i,t-2}\right) + \left(v_{it} - v_{i,t-1}\right) \quad \text{for } t = 3, ..., T,$$

with instruments selected from the orthogonality condition,

$$E\left[q_{is}\left(v_{it}-v_{i,t-1}\right)\right]=0 \quad \text{for } s=1,...,t-2 \text{ and } t=3,...,T.$$

Despite the theoretical superiority, simulation results sometimes show inaccuracy of the first-differenced GMM estimator (see Jan F. Kiviet, 1995 for example). On balance, we report both LSDV and GMM estimates of the autoregressive parameter ρ, but we leave the bias issue in LSDV estimator to the discussion section.

The estimated persistence parameters for both (4.1) and (4.2) based on least squares method are reported in Table 6. Beginning with the international estimates we see that the average autoregressive coefficient estimate under the absolute convergence hypothesis (or $\widehat{\rho}_{POOL}$) is equal to 0.88, indicating substantial persistence in real exchange rates; similar to what is found in the literature using aggregate price indices. The standard errors are small and the confidence intervals have upper-bounds typically below unity, consistent with our prior rejection of unit roots. The persistence parameters range from a low of 0.73 to a high of 0.99, across individual goods.

Turning to the estimates that allow for city-effects (or $\widehat{\rho}_{LSDV}$), the results are quite different. Now the average autoregressive coefficient is 0.43 with individual estimates ranging from 0.15 to 0.86. The standard errors are on average twice that of the pooled estimates, but the estimates under absolute and conditional convergence are statistically significantly different from one another. This point can be seen more in detail in Figure 4A which shows empirical distribution of $\widehat{\rho}_{POOL}$ and $\widehat{\rho}_{LSDV}$ with their two standard error bands. Obviously, the two estimates are economically distinct with the half-lives under absolute convergence averaging 5.4 years compared to a mere 10 months under the conditional convergence specification.

Contrary to our expectations, the persistence of Law-of-One-Price deviations across U.S. cities are comparable to what we find internationally. The autoregressive coefficients average 0.81 based on absolute price convergence estimates ($\widehat{\rho}_{POOL}$) and 0.53 based on conditional price convergence estimates ($\widehat{\rho}_{LSDV}$). Empirical distribution of both estimates with error bands for intranational case are shown in Figure 4B. The range of parameter estimates is 0.39 to 1.0 for $\widehat{\rho}_{POOL}$ and 0.05 to 0.91 for $\widehat{\rho}_{LSDV}$; standard errors for both specifications are larger than what we find in the international case, likely due to the smaller number of cities available in the U.S. panels compared to the international panels. The half lives for deviations from the Law-of-One-Price across U.S. cities are 3.3 years under absolute convergence compared to 13 months under conditional convergence.

To investigate the plausibility of two alternative specifications, we conduct the F test for the joint significance of the dummy variables used for $\widehat{\rho}_{LSDV}$. Such a test can be viewed as a test for the null of absolute convergence against the alternative of conditional convergence. The results for the F test are also reported in Table 6. The no individual effect hypothesis is significantly rejected for most international micro-prices differences, but not rejected for about half of the goods for intranational case.. This result suggests that the data is consistent with the conditional convergence hypothesis along with the theory of retail price at least for international case.

The econometric explanation for the differences in the persistence estimates under absolute and conditional convergence is straight-forward. Figure 5 sketches representative regression lines, location-by-location for a particular good with and without city-effects. The thin flat lines are what one estimates under conditional convergence, where a constant term in each regression relationship allows each good to converge to a unique, city-specific, mean. The bold line is what is estimated under the assumption of absolute convergence, which fits a single regression line to the entire panel of observations for a particular good. Stationarity of the individual real exchange rates about their location-specific means ensures that most sample points lie close to the 45° line so the slope parameter (persistence) is biased upward under absolute convergence. This argument implies that in forcing goods to satisfy the Law-of-One-Price in the long run when the sample means are far from that prediction and persistence at the micro-level is low, a severe upward bias in persistence of the deviations results.

Table 7 reports analogous results for absolute and conditional convergence based on the GMM estimator. In addition to the first-differenced GMM estimates ($\widehat{\rho}_{GMM}$) under the conditional convergence specification, we also report the GMM estimates ($\widehat{\rho}^{*}_{GMM}$) under the absolute convergence which is the GMM counterpart of $\widehat{\rho}_{POOL}$. For both $\widehat{\rho}^{*}_{GMM}$ and $\widehat{\rho}_{GMM}$, the numbers are generally very close to what we get using the least squares method. One key difference, though, is the lower average persistence measure in the intranational (U.S.) estimates under the conditional convergence specification. The consequence of this lower estimate is that the average persistence in the international and intranational data under conditional convergence is now statistically indistinguishable. Recall, that using the LSDV estimator ($\widehat{\rho}_{LSDV}$) it appeared that under conditional convergence, the U.S. deviations were actually more persistent than the international deviations. Thus in checking the robustness of our results using an alternative estimator we discount the potentially counter intuitive result that Law-of-One-Price deviations are less persistent internationally than intranationally. However, the issue of finite sample bias remains and we leave it to the next section to determine if we might be more definitive on the ranking of international and intranational price persistence. Similar to the least square case, we can conduct the test of no individual effects with GMM estimator as a test for the absolute price convergence against the conditional price convergence. The results for this test are also reported in Table 7. Just like the least squares case, the absolute convergence hypothesis is significantly rejected with international data. However, for all the case, we could not reject the same null hypothesis with intranational data with any conventional significance level.

The picture of relative price adjustment that emerges from our analysis of micro-data is that price adjustment is very rapid both across cities within countries and across cities of the world. In terms of persistence of Law-of-One-Price deviations, national borders appear not to matter. What national borders do matter for are the magnitudes of the long run deviations from the Law-of-One-Price. Across U.S. cities the deviations are economically small while across international cities the deviations are economically large. As a consequence, much of the dispersion we observe in Figures 1 and 2 is not the result of stochastic fluctuations in relative prices over time but rather reflects permanent (at least as permanent as our sample length of 10 years will allow us to estimate) deviations from the Law-of-One-Price.

5. Discussion

5.1. Aggregation

We estimate city price levels using average of log-prices: $P_{it} \equiv \frac{1}{M}\sum_{j=1}^{M} P_{it}^{j}$ and recalling that all prices were already in U.S. dollar, we construct aggregate real exchange rate in the standard way: $\overline{q}_{it} \equiv P_{it} - \overline{P}_t$ (standard except that we use the average price level in the world as the numeraire price.

Note that if the persistence parameters are the same across goods, then the aggregate real exchange rate will have the same persistence as the micro-real exchange rates. To determine if this is the case, we estimate the following regressions:

$$\begin{array}{ll} \overline{q}_{it} = \rho\overline{q}_{i,t-1} + \overline{\eta}_i + \overline{v}_{it} & \text{conditional convergence} \\ \overline{q}_{it} = \rho\overline{q}_{i,t-1} + \overline{v}_{it} & \text{absolute convergence.} \end{array}$$

where the bars over the variables indicate averages across goods, location-by-location. Estimating this specification for the aggregated real exchange rate we get parameter estimates of 0.89 and 0.34 for the international data under absolute and conditional convergence, respectively. Recall the averages of the analogous estimates at the micro-level were 0.88 and 0.43, respectively. Thus it appears that the our persistence results are preserved under this form of aggregation and by implication that the high persistence obtained elsewhere in the literature using CPI data appears not to be a consequence of simple aggregation bias.[5]

5.2. Weighting Method

Another possible explanation for the persistence of deviations found in aggregate studies is that consumption weights differ across countries i and k. Consider two goods and two locations and suppose that the price indices are:

$$\begin{aligned} P_{it} &= (P_{it}^{1})^{\alpha_i}(P_{it}^{2})^{(1-\alpha_i)} \\ \log(P_{it}/P_{kt}) &= [\alpha_i(p_{it}^{1} - p_{it}^{2}) - \alpha_k(p_{kt}^{1} - p_{kt}^{2})] + (p_{it}^{2} - p_{kt}^{2}) \end{aligned}$$

where lower case variables refer to logarithms. Now suppose prices actually obey the law of one price exactly so that $p_{it}^{j} = p_{kt}^{j}\ \forall i$, then we have:

$$\log(P_{it}/P_{kt}) = (\alpha_i - \alpha_k)(p_t^1 - p_t^2)$$

The implication is that unless the price of good 1 relative to good 2 is constant over time, national price indices will diverge from the PPP prediction even when the Law-of-One-Price holds exactly. Given that relative price levels are quite different across countries the expenditure shares might be expected to be different. Thus movements in the price of one good relative to another that is common to all countries (e.g. computer prices relative to oil prices) would move price indices by different amounts across countries due to different consumption weighting schemes. Given that these relative price movements are likely to be persistent, considerable autocorrelation might result.

[5] The issue of temporal aggregation has been discussed in Alan M. Taylor (2001).

Turning the argument on its head, we know that the goods included in the CPI are different in each country. Even if the consumption weights on various categories such as food, consumption and shelter were identical across countries the 'goods' in these bundles are often very different and the above argument goes through.

5.3. Bias

In the main section, we presented the results from LSDV and GMM estimation. For the LSDV estimator with fixed T, Nickell (1981) derived a formula for the asymptotic bias when N tends to infinity. The bias formula is given by,

$$\begin{aligned}\underset{N\to\infty}{\mathrm{plim}}\left(\widehat{\rho}_{LSDV}-\rho\right) &= \frac{-(1+\rho)}{T-1}\frac{A}{\left\{1-\frac{2\rho A}{(1-\rho)(T-1)}\right\}} \\ \text{where } A &\equiv 1-\frac{1}{T}\frac{(1-\rho^{T})}{(1-\rho)}\end{aligned}$$

Therefore, the inconsistency with fixed T results from the asymptotic bias of order $O(T^{-1})$ which is always negative given $0<\rho<1$. In contrast, the GMM estimator does not suffer from this type of asymptotic bias (and is therefore consistent) while the finite sample bias due to a fixed N may still be present. The bias formula has an approximation error of $O(N^{-1}T^{-3/2})$ and when the bias is corrected, the LSDV estimator is consistent not just for large T but also for finite T with large N. Kiviet (1995) suggests correcting the LSDV estimators with a consistent bias estimation method. Following Kiviet, we report the bias corrected LSDV estimator where the bias of original LSDV estimator $\widehat{\rho}_{LSDV}$ is evaluated using a consistent estimator of persistence, namely the GMM estimator $\widehat{\rho}_{GMM}$. Unlike Kiviet's case, it should be noted that this bias formula does not require the estimation of error variance since we do not have exogenous variables in the autoregressive model.

Table 8 reports the results of bias adjustment along with the original LSDV and GMM estimates. The bias corrected persistence parameter averages 0.58 in the international data, compared to 0.68 for the intranational data. While the bias correction for the mean of the estimated parameters is about 0.10 in both cases, the fact that the LSDV estimates start out at a higher level implies that the bias-corrected estimated remain higher in the intranational data. In terms of estimated half-lives, the average for the international cities is 15 months compared to about 21 months across U.S. cities. While there may be a statistically significant difference here, the main implication of the paper – that conditional convergence is much faster than absolute convergence in both the international and intranational context – is firmly upheld.

Figures 6A and 6B present detailed good-by-good persistence measures for the international and intranational panels. In both figures, the x-axis measures the persistence estimated under absolute convergence while the y-axis measures the persistence parameters under conditional convergence. We see that in each case the good-by-good estimates are highly positively correlated – a regression line through the two scatters yields an almost identical slope of 1.1. Thus the differences between international and intranational persistence is difficult to detect at the microeconomic level. The only substantive difference across the two figures is the center of the scatter-plots which represent the mean bias-adjusted LSDV persistence estimates which were already discussed in relation to Table 8.

5.4. Conclusions and Implications for Future Work

Our results appear to go part way toward resolving the Purchasing Power Parity Puzzle in the following sense. Viewed from the perspective of our microeconomic evidence, the persistence of real exchange rates is an artifact of mis-specification of the long run distribution of relative prices. When conditional convergence is assumed, persistence is low, easily within the range typically considered plausible for models with nominal rigidities.

The notion expressed by Obstfeld and Rogoff (in describing the PPP puzzle) that real factors induce persistence but not large changes in relative prices was another piece of the puzzle. Taken literally, our results indicate the existence of large and permanent deviations from the Law-of-One-Price in the international data. If we interpret these permanent differences as real factors – as many prominent theories, including the Balassa-Samuelson (1964) hypothesis would suggest – they are very important quantitatively. Moreover, the changes in relative prices over the business cycle due to real shocks would fall into the same category as the nominal shocks based on our empirical model. Ironically, real shocks may deliver too much persistence in the deviations!

In studying micro-economic deviations from the Law-of-One-Price we do raise an important empirical puzzle; namely the inconsistency between the persistence of good-specific deviations and deviations as measured by the CPI. We have already noted that our evidence holds up under a simple aggregation scheme (an equally weighted average of individual real exchange rates) so the difference between our finding and that in the aggregate studies does not appear to be due to aggregation.

We plan to explore a number of issues in future work. First, we plan to investigate the robustness of our aggregation results to alternative weighting methodologies – expenditure weights versus equal weighting, for example. Second, we will consider the issue of using a numeraire in the analysis as opposed to our numeraire independent measure of the real exchange rate. David Papell (1997) has emphasized that much of the existing literature has used the U.S. dollar as the numeraire and the persistence of exchange rates may be a reflection of movements in the bilateral dollar exchange rates and not necessarily movements in relative prices across other bilateral pairings. Last, we intend to conduct a more comprehensive review of the comparability of the baskets that are priced in the CPI. A major difference between the types of data found in the EIU panel (as well as that studied by Crucini, Telmer and Zachariadis (2001)) is that the CPI is intended to maximize comparability of goods over time while the cross-sectional panels attempt to maximize comparability across locations. That these two methodologies lead to different empirical implications regarding real exchange rate behavior at the disaggregate and aggregate levels may turn out to be a function of the survey design. We believe that these avenues will help us to reconcile the microeconomic evidence on Law-of-One-Prices uncovered here with the broader literature that utilizes price indices.

Appendix

A.1. Data Source

The data are collected in the *Worldwide Cost of Living Survey* compiled by the *Economist Intelligence Unit*. The intent of the survey is to provide human resource managers with data on which to base compensation decisions across countries. The survey covers 122 cities and 301 goods or services. Many goods prices are collected in different types of retail outlets. In our analysis we use only large-scale retail outlets to avoid have goods appear more than once in our analysis.

A.2. GMM Estimation

Under the assumptions in the main text, the first-differenced GMM estimator of AR coefficient ρ based on $m = (T-1)(T-2)/2$ total moment conditions can be written as

$$\widehat{\rho}_{GMM} = (\mathbf{X}'\mathbf{Z}\widehat{\mathbf{W}}_N\mathbf{Z}'\mathbf{X})^{-1}\mathbf{X}'\mathbf{Z}\widehat{\mathbf{W}}_N\mathbf{Z}'\mathbf{Y}$$

where $\mathbf{Z}' = (\mathbf{Z}_1', \mathbf{Z}_2', ..., \mathbf{Z}_N')$ is the $m \times N(T-2)$ matrix with

$$\mathbf{Z}_i = \begin{bmatrix} q_{i1} & 0 & 0 & \cdots & 0 & \cdots & 0 \\ 0 & q_{i1} & q_{i2} & \cdots & 0 & \cdots & 0 \\ \vdots & \vdots & \vdots & & \vdots & & \vdots \\ 0 & 0 & 0 & \cdots & q_{i1} & \cdots & q_{iT-2} \end{bmatrix},$$

$\mathbf{Y}' = (\mathbf{\Delta q}_1', \mathbf{\Delta q}_2', ..., \mathbf{\Delta q}_N')$ is the $N(T-2)$ vector with

$$\mathbf{\Delta q}_i = (\Delta q_{i3}, \Delta q_{i4}, ..., \Delta q_{i,T})',$$

$\mathbf{X}' = (\mathbf{\Delta q}_{1,-1}', \mathbf{\Delta q}_{2,-1}', ..., \mathbf{\Delta q}_{N,-1}')$ is the $N(T-2)$ vector with

$$\mathbf{\Delta q}_{i,-1} = (\Delta q_{i2}, \Delta q_{i3}, ..., \Delta q_{i,T-1})',$$

and $\widehat{\mathbf{W}}_N = \mathbf{S}_N^{-1}$ is an optimal weighting matrix. Following Arellano and Bond (1991), we employ

$$\mathbf{S}_N = N^{-1}\sum_{i=1}^{N}\mathbf{Z}_i'\mathbf{H}\mathbf{Z}_i^{-1}$$

where

$$\mathbf{H} = \begin{bmatrix} 2 & -1 & 0 & \cdots & 0 \\ -1 & 2 & 0 & \cdots & 0 \\ 0 & -1 & 2 & \cdots & 0 \\ \vdots & \vdots & \vdots & & \vdots \\ 0 & 0 & 0 & \cdots & 2 \end{bmatrix}$$

for the first step estimator. For the second step estimation, we use

$$\mathbf{S}_N = N^{-1}\sum_{i=1}^{N}\mathbf{Z}_i'\widehat{u}_i\widehat{u}_i'\mathbf{Z}_i^{-1}$$

where $\widehat{u}_i$ are residual vectors from the first step estimator.

For the GMM estimation without individual (city) effects, $T-1$ additional moment conditions are available since total of $m^* = T(T-1)/2$ moment conditions are implied by

$$E\left[q_{is} v_{it}\right] = 0 \quad \text{for } s = 1, ..., t-1 \text{ and } t = 2, ..., T.$$

The GMM estimator without individual effects is given by

$$\widehat{\rho}^*_{GMM} = (\mathbf{X}^{*\prime}\mathbf{Z}^*\widehat{\mathbf{W}}^*_N\mathbf{Z}^{*\prime}\mathbf{X}^*)^{-1}\mathbf{X}^{*\prime}\mathbf{Z}^*\widehat{\mathbf{W}}^*_N\mathbf{Z}^{*\prime}\mathbf{Y}^*$$

where $\mathbf{Z}^{*\prime} = (\mathbf{Z}^{*\prime}_1, \mathbf{Z}^{*\prime}_2, ..., \mathbf{Z}^{*\prime}_N)$ is the $m^* \times N(T-2) + (T-1)$ matrix with

$$\mathbf{Z}^*_i = \begin{bmatrix} \mathbf{Z}_i & 0 & \cdots & 0 \\ 0 & q_{i1} & \cdots & 0 \\ \vdots & \vdots & & \vdots \\ 0 & 0 & \cdots & q_{iT-1} \end{bmatrix},$$

$\mathbf{Y}^{*\prime} = (\mathbf{\Delta q}^{*\prime}_1, \mathbf{\Delta q}^{*\prime}_2, ..., \mathbf{\Delta q}^{*\prime}_N)$ is the $N(T-2) + (T-1)$ vector with

$$\mathbf{\Delta q}^*_i = (\mathbf{\Delta q}'_i, q_{i2}, ..., q_{i,T})',$$

$\mathbf{X}^{*\prime} = (\mathbf{\Delta q}^{*\prime}_{1,-1}, \mathbf{\Delta q}^{*\prime}_{2,-1}, ..., \mathbf{\Delta q}^{*\prime}_{N,-1})$ is the $N(T-2) + (T-1)$ vector with

$$\mathbf{\Delta q}^*_{i,-1} = (\mathbf{\Delta q}'_{i,-1}, q_{i1}, ..., q_{i,T-1})',$$

and $\widehat{\mathbf{W}}^*_N = \mathbf{S}^{*-1}_N$ is an optimal weighting matrix. Test statistic for the null hypothesis of no individual effects can be constructed based on the test of the validity of $T-1$ additional restrictions (Holtz-Eakin, 1988). Under the null hypothesis,

$$L = J^* - J$$

where J^* is the criterion function for $\widehat{\rho}^*_{GMM}$ and J is the criterion function for $\widehat{\rho}_{GMM}$ with weighting matrix obtained from the submatrix of $\mathbf{S}^*_N$, follows chi-squared distribution with $T-1$ degree of freedom.

References

[1] Arellano, Manuel, and Stephen Bond, 1991, "Some tests of specification for panel data: Monte Carlo evidence and an application to employment equations," *Review of Economic Studies* 58, 277-297.

[2] Balassa, Bela, 1964, "The purchasing power parity doctrine: A reappraisal," *Journal of Political Economy* 72, 584-596.

[3] Burstein, Ariel, Joao Neves, and Sergio Rebelo, 2001, "Distribution costs and real exchange rate dynamics during exchange-rate-based stabilizations," *Journal of Monetary Economics*, forthcoming.

[4] Cecchetti, Stephen, Nelson Mark and Robert Sonora, 1999, "Price level convergence among United States cities: lessons for the European Central Bank," Ohio State University, mimeograph.

[5] Crownover, Collin, John Pippenger and Douglas G. Steigerwald, 1996, "Testing for absolute purchasing power parity," *Journal of International Money and Finance* 15, 783-796.

[6] Crucini, Mario J., Chris I. Telmer and Marios Zachariadis, 2001, "Understanding European real exchange rates," Vanderbilt University Working Paper No. 01-20.

[7] Cumby, Robert, 1996, "Forecasting exchange rates and relative prices: Is what you want what you get with McParity?" NBER Working Paper No. 5675, July.

[8] Debreu, Gerard, 1959, The theory of value, New Haven, CT: Yale University Press.

[9] Engel, Charles, 1993, "Real exchange rates and relative prices: an empirical investigation," *Journal of Monetary Economics* 32, 35-50.

[10] Engel, Charles and John H. Rogers, 1996, "How wide is the border?" *American Economic Review* 86, 1112-1125.

[11] Froot, Kenneth A. Michael Kim and Kenneth Rogoff, 1995, "The law of one price over 700 years," National Bureau of Economic Research Working Paper No. 5132.

[12] Ghosh Atish R. and Holger C. Wolf, 1994, "Pricing in international markets: lessons from the Economist," NBER Working Paper No. 4806, July.

[13] Giovannini, Alberto, 1988, "Exchange rates and traded goods prices," *Journal of International Economics* 24, 45-68.

[14] Harris, Richard D.F. and Elias Tzavalis, 1999, "Inference for unit roots in dynamic panels where the time dimension is fixed," *Journal of Econometrics*, 91, 201-226.

[15] Haskel, Jonathan and Holger C. Wolf, 2000, "The law of one price - a case study," mimeograph, August.

[16] Holtz-Eakin, Douglas, 1988, "Testing for individual effects in autoregressive models," *Journal of Econometrics* 39, 297-307.

[17] Isard, Peter, 1977, "How far can we push the Law of One Price," *American Economic Review* 67(3), 942-948.

[18] Kiviet, Jan. F., 1995, "On bias, inconsistency, and efficiency of various estimators in dynamic panel data models," *Journal of Econometrics* 68, 53-78.

[19] Kravis, Irving B. and Robert E. Lipsey, 1983, "Toward an explanation of national price levels," Princeton Studies in International Finance, No. 52, Princeton: Princeton University.

[20] Lutz, Mattias, 2001, "Pricing in segmented markets, arbitrage barriers and the law of one price: evidence from the European car market," University of St. Gallen, mimeograph, January.

[21] Obstfeld, Maurice and Kenneth Rogoff, 1998, "Open-economy macroeconomics: Developments in theory and policy, *Scandinavian Journal of Economics* 100, 247-275.

[22] Obstfeld, Maurice and Kenneth Rogoff, 2000, "New directions for stochastic open economy models," *Journal of International Economics* 50, 117-153.

[23] O'Connell, Paul G. J., 1998. "The overvaluation of purchasing power parity." *Journal of International Economics* 44, 1-19.

[24] Nickell, Stephen, 1981, "Biases in dynamic models with fixed effects," *Econometrica* 49, 1417-1426.

[25] Papell, David 1997, "Searching for stationarity: Purchasing power parity under the current float," *Journal of International Economics* 43, 313-332.

[26] Parsley, David and Shang-Jin Wei, 1996, "Convergence to the law of one price without trade barriers or currency fluctuations," *Quarterly Journal of Economics* 61, 1211-1236.

[27] Rogers, John H., 2001, "Price level convergence, relative prices, and inflation in Europe," *Federal Reserve Board International Finance Discussion Paper*, No. 699.

[28] Rogers, John H. and Michael Jenkins, 1995, "Haircuts or hysteresis? Sources of movements in real exchange rates," *Journal of International Economics* 38, 339-360.

[29] Rogoff, Kenneth 1996, "The purchasing power parity puzzle," *Journal of Economic Literature* 34(2), 647-668.

[30] Samuelson, Paul A., 1964, "Theoretical notes on trade problems," *Review of Economics and Statistics* 46, 145-154.

[31] Taylor, Alan M., 2001. "Potential pitfalls for the purchasing-power-parity puzzle? Sampling and specification biases in mean-reversion tests of the law of one price." *Econometrica* 69, 473-98.

TABLE 1 – LOCATIONS

City, Country (No. of goods)	City, Country (No. of goods)	City, Country (No. of goods)
(1) International Data	Vienna, Austria (263)	Adelaide, Australia (251)
Bahrain, Bahrain (230)	Brussels, Belgium (263)	Brisbane, Australia (12)
Dhaka, Bangladesh (133)	Prague, Czech (188)	Melbourne, Australia (2)
Beijing, China (144)	Copenhagen, Denmark (264)	Perth, Australia (2)
Hong Kong, Hong Kong (242)	Helsinki, Finland (255)	Sydney, Australia (2)
New Delhi, India (57)	Lyon, France (261)	Auckland, New Zealand (257)
Mumbai, India (146)	Paris, France (7)	Wellington, New Zealand (5)
Jakarta, Indonesia (183)	Berlin, Germany (265)	
Tehran, Iran (181)	Dusseldorf, Germany (5)	San Jose, Costa Rica (230)
Tel Aviv, Israel (255)	Athens, Greece (247)	Guatemala City, Guatemala (221)
Osaka Kobe, Japan (244)	Budapest, Hungary (255)	Mexico City, Mexico (238)
Tokyo, Japan (7)	Dublin, Ireland (248)	Panama City, Panama (242)
Amman, Jordan (137)	Milan, Italy (263)	
Seoul, Korea (167)	Rome, Italy (5)	Calgary, Canada (250)
Kuala Lumpur,Malaysia (244)	Luxembourg, Luxembourg (260)	Montreal, Canada (15)
Karachi, Pakistan (192)	Amsterdam, Netherlands (260)	Toronto, Canada (3)
Manila,Philippines (211)	Oslo, Norway (233)	Atlanta, USA (249)
Al Khobar, Saudi Arabia (203)	Warsaw, Poland (232)	Boston, USA (11)
Jeddah, Saudi Arabia (17)	Lisbon, Portugal (267)	Chicago, USA (5)
Singapore, Singapore (256)	Bucharest, Romania (1)	Cleveland, USA (3)
Colombo, Sri Lanka (212)	Moscow, Russia (116)	New York, USA (1)
Taipei, Taiwan (215)	Barcelona, Spain (268)	
Bangkok, Thailand (257)	Stockholm, Sweden (252)	
Abu Dhabi, UAE (238)	Geneva, Switzerland (262)	*(2) Intranational Data*
Dubai, UAE (11)	Zurich, Switzerland (6)	Atlanta, USA (248)
	Istanbul, Turkey (253)	Boston, USA (257)
Abidjan, Cote dIvoire (242)	London, UK (261)	Chicago, USA (251)
Cairo, Egypt (197)	Belgrade, Yugoslavia (105)	Cleveland, USA (249)
Nairobi, Kenya (233)		Detroit, USA (260)
Tripoli, Libya (51)	Buenos Aires, Argentina (253)	Houston, USA (250)
Casa Blanca, Morocco (199)	Sao Paulo, Brazil (255)	Los Angeles, USA (248)
Lagos, Nigeria (204)	Santiago, Chile (257)	Miami, USA (253)
Dakar, Senegal (197)	Bogota, Columbia (235)	New York, USA (234)
Johannesburg, South Africa (253)	Quito, Ecuador (177)	Pittsburgh, USA (235)
Tunis, Tunisia (186)	Asuncion, Paraguay (250)	San Francisco, USA (230)
Harare, Zimbabwe (200)	Lima, Peru (1)	Seattle, USA (252)
	Monte Video, Uruguay (257)	Washington DC, USA (255)
	Caracas, Venezuela (238)	

Note: Entries are the city in which the price data are collected, the country to which the city belongs and the number of goods in the analysis for which that city is used

TABLE 2 – GOODS BY CONSUMPTION CATEGORY

EIU Category	No. of goods	BLS Category	No. of goods
Food and beverages	112		
Alcohol	20		
Tobacco	5		
Category total	137	Food and beverages	73
Utilities	6	Fuel and Utilities	12
Household goods	26		
Shelter	19		
Category total	45	Household services and furnishings	66
Clothing	32	Apparel and upkeep	47
Transportation	18	Transportation	34
		Medical care	18
Recreation	12	Entertainment	27
Domestic help	3		
Personal care	16		
Education	22		
Business related	10		
Salaries	3		
Category total	54	Other commodities and services	21
Grand total	304	Grand total	298

TABLE 3 – SUMMARY RESULTS OF UNIT ROOT TESTS

Significance levels	No constant	Constant	Constant and trend
	Panel A: International estimates		
0.01	268 (99%)	270 (100%)	247 (91%)
0.05	268 (99%)	270 (100%)	259 (96%)
0.10	268 (99%)	270 (100%)	259 (96%)
	Panel B: Intranational estimates		
0.01	199 (81%)	181 (74%)	76 (31%)
0.05	217 (88%)	206 (84%)	113 (46%)
0.10	229 (93%)	221 (90%)	136 (55%)

Notes: Number of goods is 270 in Panel A and 245 in Panel B. Number in the table are the number of goods we reject the null hypothesis of a unit root. The numbers in parentheses are the numbers of rejection expressed as a percentage of the total number of goods, 270 or 245. Models with no constant, constant and constant/trend for individual (city) effects are estimated by least squares pooled estimator, least squares dummy variable estimator, and least squares dummy variable with trend estimator, respectively. See Harris and Tzavalis (1999).

TABLE 4 – MEAN REAL EXCHANGE RATES

	All goods		Traded goods		Non-traded goods	
Region (No. of cities)	1990	2000	1990	2000	1990	2000
All regions (90)	0.00	0.00	0.00	0.00	0.00	0.00
Africa (10)	-0.02	-0.21	0.01	-0.15	-0.13	-0.42
Asia (24)	-0.06	0.04	-0.04	0.07	-0.11	-0.05
Central/South America (13)	-0.33	-0.02	-0.32	-0.05	-0.35	0.05
Europe (28)	0.21	0.06	0.18	0.03	0.30	0.17
North America (8)	-0.17	0.09	-0.16	0.07	-0.20	0.13
Oceania (7)	0.07	-0.13	0.05	-0.12	0.12	-0.14
High income regions (42)	0.33	0.20	0.30	0.16	0.45	0.32
Africa (0)	n.a.	n.a.	n.a.	n.a.	n.a.	n.a.
Asia (8)	0.31	0.48	0.28	0.45	0.42	0.60
Central/South America (0)	n.a.	n.a.	n.a.	n.a.	n.a.	n.a.
Europe (19)	0.41	0.15	0.37	0.11	0.57	0.30
North America (8)	-0.17	0.09	-0.16	0.07	-0.20	0.13
Oceania (7)	0.07	-0.13	0.05	-0.12	0.12	-0.14
Low/middle income regions (48)	-0.23	-0.14	-0.21	-0.12	-0.30	-0.21
Africa (10)	-0.02	-0.21	0.01	-0.15	-0.13	-0.42
Asia (16)	-0.26	-0.20	-0.22	-0.14	-0.37	-0.38
Central/South America (13)	-0.33	-0.02	-0.32	-0.05	-0.35	0.05
Europe (9)	-0.27	-0.16	-0.26	-0.16	-0.31	-0.13
North America (0)	n.a.	n.a.	n.a.	n.a.	n.a.	n.a.
Oceania (0)	n.a.	n.a.	n.a.	n.a.	n.a.	n.a.

Note: These real exchange rates use the world average price.

TABLE 5 – DISPERSION IN REAL EXCHANGE RATES

Region (No. of cities)	All goods 1990	All goods 2000	Traded goods 1990	Traded goods 2000	Non-traded goods 1990	Non-traded goods 2000
All regions (90)	0.685	0.594	0.626	0.527	0.855	0.777
Africa (10)	0.659	0.596	0.634	0.533	0.724	0.727
Asia (24)	0.724	0.706	0.675	0.656	0.861	0.844
Central/South America (13)	0.637	0.540	0.540	0.479	0.893	0.704
Europe (28)	0.648	0.526	0.595	0.448	0.800	0.726
North America (8)	0.479	0.500	0.377	0.389	0.737	0.776
Oceania (7)	0.449	0.408	0.360	0.343	0.675	0.582
High income regions (42)	0.525	0.499	0.467	0.450	0.680	0.630
Africa (0)						
Asia (8)	0.568	0.521	0.520	0.490	0.699	0.602
Central/South America (0)	n.a.	n.a.	n.a.	n.a.	n.a.	n.a.
Europe (19)	0.500	0.452	0.367	0.107	0.625	0.572
North America (8)	0.479	0.500	0.377	0.389	0.737	0.776
Oceania (7)	0.449	0.408	0.360	0.343	0.675	0.582
Low/middle income regions (48)	0.689	0.614	0.640	0.547	0.826	0.791
Africa (10)	0.659	0.596	0.634	0.533	0.724	0.727
Asia (16)	0.721	0.679	0.685	0.642	0.812	0.754
Central/South America (13)	0.637	0.540	0.540	0.479	0.893	0.704
Europe (9)	0.703	0.617	0.664	0.493	0.821	0.923
North America (0)	n.a.	n.a.	n.a.	n.a.	n.a.	n.a.
Oceania (0)	n.a.	n.a.	n.a.	n.a.	n.a.	n.a.
U.S. cities (13)	0.281	0.254	0.262	0.234	0.344	0.318

TABLE 6 – SUMMARY RESULTS OF POOLED/LSDV ESTIMATES

	No. Cities	$\widehat{\rho}_{POOL}$	se	$\widehat{\rho}_{LSDV}$	se	F	*p-value*
			Panel A: International estimates				
median	60.000	0.889	0.014	0.430	0.033	3.082	0.000
mean	57.893	0.880	0.017	0.433	0.040	3.168	0.001
std	8.881	0.055	0.013	0.103	0.031	0.763	0.020
min	23.000	0.735	0.002	0.153	0.005	1.095	0.000
max	69.000	0.987	0.135	0.864	0.290	5.286	0.330
			Panel B: Intranational estimates				
median	13.000	0.833	0.034	0.542	0.051	2.003	0.032
mean	12.689	0.814	0.042	0.531	0.066	2.218	0.106
std	0.732	0.114	0.031	0.170	0.053	1.335	0.161
min	10.000	0.392	0.004	0.051	0.002	0.613	0.000
max	13.000	1.016	0.180	0.914	0.401	17.469	0.816

Notes: Number of goods is 270 in Panel A and 245 in Panel B. Pooled estimates ($\widehat{\rho}_{POOL}$) and LSDV estimates ($\widehat{\rho}_{LSDV}$) are reported with heterosckedasticity-consistent standard errors (se). F and *p-value* are the F test statistics for no individual (city) effects and their p-values, respectively.

TABLE 7 – SUMMARY RESULTS OF GMM ESTIMATES

	No. Cities	$\widehat{\rho}^*_{GMM}$	se	$\widehat{\rho}_{GMM}$	se	L	*p-value*
			Panel A: International estimates				
median	60.000	0.819	0.003	0.406	0.011	51.350	0.000
mean	57.893	0.811	0.004	0.404	0.012	49.613	0.005
std	8.881	0.077	0.002	0.137	0.005	8.181	0.064
min	23.000	0.567	0.000	-0.123	0.000	3.269	0.000
max	69.000	0.979	0.017	0.734	0.042	61.671	0.974
			Panel B: Intranational estimates				
median	13.000	0.790	0.050	0.445	0.114	11.451	0.323
mean	12.661	0.754	0.053	0.405	0.124	11.096	0.359
std	0.732	0.158	0.029	0.231	0.066	1.503	0.123
min	10.000	0.179	0.003	-0.251	0.011	3.185	0.229
max	13.000	1.019	0.174	0.910	0.383	12.911	0.977

Notes: Number of goods is 270 in Panel A and 245 in Panel B. Two-step GMM estimates without individual (city) effect ($\widehat{\rho}^*_{GMM}$) and two-step GMM estimates based on first difference ($\widehat{\rho}_{GMM}$) are reported with heterosckedasticity-consistent standard errors (se). L and *p-value* are the GMM-based test statistics for no individual (city) effects and their p-values, respectively.

TABLE 8 – SUMMARY RESULTS OF BIAS CORRECTIONS

	$\widehat{\rho}_{POOL}$	$\widehat{\rho}_{LSDV}$	$\widehat{\rho}_{GMM}$	$\widehat{bias}$	bias corrected $\widehat{\rho}_{LSDV}$
Panel A: International estimates					
median	0.889	0.430	0.406	-0.148	0.573
mean	0.880	0.433	0.404	-0.148	0.581
std	0.055	0.103	0.137	0.014	0.114
min	0.735	0.153	-0.123	-0.182	0.265
max	0.987	0.864	0.734	-0.092	1.012
Panel B: Intranational estimates					
median	0.833	0.540	0.445	-0.150	0.683
mean	0.814	0.532	0.405	-0.146	0.678
std	0.114	0.170	0.231	0.024	0.183
min	0.392	0.051	-0.251	-0.199	0.174
max	1.016	0.914	0.910	-0.078	1.109

Notes: Number of goods is 270 in Panel A and 245 in Panel B. Bias of LSDV estimator is estimated by GMM estimator.

Figure 1
International and Intranational Price (Individual Goods)

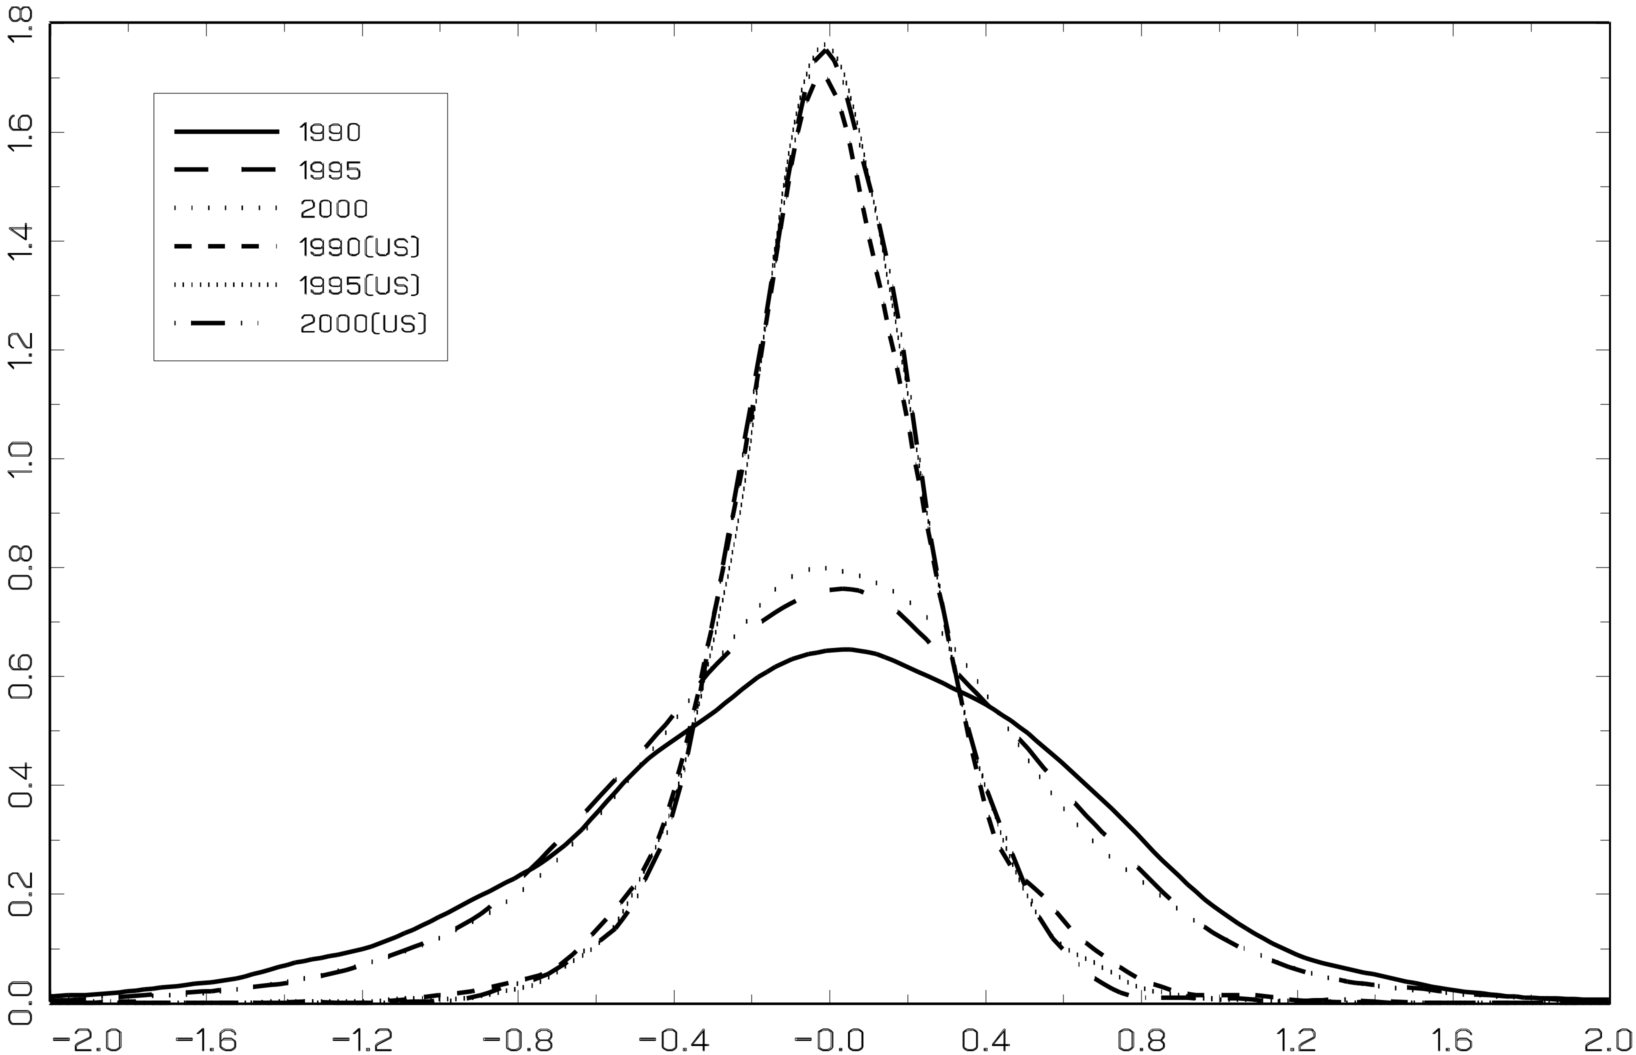

Figure 2
International and Intranational Price (Index)

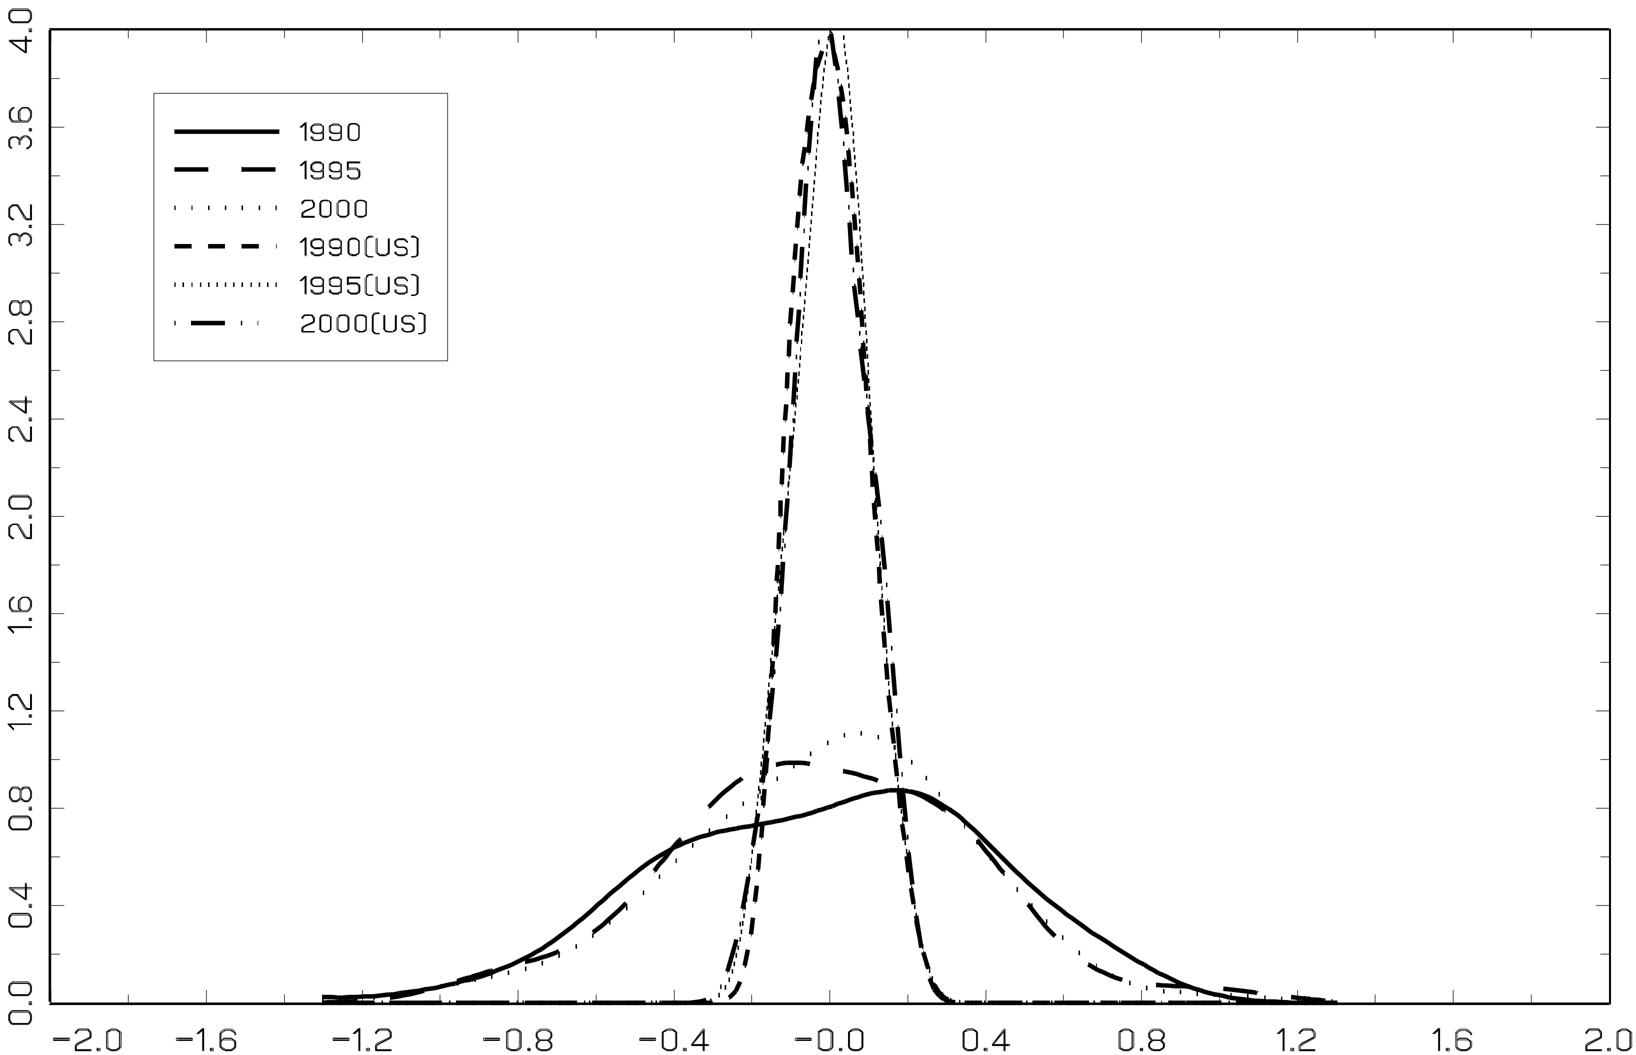

Figure 3
International and Intranational Price in 2000

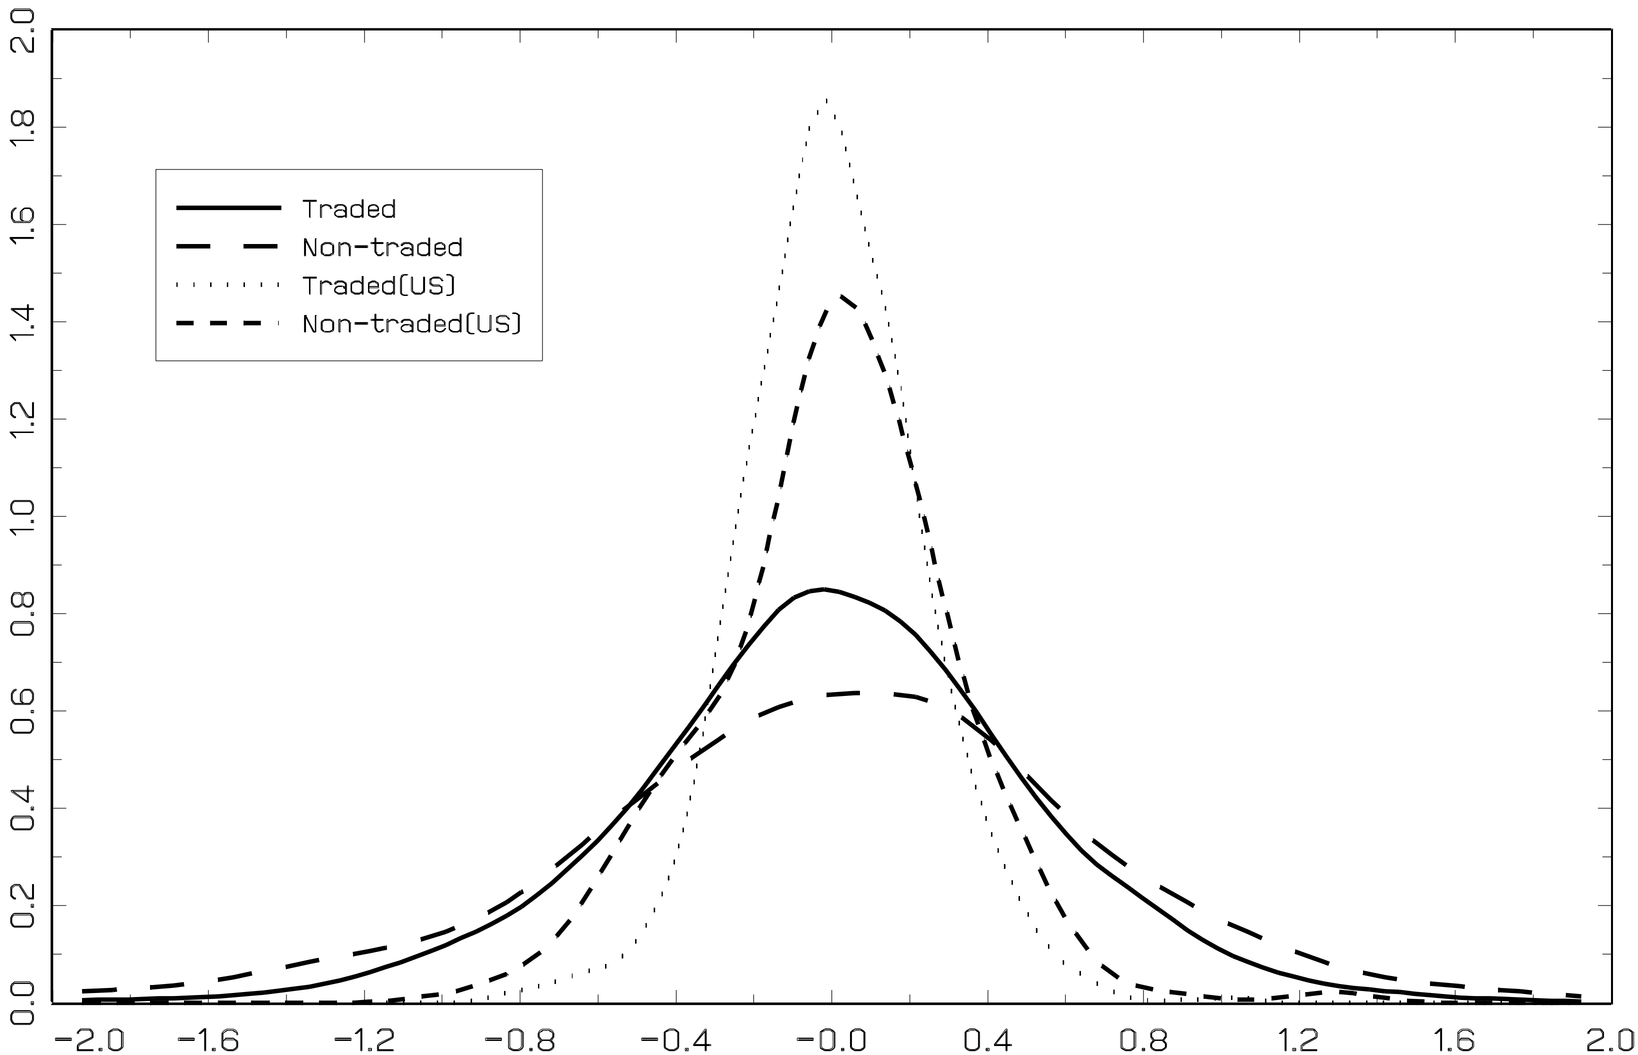

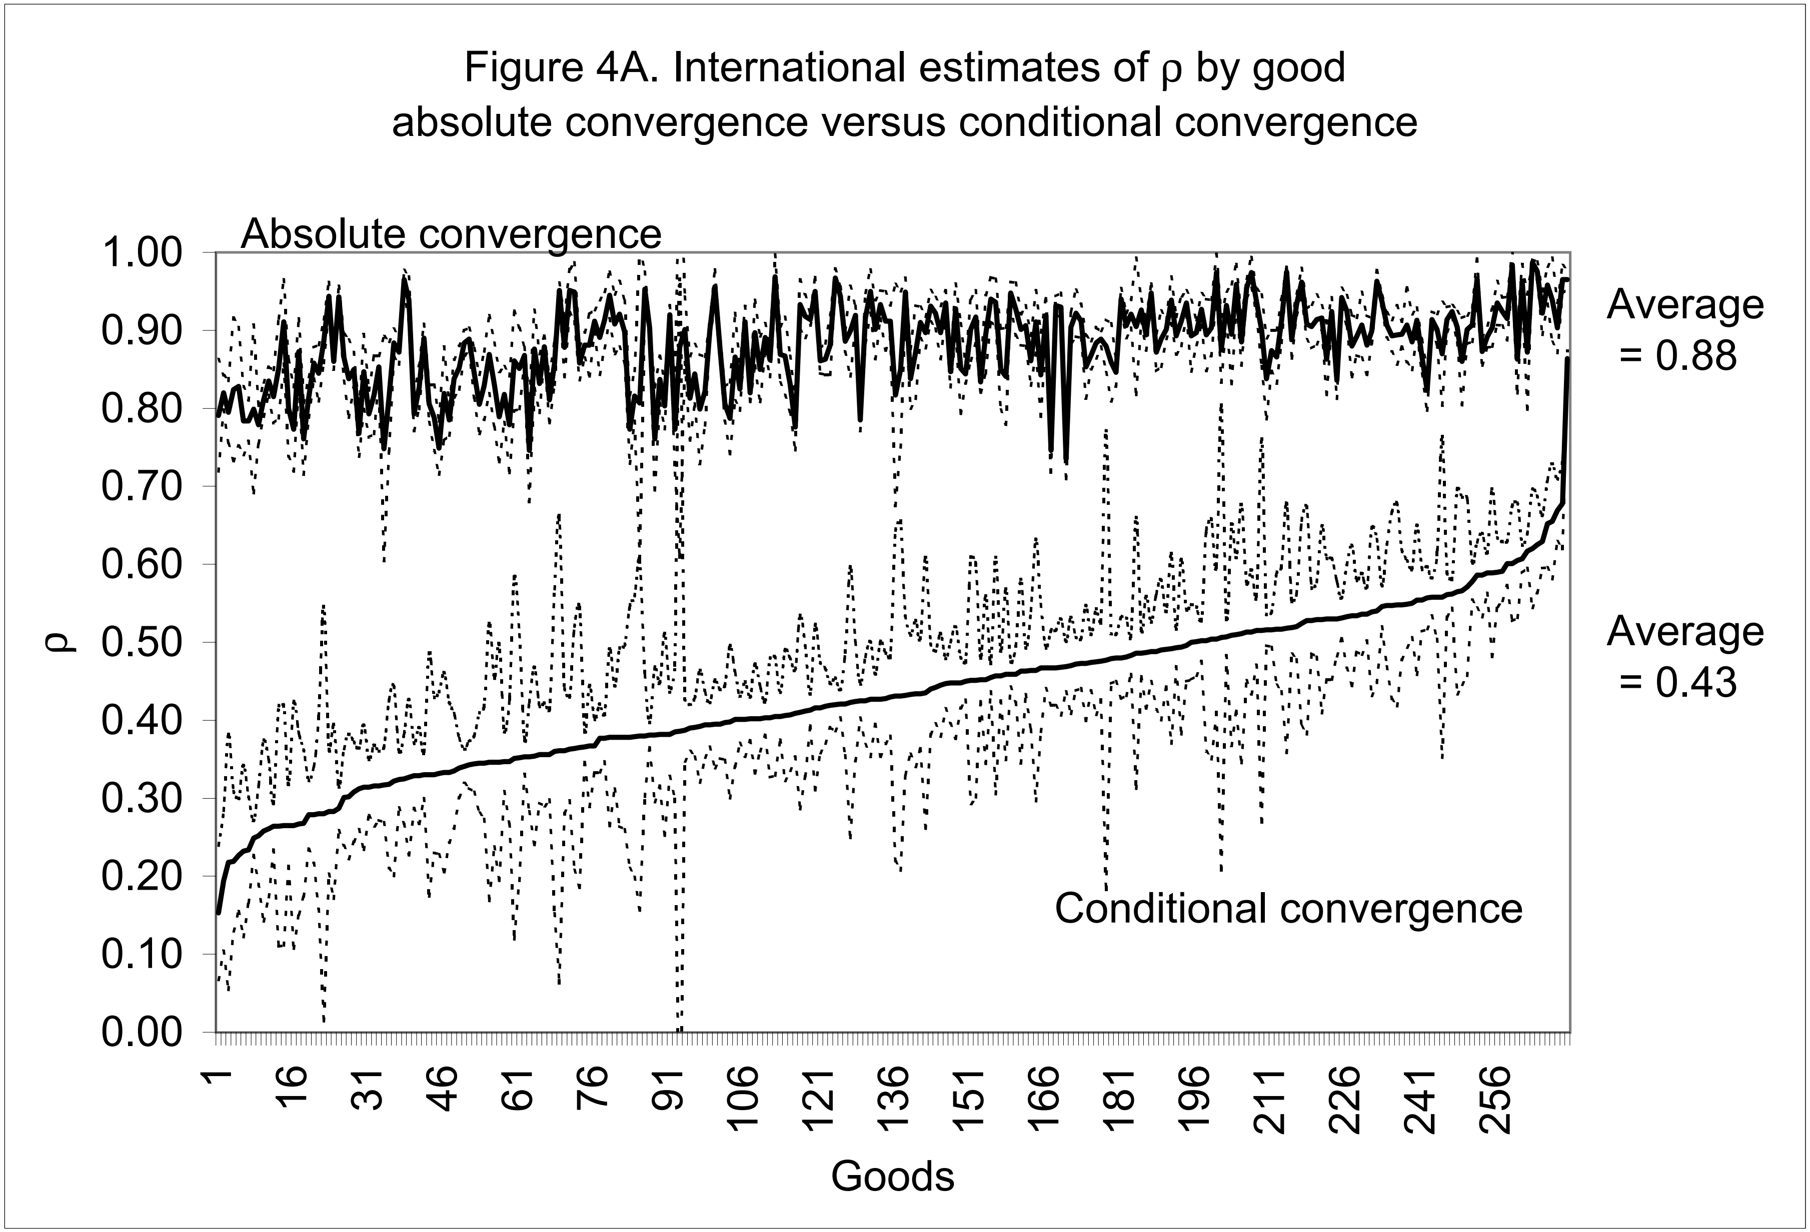
Figure 4A. International estimates of ρ by good
absolute convergence versus conditional convergence
Absolute convergence
Average = 0.88
Average = 0.43
Conditional convergence
ρ
1.00
0.90
0.80
0.70
0.60
0.50
0.40
0.30
0.20
0.10
0.00
1
16
31
46
61
76
91
106
121
136
151
166
181
196
211
226
241
256
Goods

Figure 4B. Intranational estimates of ρ by good
absolute convergence versus conditional convergence

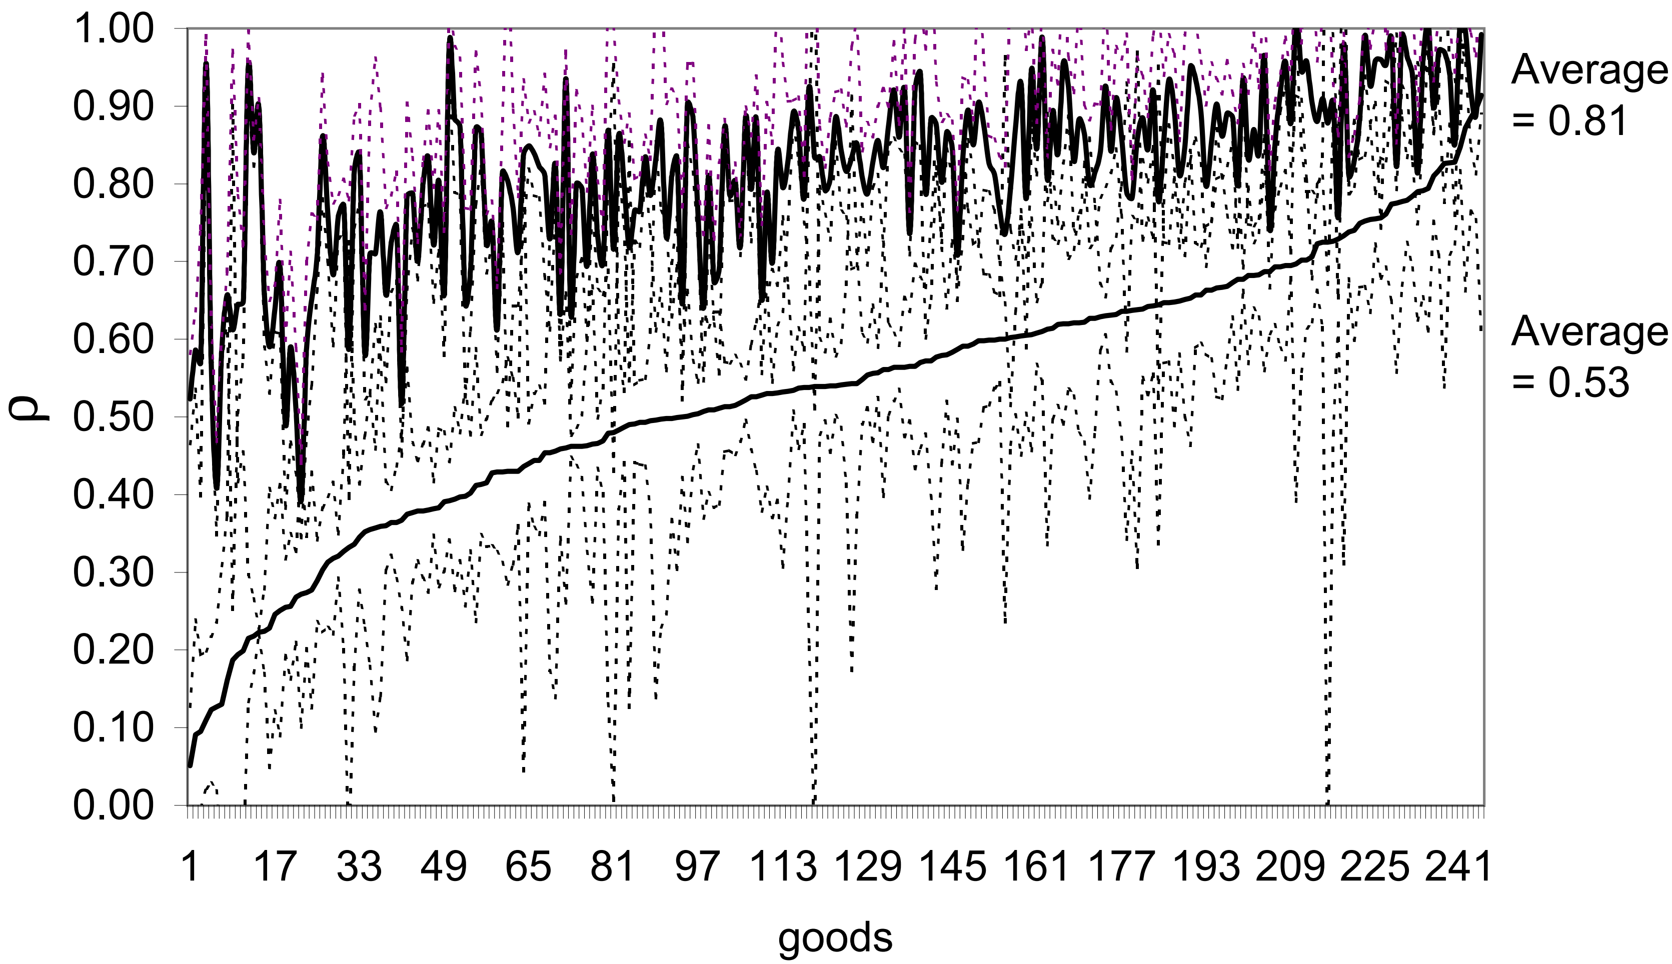

Figure 5. Upward bias in real exchange rate persistence

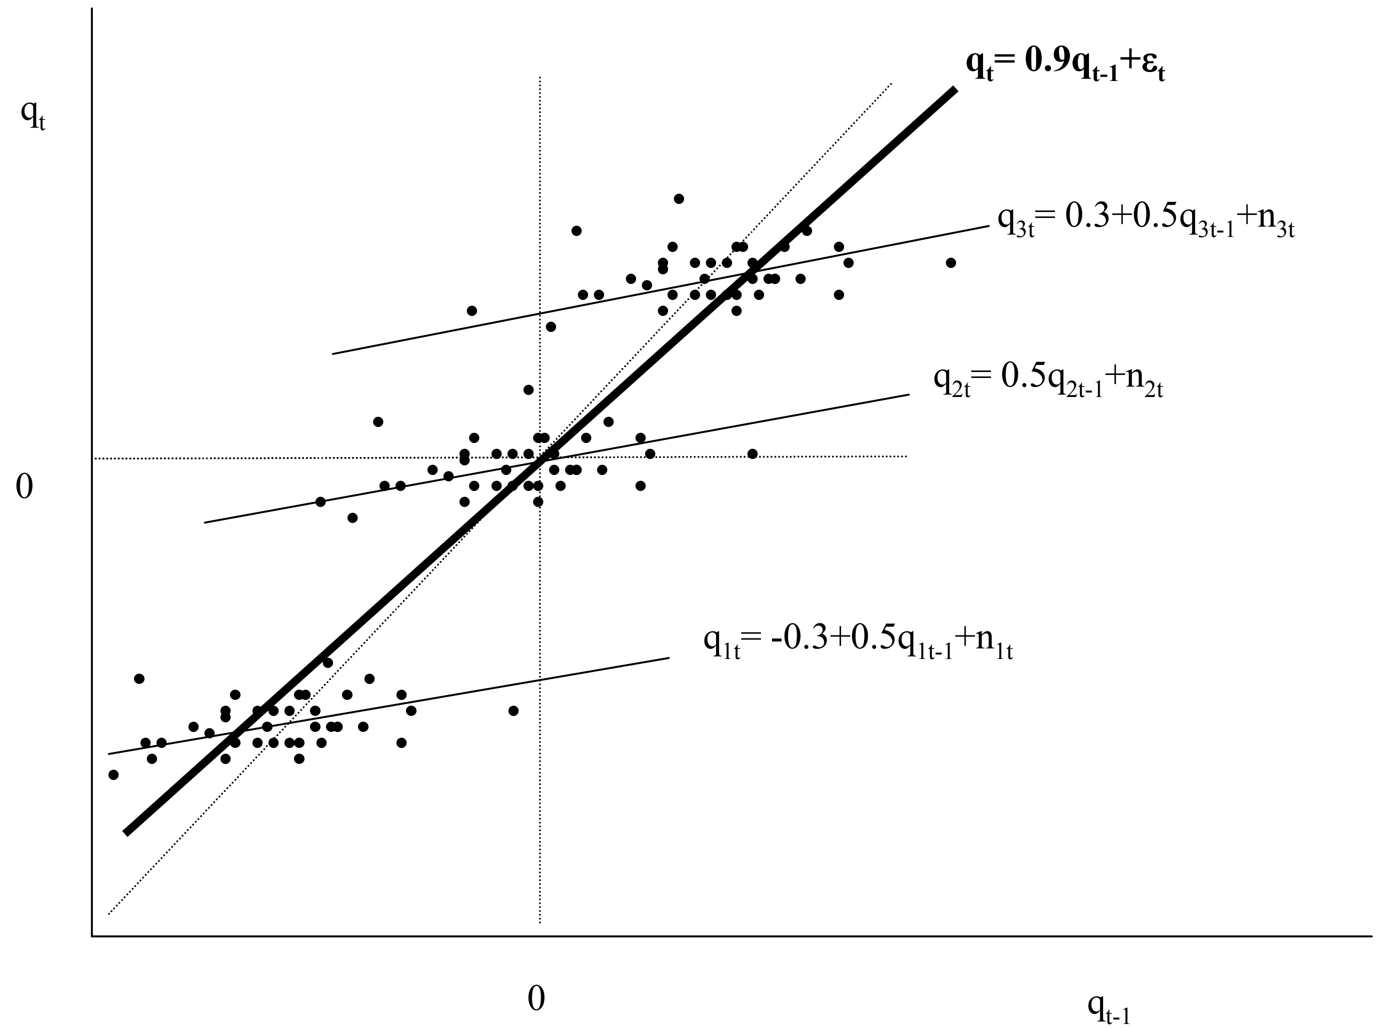

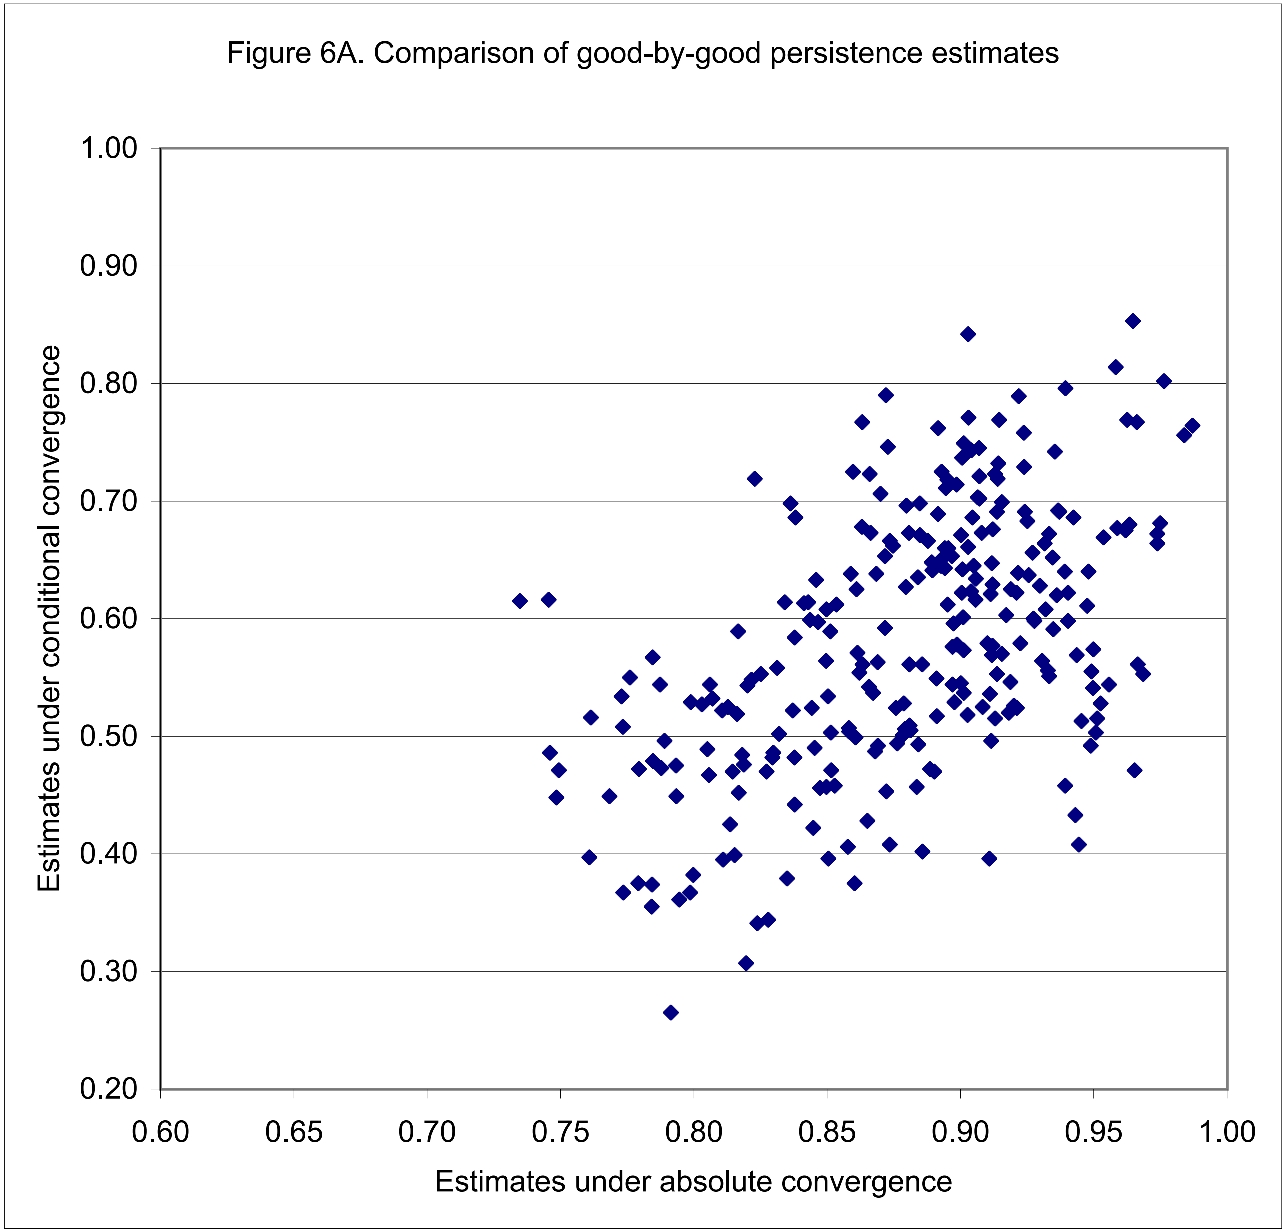
Figure 6A. Comparison of good-by-good persistence estimates
1.00
0.90
0.80
0.70
0.60
0.50
0.40
0.30
0.20
Estimates under conditional convergence
0.60
0.65
0.70
0.75
0.80
0.85
0.90
0.95
1.00
Estimates under absolute convergence

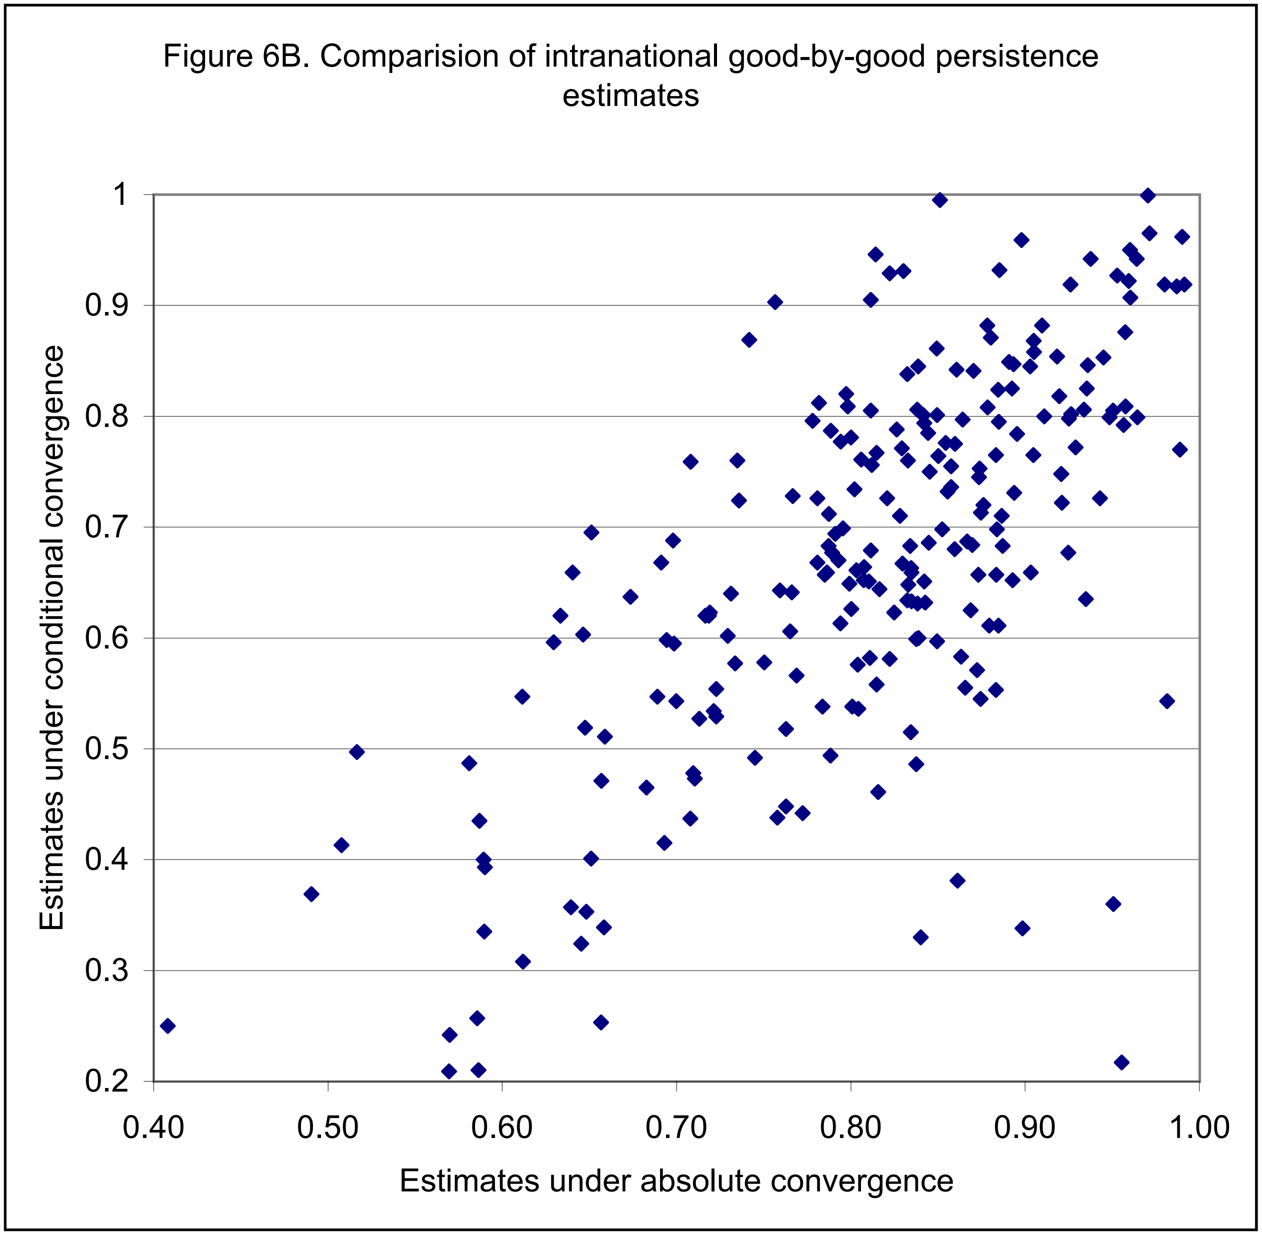
Figure 6B. Comparision of intranational good-by-good persistence estimates
Estimates under conditional convergence
Estimates under absolute convergence
1
0.9
0.8
0.7
0.6
0.5
0.4
0.3
0.2
0.40
0.50
0.60
0.70
0.80
0.90
1.00