DEPARTMENT OF ECONOMICS
UNIVERSITY OF CYPRUS

TRENDS IN INTERNATIONAL PRICES

Philippe Andrade and Marios Zachariadis

Discussion Paper 2010-02

P.O. Box 20537, 1678 Nicosia, CYPRUS Tel.: ++357-2-892430, Fax: ++357-2-892432
Web site: http://www.econ.ucy.ac.cy

Trends in international prices *

Philippe Andrade †
Banque de France & CREM

Marios Zachariadis ‡
University of Cyprus

December 2009

Abstract

We exploit the panel dimension of a price levels dataset for more than one hundred product items across 140 cities in 90 countries for the period from 1990 to 2009 in order to improve our understanding of international price dispersion and the evolution of prices over time. We consider a panel data model with exchangeable units that allows for the possibility of common components for different dimensions of the panel. This allows one to gauge the contribution of each dimension of the data to total variation and to disentangle the sources of potential non-stationarity. It also allows us to identify differences in the speed of convergence for different time-varying components in response to location-specific, product-specific, and idiosyncratic shocks. Finally, we proceed to identify the economic determinants of different components to show that particular dimensions of the data are more suited for examining particular theories.

Keywords: prive levels, variance decomposition, convergence, non-stationarity, international price dispersion.

JEL Classification: E31, F4, C23

*We would like to thank Jon Copestake, Editor of the Worldwide Cost of Living Surveys, for the insights and detailed explanations offered to us regarding the price data utilized in our application.

†Philippe Andrade, Banque de France, Monetary Policy Research Division, 31 rue Croix des Petits-Champs, 75049 Paris Cedex 1, France. Phone#:+33-142924995. Fax#: +33-149924818. E-mail: philippe.andrade@banque-france.fr

‡Marios Zachariadis, Department of Economics, University of Cyprus, 1678 Nicosia, Cyprus. Phone#: 357-22892454, Fax#: 357-22892432. E-mail: zachariadis@ucy.ac.cy

1 Introduction

What are the forces driving international prices in a globalized economy? To answer the question we study how individual price levels for comparable goods across a large sample of countries have evolved over the last two decades. More precisely, we exploit the panel dimension of the Economist Intelligence Unit (EIU) price levels dataset, which reports prices for over one hundred product items across 140 cities spanning 90 countries over the period 1990-2009, and we implement a panel methodology that allows us to split the variation in prices between global, (time-varying and time-invariant) location-specific, (time-varying and time-invariant) good-specific, and idiosyncratic good-location components.

We find the following results. First, nearly 95% of total international price variability comes from good-specific characteristics that do not evolve over time. The location-specific component of international price dispersion accounts for less than 2%, with time-invariant location characteristics explaining roughly 65% of this. This is in sharp contrast with the fact that the literature has mostly focused on the time evolution of cross-country international price differences. Second, we show that these prices all share the same global stochastic trend. This accounts for a very tiny fraction of total variance which is less than 0.5%. At the same time, this global trend is behind the observed persistence in price levels. Third, we document that shocks to city relative prices (i.e. to price differences for the same good across space), and shocks to goods relative prices (i.e. to price differences across goods within a location) have a short duration. Focusing on city-relative prices, location-specific shocks appear to be more rapidly corrected than idiosyncratic good-and-location-specific ones, whereas looking at goods relative prices, idiosyncratic good-and-location-specific shocks appear to be more rapidly corrected than good-specific ones. Moreover, the reaction to these shocks differs depending on a country's development level and goods' tradeability, with shocks being more rapidly corrected for traded goods and in less developed countries. Fourth, we relate the location-specific and good-specific components to economic variables, and examine whether competing theory-implied variables may be more adapted to describing one among these dimensions. We find significant explanatory power of variables related to production and distribution costs, monetary policy, and trade costs. We also show that some variables may significantly account for one component of price dispersion but not an other, so that one should not be dismissive of a theory simply based on a single dimension or component of the data.

The broad picture that emerges from these results is that of a fairly integrated global economy where the distribution of prices across different goods and locations is relatively stable over time, but slightly evolving in a common movement along with comparatively small and temporary location-specific churns. Consequently, our empirical analysis underlines that in order to understand price persistence one has to analyze global trends, while in order to understand international price dispersion one has to focus on goods characteristics. Finally, in order to understand medium-run fluctuations in international prices one has to turn to differences across locations.

Our paper relates to the vast literature on the persistence of international deviations from the Law of One Price (LOP). Until recently, these were considered to be very persistent with a half-life of several years (as documented in the surveys by Goldberg and Knetter, 1997, and Obstfeld and Rogoff, 2000), a result conveying a lack of integration in international markets. Instead, we confirm the recent evidence of Crucini & Shintani (2008) using EIU data, Goldberg & Verboven (2008) using European car market data, and Broda & Weinstein (2008) using barcode data, showing that persistence of the deviations from the LOP is reduced sharply when one increases the comparability of goods being studied. We find a half-life of roughly 1.5 year. Our contribution here, is to implement a panel unit root test and a convergence rate estimation method which takes into account that the units of the panel under study are cross-correlated because of global, location-specific, or good-specific shocks. Using standard panel procedures in the presence of such presumable cross-dependencies among units of the panel, may lead to downward bias in persistence estimates and to tests favoring the conclusion that prices are mean-reverting even if they are in fact affected by stochastic trends.[1] In addition to dealing with dependencies between units of the panel, an other interesting feature of our approach compared to the existing literature on the convergence to the LOP, is to allow for potential trends in international prices other than location-specific ones; namely, global (worldwide), good-specific and idiosyncratic good-location ones. This also allows us to estimate convergence rates for those among these components that are shown to be stationary. That is, we can distinguish among mean-reverting reactions to location-specific, product-specific, and idiosyncratic shocks.

[1] For instance, O'Connell (1998) showed how neglecting to correct for cross-sectional dependence between real exchange rates (due to common macro shocks) leads to wrongly conclude in favor of long-run PPP.

The paper is also linked to the literature that aims to explain the mechanisms behind these deviations from the LOP. Similar to Hellerstein (2008) for the US beer market or Nakamura & Zerom (2009) for the US coffee market, we find evidence that variables affecting production costs, distribution costs, mark-ups, and trade costs can be useful for understanding international price differences. The latter results are consistent with Bergin and Glick (2003) and Atkeson and Burstein (2007, 2009) who assign a central importance to trade costs. Our results also point to the importance of goods characteristics to understand international prices, in line with Crucini, Telmer & Zachariadis (2005), and Crucini & Shintani (2008). This crucial role of goods features also connects our work with papers that aim to fill the gap between microdata-based results and results based on aggregate macro price indices. Imbs *et al.* (2005) argue this comes from an aggregation of heterogenous individuals or sectoral price dynamics. Given the huge part of price heterogeneity that comes from goods fixed effects, our analysis also suggests the importance of a composition bias in the indices being compared as another culprit responsible for the differences between macro and micro studies.

Our results on good-specific inflation rates convergence within a country can be compared to studies that attempt to disentangle price fluctuations into different components to account for country-specific, industry-specific, and common components. For example, Ciccarelli & Mojon (2008) show that inflation across 22 OECD countries shares a common stochastic factor that accounts for 70% of total inflation variance and drives the worldwide inflation trend. Like them, we find that international prices share a common trend and that their result can be extended to a larger group of countries than just OECD ones. We note that as they work with price indices, they cannot assess the importance of good-specific variability as we do. Our results can also be seen as extensions to an international environment of the results in Clark (2006), Boivin, Giannoni & Mihov (2008), Maćkowiak, Moench & Wiederholt (2009) and Reis & Watson (2009). Using monthly sectoral data, these papers underline that the persistence in disaggregated US sectoral price indices is due to a common stochastic trend whereas the majority of variance comes from transitory US sectoral shocks.[2]

Finally, we note that Engel & Rogers (2004), Crucini, Telmer, and Zachariadis (2005b), Bergin and Glick (2007), Crucini & Shintani (2008), and Crucini and Yilmazkuday (2009) also exploited sub-samples of the same dataset being utilized in its entirety here. The first paper focuses on a

[2]See Altissimo, Mojon & Zaffaroni (2008) for comparable results on Euro area sectoral price indexes.

sample of prices in 18 European cities for 101 traded and 38 non-traded products for the period from 1990 to 2003, to ask how much more integrated the EU has become after the introduction of the euro. The second paper utilizes the EIU data averaged over the period from 1990 to 2000, and focuses on the first and second moments of the cross-sectional distribution of bilateral country prices across goods, to assign a role to geographic variables. The paper by Bergin and Glick focuses on a sample of 101 tradeable goods in 108 cities in 70 countries for the period from 1990 to 2005, to assess global price convergence. Crucini & Shintani (2008) focus on a sample of 90 cities in 63 countries for the period from 1990 to 2005, to assess the rate of price convergence for the relative price of each good. Crucini and Yilmazkuday (2009) average the data over the period from 1990 to 2005 and explain this cross-sectional dimension with trade and distribution costs. As compared to these papers, we use the complete sample of EIU prices that are available across all cities for the period from 1990 to 2009, and implement an exchangeable units panel model that allows us to exploit all the dimensions of the dataset to examine the different components responsible for the presence of stochastic trends in non-stationary price processes or for the rate of mean reversion in stationary relative prices.

The rest of the paper is organized as follows. In Section 2 we provide a detailed presentation of the EIU data, underlining goods comparability across locations and time. Section 3 presents the statistical model on which we rely to decompose the international price dynamics into its global, country-specific, and good-specific components. Section 4 relies on this model and a variance decomposition exercise to assess the relative importance of each of these terms for total variability of international prices. In Section 5, we implement an original strategy to test for unit roots in a panel with cross-section dependencies and a small number of time periods. Section 6 is devoted to the estimation of the speed of convergence to the LOP and of goods relative prices, in response to various types of shocks. Section 7 presents the mapping of location-specific and good-specific components of prices onto economic variables drawn from theoretical models of international prices.

2 Data

The main source of data utilized in our application comes from the EIU. These data is available for a sample of 327 items for 140 cities in 90 countries for the period from 1990 to 2009. Some summary statistics regarding these prices are presented in Table (1). As can be seen there, this sample includes vastly different priced items, with the standard deviation much greater for traded products as compared to non-traded items, and across LDC's as compared to across developed economies. We also note that there is a much lower number of non-traded items available as compared to traded products.

A number of explanatory variables was obtained directly from the EIU dataset at the city level. These include electricity cost (entry 185 in the EIU data), regular unleaded petrol (entry 250), and residential rent for two-bedroom unfurnished apartment (entry 262) chosen among other rents since it contained the least missing observations.[3] We also utilize the country's exchange rate measured as the number of national money units for one US dollar and assembled by the EIU to match the sampling periods of the city price levels data.

We downloaded disaggregated MFN tariffs data for the items in our price database from the Trade Analysis Information System (TRAINS) dataset, which contains tariffs and import values data for 119 countries available at the most detailed commodity level of the national tariffs. We also obtained city-specific population from the Henderson cities dataset that includes population sizes for each decade between 1960 and 2000 for about 3,000 cities across the world.[4] Money supply as given by M1 in billion $US was obtained from the Economic Indicators dataset of the Country Data made available by the PRS group. In addition, we obtained real GDP per worker from the Penn World Tables. Finally, services GDP share and imports shares were obtained from the World Development Indicators.

Below, we undertake a detailed presentation of how these prices are collected and put together, meant to help the reader understand the potential advantages and disadvantages of using this dataset to study international prices.[5] Although subsamples of these data have been used previously as described above, the information provided below is largely new. Our intent is to make this available in order to assist future researchers in appropriately handling these data.[6]

---

[3] Wages in current dollars are available from the EIU dataset at the country-level. Including these country-level input cost would reduce the number of cities being considered in Table 6 to 61, without changing our inference regarding the importance of other input costs.

[4] The revised Henderson dataset was provided to us by Yiannis Ioannides to whom we are grateful for this.

[5] This discussion has benefitted greatly from systematic direct communication with the EIU office over the past few years, and in particular, from the insights and detailed explanations offered to us by the Editor of the Cost of Living surveys of the EIU, Jon Copestake.

[6] We would thus include this information in an appendix upon publication of this paper.

Selection of stores and goods

Considerable care is taken by the EIU team to assess accurately the normal or average prices international executives and their families can expect to encounter in the cities surveyed. Survey prices are gathered from three types of stores: supermarkets, medium-priced retailers and more expensive specialty shops. Only outlets where items of internationally comparable quality are available for normal sale are visited. While the majority of cities provide a wide selection of goods and stores at different price levels, this range narrows considerably at several locations. In some cities the entire range of prices has to be collected at the few stores where goods of internationally comparable quality are found. Local markets and bazaars are visited only if the goods available are of standard quality and if shopping in these areas does not present any danger. For certain items like monthly rent and clothing, there are many subjective factors, questions of personal preferences and taste at play, as well as a wide variety of choice. Therefore, the price data given for certain items should be considered to be merely an indication of the general level of prices in these categories. As a result, we felt the need to create a sub-sample of goods that are more likely to be comparable across locations. All results currently presented in the paper, are being replicated using this more highly comparable sub-sample of goods.

The price range presented in the survey utilized in the current study is for supermarkets and mid price outlets. The EIU takes one representative price per store, sampling only one price from each type of store, and generally surveys two stores per item for most products. This represents a market snapshot during the survey timing. In all cases, the EIU aims to keep the same stores and the same brands and sizes in obtaining the price for each item, so as to ensure ongoing consistency between surveys in each location. Store and product consistency has been an aim of the survey since its inception. The aim of sampling the same stores has remained consistent and the ability to do so has varied based on specific events in certain years relating to availability or specific situations affecting correspondents, like being refused entry to a store under new management. There would therefore be no typicality of this kind specific to earlier or later surveys.

However, such consistency depends on and varies within individual markets. The surveyors seek to keep to the same stores, brands and weights between surveys. However, given that the survey takes place simultaneously in 140 cities over a period of twenty years, there may be substitutions or changes. This can occur in an evolutionary sense as certain brands or stores or sizes overtake others as the popular interpretation of a particular item changes over time. Alternatively, there may be sudden changes in brand, store or item based on availability in the market during a particular period. For example, a store may close and a certain brand may become temporarily or permanently unavailable. In these cases, substitutes are sought to reflect the price of obtaining the item in question at that particular time. This is more common in less developed markets where availability and price can fluctuate on a day to day basis, but even mature markets are prone to pricing or availability shocks and other changes of this kind especially over longer periods. We note than while the BLS adapts its basket of goods regularly and also changes the weighting system based on consumption trends, the EIU seeks to be more generally representative and has for the most part not changed in this manner, in an attempt to ensure a consistent dataset of like for like products going back over time.

The general conclusion from the discussion in this sub-section is that the EIU city-level prices are highly comparable across both space[7] and time, and are thus suitable for the study of LOP deviations and their evolution over time. That is, one can use these prices to understand both the degree of market segmentaton at any given point in time, and the process of market integration over time. However, these prices appear less suitable for overall cost of living comparisons across locations since the goods sampled do not necessarily reflect local preferences as much as the shopping basket of executives and other multinational employees and their families.

Sampling, seasonality, and sales

The fieldwork for the Worldwide Cost of Living Surveys is carried out on location by the EIU researchers during the first week of March for the Spring edition and during the first week of September for the Autumn edition. For the historic data in the "citydata" publication the majority of prices have been gathered in September. That is, since the data overwrites old data each year, most of the price data made available historically by the EIU is September data. There are two types of exception to this. First, are cities surveyed annually and only in March. These are: Baku,

[7] The degree of comparability across locations is high but varies with the general availability of goods in a given city. Given that the survey takes place in 140 cities worldwide, it is not always the case that an identical product is taken in all cities for all items. For example, it is more likely that while London has a quality Burberry raincoat available, Brussels does not have the same item or brand and the correspondent has taken a price based on the designer raincoats that are available. For such products, prices will reflect the general availability and local demand conditions in a location. Given these concerns, one should consider subsamples that exclude products likely to be less homogeneous across locations. The latter category includes pretty much all clothing items, automobiles, and a number of other products.

Bratislava, Calgary, Douala, Harare, Port Moresby, San Juan, and Tunis. For these cities, data is gathered since 2001 during the first week of March. Second, are cities where there are problems or delays in gathering data. These are individual cases and are not tracked, but it would generally be the case that such data is still gathered within a month or two, so that prices can still be relevant and comparable to other cities. Moreover, no such lags are allowed in high inflation locations.

The March and September dates for gathering data are specifically designed to avoid standard sales seasons, like traditional sales in December, January, May and June which take place in many countries. Correspondents are instructed not to take sale prices for items, but to take standard recommended retail prices. There is an element of common sense here as well though. That is, correspondents may take sales prices for general promotions if they feel the price reflects the "true worth" of an item. This might be the case for some items since retailers commonly use tactics of promoting an item by describing it as on "sale" when in fact they have previously artificially inflated the retail price of the item in order to later reduce it to a more reasonable price and make consumers think they are purchasing a bargain. This is true of items like CDs, wine, certain fresh food items, and other consumer goods. A few adjustments of the survey prices have been made in some cases where seasonal discount sales and changes in brand names, package sizes, and quality would have unduly distorted the index results. This procedure is limited to cases where it would not entail misrepresentation of actual prices in the EIU team's judgement.

The conclusion from the above paragraph is that the astonishing price differences for specific items across cities observed by the EIU team, are not due to sales or discounting. For example, the difference between a Burberry raincoat in Brussels and London is not due to sales prices or discounting as the EIU does not seek to include such seasonal data in the price survey.

Reliability of data

We have opted to be extremely conservative in removing entries that at first might appear to be price outliers. Moreover, we never opt to adjust prices for what might at first appear to be "obvious" mistakes, like misplacing a digit or otherwise using a wrong unit, or misplacing part of a price entry in previous or subsequent entries. In this respect, our treatment of the data is very different than Crucini and Shintani (2008).

We opted to treat the data as a rather reliable representation of actual prices since in our discussions with the EIU office it was convincingly explained to us that specifying for instance the price variance between surveys not to be less than half or more than twice the CPI rate would be an extremely narrow margin for highlighting outliers, as the EIU team has historically observed prices that regularly change by as much as four times or more the CPI rate, while other prices remain unchanged year after year or even move down. It was also explained to us by the EIU office, that every survey price is "sense checked" as it comes in compared to those returned six months ago and those returned one year ago. Sense checking is simply to ensure that prices look broadly comparable to those returned previously. However, the final prices reported in the EIU surveys are based on actual ones as returned from field correspondents in each city, and are never a calculation based on a ratio of expected price movement to reported inflation levels. As a result, prices of individual items in the basket the EIU surveys can fluctuate wildly based on the basket snapshot that is taken.

Where a user has serious concerns, the EIU recommends removing a price rather than guessing at its original value. For instance, if we suspect that certain prices were simply misinput in error then this price would need to be removed from consideration as an outlier rather than tweaked into something resembling what it "should be". While it is completely valid that a tiny proportion of the reported prices may include errors, the vast majority of prices are arguably valid snapshots at the time of the survey and most prices that vary disproportionately with the CPI can be explained simply by looking at the context in which the prices were taken. Finally, even if all prices that move very differently than the CPI were assumed to be errors, these would represent a proportion below 0.5% of the available datapoints.[8]

Errors that emerge may be a currency issue where back-rates are recalculated to cater to currency redenominations caused by inflationary spikes, or where devalued/alternative exchange rates are in operation. It is possible that some prices might be entered in a sub-unit of currency (e.g. in pence or cents) then reported in standard units (e.g. in pounds, euros or dollars). However, this is something the EIU generally seeks to rectify on a rolling basis. Still, the EIU cannot double-check many of the prices since the citydata feed automatically takes from the source files. These are taken from surveys based on manually collected data by correspondents in each location. The price

[8] This might be one reason for the robustness of the Crucini and Shintani (2008) convergernce rate estimates in this altered sample.

dataset is built as the accumulation of decades of data submitted from a variety of sources in a variety of formats. Any data collected before 1998, for example, would have been returned in paper format and manually input into the base files eventually used, and the original paper versions have long since been disposed of. Thus, the EIU may only be able to check sources for items after 1998 but such a process would be time-consuming and unnecessary according to the EIU office, since most of the price entries that appear at first to be errors are actually valid price entries.

For instance, a seemingly wrong but actually correct price entry comes from Casablanca in the case of bread. The figures for years 1992 to 1995 seem to be missing the initial "1". This example of bread in Casablanca between 1992 and 1996 is a prime example of how EIU prices should be considered valid even if they look peculiar relative to general price trends. Between 1992 and 1995, Morocco suffered from a period of drought which caused three harvests to fail (1992, 1993, and 1995). This had an impact on economic growth and prompted a recession. In response, the government will have extended price controls on staples. In the Moroccan diet, bread is considered to be the staple food of the poor and would have been the first and most heavily price-regulated item. Upon recovery and under external pressure the government pledged to relax such controls in 1996. In the case of the survey, we can clearly see this reflected. Lower priced bread in line with the 1992-1995 prices may have been widely available before and after this period, but during this period shortages, economic stagnation, suppressed demand for more expensive consumer goods, and price controls may have meant that these were the only prices available for bread. This situation was rectified as Morocco emerged from this period. Similarly, many prices could be flagged in developing countries during times of instability as these experience massive fluctuations in prices dependent on localised supply and demand factors. Thus, the EIU suggests that users consider reasons why a particular price may deviate from expectation based on the political, social and economic market context, globally, nationally or at city level before removing a price entry.

Nominal exchange rate issues

Spot exchange rates are applied to the citydata surveyed by the EIU, and are available along with the price data for each year. The post rates are FT rates taken on the Friday of the first week of each month of the survey. Since the data overwrites old data each year, most of the exchange rate data (and price data) supplied historically is September data except in a few instances where a city is only surveyed every March - in which case all price and exchange rate data is from the first week of March. Thus, the exchange rate reported is the spot rate for the survey when the published data was gathered.

For pre-1999 price series, the conversion from legacy currencies to euros is made using the appropriate legacy currency, i.e. Ecu exchange rates prevailing at the time. Like Eurostat, the EIU has chosen to use the Ecu exchange rates because there is no universally agreed methodology for calculating a synthetic euro exchange rate. One Ecu was worth exactly one euro when the euro was launched at the beginning of 1999. The EIU used the September end-period rate from Eurostat to convert the legacy prices. Although surveys were completed for Euro cities at slightly different times in September, the EIU wanted to apply a standard rate to maintain relative prices between cities and also maintain distances between published Cost of Living indices.

3 A statistical model of international prices

Let p_{ilt} denote the (log) price for good i sold in location l at date t and expressed in US dollars. We observe n_i different goods and services prices across a set of n_l locations. With these three dimensions of our data set, the evolution of international prices can then be decomposed into four components: a trend that is common to all locations, a trend specific to locations l, a trend specific to product i, and a trend specific to each product and location, so that:

$$p_{ilt} = m_t + m_{it} + m_{lt} + m_{ilt} \tag{1}$$

These components can be understood as trends resulting from shocks common to all locations and products, ϵ_t, from individual product shocks common to all locations, ϵ_{it}, from city shocks common to all products, ϵ_{lt}, and finally, from shocks specific to each product and location, ϵ_{ilt}. We also allow for a non-zero time-mean in each of these components, respectively denoted as μ, μ_i, μ_l and μ_{il}.

We consider a panel model of international prices that has a weakly-linearly exchangeable structure (see Gregoir, 2003). Weak-linear or covariance exchangeability implies that the dependencies across time and units in a panel can be described by a covariance function that does not depend on any ordering of the units. This translates into the following covariance structure across dates and units:

$$\text{cov}\left(\Delta p_{ilt}, \Delta p_{jft-h}\right) = \gamma_1\left(h\right) + \gamma_2\left(h\right)1_{ij} + \gamma_3\left(h\right)1_{lf} + \gamma_4\left(h\right)1_{iljf}$$

with 1_{ij}, 1_{lf} and 1_{iljf} indicator functions that equal 1 when, respectively, good $i = j$, location $l = f$, and both $i = j$ and $l = f$, and 0 otherwise. This covariance structure amounts to imposing that each of the preceding shocks – ϵ_t, ϵ_{lt}, ϵ_{it}, and ϵ_{ilt} – are statistical innovations to each of the related trends, and uncorrelated with each other.

The literature on international prices traditionally focuses on the behavior of deviations from the LOP. Namely, it investigates the determinants of differences in the price of a specific good i observed in different locations l and f, $q_{ilft} = p_{ilt} - p_{ift}$. With the postulated structure for international prices, this implies that

$$q_{ilft} = (m_{lt} - m_{ft}) + (m_{ilt} - m_{ift})$$

The usual practice is then to construct a panel by stacking observations for every possible pair of locations $\{l, f\}$ in the sample. By construction, this approach will create dependencies across the units of this panel and therefore raise OLS estimation efficiency issues. Indeed, under our assumptions it holds that for any three locations l, f, k

$$\text{cov}(q_{ilft}, q_{ilkt}) = \text{V}(m_{lt}) + \text{V}(m_{ilt}), \quad \text{cov}(q_{ilft}, q_{jlkt}) = \text{V}(m_{lt})$$

Moreover, as explained in the next section, cross-correlation between individual panel units is problematic when one tries to assess the stationarity of a variable using a panel testing approach. It can be noted that the cross-correlation problem does not come only from the hypothesis of a location-specific component, m_{lt}. Even without assuming the existence of a location-specific component, the simple replication of the same price across units of the panel correlates units through repeated m_{ilt}. Much of the previous literature has instead opted to use all unique bilateral price comparisons between every possible pair of locations in the data, introducing such cross-dependencies among panel units.

We thus rather look at a location l's price for good i relative to the average across cities

$$q_{ilt} = p_{ilt} - \overline{p}_{it}, \quad \overline{p}_{it} = (1/n_l)\sum_{l=1}^{n_l} p_{ilt}$$

By definition, the preceding model implies that some components in prices are common to every good and every location, so that deviations from the LOP have the following structure:

$$q_{ilt} = m_{lt} + m^q_{ilt}$$

with $m^q_{ilt} = m_{ilt} - (1/n_l)\sum_{l=1}^{n_l} m_{ilt}$. It can be noted that the correlation among the units of the panel is then given by

$$\mathrm{cov}(q_{ilt}, q_{ift}) = (1/n_l)\mathrm{V}(m_{ilt}), \quad \mathrm{cov}(q_{ilt}, q_{jlt}) = \mathrm{V}(m_{lt})$$

The cross-correlation issue disappears if one is able to estimate and remove the common component, m_{lt} and if the number of locations in the sample is sufficiently large.

Aside to the usual investigation of the differences between cities of the price for the same good, one may also be interested in looking at the evolution of goods relative prices within a location, namely

$$r_{ilt} = p_{ilt} - \overline{p}_{lt}, \quad \overline{p}_{lt} = (1/n_i)\sum_{i=1}^{n_i} p_{ilt}$$

Again, the exchangeable structure we consider implies that some movements in prices are common to every good and every location, so that good relative prices evolve according to

$$r_{ilt} = m_{it} + m^r_{ilt}$$

with $m^r_{ilt} = m_{ilt} - (1/n_i)\sum_{i=1}^{n_i} m_{ilt}$.

Estimating the unobservable common factors, m_t, m_{it}, m_{lt}, is needed if one wants to correct for cross-correlation among the units of the panel. Estimates of these common factors may also be wanted to understand what are the key variables driving them. It turns out that an interesting feature of our approach is that these unobservable trend components have simple natural estimators. Namely, let $n = n_i + n_l$, (with n_i the number of goods in the sample and n_l the number of locations) and define

$$\overline{p}_t = \frac{1}{n}\sum_i \sum_l p_{ilt}, \quad \overline{p}_{it} = \frac{1}{nl}\sum_l p_{ilt}, \quad \overline{p}_{lt} = \frac{1}{ni}\sum_i p_{ilt}$$

then

$$\widehat{m}_t = \overline{p}_t, \quad \widehat{m}_{it} = \overline{p}_{it} - \overline{p}_t, \quad \widehat{m}_{lt} = \overline{p}_{lt} - \overline{p}_t, \quad \widehat{m}_{ilt} = p_{ilt} - \overline{p}_{it} - \overline{p}_{lt} + \overline{p}_t$$

Thus, this gives us a simple way to identify how each of these components matters for the dynamics of international price levels and differences.[9]

4 The relative weight of global, location, and goods components

The literature on international prices stresses different dimensions of their differences. For example, the macroeconomic literature typified in the work of Charles Engel (see, for example, Engel, 1993, and Engel and Rogers, 1996) had until the recent past been focusing on time variation of international relative prices. This framework very often favored nominal considerations and sticky price explanations of international relative prices. Crucini, Telmer and Zachariadis (2005) focus instead on the variation of prices of individual goods across locations, emphasizing economic explanations related to product characteristics, while Crucini, Telmer and Zachariadis (2005b) emphasize variation of bilateral LOP deviations across goods allowing economic geography considerations to enter the analysis.

The decomposition of the price process we postulated in equation (1) provides a simple way to assess the relative importance of each component through a variance decomposition exercise. Indeed, relying on the orthogonality between the components of (1) allows us to decompose the total price variance into four components

$$\mathrm{V}(p_{ilt}) = \mathrm{V}(m_t) + \mathrm{V}(m_{it}) + \mathrm{V}(m_{lt}) + \mathrm{V}(m_{ilt}).$$

We also decompose the total variance in locations and goods relative prices into

$$\mathrm{V}(q_{ilt}) = \mathrm{V}(m_{lt}) + \mathrm{V}(m^q_{ilt}), \quad \mathrm{V}(r_{ilt}) = \mathrm{V}(m_{it}) + \mathrm{V}(m^r_{ilt}).$$

Moreover, for each of these components, one can also compare the variance that comes from the part of these components that is fixed over time relative to the total variance. More specifically, one can decompose the total variance of the good-specific component $\mathrm{V}(m_{it})$ and of the location-specific component $\mathrm{V}(m_{lt})$, into the variance of good-specific and location-specific fixed effects $\mathrm{V}(\mu_i)$ and $\mathrm{V}(\mu_l)$, and their respective time-varying components $\mathrm{V}(m_{it} - \mu_i)$ and $\mathrm{V}(m_{lt} - \mu_l)$.

Table (2) shows the results. The global component represents a very tiny fraction of total variance in prices with less than 0.5%. As shown in Figure 1, this common trend moves in accordance

[9] The annual frequency of the data set we use limits the number of dates in our sample and therefore the feasability of altternative approaches using common factor models.

with common wisdom, with global inflation rising until the mid 90's, then declining until 2000, and then rising again since 2001 up until 2008, with a sharp fall in 2009.

The location-specific component represents only roughly 1.5% of international prices total variance with time invariant factors accounting for about 65% of the variation of this particular component. Finally and strikingly, almost 95% of total price differences across dates, locations, and goods comes from the good-specific dimension, m_{it}, where within this component 99.9% comes from differences between time-invariant goods effects, μ_i.

In Figure 2, we present the evolution of relative prices for eight representative cities over time (Budapest, Cairo, Nairobi, New Dehli, New York, Paris, Rio, and Tokyo), and in Figure 3 we do the same for eight representative goods (aspirin, bread, Coca-Cola, Cointreau liqueur, cinema seat, haircut, lettuce, and light bulbs.) As shown in those figures, there are no clear trends in cities or goods relative prices over time. Moreover, comparing these two figures, it becomes evident that the time varying part of cities relative prices is relatively more important than for goods relative prices, at least for the representative cities and goods shown there. Figure 4 presents two panels with the complete cross-sectional distribution for the share of the time varying component in total variance for each relative price for all cities and goods in the top and bottom panels respectively. As illustrated there, and consistent with the results of Table (2) and with what was shown in Figures 2 and 3 for a small number of representative cities and goods respectively, the time-varying component is much less important for goods relative prices as compared to city relative prices.

5 Testing for stochastic trends in international prices

A vast literature in international macroeconomics has investigated whether international price differences have a tendency to disappear over time, and at which pace this convergence, if any, takes place. These studies test for unit roots (UR) leading to stochastic trends and non-convergence in the autoregressive dynamics of price differences at date t for the same good or basket i between two countries, l and f, given as $q_{ilft} = p_{ilt} - p_{ift}$, where p is the common currency price in each location. Specifically, these studies test for $\{\rho^q_{ilf} = 1\}$ in a regression of the following kind:

$$\Delta q_{ilft} = c_{ilf} + (\rho^q_{ilf} - 1)q_{ilft-1} + \sum_{h=1}^{H_q} \phi^q_{ilfh} \Delta q_{ilft-h} + \varepsilon^q_{ilft},$$

where ε^q_{ilft} is a white noise process. A lot of the earlier work that relied on aggregate price indices found evidence in favor of the null of non-convergence (see Rogoff, 1996). However, UR tests are known for their low in sample power. In order to increase the power of these tests, one can stack individual price processes to increase the sample size by increasing n instead of T, since in many cases it is not possible to increase the time dimension. Modelers have developed so-called panel UR tests (see Levin *et al.* 2002, Im *et al.* 2003) and several recent studies investigate the convergence of international (log) relative prices by implementing these tests. Examples are Goldberg & Verboven (2005) using European car market data, Imbs *et al.* (2005) using European sectoral data, Broda & Weinstein (2008) using US and Canadian Barcode data, and Crucini & Shintani (2008) using a shorter time sample of the EIU data. They all reject the null of a unit root implying that international price differentials are stationary and therefore providing evidence in favor of long-run convergence. They also find a much more rapid convergence rate than studies using aggregate price indices.

Testing for unit roots in such a panel model is usually done by implementing procedures that postulate some homogeneity across individual parameters (for instance $\rho_{ilf} = \rho_i$, $\forall\ l, f$), i.e. the stacked units are assumed to be comparable, and most importantly, it is also often assumed that there is no cross-individual dependence in the error term ε_{ilft}. Using standard panel UR tests while there is such presumable cross-dependencies among units of the panel is problematic as it induces severe test size distortions as reported in Maddala and Wu (1999). That is, one may reject the null of non-stationarity too frequently, concluding that prices are mean-reverting even though they are in fact affected by stochastic trends. For instance, O'Connell (1998) shows how neglecting to correct for cross-sectional dependence between real exchange rates (due to common macro shocks) leads to wrongly conclude in favor of long-run PPP.[10] It therefore seems important to tackle cross-dependence when assessing long-run convergence in international prices, an issue that is not dealt with in any of the above mentioned studies utilizing microeconomic price data.

Instead of the usual UR tests regressions presented above, we rely on an exchangeable model of international prices to assess the issue of international price convergence. Following Gregoir (2003),

[10] We note that cross-sectional dependence also raises estimation efficiency issues so that GLS should be preferred.

the dynamics of the exchangeable price process given in expression (1) can be rewritten as

$$\begin{aligned}\Delta p_{ilt} &= c_{il} + (\rho_1 - 1)\widehat{m}_{t-1} + (\rho_2 - 1)\widehat{m}_{it-1} + (\rho_3 - 1)\widehat{m}_{lt-1} + (\rho_4 - 1)\widehat{m}_{ilt-1} \\ &+ \sum_{h=1}^{H} \left(\phi_{1h}\Delta\overline{p}_{t-h} + \phi_{2h}\Delta\overline{p}_{it-h} + \phi_{3h}\Delta\overline{p}_{lt-h} + \phi_{4h}\Delta p_{ilt-h}\right) + \widetilde{u}_{ilt}, \end{aligned} \tag{2}$$

where $\widetilde{u}_{ilt}$ is a white noise process satisfying $\widetilde{u}_{ilt} \to u_{ilt}$ as $n \to \infty$, with $u_{ilt} = \epsilon_t + \epsilon_{it} + \epsilon_{lt} + \epsilon_{ilt}$ defined as the sum of the innovations associated with each of the components m_t, m_{it}, m_{lt} and m_{ilt} of the price process. Several sources of stochastic trends (and their combinations) can be encountered. When $\{\rho_1 = 1\}$, the global component is non-stationary, when $\{\rho_2 = 1\}$, the location-specific component is non-stationary, when $\{\rho_3 = 1\}$ the good-specific component is non-stationary, and lastly, when $\{\rho_4 = 1\}$ the good-and-location-specific component is non-stationary. Note that the widely debated convergence to the LOP holds when $\{\rho_2 < 1\}$ and $\{\rho_4 < 1\}$, no matter what are the values of ρ_1 and ρ_3. Putting it differently, in addition to dealing with dependencies between units of the panel, an interesting feature of our approach compared to the existing literature on the convergence to the LOP, is to allow for potential trends in international prices other than location-specific ones; namely, global (worldwide) and good-specific ones.

Gregoir (2003) derives the asymptotic distribution of the usual test statistics obtained after a Least Squares Dummy (LSD) estimation of the test regression (2). The asymptotics are obtained for both n and T going to infinity. In our sample, we only have a small number of time periods $T = 20$. As n is rather large (greater than 4000), one solution is to consider that one can approximate the characteristics of an infinite sample where $n, T \to \infty$ with $T/n \to 0$. Indeed, the asymptotic properties of the test statistics hold for both situations where $n, T \to \infty$ with $n/T \to 0$ (T goes to infinity "first") or with $T/n \to 0$ (n goes to infinity "first"). Therefore, relying on the asymptotic distribution of the test statistic could be considered as acceptable. Given the small number of time periods available, we prefer to rely on an asymptotic analysis where $n \to \infty$ and T is given. Along the lines of Harris & Tzavalis (1999), an asymptotic distribution for given T, and n going to infinity, can be applied to the test statistics developed for the case of common shocks. Of course this approach cannot be implemented for the global trend process, $\widehat{m}_t$, for which we can only observe one realization. Gregoir (2003) shows that the test statistic for ρ_1 has the same non-standard limiting distribution as in usual UR tests. For that component, we therefore based our analysis on a small sample approximation of this asymptotic distribution given by McKinnon (1996).

Table (3) provides the results for two different samples. The first one is made of every good in every location, so that the resulting panel is unbalanced. To deal with missing values in this case, one needs to assume these to be randomly chosen across goods and locations. However, as missing values in the EIU survey are not likely to be purely random but probably related to a country's development level and the life-cycle of each product, we also implement an analysis for a sample made of goods that are always observed in every location resulting in a balanced panel. This allows us to compare results for the unbalanced versus the balanced panel samples, which also serves to indicate the robustness of our results to missing values issues.

Table (3) shows that only the global component has a stochastic trend. The good-specific component is relatively persistent but stationary. Lastly, there are no stochastic trends in the location-specific and in the good-and-location-specific components. The fact that non-stationary components appear only in the global trend implies that international price differences for the same good across locations are stationary. Therefore, the preceding results confirm the recent ones of Goldberg & Verboven (2005), Imbs *et al.*, (2005), Broda & Weinstein (2008), and Crucini & Shintani (2008), who all reject the null of a unit root in LOP deviations. Unlike previous work, we have dealt with the problems induced by dependence across units of the panel, so that our results suggest that the stationarity of deviations from the LOP is not due to this potential concern. At the same time, our results suggest the existence of a trend in international prices that is common to every good and location. Since the sole source of non-stationarity appears to be this global component, standard panel UR test procedures applied to country relative prices as in much of the literature, would by construction fail to detect this stochastic trend.

6 The pace of convergence in international prices

6.1 Convergence to the LOP

The results from the preceding section imply that the price for the same good i across locations all share the same stochastic trend. Consequently, the price for a particular good in a particular location is cointegrated with the average price for that good across locations. In other words, deviations from the LOP are transitory up to an average gap that is constant over time.

Once the stationarity of these deviations has been established, the question of the pace of this convergence to the LOP can be investigated by estimating the following dynamic model for the product-level relative price across international locations

$$\Delta q_{ilt} = c_{il}^{q} + (\rho_1^q - 1)\widehat{m}_{lt-1} + (\rho_2^q - 1)\widehat{m}_{ilt-1} + \sum_{h=1}^{H_q}(\phi_{1h}^q \Delta\widehat{m}_{lt-h} + \phi_{2h}^q \Delta\widehat{m}_{ilt-h}) + \widetilde{u}_{ilt}^q \qquad (3)$$

where $\widetilde{u}_{ilt}^q$ is a white noise process satisfying $\widetilde{u}_{ilt}^q \to u_{ilt}^q$ when $n \to \infty$ with $u_{ilt}^q = \epsilon_{lt} + \epsilon_{ilt}$.

The literature mostly focuses on the value of the first-order autoregressive coefficient in that equation, ρ, and derives the long-term reaction from power functions of this parameter. It is therefore implicitly assumed there that the relative price process is a pure AR(1). For more complex dynamics of the kind postulated here, the transmission and correction of the shock is related to the whole set of parameters describing the dynamics, i.e. ρ and ϕ_h, (see Murray and Papell, 2002) through a complex function that satisfies $d(H) = (\rho + \phi_1)d(H-1) + \sum_{h=2}^{H}(\phi_h - \phi_{h-1})d(H-h)$, where $d(H)$ denotes the reaction at horizon H and with $d(0) = 1$. Yet another difference with what is done in previous work, is that our decomposition calls for a distinction of the reaction to economy-wide shocks (ϵ_{lt}) that will be given by

$$d_l^q(H) = (\rho_1^q + \phi_1^q)d(H-1) + \sum_{h=2}^{H}(\phi_{1h}^q - \phi_{1h-1}^q)d_l^q(H-h)$$

and the reaction to purely idiosyncratic (location-and-good-specific) shocks (ϵ_{ilt}) given by

$$d_{il}^q(H) = (\rho_2^r + \phi_2^q)d(H-1) + \sum_{h=2}^{H}(\phi_{2h}^q - \phi_{2h-1}^q)d_{il}^q(H-h)$$

with $d_l^q(0) = d_{il}^q(0) = 1$.

Table (4) provides the results. The speed of convergence to the LOP is of comparable magnitude to the one found in Crucini and Shintani (2008), namely a persistence (first-order autoregressive) parameter of 0.65. This is the case, even though they handle potential price outliers entirely different as explained in the data section, and in spite of the fact that the former paper does not handle cross-dependencies among panel units. In the first panel of Figure 5, we present graphically the speeds of convergence for our sample relative to Crucini and Shintani (2008).

An interesting aspect of our methodology is that it can be applied to show that persistence differs depending on whether the initial shock is idiosyncratic or location-specific. As shown in Table (4), whatever the sample, location-specific shocks are more rapidly corrected than good-and-location-specific ones. Figure 5 demonstrates this graphically, where the reaction to location-specific shocks is shown in the three figures on the left panel of Figure 5 and the reaction to good-and-location specific shocks is shown in the three figures on the right panel of Figure 5. This result might relate to the fact that location-specific shocks can often be related to transitory effects such us local weather conditions in a city at one point in time, which typically revert back to their temporal mean levels soon after a shock occurs.

We proceed by considering an exercise that splits the analysis of the speed of convergence according to subsamples of goods that can be classified as traded (TR) and goods that can be classified as mostly non-traded (NT) in international markets. This allows us to take a first glance at the role of trade costs in determining the speed of convergence. Likewise, we can compare the behavior of cities in developed economies (DEV) versus those in less developed ones (LDC), to allow a first glance at the potential role of income levels in determining the speed of convergence.

Indeed, the dispersion and convergence in international prices depend on whether the goods considered are traded or non-traded, or the cities considered are in developed or less developed countries. For instance, convergence to the LOP is more rapid across LDCs than across developed economies. At the same time, the dispersion of international relative prices is greater among LDCs, with a standard deviation equal to 12,100 USD, than among developed economics which have a standard deviation equal to 9,700 USD as shown in Table (1). These results are consistent with more dispersion in fixed effects and more volatile shocks for LDCs. Likewise, Table (4) shows that traded goods adjust somewhat more rapidly to shocks as compared to non-traded ones. In Figure 5, we present graphically the different speeds of convergence for developed versus less developed economies and for traded versus non-traded goods. The visual evidence confirms the small differences in convergence rates for traded versus non-traded goods, and somewhat larger differences in convergence across LDC locations as compared to convergence across cities of developed economies.

6.2 Convergence across goods inflation rates

The result that there exists only one stochastic trend common to every international price also implies that product-level relative prices inside the same country converge to constant terms, leading to convergence of product-level inflation rates. The speed of convergence to these constants can be assessed by estimating the following dynamic regression model

$$\Delta r_{ilt} = c^r_{il} + (\rho^r_1 - 1)\widehat{m}_{it-1} + (\rho^r_2 - 1)\widehat{m}_{ilt-1} + \sum_{h=1}^{H_r}(\phi^r_{1h}\Delta\widehat{m}_{it-h} + \phi^r_{2h}\Delta\widehat{m}_{ilt-h}) + u^r_{ilt}, \qquad (4)$$

where $\widetilde{u}^r_{ilt}$ is a white noise process satisfying $\widetilde{u}^r_{ilt} \to u^r_{ilt}$ when $n \to \infty$ with $u^r_{ilt} = \epsilon_{it} + \epsilon_{ilt}$. The reaction to a worldwide good-specific shock (ϵ_{it}) and to a purely idiosyncratic shock (ϵ_{ilt}) will be given respectively by $d^r_l(H)$ and $d^r_{il}(H)$, which are determined by the same type of recursion as in the preceding section, with $d^r_l(0) = d^r_{il}(0) = 1$, and parameters ρ^r and ϕ^r replacing the ρ^q's and ϕ^q's.

Table (5) provides the results. We see that convergence in inflation rates across goods is faster after an idiosyncratic (good-and-location) shock as compared to the reaction to good-specific shocks. Figure 6 demonstrates this result graphically, where the reaction to the former type of shock is shown in the three figures on the left panel of Figure 6 and the reaction to idiosyncratic shocks is shown in the three figures on the right panel of Figure 6. This result might relate to the fact that good-specific shocks are more likely to be related to the currently available production technology for a good which changes only slowly over time so that such shocks are less likely to be transitory. In Table (5), we can also see that convergence in goods inflation rates is somewhat faster for cities in the LDC sample as compared to the developing economies sample, and for internationally traded items as compared to non-traded ones. This comparison is shown graphically in Figure 6.

7 Factors that drive international prices

As argued earlier, one advantage of our approach is that it can be used to assess how important each component is relative to the others in explaining total variation and persistence. In addition, having figured out which components are important for dispersion or persistence, our methodology also enables us to proceed with mapping the different components onto economic variables related to different theories. This allows us to consider the possibility that different theory-related variables and thus different theories might be more likely to match one or another dimension of the data.

Thus, it seems interesting to provide an assessment of what economic factors drive each component in the exchangeable price dynamics. This can be achieved by regressing $\hat{m}_t$, $\hat{m}_{it}$ and $\hat{m}_{lt}$ over potential determinants such as productivity, openness, development levels, and the money supply.[11]

[11] The measurement error implied by the fact that we work with estimates rather than the true unobserved com-

We abstain from considering regressions based on the global trend, $\hat{m}_t$, since this would be based on merely twenty observations at best. In practice, and in the presence of non-stationarity in this component, the sample would be reduced even more once we consider appropriate lags and add explanatory variables that are typically not available for the last few years.

The importance of our approach here relative to prior work is that we do not restrict ourselves a priori to studying a specific dimension of our dataset. As alluded to earlier, the macroeconomic literature going back to Engel (1993) and Engel and Rogers (1996) had until recently been focusing on time variation of international relative prices, often favoring nominal or sticky price explanations of these. More recently, Crucini, Telmer, and Zachariadis (2005) focus instead on the variation of prices of particular goods across countries, emphasizing economic explanations related to product characteristics, while Crucini, Telmer, and Zachariadis (2005b) emphasize variation of relative prices across goods between country pairs, allowing for the importance of economic geography. Crucini and Yilmazkuday (2009) average the data over time and largely explain the cross-sectional component with a theory that encompasses trade and distribution costs. We argue, that these approaches might be more suitable for investigating one or another theory, and that focusing on only a certain dimension of the dataset can thus bias the results in favor or against one or another theory. At the same time, one has to recognize that certain dimensions of the data are more important in tems of total variance or persistence so that one might want to consider theories that best match those dimensions of the panel data. We proceed to explain the different components contributing to total variation in prices, in an attempt to compare the importance of different theories.

To discipline the discussion, we present a broad empirical model that considers the relative price, P_i/P, of a good i in a specific location at a specific date (with P the aggregate price level in that location at that date) as resulting from a markup, ν_i, and a trade cost (including distribution and transport costs), D_i, over the (real) marginal cost of producing the good, MC_i,

$$P_i/P = \nu_i \times D_i \times MC_i$$

We note that each of these three terms can be split into global, country, good, and idiosyncratic components. Let's assume that $Y_i = AL_i$. Then, $C(Y_i) = (W/P)(Y_i/A)$ and $MC_i = MC =$

ponent becomes smaller when the number of individuals in the sample is high.

$(W/P)(1/A)\ \forall i$, with W/P representing a real input cost. Approximating A with Y/L, then

$$P_i/P = \nu_i \times D_i \times (W/P)/(Y/L) \Leftrightarrow P_i = \nu_i \times D_i \times W/(Y/L).$$

Converting the price of good i in different countries into US dollars leads to

$$P_i/S = \nu_i \times D_i \times W/(Y/L)(1/S).$$

Supposing that $\nu_i = \nu(S, X_i)$ and $D_i = D(Z_i)$ with X_i, Z_i variables that can be both either location or product specific, then international price differences can be analyzed in a log-linear regression model where $(p_i - s)$ is regressed over $y - l$, w, s and variables x_i and z_i chosen to account for $\log \nu_i$ and d_i (denoting the log values in lowercase letters.)

7.1 City-specific fluctuations

We define the city-specific component we set out to explain here, as $\overline{q}_{lt} = (1/n_i)\sum_i q_{ilt}$. This is the average relative price in city l. We consider a number of economic variables in addition to the city-specific lagged average price, in an attempt to explain this component of the data. We present these below in the order in which they appear in Table (6). We consider in turn the country-level log of real gdp per worker related to productivity and development levels, and then the service sector share of GDP as a measure of the distribution sector. Moreover, we consider a number of city-specific variables: electricity cost, regular unleaded petrol, and city-specific residential rent for a two-bedroom unfurnished apartment. These are meant to capture local production and distribution costs. Moreover, we consider city-specific population as a proxy for local scale economies in distribution and as a proxy of the degree of local competition. We then add a couple of nominal variables related to a country's (log) money supply (M1 in billion $US), and the log of a country's nominal exchange rate relative to the US dollar. Finally, we consider two measures of economy-wide openness as given by a country's import share in GDP, and a country's average tariff on imports.

Table (6) gives the results for regressions of $\overline{q}_{lt}$ over the preceding city-specific and country-specific macro-variables. The results are for the sample restricted to goods and cities always present in the survey. Input costs, relating to the cost of electricity, petrol, and residential rents in each city, appear to be the most robust set of determinants acting positively on prices. City population systematically has a negative impact, consistent with economies of scale in distribution and production as well as higher competition and lower markups for bigger cities. Finally, monetary policy has a positive significant impact on prices. The impact of real GDP per worker is positive but becomes insignificant as soon as we include a measure of the relative size of the service sector in column (2) of Table (6), and remains insignificant throughout. Finally and surprisingly, we also find that country differences in terms of import tariff rates have no impact on international LOP deviations across cities.

7.2 Good-specific fluctuations

The good-specific component we set out to explain here, is defined by $\overline{r}_{it} = (1/n_l)\sum_l r_{ilt}$. This is the average relative (to the average-priced good) price over cities for good i, referred to as *rlprice*. We consider the following regressors, in addition to the lagged relative price, to explain this component: the standard deviation (across locations) of the goods' price, the average (across locations) import tariff for the good, and the log of the goods' average import value. The standard deviation across cities is a measure of the dispersion for each good which captures a variety of factors that might inhibit trade and thus result in higher dispersion. The cross-country average tariff for each good is a direct measure of trade costs. Finally, the import value captures the degree of realized trade specific to each good. One would desire a number of additional good-specific measures of theoretical variables relating to trade costs, distribution costs, and mark-ups. However, such measures are hard to come by, so we defer from doing so for the purposes of this paper. As such, the results being discussed below are merely indicative. Table (7) presents the resulting estimates. The interesting finding here is the significant role of tariffs in determining final prices. This is in accordance with common wisdom, but, interestingly, not evident when considering deviations from the LOP across cities as shown previously in Figure (6).

8 Conclusion

This paper set out to decompose the variance present in a panel of international price data into different components, to disentangle the sources of non-stationarity present in these international prices, and to assess the convergence rates for the remaining time-varying components. In the first instance, we have shown that the time-invariant component of the good-specific dimension accounts for the great majority of total variance while the global component accounts for less than one percent of this. Relating to the second question, we have shown that, nevertheless, the sole source of non-stationarity appears to be this global component so that the presence of non-stationariy would go undetected by construction when one builds country relative prices as in much of the literature. In the third instance, we have shown that convergence as a reaction to a location-specific shock is faster than after an idiosynratic shock. This result might relate to the fact that location-specific shocks can realistically be transitory effects related to local weather or other local conditions likely to change over time, whereas good-location-specific shocks might relate to the currently available production or distribution technologies for certain goods in certain locations which can be slower to change over time. Similarly, good-specific shocks are more likely to be related to the available production technology for a good which changes only slowly over time, relative to location-related shocks.

Our ultimate goal was to relate the different components of this panel to economic theory explanations. To this end, we attempted to explain the location-specific and good-specific components of prices, and found a robust impact of input costs (related to production and distribution costs), city population (related to distribution costs and markups), and nominal factors such us the money supply on average LOP deviations across cities, and a robust impact of trade costs in the form of tariff rates on goods relative prices.

By considering different dimensions and components of a panel of prices, this paper also serves the goal of placing previous work on this topic that focused on particular dimensions or components of prices, within a more general framework that allows gauging the relative significance of different components and different economic-theory explanatory variables. Additional work, would be needed in this last direction in order to properly assess what appears to be the most important source of variation in the data; that is, to explain the time-invariant good-specific dimension using a number of additional good-specific economic-theory factors. Crucini, Telmer, and Zachariadis (2005) take a first step towards this direction, but further work with more detailed data is called for in order to explain what is evidently the richest source of variation in the data.

References

[1] Altissimo, Filippo and Mojon, Benoît and Zaffaroni, Paolo (2008) "Can aggregation explain the persistence of inlfation?" *Journal of Monetary Economics*.

[2] Atkeson, Andrew, and Ariel Burstein (2007) "Pricing-to-market in a Ricardian Model of International Trade" *The American Economic Review Papers and Proceedings*.

[3] Atkeson, Andrew, and Ariel Burstein (2009) "Pricing-to-Market, Trade Costs, and International Relative Prices," *The American Economic Review*.

[4] Bergin, R. Paul and Reuven Glick, Endogenous Nontradability and Macroeconomic Implications, NBER Working Paper No. W9739, June 2003.

[5] Bergin, R. Paul and Reuven Glick (2007) "Global price dispersion: Are prices converging or diverging?" *Journal of International Money and Finance* 26(5) 703-729.

[6] Boivin, Jean, Giannoni, Marc P. and Mihov, Ilian (2008) "Sticky prices and monetary policy: Evidence from disaggregated data" *The American Economic Review*.

[7] Christian Broda David E. Weinstein (2008) "Understanding International Price Differences using barcode data," NBER Working Paper 14017.

[8] Carlos Carvalho and Fernanda Nechio (2008) "Aggregation and the PPP Puzzle in a Sticky Price Model" unpublished manuscript, FRB New York and Princeton University.

[9] Crucini, Mario, Chris Telmer, and Marios Zachariadis, (2005) "Understanding European Real Exchange Rates," *The American Economic Review*, 95:3, 724-738.

[10] Crucini, Mario, Chris Telmer, and Marios Zachariadis, (2005b) "Price Dispersion: The role of Borders, Distance, and Location," unpublished manuscript.

[11] Crucini, M.J. and Shintani, M. (2008). Persistence in the Law of One Price deviations: Evidence from micro-data. Journal of Monetary Economics, 55, 629-644.

[12] Crucini, Mario J. and Hakan Yilmazkuday (2009) "A Model of International Cities: Implications for Real Exchange Rates" NBER Working Paper No. 14834.

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

[14] Engel, Charles and John Rogers (1996) 'How wide is the border,' *The American Economic Review* 86, 1112-1125.

[15] Engel C. and J. Rogers (2004). European product market integration after the euro, Economic Policy, 39:347-384.

[16] Gregoir S. (2003). Representation and statistical analysis of weakly linearly exchangeable dynamic panels. mimeo

[17] Goldberg, P.K. and M. Knetter (1997) "Goods Prices and Exchange Rates: What Have We Learned?" Journal of Economic Literature, 35: 1243-1272.

[18] Goldberg, P.K. and Verboven, F. (2005). Market integration and convergence to the Law of One Price: Evidence from the European car market. Journal of International Economics 65:49–73.

[19] Imbs, J., Mumtaz, A., Ravn, M.O. and Rey, H. (2005). PPP strikes back: Aggregation and the real exchange rate. Quarterly Journal of Economics 120:1–43.

[20] Im, K., M.H. Pesaran and Y. Shin (2003). Testing for unit roots in heterogeneous panels. Journal of Econometrics, 115:53–74.

[21] Keane, M. P. and Runkle, D. E. (1990). Testing the rationality of price forecasts: New evidence from panel data. American Economic Review, 80:714–35.

[22] Kwiatkowski, D., Phillips, P. C. B., Schmidt, P. and Shin, Y. (1992). Testing the null hypothesis of stationarity against the alternative of a unit root. Journal of Econometrics, 54:159-178.

[23] Levin, A., Lin, C. F. and Chu, J. (2002). unit root tests in panel data: asymptotic and finite sample properties. Journal of Econometrics, 108:1–24.

[24] Maddala, G. S. and Wu, S. (1999). A comparative study of unit root tests with panel data and a new simple test: evidence from simulations and bootstrap Oxford Bulletin of Economics and Statistics, 61.

[25] Moon, R. and Perron, B. (2004). Testing for a unit root in panels with dynamic factors. Journal of Econometrics, 122:81–126.

[26] Murray, C.J. and Pappell, D. (2002). The purchasing parity persistence paradigm. Journal of International Economics 56:1–19.

[27] Maćkowiak, Bartosz, Moench, Emanuel and Wiederholt, Mirko (2009). Sectoral data and models of price setting. Journal of Monetary Economics.

[28] Emi Nakamura and Dawit Zerom (2009) "Accounting for incomplete pass-through." NBER Working Paper 15255.

[29] Obstfeld, Maurice and Kenneth Rogoff (2000) "The Six Major Puzzles in International Macroeconomics: Is there a common cause?" NBER working paper No. 7777

[30] O'Connel, P.G.J. (1998). The overvaluation of purchasing power parity. Journal of International Economics 44:1–19.

[31] Pesaran, M. H. and Smith, R. P. (1995). Estimating long-run relationships from dynamic heterogeneous panels. Journal of Econometrics, 68:79–113.

[32] Reis, Ricardo and Watson, Mark W., (2009) "Relative goods' prices, pure inflation and the Phillips correlation," *Mimeo.*

Table 1: Descriptive Statistics

SAMPLE PERIOD: 1990–2009[a]	WHS	LDC	DEV	NT	TR
PRICE LEVEL (USD), P_{ilt}					
Mean	1621	1750	1529	1328	1707
Med	7.7	5.8	9.2	224	4.8
Max	534896	534896	452783	534896	436190
Min	$3e^{-4}$	$3e^{-4}$	$2e^{-2}$	$1e^{-2}$	$3e^{-4}$
Std Dev	10851	12269	9703	3499	12176
# of obs	678153	283483	394670	15340	525813

[a]WHS = Whole sample; LDC = less developed countries (income per capita < 12000$ per year); DEV = developed countries; NT = non-traded goods; TR = traded goods (see classification in Appendix)

Table 2: Variance decomposition

SAMPLE PERIOD: 1990–2009[a]	WHS	LDC	DEV	NT	TR
LOG PRICE, p_{ilt}					
$\mathrm{V}(p_{ilt})$	7.46	7.79	7.16	7.18	5.59
$\widehat{\mathrm{V}}(m_t)$	.03	.03	.03	.03	.03
$\widehat{\mathrm{V}}(m_{lt})$	.12	.11	.06	.12	.12
$\widehat{\mathrm{V}}(m_{it})$	7.10	7.11	7.08	6.58	5.26
LOG DEVIATIONS OF LOP, q_{ilt}					
$\widehat{\mathrm{V}}(q_{ilt})$	.35	.42	.24	.54	.29
$\widehat{\mathrm{V}}(m_{lt})$	.12	.11	.05	.12	.12
$\widehat{\mathrm{V}}(\mu_l)$	.09	.08	.03	.08	.08
LOG GOODS RELATIVE PRICES, r_{ilt}					
$\widehat{\mathrm{V}}(r_{ilt})$	7.32	7.66	7.09	6.97	5.48
$\widehat{\mathrm{V}}(m_{it})$	7.10	7.11	7.08	6.58	5.26
$\widehat{\mathrm{V}}(\mu_i)$	7.08	7.10	7.07	6.57	5.25

[a]WHS = Whole sample; LDC = less developed countries (income per capita < 12000$ per year); DEV = developed countries; NT = non-traded goods; TR = traded goods (see classification in Appendix)

Table 3: Testing for unit-roots in the price levels

SAMPLE PERIOD: 1990–2009[a]	**(1)**	**(2)**
EXCHANGEABLE PANEL		
$\widehat{m}_{t-1}$	−1.25	−1.25
	(.627)	(.627)
$\widehat{m}_{lt-1}$	−213	−67
	(.000)	(.000)
$\widehat{m}_{it-1}$	−101	−31
	(.000)	(.000)
$\widehat{m}_{ilt-1}$	−804	−278
	(.000)	(.000)
# of obs	504025	357792

[a]Column (1) = Whole sample; Column (2) = Restricted sample (Goods and cities always present). Results are DF unit-root test statistics for panel model corrected for small T. Number in brackets are the significance level of the test statistic.

Table 4: Convergence to the LOOP

SAMPLE PERIOD: 1990–2009[a]	**WHS**	**LDC**	**DEV**	**NT**	**TR**
RESTRICTED SAMPLE (Goods and cities always present)					
# of obs	357792	110160	247632	68160	289632
1. *Standard regressions*					
q_{ilt-1}	−.34	−.39	−.31	−.31	−.36
	(.003)	(.006)	(.003)	(.011)	(.003)
2. *Exchangeable regressions*					
$\widehat{m}_{lt-1}$	−.39	−.44	−.36	−.36	−.40
	(.005)	(.007)	(.004)	(.011)	(.005)
$\widehat{m}_{ilt-1}$	−.33	−.38	−.30	−.30	−.34
	(.004)	(.008)	(.003)	(.015)	(.003)

[a]WHS = Whole sample; LDC = less developed countries (income per capita < 12000$ per year); DEV = developed countries; NT = non-traded goods; TR = traded goods (see classification in Appendix). Results are first-order autoregressive parameter estimates in equation (3). Additional regressors (not shown) are a constant and 3 lags of Δq_{ilt}. Estimation is achieved by the fixed effect (within) method. Standard errors below the coefficient estimates are White's robust estimators.

Table 5: Convergence across goods inflation rates

SAMPLE PERIOD: 1990–2009[a]	**WHS**	**LDC**	**DEV**	**TR**	**NT**
RESTRICTED SAMPLE (Goods and cities always present)					
# of obs	357792	1100160	247632	68160	289632
1. *Standard regressions*					
r_{ilt-1}	$-.30$	$-.31$	$-.27$	$-.30$	$-.27$
	(.003)	(.006)	(.003)	(.011)	(.003)
2. *Exchangeable regressions*					
$\widehat{m}_{it-1}$	$-.25$	$-.26$	$-.23$	$-.25$	$-.21$
	(.005)	(.007)	(.004)	(.011)	(.005)
$\widehat{m}_{ilt-1}$	$-.31$	$-.32$	$-.28$	$-.31$	$-.28$
	(.004)	(.008)	(.003)	(.015)	(.003)

[a] WHS = Whole sample; LDC = less developed countries (income per capita < 12000$ per year); DEV = developed countries; NT = non-traded goods; TR = traded goods (see classification in Appendix). Results are first-order autoregressive parameter estimates in equation (4). Additional regressors (not shown) are a constant and 3 lags of Δr_{ilt}. Estimation is achieved by the fixed effect (within) method. Standard errors below the coefficient estimates are White's robust estimators.

Table 6: Deviation to the LOP across cities

	(1)	(2)	(3)	(4)	(5)	(6)
lagged price	0.622***	0.623***	0.396***	0.391***	0.376***	0.331***
	(0.049)	(0.046)	(0.048)	(0.046)	(0.047)	(0.055)
real gdp per worker	0.141**	0.078	0.066	0.046	0.027	0.058
	(0.054)	(0.093)	(0.050)	(0.053)	(0.056)	(0.064)
services share		0.002	0.001	0.001	0.000	0.000
		(0.002)	(0.001)	(0.001)	(0.001)	(0.001)
electricity cost			0.090***	0.092***	0.086***	0.110***
			(0.025)	(0.025)	(0.025)	(0.027)
petrol cost			0.149***	0.150***	0.149***	0.140***
			(0.020)	(0.020)	(0.020)	(0.022)
rental cost			0.136***	0.136***	0.141***	0.140***
			(0.018)	(0.019)	(0.018)	(0.017)
city population size				-0.111**	-0.136**	-0.171**
				(0.052)	(0.054)	(0.071)
money supply					0.047**	0.061***
					(0.020)	(0.019)
nominal exchange rate					-0.003	-0.008
					(0.007)	(0.009)
imports share						-0.001**
						(0.001)
average tariff rate						0.021
						(0.080)
constant	0.001	-0.002	-0.026***	-0.021***	-0.018**	-0.020*
	(0.005)	(0.010)	(0.006)	(0.006)	(0.007)	(0.010)
Observations	1598	1416	1287	1270	1230	1132
Number of cities	95	89	81	80	76	76
adjusted R^2	0.403	0.423	0.687	0.691	0.705	0.717

Robust standard errors in parentheses

*** p<0.01, ** p<0.05, * p<0.1

Table 7: Differences in inflation rates across goods

	(1)	(2)	(3)	(4)	(5)	(6)	(7) drlprice
price lag	-0.127***	-0.125***	-0.133***	-0.125***	-0.132***	-0.132***	-0.132***
	(0.015)	(0.018)	(0.010)	(0.018)	(0.010)	(0.010)	(0.010)
standard deviation	-0.038***			-0.025	-0.015		-0.011
	(0.012)			(0.016)	(0.015)		(0.015)
tariff		0.052***		0.050***		0.046***	0.046***
		(0.013)		(0.013)		(0.013)	(0.013)
import value			-0.001		-0.001	-0.000	-0.001
			(0.001)		(0.001)	(0.001)	(0.001)
Constant	0.028***	-0.057***	-0.063***	-0.045***	-0.058***	-0.066***	-0.062***
	(0.006)	(0.007)	(0.017)	(0.013)	(0.018)	(0.017)	(0.018)
Observations	5092	3918	3445	3918	3445	3445	3445
Number of goods	268	214	214	214	214	214	214
adjusted R^2	0.0842	0.0826	0.0819	0.0837	0.0822	0.0864	0.0865

Robust standard errors in parentheses

*** p<0.01, ** p<0.05, * p<0.1

Figure 1: Evolution of world-average prices (Global trend)

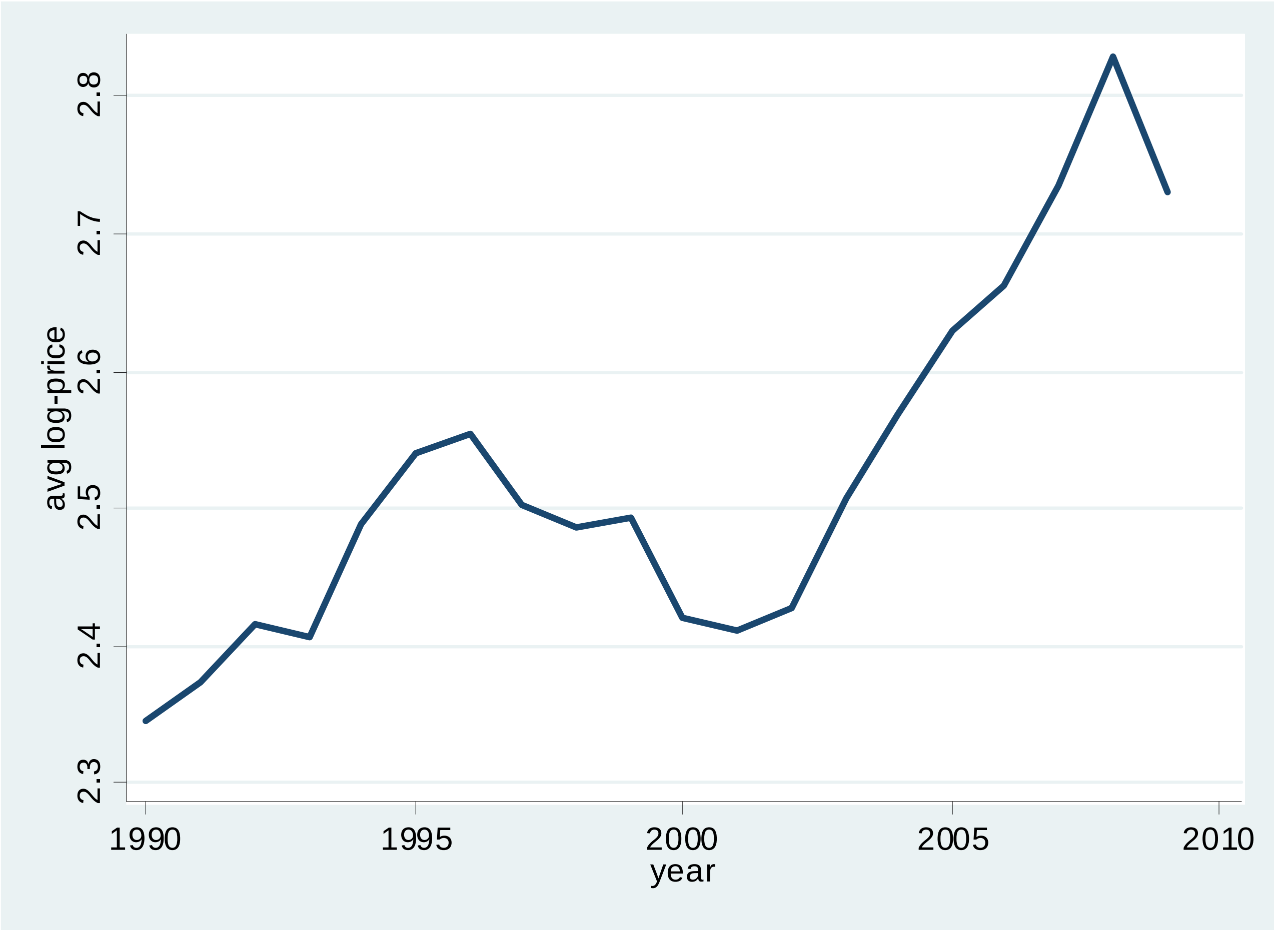

Figure 2: Evolution of cities relative prices

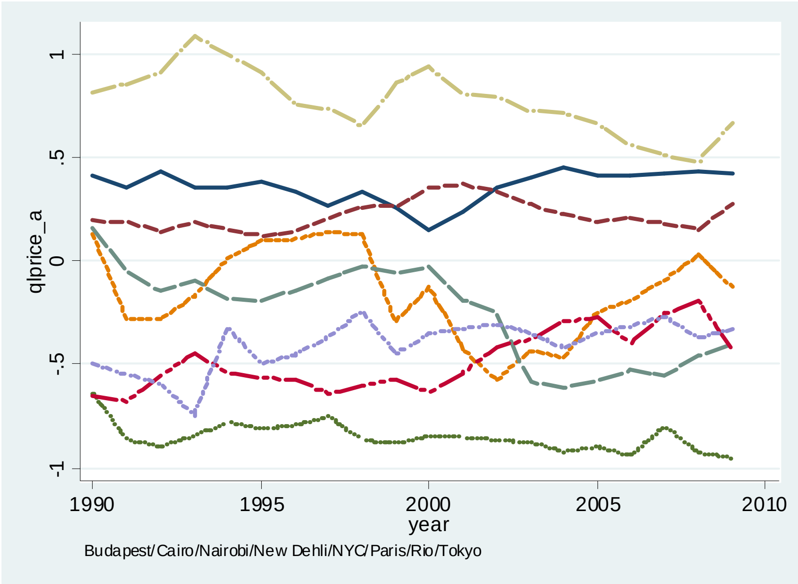

Figure 3: Evolution of goods relative prices

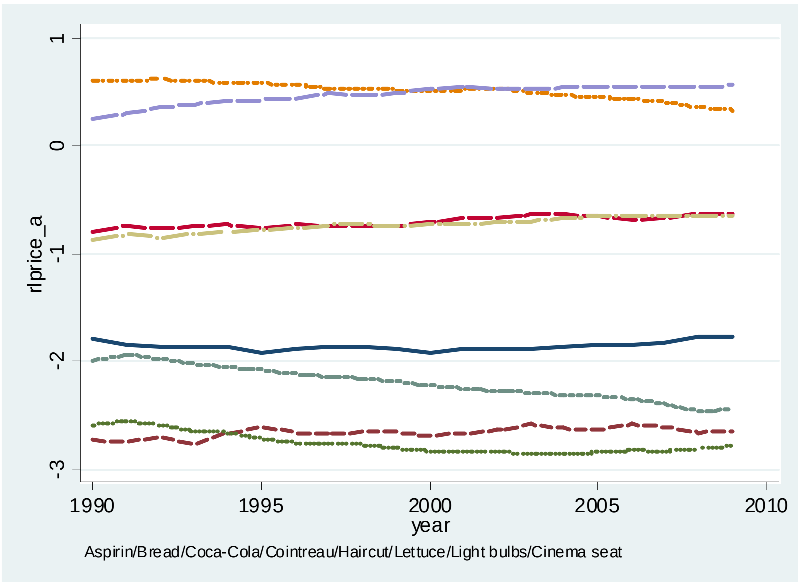

Figure 4: Distribution of the share of the time-varying component in total variance

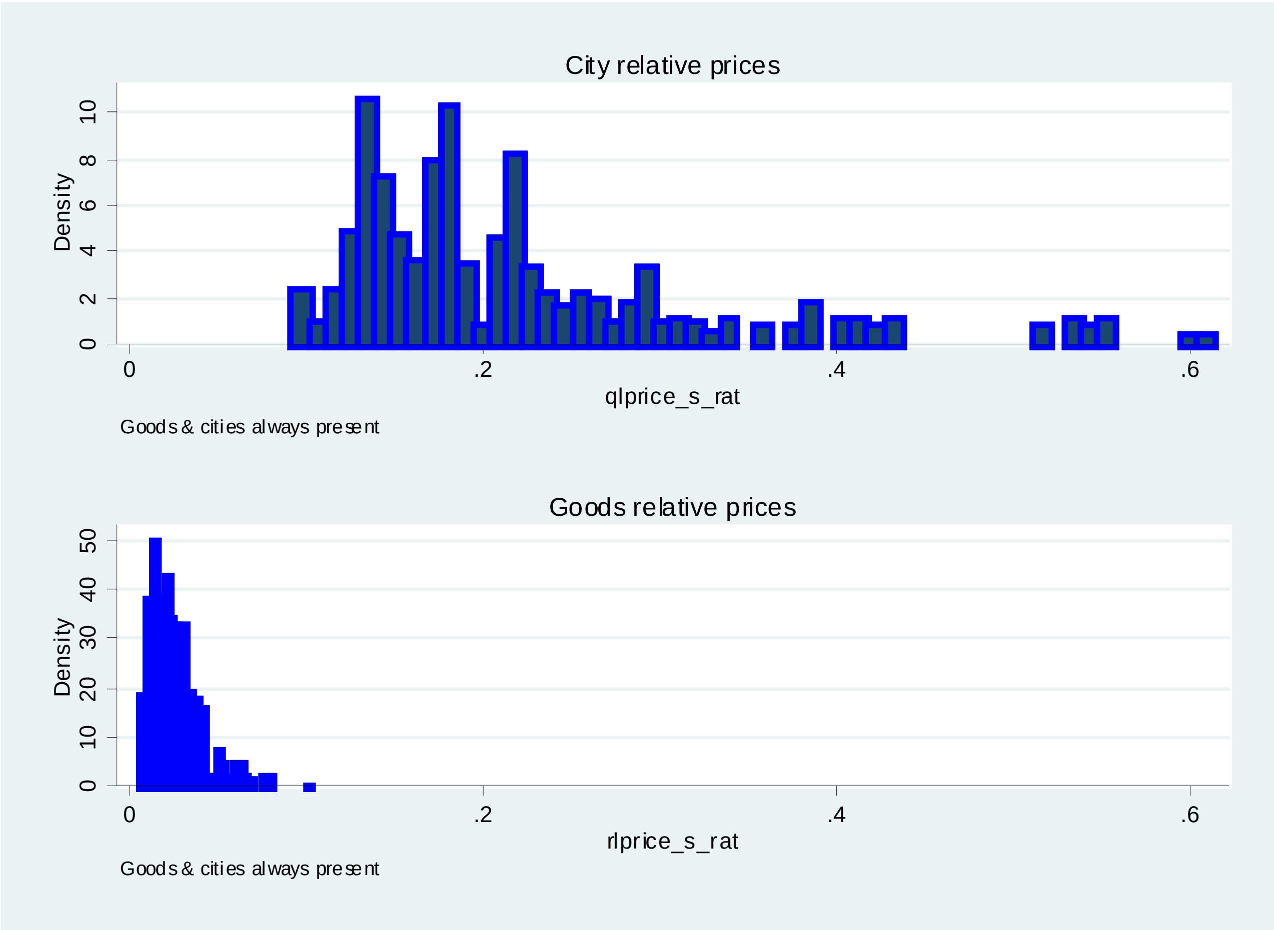

Figure 5: Convergence to the LOP across cities

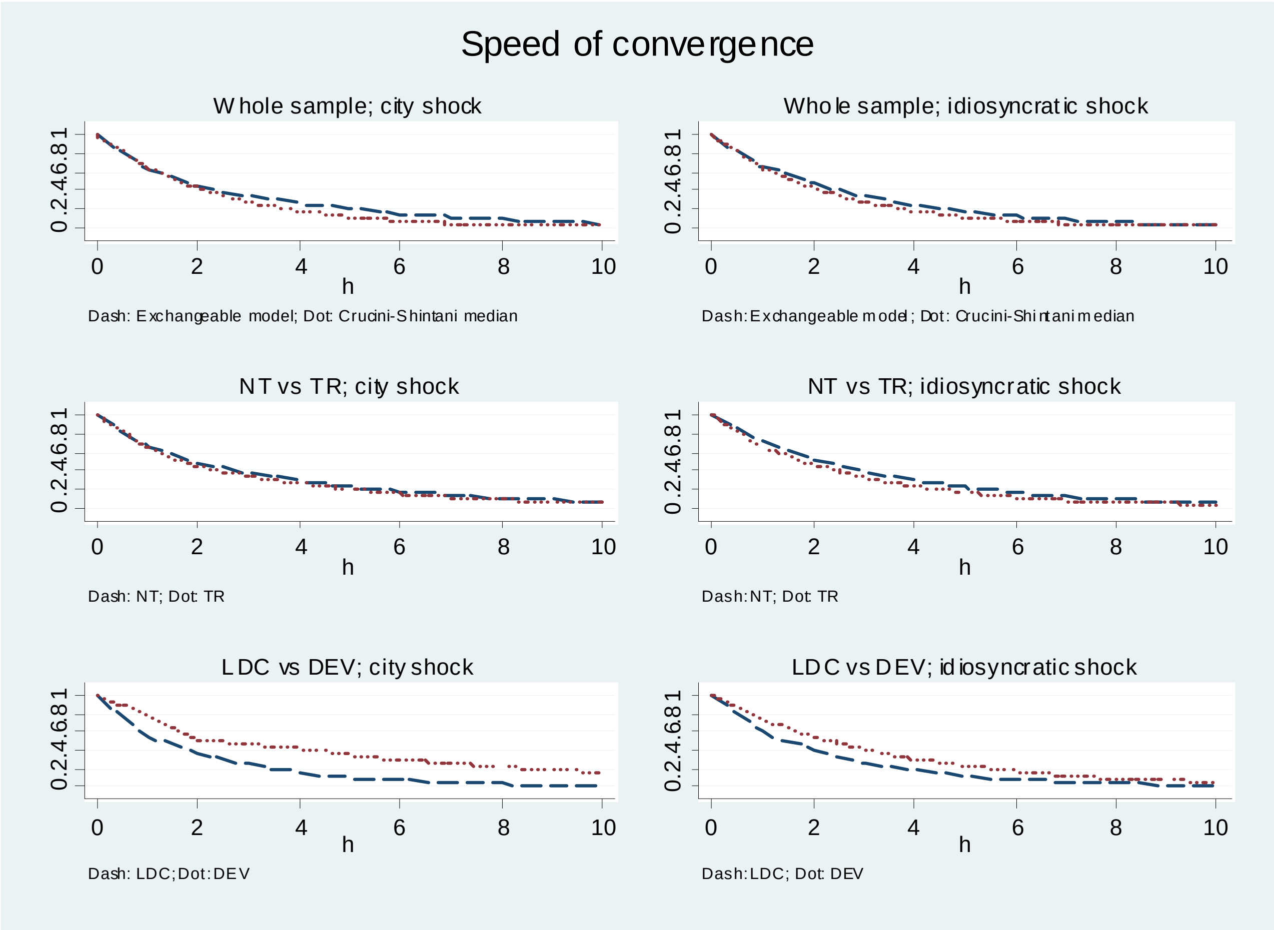

Figure 6: Convergence to the same inflation rate within cities