

NBER WORKING PAPER SERIES

UNDERSTANDING LONG-RUN PRICE DISPERSION

Mario J. Crucini
Hakan Yilmazkuday

Working Paper 18811
<http://www.nber.org/papers/w18811>

NATIONAL BUREAU OF ECONOMIC RESEARCH
1050 Massachusetts Avenue
Cambridge, MA 02138
February 2013

We especially thank Mary Amiti, Eric Bond, Paolo Pesenti and Robert Vigfusson for valuable comments and criticisms as well as participants at numerous conferences and seminars. We gratefully acknowledge the financial support of the National Science Foundation (SES-1030164). The views expressed herein are those of the authors and do not necessarily reflect the views of the National Bureau of Economic Research.

NBER working papers are circulated for discussion and comment purposes. They have not been peer-reviewed or been subject to the review by the NBER Board of Directors that accompanies official NBER publications.

© 2013 by Mario J. Crucini and Hakan Yilmazkuday. All rights reserved. Short sections of text, not to exceed two paragraphs, may be quoted without explicit permission provided that full credit, including © notice, is given to the source.

Understanding Long-run Price Dispersion
Mario J. Crucini and Hakan Yilmazkuday
NBER Working Paper No. 18811
February 2013
JEL No. F0,F11,F15,F63

ABSTRACT

We use a unique panel of retail prices spanning 123 cities in 79 countries from 1990 to 2005, to uncover six novel properties of long-run international price dispersion. First, at the PPP level, virtually all (91.6%) of price dispersion is attributed to service-sector wages, consistent with a dominant role of the retail distribution margin. Second, at the level of individual goods and services, the average contribution of service-sector wages is significantly reduced, one-third as large (31.9%). This reflects the fact that good-specific sources of price dispersion, such as trade costs and good-specific markups, tend to average out across goods. Third, at the LOP level, borders and distance contribute about equally to price dispersion with distance elasticities consistent with the existing trade gravity literature which links trade volumes (rather than relative prices) to borders and distance. Fourth, in the cross-section, price dispersion is rising in the distribution share consistent with the notion that baby-sitting services and haircuts embody local wages to a far greater extent than highly traded manufactured goods. Fifth, we provide the first estimates of distribution margins at the micro-level and show them to be very different across goods and substantial in the aggregate, where they account for about 55% of consumption expenditure. Sixth, these estimates are broadly consistent with more aggregated U.S. NIPA measures currently used in the literature.

Mario J. Crucini
Department of Economics
Vanderbilt University
Box 1819 Station B
Nashville, TN 37235-1819
and NBER
mario.j.crucini@vanderbilt.edu

Hakan Yilmazkuday
Department of Economics
Florida International University
Miami, FL 33199
skuday@gmail.com

1. Introduction

Figure 1 is what we seek to explain, it plots long-run levels of retail prices in U.S. dollars of about 300 goods and services across 123 cities of the world against long-run levels of *hourly wages of domestic cleaning help*, a proxy for wages in each city’s service sector. Long-run in this context means time-averaged deviations computed over the period 1990 to 2005. The prices and wages used to construct these averages are from the Economist Intelligence Unit (EIU) World Cost of Living Survey which spans 79 countries. As far we know, this is the first study to use time-averaged data to study long-run deviations from the Law of One Price (LOP) and Purchasing Power Parity (PPP).

The use of time-averaged data is designed to eliminate the transitory deviations from the LOP in these data documented by Crucini and Shintani (2008). The theoretical focus thus shifts from nominal frictions and transitory dynamics of relative prices emphasized in the business cycle literature to real frictions emphasized in the trade literature. Following the excellent survey by Anderson and van Wincoop (2004), we adopt a broad definition of trade costs, including all of the costs associated with the movement of goods from the factory gate to the final consumer.¹

Our approach boils down to measuring the components of the unit cost of a good sold in a retail outlet in a particular city under the assumption of perfect competition. In logarithms, the retail price of good i , sold in city j relative to city k , p_{ijk} , is assumed to be given by:

$$p_{ijk} = -z_{jk} + \alpha_i w_{jk} + (1 - \alpha_i) q_{ijk} \quad (1.1)$$

where z_{jk} is relative TFP in the distribution (retailing) sector and w_{jk} is the relative wage paid to labor in the retail sector.² The variable, q_{ijk} , denotes the relative import price of the good being sold and is assumed to differ across cities due to iceberg shipping costs. The parameter α_i is the good-specific share of unit cost contributed by distribution costs. Our data includes, p_{ijk} and w_{jk} , but not α_i , z_{jk} or q_{ijk} . Identification of α_i is achieved by assuming wages are uncorrelated with either trade costs or retail productivity. The notion

¹Burstein, Eichenbaum and Rebelo (2007) coined the phrase *distribution costs* to encompass both traditional trade costs (tariffs and shipping costs) and retail costs (labor and capital employed in the retail sector).

²An earlier version of the paper included capital as an input in the production function for retail services. Since the cost of capital is more challenging to measure than hourly wages we subsume it in the z_j term.

that traditional trade costs are orthogonal to wage differences is based on the fact that there is little evidence in the data to indicate that either tariff barriers or iceberg shipping costs are correlated with wage levels. The assumption that wages and retail productivity are uncorrelated is what defines retail productivity.

To see the implication of this, consider the same restaurant meal sold in Istanbul and New York city. According to the US NIPA, a restaurant meal has a distribution share of about 0.75. Suppose that restaurants in Istanbul and New York happen to pay the same common currency price for the ingredients used to produce this meal and that the New York restaurateur must pay a 100% higher wage for cooks and waiters compared to his counterpart in Istanbul. Further, suppose the EIU data indicates that the price of the meal is only 50% higher in New York compared to Istanbul; by our definition, New York is 25% more productive than Istanbul in the retail sector ($50\% = -25\% + 0.75(100\%)$).

In reality, these retail productivity gaps may arise from differences in the physical capital employed in the retail sectors across the two cities (including both public infrastructure and private structures) or human (managerial) capital, to name just two. As a practical matter, our method also does not allow us to distinguish retail productivity differences from differences in markups across cities. Thus, the price may be higher in New York because restaurants are charging higher markups over cost than their counterparts in Istanbul.

Aggregating across goods, our model predicts the long-run deviation of the price level, across cities j and k , will follow:

$$p_{jk} = -z_{jk} + \alpha w_{jk} + (1 - \alpha)q_{jk} \tag{1.2}$$

Here, p_{jk} , are constructed as expenditure-weighted averages of the p_{ijk} . By construction, the parameter α is the cost share of distribution in aggregate consumption ($\alpha \equiv \sum_i \alpha_i \omega_i$, where ω_i is the consumption expenditure share of good i). The price level deviations arising from the traditional trade cost term, q_{jk} , tend to be small reflecting the common-sense notion that trade costs tend to average out across goods.³

In Figure 1, the estimated line through the scatter of price levels has a slope of 0.52 and an R^2 value of 0.37.⁴ In words: a doubling of wages is associated with a 52 percent higher

³The traded input index is a bit subtle since the prices are both good and destination specific so we normalize the weights as follows for expositional convenience $(1 - \alpha)q_j \equiv (1 - \alpha) \sum_i \omega_i \frac{\alpha_i}{1 - \alpha} q_{ij}$.

⁴The estimation is by geometric mean regression to consider for possible measurement errors in both the price and wage data. A common set of consumption expenditure weights are used for all cities. These

price level. This finding is typically associated with the seminal work of Harrod (1933), Balassa (1964), and Samuelson (1964). The HBS theory assumes that the Law of One Price (LOP) holds for traded goods but not for non-traded goods. To the extent that the variable plotted on the x-axis, wages for domestic help, proxies for the unit costs of producing purely non-traded goods, the slope of 0.52 estimated using price levels would be the share of non-traded goods in total consumption expenditure, α . As it turns out, the expenditure weighted distribution share, α , is 0.55, lending credibility to our explanation.

The success of HBS, however, is deceptive as it completely abstracts from two sources of deviations from the LOP in final traded goods. The first source of international price variation comes from traditional trade costs. While we find these trade costs aggregate toward a modest number at the PPP level, they can be very substantial at the microeconomic level. Second, the classical dichotomy upon which the HBS theory is based is patently false. There are virtually no goods for which $\alpha_i = 0$ and thus no goods with a zero sloped line in Figure 1.⁵ We fill this void by explicitly estimating the good-level distribution share parameters while using a regression technique to recover the role of geographic distance and borders in determining LOP deviations in the traded inputs, q_{ijk} . That this source of cross-sectional variance is important is evident from the fact that repeating the same regression with wages on the right-hand-side, but LOP deviations on the left-hand-side, the slope is virtually identical, equaling 0.53, but the R^2 value drops by a factor of three to 0.11. The lower explanatory power of the distribution wedge at the microeconomic level suggests the need for good-specific trade costs and other factors that prevent the LOP from holding in the long-run.

Turning to the details of our findings, we have two sets of results, one for relative price levels (PPP) and the other at the level of individual goods (LOP). The variance of price levels for international city pairs is almost entirely explained by international wage differences, 92% by our estimate. Restricting attention to OECD international pairs significantly elevates the retail productivity effect such that it accounts for 61% of PPP deviations. Because these countries have similar wage levels, it is not surprising that the HBS theory is not as helpful in accounting for what limited price dispersion does exist. The role of cross-city wage differences is limited even further when the sample is restricted to city pairs within

consumption expenditure weights are taken from the PWT, averaged across all OECD nations.

⁵Sposi (2013) uses price indices from the Penn World Tables and finds a slope coefficient of 0.02 when the price index for machinery and equipment is regressed on income per worker.

the same country with wage differences now accounting for a mere 8% of price level (PPP) dispersion. It is important to keep in mind that the amount of price level dispersion across cities that are located in the same country is a trivial 3-5%, it should be clear that a modest amount of wage or retail productivity variance goes a long way in terms of accounting for the lion's share of the variance. The thrust of the PPP analysis is that when long run price level differences are consequential, the differences are attributable to the level of economic development, not traditional trade frictions or retail productivity.

Turning to LOP deviations the table turns dramatically in favor of borders and trade costs and away from the wages and retail productivity as explanatory factors. For international city pairs, the explanatory power of the HBS theory (wage dispersion) falls by a factor of three, to about 32%. Traditional theories of trade that emphasis distance and borders account for the lion's share, about 41%. City effects account for almost none of the international LOP variation. Essentially, this is because international LOP deviations are both large and idiosyncratic to the good once we condition on the wage level. The remainder is a residual term, which may reflect good and location-specific markups as well as other variables omitted from the model.

It is important to be up-front about what our measurements leave out and what implication this has for the interpretation of the results. Perhaps the most obvious omission is the neglect of markups. There are two places where it seems natural to allow for markups over marginal cost. The first is at the retail level, a markup of retail prices over the unit cost of retailing (including both distribution and traded input costs) that differs across city pairs. A city may have a relatively high price level than we predict based on relative wage costs and relative traded input costs (i.e., $p_{jk} - \hat{\alpha}w_{jk} - q_{jk} = -z_{jk} > 0$) either because it has a relatively inefficient retail sector ($z_{jk} < 0$, perhaps due to poor public infrastructure) or because retailers have market power and choose to set different markups over marginal cost across cities. In the latter case, the $z_{jk} < 0$ term is in fact picking up a relatively high markup in city j compared to k , not relative inefficiency in retail in city j compared to k . Arguably, cities with inefficient infrastructure may also tend to be those with less competitive retailing. Consider a very efficient big-box retailer such as Walmart in a U.S. suburb compared to an open market in the center of Istanbul. One would expect the average markup over marginal cost to fall upon Walmart's entry into the Istanbul market in the same fashion that Walmart has lowered retail prices in the United States over time. While this is plausible,

our data are simply not up to the task of exploring these alternative interpretations of the retail productivity term and these issues are thus left to future work.

The second place a markup would naturally appear is between the factory gate price and the price the retailer pays for the traded input, q_{ijk} . Thus some of the variation in traded prices we attribute to borders, distance and a residual term could be due to markups that exporters charge to the various destination markets. Where such a markup over cost on imported goods gets allocated in our variance decomposition depends crucially on its covariance with other location-specific variables used in our regression framework. If the markup is not correlated with wages, distance or borders, it will be relegated to the residual term.

More problematic is the idea that markups of retail prices over factory gate prices are positively correlated with wages in the destination market since this is inconsistent with the orthogonality assumption used to estimate the distribution share. The notion that markups are increasing in the wage levels in the destination market was first theoretical developed and empirical documented in a very interesting paper by Alessandria and Kaboski (2011). They estimated a robust positive correlation between good-level U.S. export unit values and aggregate wage levels in the destination market. They attribute this correlation to costly search by consumers in retail markets. The implication of this line of reasoning for our work is that our regression coefficient would overestimate the distribution cost share based on a standard omitted regressor bias argument.⁶

Identification of markups, of course, is a challenge confronting many sub-fields of macroeconomics. The model that we develop and estimate is based on the premise that markets are perfectly competitive. The NIPA is constructed along these same lines in the sense that payments to labor and capital exhaust value added, there is no separate line item for markups over unit cost. An important task going forward is to consider how bargaining power and imperfect competition leads shares of payments to factors of production that exceed their shares in the production function. All that may be claimed at this point is that both our estimating approach and variance accounting of LOP and PPP deviations are consistent with the NIPA constructs and perfect competition.

⁶That is, if an omitted regressor (markup) is positively correlated with an included regressor (wages), the coefficient on the included regressor (the distribution margin) is biased upward.

2. The Model

The model consists of an arbitrary number of cities, each inhabited by two representative agents. One representative agent is a manufacturer who specializes in the production of a single good and exports this good to all other cities of the world. The second representative agent is a retailer who imports all of the manufactured goods and makes the goods available in retail outlets in her city of operation. To import a good, the retailer must pay an iceberg shipping cost over the factory-gate price in the producer's location. The shipping cost is hypothesized to be increasing in the distance shipped and makes a discrete jump if a national border is crossed. The retailing activity is labor intensive with the retailers allocating their non-leisure time across all of the goods they sell. Some cities have more productive retailers than others which is captured by total-factor-productivity (TFP) at the retail level, specific to the city, common to all goods the retailer sells. Part of the TFP effect might be local public infrastructure and private capital, neither of which are modeled here.⁷

Turning to the details, the retailing technology for each good is Cobb-Douglas in retailer hours, N_{ij} , and the quantity of the imported manufactured good, G_{ij} , with TFP level, Z_j :

$$R_{ij} = Z_j N_{ij}^{\alpha_i} G_{ij}^{1-\alpha_i} . \quad (2.1)$$

While the production function is restricted to be common to all locations, it is very flexible across goods. It captures pure labor services (e.g., baby-sitting services) with α_i equal to one and internet purchases (e.g., Amazon.com book purchases), α_i equal to zero, and all points in between.

The retailer in city, j , minimizes cost of each good, i , by optimally choosing the two inputs needed to produce the good: i) the amount of the traded input, G_{ij} , to import and ii) the fraction of her time devoted to the good, N_{ij} :

$$\min_{N_{ij}, G_{ij}} (W_j N_{ij} + Q_{ij} G_{ij}) \quad (2.2)$$

Note that the W_j reflects the single opportunity cost of time relevant to the problem, that of the retailer. The two constraints on this minimization problem are the production function,

⁷In an earlier version of the paper, we included capital as a factor of production. The simpler formulation here focuses on aggregate retail efficiency, labor in the retail sector and trade costs. To the extent public and private infrastructure capital alter efficiency, these would be allocated to the TFP term.

(2.1) and that total hours available in the period are exhausted between leisure hours and total time allocated to all retail goods.

The resulting retail price is a Cobb-Douglas aggregate of the price (inclusive of trade cost) that the retailer paid to acquire the traded input, Q_{ij} , and the retailer's opportunity cost of time, W_j :

$$P_{ij} = \frac{W_j^{\alpha_i} Q_{ij}^{1-\alpha_i}}{Z_j}. \quad (2.3)$$

It is important to note that the weights on the two inputs are good specific. Not surprisingly, all retail prices decrease in proportion to total factor productivity in retailing, Z_j and increase in proportion to retail input prices, W_j and Q_{ij} (with the factor of proportionality being their respective cost shares).

The factory-gate price of the traded input is determined as follows. A manufacturer in each city operates a simple linear technology, $Y_i = A_i N_i^m$, and maximizes profits from world-wide sales. She charges the same factory-gate price to all destination markets, $Q_{ii} = \widetilde{W}_i/A_i$ where \widetilde{W}_i is the manufacturing wage in city i and A_i is productivity in manufacturing. The complete equilibrium solution to the model is presented in a separate technical appendix (see Crucini and Yilmazkuday, 2009).

For the purposes of studying relative prices, the only remaining piece of information needed is the relationship between the factory gate price and the destination price. We assume a proportional, good and location-specific shipping cost: $Q_{ij} = (1 + \tau_{ij}) Q_{ii}$.

The prediction of this model for the common-currency relative price of good i in city j relative to city k is:

$$\frac{P_{ij}}{P_{ik}} = \frac{Z_k}{Z_j} \left(\frac{W_j}{W_k} \right)^{\alpha_i} \left(\frac{Q_{ij}}{Q_{ik}} \right)^{1-\alpha_i}. \quad (2.4)$$

Taking logs gives the object of interest:

$$p_{ijk} = -z_{jk} + \alpha_i w_{jk} + (1 - \alpha_i) q_{ijk} \quad (2.5)$$

where $p_{ijk} = \log(P_{ij}/P_{ik})$, $z_{jk} = \log(Z_j/Z_k)$, $w_{jk} = \log(W_j/W_k)$, and $q_{ijk} = \log(Q_{ij}/Q_{ik})$.

3. The Data

We use city-level data on retail prices, wages, and the greater-circle distance in our empirical work. The prices and wages are from the World Cost of Living Survey conducted by the Economist Intelligence Unit (EIU). The surveys took place in 123 cities, located in 79

countries. The vast majority of the cities in the survey are national capitals and since urban areas are typically densely populated with higher per capita income than rural areas, these cities account for a significant fraction of global consumption and production; they are also typically major ports and centralized trading locations (see Figure 2). The larger number of cities than countries is due to the fact that the survey includes multiple cities in a few countries. Noteworthy are the 16 U.S. cities included in the survey; the next largest number of cities surveyed equals 5 in Australia, China and Germany. Our sample is annual from 1990 to 2005. Up to data availability for particular years and cities, the number of goods and services surveyed by EIU staff is 300. Each price observation is collected from the same retail outlet. Examples of goods found in the survey are: *Butter (500 grams)*, *Compact disc album*, *Light bulbs (two, 60 watts)*. Typical examples of services are: *Dry cleaning*, *mans suit (standard high-street outlet)*, *baby-sitters rate per hour (average)*, *Hilton-type hotel, single room, one night including breakfast (average)*.

Let $P_{ij,t}$ be the price of good i , in city j and year t in U.S. dollars. The object of interest is the long-run bilateral price deviation across city pair j and k , computed as the time-averaged log-relative price:

$$p_{ijk} \equiv T^{-1} \sum_t p_{ijk,t}$$

where $p_{ijk,t} \equiv \log(P_{ij,t}/P_{ik,t})$.

To gain an appreciation of the relative importance of long-run price dispersion compared to time series price variation, Figure 3 presents kernel density estimates of relative prices. The solid lines are the distributions of the time-averaged prices, p_{ijk} , while the dashed lines are the distributions of the annual deviations of relative prices from these long-run levels, $p_{ijk,t} - p_{ijk}$. The two charts on the left are distributions for U.S. city pairs and the two charts on the right are all international cross-border city pairs. The charts in the top row include only non-traded goods prices while the charts in the bottom row include only traded goods prices.

If the LOP held *always and everywhere*, all the distributions would be degenerate at zero. Such a situation would describe a world of frictionless trade in goods markets and instantaneous arbitrage. Given the continuous and often large movements of nominal exchange rates and what is known about the infrequency of local currency price changes, it is not surprising that the dashed lines reveal transitory deviations of relative prices from their long-run means. What is very surprising is the distribution of the long-run means themselves. In each case,

with the possible exception of traded goods across U.S. cities, the dispersion of the long-run price distribution is greater than the variation of the time series deviations around these long-run means. Put differently, the time series movements seem less puzzling in light of the size of the long-run deviations. This motivates our focus on estimating the sources of long-run relative price deviations.

Table 1 presents summary statistics related to the data in Figure 3. The least amount of price dispersion is found in U.S. traded goods, 0.29 and the greatest amount is found in the case of non-traded goods involving international border crossings, 1.07. Remarkably, non-traded goods in the U.S. actually have less price dispersion than do traded goods internationally, 0.54 compared to 0.68. Inter-quartile differences yield similar measures of price dispersion. As originally discovered by Crucini and Telmer (2012), time series variation is typically less than the long-run variance, with the possible exception of traded goods across U.S. cities, and even in this case, one of the two measures (inter-quartile difference) also gives a ranking consistent with the broader samples. Notice also that the distinction between traded and non-traded goods is obvious in the long-run measure. This contrasts with the existing international finance literature where the time series variance of non-traded and traded real exchange rates are found to be very comparable (Engel (1999) and Crucini and Landry (2012)).

The remaining data utilized are wages, measures of distribution costs and distance. Directly measuring trade costs is a significant challenge in the literature. Hummels (2001) provides the most comprehensive estimates of sectoral trade costs using import unit values, a more direct method than employed here. Unfortunately, these estimates are available for a very limited number of countries and are more aggregated than our retail price data. Instead, we follow the gravity literature in trade and use the greater circle distance between cities in the EIU sample to estimate trade costs in LOP deviations at the retail level. The implied trade costs are consistent with Hummels estimates.

The wage measure is *hourly rate for domestic cleaning help (average)* from the EIU survey. This wage measure is chosen for a number of reasons: (i) it spans the entire 1990-2005 sample period and (ii) the number of missing observations is substantially lower than the alternative available wage series in the EIU survey (i.e., only 269 missing observations out of 7,503 city pairs, less than 4% of the sample), and (iii) it has a high cross-sectional correlation with alternative sources of wage data at the country level. In particular, we

considered an alternative hourly wage measure found in the EIU sample and the per capita income measure from the PWT, which is often used as a proxy variable in studies of HBS. The alternative hourly wage in the EIU is available for far fewer cities and the PWT is a national measure, not a city-level measure. We also considered wage measures from a survey done by the Union Bank of Switzerland (UBS) because it is available at the city level and has some occupational disaggregation, but it has 4,947 missing observations (out of 7,503 city pairs), and it is a single cross-section for the year 2006 and thus is unlikely to be representative of the long-run wage distribution over the period of 1990-2005. For the wage series that could be compared, the cross-sectional correlations were quite high.

Sectoral U.S. NIPA data and U.S. input-output tables are used to cross-validate the distribution share measures estimated from our regression model. Our model recovers close to 300 good-specific distribution shares while the U.S. data provide 57 sectoral distribution shares and the input-output data span 33 sectors. The NIPA shares are computed as the value the producers receive relative to the value consumers pay for the output of a particular sector. For the typical traded good, the distribution margin computed using the consumer value less the producer value relative to the consumer value is about 50%. That is, the retail price is about twice the producer price. However, for services, the same NIPA data would produce an estimate of the distribution margin close to zero. Consider a visit to the doctor's office to receive an expensive vaccine injection by a nurse. Because of the arms length nature of the transaction, it appears as though what the consumer pays, the producer gets. Most existing studies record the distribution margin to be zero in these situations. However, the economic concept that the distribution margin is intended to capture in our model is the distinction between retail prices and traded inputs. The goal is to treat the labor services of the nurse at the doctor's office in a consistent manner with the labor services of the salesperson at Walmart. To our knowledge, this issue has not been dealt with in the existing literature because the focus has been mostly on traded goods.⁸ Since the CPI consists of a large and growing fraction of services, measuring the distribution share for the service sector is an important facet of our work.

⁸See, for example, Burstein, Neves and Rebelo (2003). Following a conversation between Crucini and Rebelo, the distribution margin in Burstein, Eichenbaum and Rebelo (2005, 2007) makes an approximate correction for this effect.

4. The Estimation Approach

Taking the model to the data involves a number of empirical challenges. We have rich data with which to measure LOP deviations and reliable measures of wages across cities in our panel, but we lack both measures of retail productivity and micro-level data on traded input prices. This section describes the three-step approach taken to identify distribution shares, the α_i 's for each good, city-level retail productivities, the z_j 's, and trade costs between any city pair at the good level.

The first stage utilizes the available data on price and wages across cities to estimate a good-specific distribution share α_i , by regressing LOP deviations on the wage ratio:

$$p_{ijk} = \alpha_i w_{jk} + \theta_{ijk} \quad . \quad (4.1)$$

where, according to the model, the residual is $\theta_{ijk} = -z_{jk} + (1 - \alpha_i) q_{ijk}$. The slope parameter in the relationship between prices and relative wages in the HBS scatterplot of Figure 1 is the empirical counterpart to α_i .

The parameter, α_i , is estimated by geometric mean regression (GMR). The GMR estimate is the geometric average of the coefficient from an ordinary least squares (OLS) regression,

$$p_{ijk} = \hat{\alpha}_i w_{jk} + \epsilon_{ijk}$$

and an inverse least squares (ILS) regression

$$w_{jk} = \hat{\gamma}_i p_{ijk} + \varepsilon_{ijk} \quad ,$$

$\hat{\alpha}_i^G = \sqrt{\hat{\alpha}_i / \hat{\gamma}_i}$. As described in Kennedy (2003), this estimator is consistent when the two variables have comparable measurement error variances relative to the variance of the true underlying economic variables.

The second stage regression uses the residuals from the first stage regression,

$$\hat{\theta}_{ijk} = p_{ijk} - \hat{\alpha}_i^G w_{jk}$$

and pools all goods and bilateral city pairs to estimate city fixed-effects ($-z_{jk}$):

$$\hat{\theta}_{ijk} = -z_{jk} + \varphi_{ijk}$$

where $-z_{jk} \equiv -\mu_j + \mu_k$ is the retail productivity differential and the residual, $\varphi_{ijk} = (1 - \alpha_i) q_{ijk}$, represents the traded input cost ratio, which is assumed to be mean zero across

goods, for each bilateral pair. Note that while the z 's capture retail productivities under the assumption of perfect competition, they would also capture city-specific markups if retailers in each city have market power. Since we cannot separately identify retail productivity and retail markups, we will call z_j 's (log) retail productivities simply to be consistent with our competitive equilibrium model of retailing and trade. Estimation at this stages is by OLS.

The third stage considers relative prices of traded inputs, q_{ijk} .⁹ To place some structure on these trade costs, consider the no-arbitrage condition for good i , across city pair j and k :

$$-\delta_i d_{jk} - \rho_i B_{jk} \leq q_{ijk} \leq \delta_i d_{jk} + \rho_i B_{jk} \quad (4.2)$$

where $\delta_i d_{jk} + \rho_i B_{jk}$ is the trade cost between city j and city k for good i ; $d_{jk} > 0$ is the log distance between cities j and k , $\delta_i > 0$ is the elasticity of trade costs with respect to distance for traded good i , B_{jk} is a border dummy taking a value of 1 if cities k and j are in different countries (and 0 otherwise), and $\rho_i > 0$ is the logarithm of the additional cost of crossing the border between city j and k (if one exists) with traded-input i . While the inclusion of distance is to capture geographical barriers to trade, the inclusion of a border dummy is to capture official barriers to international trade. As is evident from the specification above, the border effect estimated here differs across goods, but is common to all border crossings.

Equation 4.2 is a standard arbitrage condition showing that arbitrage is profitable in the sense of shipping goods from city k to city j only if the price in city j is high enough relative to the price in city k to cover the arbitrage costs in that direction: $q_{ijk} > \delta_i d_{jk} + \rho_i B_{jk}$. Conversely, goods should be shipped from city j to city k when the price in city j is sufficiently low: $-q_{ijk} > \delta_i d_{kj} + \rho_i B_{kj}$.

Since the traded input cost ratios, $(1 - \alpha_i) q_{ijk}$, are estimated by the $\widehat{\varphi}_{ijk}$'s, the sign of $\widehat{\varphi}_{ijk}$ provides useful information on the profitable direction of arbitrage. Since trade costs measured by $\delta_i d_{jk} + \rho_i B_{jk}$ are positive, positive values of estimated φ_{ijk} 's will be associated with city j importing from city k , and negative values of φ_{ijk} 's will be associated with city j exporting to city k , according to our model. Thus, the direction-of-arbitrage indicator function \mathbf{I}_{ijk} is set to a value of 1 or -1 according to:

$$\widehat{\mathbf{I}}_{ijk} = \begin{cases} 1 & \text{if } \widehat{\varphi}_{ijk} > 0 & j \text{ imports from } k \\ -1 & \text{if } \widehat{\varphi}_{ijk} < 0 & k \text{ imports from } j \end{cases} \quad (4.3)$$

⁹Note, the trade costs are estimated as $\widehat{q}_{ijk} = \widehat{\varphi}_{ijk} / (1 - \widehat{\alpha}_i)$.

The importance of controlling for local distribution costs by using $\widehat{\varphi}_{ijk}$ rather than the real exchange rates themselves should be evident: if we did not do this, our indicator function would suggest that all goods are imported by rich nations from poor ones due to the lower price levels in poor countries at the retail level (recall Figure 1). Clearly, the first and second stages of the regression are crucial prior steps since they remove the influence of local distribution costs.

Consider, now, choosing the variables, κ_{ijk} and η_{ijk} , such that Equation 4.2 holds with equality:

$$-\delta_i d_{jk} - \rho_i B_{jk} - \kappa_{ijk} = q_{ijk} = \delta_i d_{jk} + \rho_i B_{jk} + \eta_{ijk} .$$

Note that given the sign conventions for these plug-in values, they satisfy $\kappa_{ijk} \leq 0$ and $\eta_{ijk} \leq 0$. By using the estimated direction-of-arbitrage indicator $\widehat{\mathbf{I}}_{ijk}$, these two equalities can be combined in the following expression:

$$q_{ijk} = \widehat{\mathbf{I}}_{ijk} (\delta_i d_{jk} + \rho_i B_{jk}) + c_i + \varepsilon_{ijk} \quad (4.4)$$

where $\varepsilon_{ijk} = \widehat{\mathbf{t}}_{ikj} \eta_{ijk} + \widehat{\mathbf{t}}_{ijk} \kappa_{ijk} - c_i$. The presence of c_i is to ensure that $E_{jk}(\varepsilon_{ijk}) = 0$. The indicator $\widehat{\mathbf{t}}_{ikj}$ takes a value of 1 if the direction of trade is from city k to city j (and 0 otherwise), $\widehat{\mathbf{t}}_{ijk}$ takes a value of -1 if the direction of trade is from city j to city k (and 0 otherwise). The two indicators add up to the original one: $\widehat{\mathbf{I}}_{ijk} = \widehat{\mathbf{t}}_{ikj} + \widehat{\mathbf{t}}_{ijk}$. Using $\widehat{\varphi}_{ijk} = (1 - \widehat{\alpha}_i) q_{ijk}$ (the fitted residuals from the second stage regression) and estimated $\widehat{\alpha}_i$'s from the first stage regression: $\widehat{q}_{ijk} = \widehat{\varphi}_{ijk} / (1 - \widehat{\alpha}_i)$. Thus, all the variables necessary to estimate equation 4.4 are available: the relative input cost, \widehat{q}_{ijk} , the direction of trade indicator, $\widehat{\mathbf{I}}_{ijk}$, greater circle distance and border dummies (d_{jk} and B_{jk}).

Although we have confidence about what the good-specific intercept and residuals ($c_i + \varepsilon_{ijk}$) do not represent, namely relative distribution costs (wage and retail productivity components), border-related costs, or distance-related costs, there are a number of plausible alternative explanations for what they do represent. The sources of these deviations from LOP unexplained by the model include: measurement error in retail prices, deviations from LOP that are below the threshold arbitrage value, trade costs that do not depend on distance or borders (e.g., trade finance) and markups specific to goods rather than common to bilateral locations.

5. The Results

This section reports our findings, beginning with a careful review of the parameter estimates obtained by the three-stage regression approach using the full sample of international cities. Next, the implications of these estimates for geographic price dispersion at the good level (i.e., deviations from LOP) are reported using variance decompositions. Results that serve to contrast heterogeneity across goods and the role of national borders are highlighted. Finally, the LOP deviations are aggregated and a variance decomposition of PPP is conducted.

5.1. Parameter estimates

The distribution share of each good α_i is estimated by the first-stage regression using the GMR estimator described in the previous section. The mean $\hat{\alpha}_i$ estimate is 0.48, while the median adjusted R^2 of the first-stage regression is 0.34 (both taken across goods). Recall that using price levels computed from the same data and the same wage measure (Figure 1), the slope coefficient was 0.52 and the R^2 was 0.37. As we shall see below, the higher slope coefficient in the aggregate for PPP is to be expected since, according to the model, it represents a consumption-expenditure-weighted average of the microeconomic distribution shares. Since non-traded goods tend to carry the largest consumption shares and involve more distribution costs (non-traded inputs), it is expected that the PPP slope exceeds the simple average LOP slope in the cross-section.

Given that a considerable number of variables implied by our theory that are omitted in the first stage regression, it is natural to ask how our distribution share estimates compare to distribution shares in the U.S. NIPA accounts and input-output tables. Since the latter are more aggregated than our estimates, we average our good-level estimates within each NIPA sector. Table 2 compares our microeconomic estimates to the more aggregated NIPA values at six points in the distribution. The median distribution share is estimated to be 0.45 compared to 0.41 using US NIPA data. The estimated values match remarkably closely throughout the distribution except at the very high-end of the distribution: at the third quartile the estimated distribution share is 0.55 compared to 0.75 in the US NIPA.

In summary, our estimates are broadly consistent with direct U.S. NIPA measures, but our estimates are preferred in the context of our study for three reasons. First, they are good

specific allowing our subsequent analysis and variance decompositions to exploit the richness of our micro-price data. Second, the prices are consistent with the wage data since both are taken from the same EIU survey. By consistent, we mean not averaged across occupations and covering the same 123 cities as the price survey. Third, the U.S. distribution shares are not necessarily representative of those in other nations. Our estimates are literally global estimates.

Turning to the second stage of the estimation process which recovers the retail productivity of each city (Z_j), the median adjusted R^2 is 0.40. Table 3 provides summary statistics for the Z_j 's, which are all statistically significant at 5% level. The mean and median are not informative because they reflect an arbitrary choice of units in which to measure productivity. What is interesting are the large differences in productivity across cities, the third quartile city is 65% more productive than the first quartile city. In terms of the retail sector, the most productive city is 6.6 (3.32/0.50) times more productive than the least productive city.

To cross-validate our inference about retail productivity, the estimated Z_j 's are compared to distribution sector productivities from the GGDC Productivity Level Database. This database covers only 18 of the 79 countries in the EIU sample. For countries in the EIU sample with more than one city in the price survey, we take the simple average of city-level retail productivities as our estimate of national retail productivity. The correlation between the two estimates of retail productivity across the 18 countries common to both samples is 0.48.

The effects of trade costs are estimated in the third-stage regression described earlier, which have a median adjusted R^2 of 0.61. The summary statistics for distance elasticities (δ_i 's), which are all statistically significant at the 5% level, are reported in Table 3. The median distance elasticity is 0.05 which implies that price deviations increase by 40% per 1,000 miles of distance between the source and destination. Taking into account that traded inputs account for only a fraction of the retail price, the average effect of 1,000 miles of distance on retail prices is given by $\exp\left(\widehat{\delta}_i(1 - \widehat{\alpha}_i)d_{jk}\right) - 1$. Evaluated at the median values of both the distance elasticities and distribution share parameters implies about a 21% increase in the retail price from source to destination from the trade cost channel.

The calculation above does not take into account the possible impact of market segmentation associated with national borders. The border effects are statistically significant at the

5% level. Borders are also economically significant. For the median good, the border adds 33% to the traded-input price (see Table 3). Taking into account the share of traded inputs in the production of the median retail good using, $\hat{\rho}_i(1 - \hat{\alpha}_i)$, the average border wedge is 18%. Engel and Rogers (1996) have popularized the transformation of price deviations into distance equivalent measures. The border effect in distance equivalent units is computed as $\exp\left(\frac{\hat{\rho}_i}{\hat{\delta}_i}\right)$. The average border effect across goods, pooling all location pairs, is 735 miles. The median border width is 522 miles. As one might expect, the differences across goods are large. Moving from the first quartile of the distribution of the border effect to the third quartile, the border width increases from a mere 38 miles to an astounding 55,322 miles.¹⁰ Interestingly, the first quartile contains the traded good, *banana*, while the third quartile contains the non-traded good, *a three course dinner for four people*. What this suggests to us is that extrapolating the distance metric becomes less useful as the item in question becomes inherently less traded in the sense of being produced mostly with local inputs. In such cases it is preferable to report price dispersion and attempt to account for that dispersion with something other than iceberg shipping costs.

The cumulative explanatory power of the three-stage estimation is a useful metric for summarizing the completeness of our model in accounting for price dispersion. The median R-bar squared value across goods is 0.72. In other words, the parsimonious set of controls account for the bulk of good-level heterogeneity in long-run price deviations. Having described the parameter estimates and their economic interpretations, it is now possible to decompose the variance of LOP and PPP deviations into the contributions of retail productivity, wages, distance and borders.

¹⁰The baseline border width estimate of Engel and Rogers (1996) was 75,000 miles. It is not possible to make a direct comparison between our estimates and theirs for a number of reasons. First, they use the time series variance of changes in relative prices. Second, they use CPI data aggregated to roughly two-digit categories. Third, they focus on U.S.-Canada city pairs.

5.2. Variance decomposition of LOP deviations

The goal is to explain the variance of relative prices across all unique city pairs, good-by-good¹¹:

$$V_i = \text{var}_{jk}(p_{ijk}).$$

Intuitively, the larger is V_i , the greater are the deviations from LOP over the geography of bilateral city pairs indexed by j and k . The natural economic benchmark for the lower bound is when the LOP holds across all bilateral pairs used in the calculation, in which case, $V_i = 0$. The model is designed to elucidate the sources of price dispersion emanating from plausible real frictions related to economic geography.

Recall that the log-relative price consists of three main components, two related to retailing (relative TFP in retailing and the relative wages of the retailers) and the relative cost of acquiring the traded input:

$$p_{ijk} = -z_{jk} + \alpha_i w_{jk} + (1 - \alpha_i) q_{ijk} .$$

Substituting our model of the traded input component into this equation gives the rather intimidating expression,

$$p_{ijk} = -z_{jk} + \alpha_i w_{jk} + (1 - \alpha_i) [\widehat{\mathbf{I}}_{ijk} (\delta_i d_{jk} + \rho_i B_{jk}) + c_i + \varepsilon_{ijk}] . \quad (5.1)$$

The variance decomposition will be computed using the fact that $\text{var}(p_{ijk}) = \text{cov}(p_{ijk}, p_{ijk})$ with the second p_{ijk} replaced by all of the terms on the right-hand-side of (5.1). The resulting variance decomposition is:

$$1 = \underbrace{\beta_{iz} + \alpha_i \beta_{iw}}_{\text{Non-traded inputs}} + \underbrace{(1 - \alpha_i) \delta_i \beta_{id} + (1 - \alpha_i) \rho_i \beta_{iB}}_{\text{Traded inputs}} + \underbrace{(1 - \alpha_i) \beta_{i\varepsilon}}_{\text{Unexplained}} .$$

The use of the notation β is natural here since the contribution of each component to the variance is effectively a regression coefficient. For example, $\beta_{iw} = \text{cov}_{jk}(w_{jk}, p_{ijk}) / \text{var}_{jk}(p_{ijk})$ is the coefficient from a regression of relative wages on relative prices of good i across all city-pairs. Essentially this component tells us the role of international wage dispersion on

¹¹The variance metric is different from the one used by Crucini, Telmer and Zachariadis (2005) to study price dispersion across European capital cities. They normalized prices to their cross-city means, $\text{var}_j(p_{ij} - \bar{p}_i)$, which does not allow for the role of bilateral distance in the trade cost component of our structural model.

retail price dispersion of good i , across locations. The fact that this β has a coefficient α_i means that the international wage dispersion has an effect on retail price dispersion that is increasing in the share of distribution in cost.

The contribution of traded inputs is more nuanced, involving a traditional trade cost component and a border effect. As one would expect, both terms are weighted by the cost share of traded inputs, $(1 - \alpha_i)$. The first component which has a traditional shipping cost formulation involves the product of δ_i and β_{id} . Recall that δ_i is the good-specific elasticity of trade cost with respect to distance; ceteris paribus, a good that is more costly to transport will contribute more to the variance of retail prices. The more subtle part of the expression is the covariance between relative prices at the retail level and interaction of the direction of trade indicator and distance, the $\beta_{id} = cov(\widehat{\mathbf{I}}_{ijk}d_{jk}, p_{ijk})/var_{jk}(p_{ijk})$ term. Due to the presence of the indicator function and how it was defined, this is the coefficient of a regression of the absolute value of the relative price of traded inputs on the retail price. Intuitively, we want to relate trade costs to distance in a symmetric fashion in the sense that the distance matters, not the direction of trade. The absolute value ensures that trade costs are non-negative in the estimation equation. Basically, if retail price deviations are increasing in the estimated trade cost, which themselves are rising in distance, then trade costs contribute positively to long-run price dispersion. The contribution is greater for locations separated by greater distances since $\delta_i > 0$. Holding distance fixed, goods that are more costly to ship, greater δ_i , will exhibit more long-run price dispersion. The border effect contributes only to the variance of cross-border city pairs and does so as a level effect, not as a function of distance. The level effect on volatility is good specific due to the presence of ρ_i .

Figures 4 and 5 present the variance decomposition in the cross-section of goods, sorted by the estimated distribution share, $\widehat{\alpha}_i$. Thus, at the left-hand boundary are goods with 0.20 of their cost attributed to distribution inputs and the rest of the cost attributed to the traded good itself. Unleaded gasoline is an example of such a good. The right-hand boundary is a pure service, such as the hourly wage of baby-sitters. The variance of retail prices (V_i) is the upper contour. We clearly see a positive relationship between the distribution share and price dispersion. The increases are substantial. For example, price dispersion in goods involving the lowest distribution share practically satisfy the LOP in the case of Canada and the United States and among OECD cities, whereas the deviations for services approach 80%. The countries included in the sample matter for both the absolute level of price dispersion

and the relative contributions of various components. Price dispersion is uniformly greater across LDC city pairs and World city pairs than across cities of the OECD or North America.

It is important to point out that the obvious contribution of non-traded inputs in these figures contrasts sharply with much of the existing international finance literature where skepticism regarding the value of the HBS theory originated. In that literature, the metric is the time series variance of the real exchange rates of two sub-indices of the CPI typically used to elucidate the HBS theory (e.g., Engel (1999)). Recall, however, that in Table 1 the short-run variance (a measure of time series variance) failed to reveal a sharp difference in variability of real exchange rates across traded and non-traded goods. That classification, however, is based on applying the HBS theory to final goods, not intermediate inputs. Crucini and Landry (2012) show that the time series variance of LOP deviations are in fact rising in the distribution share. In other words, when the HBS theory is applied to intermediate inputs, it is successful in accounting for differences in both the long-run and the short-run properties of price dispersion.

Turning to the details, intranational price dispersion, displayed in the left-hand charts, is largely accounted for by distance. International price dispersion is driven significantly by three components: wages, distance and borders. The relative importance of the three depends on the set of locations under examination and differs across goods. As expected, the wage component becomes more important as we move from, say unleaded gasoline to baby-sitting services (left-to-right along the x-axis). This is true for all location groups. However, the absolute contribution of wages is greatest in the case of the world panel. This makes sense since this is the panel in which wage differences across the rich and poor countries enters into the picture. Thus when international price dispersion is economically significant, it is driven largely by wage differences in the service sector. When the locations are restricted to the OECD, however, the wage differences are not large enough to account for the observed price dispersion and the U.S. and Canada cross-border pairs are no exception to this result. Simply put, when cities are in countries with similar wealth levels and labor productivity levels, price dispersion is lower and more of the observed dispersion is accounted for by trade costs: distance in the case of intranational price dispersion, distance and borders in the case of international price dispersion.

The figures provide a very complete picture of sources of price dispersion across goods and locations, but it is difficult to get an overall sense of the role of various factors without

averaging across goods. Table 4 reports the average long-run price dispersion and the contribution of each factor to that average. Thus, the averages reported in Table 4 are influenced by the fact that goods are not equally distributed over the distribution shares of 0.2 to 1.0 displayed in the figures. For example, the modal good is a food item with a distribution share of about 0.37 (the median is 0.45). The columns labelled NB report results that use only within country city pairs in the analysis whereas the columns labelled B report results that use only city pairs separated by a national border. The main column headings, LDC, OECD and CAN-US describe the regions in which these cities are located while WORLD is the entire pooled sample of 123 cities. The upper panel of the table reports the absolute geographic variance coming from each economic channel while the lower panel reports the proportion of the total variance accounted for by each economic channel.

Table 4 shows the dominant factor accounting of geographic microeconomic price dispersion is the traded input component. Recall that the traded input component is the sum of traditional trade costs (distance) and the border effect. The border effect for the average good is between 7.5% and 10.4%. Interestingly the U.S.-Canada border is not as wide as the typical international border, though the differences are not large. Distance contributes between 6.4% to 14.7%. These differences are simply a reflection of the differences in the average distance separating bilateral city pairs in the respective columns. As it turns out, intranational city pairs are closest together in the LDC sub-sample with an average greater circle distance of 564 miles compared to over 4,000 miles for international city pairs in the entire world sample. Understandably, distance tends to generate much more price dispersion in the latter geography than in the former. North America is interesting because it has the exceptional property that intranational city pairs are almost as far apart as international city pairs, 1,083 versus 1,134. Thus distance is estimated to account for a comparable amount of absolute price dispersion in the last two columns, 7.04 versus 8.18.

The second most important contributor to price dispersion depends on whether the focus is intranational price dispersion or international price dispersion. Since labor markets tend to be nationally segmented, services wages are much more important sources of price dispersion for cross-border (B) city pairs than for with-country city pairs (NB) holding fixed the set of locations in the comparison. For example, for the pooled sample service wages contribute 0.42% to price dispersion across cities within the same country, less than one-fortieth their contribution to price dispersion across cities located in different countries. Obviously when

the countries used in the comparison are all at comparable levels of development, such as the OECD and particularly North America, services wages are sufficiently close across locations that for even cross-border city pairs they play a trivial role. For example, across the OECD service wages contribute only 2.22% to price dispersion, in North America, 2.77%.

Retail productivity variance is highest for LDC cross-border city pairs, as one might expect, but the variance is not great at 7.37%. While it might be expected that border crossings should result in higher retail productivity variance, this is true in only two of the four regions of Table 4. It could be that since most of the international cities are central hubs of trade on the coast the productivity is more similar for these city pairs than when comparing a city on the ocean to one in the interior. It may also be sensible to aggregate the retail productivity effect with the border effect to more accurately contrast international and intranational distribution costs that are not related directly to distance (shipping costs). That is, part of what makes traded inputs more expensive may be a fixed border cost of sourcing goods from outside the country.

Once the absolute variance patterns are understood, their percentage contributions directly follow. As the lower panel of Table 4 indicates, trade costs are the most important contributor to microeconomic price dispersion, accounting for between 23.4% (World, B) and 49.2% (OECD, NB) of total price dispersion. It is important to note that despite shorter distance and thus lower trade costs between cities within countries (compared to internationally), trade costs actually account for a larger fraction of total price dispersion intranationally than internationally because the absolute variance is smaller intranationally. Service wages, however, rank first for the overall international sample (World, B) accounting for 31.9% of total international price dispersion, with trade costs (distance) a distant second, at 23.4%.

To summarize, the model has provided a useful conceptual framework with which to conduct a variance decomposition of long-run LOP deviations into the role of distance, borders, wages, retail productivity and a residual. The relative importance of the factors depends on the good and set of locations in a fashion that makes intuitive sense. Distance and borders loom large in general and wages differences (consistence with HBS) seem particularly important once the analysis moves outside the OECD (i.e., once comparisons involve countries at very different income levels). We turn now to the aggregate implications of our model for PPP.

5.3. Variance decomposition of PPP deviations

Our ability to account for the heterogeneity in the sources of price dispersion in the micro-price data is reassuring as it suggests a plausible structural interpretation of retail productivity (or markups), distance and borders. The question addressed in this section is what lessons should be drawn for international macroeconomics. In this regard, the tables reporting the cross-good averages may not be the appropriate metric for two reasons. First, they are averages of variances, not variances of averages (or aggregates): to the extent LOP deviations average out across goods, these LOP variance decompositions may be misleading in terms of the level and sources of international price level dispersion.¹² Second, the economic characteristics of the median good or the average good in the EIU sample may not be representative of the economic characteristics embodied in the aggregate consumption basket, which is more relevant for aggregate real exchange rates and consumption levels.

The appropriate way to draw implications for international macroeconomics is to aggregate the data and the model in a consistent fashion and perform the same type of variance decomposition as was performed good-by-good. The starting point is the cost-of-living index for city j motivated by reference to a Cobb-Douglas aggregator function:

$$P_j = \prod_i (P_{ij})^{\omega_i}$$

where ω_i is the consumption-expenditure-share of good i .

According to Equations 2.5 and 4.4, the log deviations from PPP may be written as:

$$\begin{aligned} p_{jk} = z_k - z_j + & \left(\sum_i \omega_i \alpha_i \right) w_{jk} + \sum_i \omega_i (1 - \alpha_i) \widehat{\mathbf{I}}_{ijk} \delta_i d_{jk} \\ & + \sum_i \omega_i (1 - \alpha_i) \widehat{\mathbf{I}}_{ijk} \rho_i B_{jk} + \sum_i \omega_i (1 - \alpha_i) c_i + \sum_i \omega_i (1 - \alpha_i) \varepsilon_{ijk} \end{aligned} \quad (5.2)$$

where $p_{jk} = \log(P_j/P_k)$ is the aggregate real exchange rate across city pair j and k .

The variance decomposition of real exchange rates (using Equation 5.2) is:

$$1 = \underbrace{\beta_z + \omega \beta_w}_{\text{Non-traded inputs}} + \underbrace{\sum_i \omega_i (1 - \alpha_i) [\delta_i \beta_{id} + \rho_i \beta_{iB}]}_{\text{Traded inputs}} + \underbrace{\sum_i \omega_i (1 - \alpha_i) \beta_\varepsilon}_{\text{Unexplained}}$$

¹²The averaging-out property of LOP deviations was first documented by Crucini, Telmer and Zachariadis (2005) in the context of mostly EU capital cities over the period 1975 to 1990 at five-year intervals.

where $\omega = \sum_i \omega_i \alpha_i$ and the β 's are now covariances using the same right-hand-variables as before, but with the aggregate real exchange rate, p_{jk} , replacing the LOP deviations.

The results are reported in Table 5. As one would expect, the aggregate real exchange rate is much less volatile than LOP deviations since the LOP deviations tend to average out across goods. The variance falls by a factor of between 4 and 5. The impact of aggregation on the sources of price dispersion is most dramatic in the case of international price level dispersion across the world. Relative wages now account for virtually the entire variance (91.5%). Note that the coefficient on w_{jk} in Equation 5.2 can be estimated from a regression of log relative price level on the log relative wages. Such a regression produces a coefficient of 0.55 on wages and an R^2 of 0.85. The same coefficient can be estimated from the micro-price data, by calculating the consumption-expenditure-weighted average of distribution shares, $\sum_i \omega_i \alpha_i$. And what is this sum? Exactly 0.55! Note that this is higher than the average distribution share across goods (0.48) since consumption expenditure weights are positively correlated with the distribution share. In words: goods with higher non-traded factor content tend to account for a greater share of consumption expenditure.

That the slope coefficient recovers the product of these two structural share parameters without other controls in the regression equation, is more subtle. What is sufficient for this to occur is that trade costs tend to average out across goods for a given bilateral pair. This is what happens, as is evident in Table 5, the significant effects of trade costs and productivities in LOP deviations have basically aggregated away leaving relative wages to account for virtually all of the price level dispersion. While the results are not quite so stark for more narrowly defined international comparisons, wages are still by far dominant factor. For Canada-U.S. cross-border pairs, wages are about twice as important as distance, but distance clearly matters.

The results for *intranational* city pairs provide an interesting contrast. As one would expect, labor mobility tends to eliminate the role of wages in accounting for deviations from PPP. To a first approximation, the proportion of variance shifts from wages to retail productivity. Since retailing involves time-to-build and an immobile factor (land), it is arguably more difficult to arbitrage across locations and continues to play a role in accounting for price differences even within countries. An evolving literature asks how big-box retailers such as Walmart may alter the distribution margin across locations. These are interesting questions for future research but are beyond the scope of what can be measured given

available data in our panel.

6. Conclusion

The growth of international trade and financial integration has moved the fields of international trade and international finance ever closer. This is a healthy development, bringing together the microeconomic details of trade theory such as patterns of specialization and the extensive margin of trade with the central facets of macroeconomic theory, dynamic equilibrium concepts and expectation formation. The fact that trade focuses on absolute deviations from the LOP while finance focuses on relative deviations from PPP has been an impediment to progress. Recent efforts to expand the availability of large-scale micro-price panels allows absolute deviations from PPP to be traced back to LOP deviations as was done in this paper.

The amount of price dispersion at the good level is 3 to 5 times larger than at the aggregate, PPP level. The main cause of this averaging out of deviations specific to the good is attributed to trade costs and a residual term. The residual term may involve markups specific to the good and bilateral city pair as well as measurement error. What remains at the PPP level is largely a wage effect consistent with the HBS theory, though how dominant this factor is depends on the set of countries studied. Retail productivities and trade costs remain substantial at the PPP level for OECD countries, while they are dwarfed by wage effects when all international cities are pooled, consistent with the impression given by Figure 1.

Effectively, this means that theories intended to match relative prices both within the OECD and between the OECD and the rest of the world will need to include both a rich trade structure drawn from trade theory and a role for the distribution margin and markups. The importance of markups relative to real costs of wholesale and distribution remains an open question.

References

- [1] **Alessandria, George and Joseph Kaboski.** 2011. "Pricing-to-Market and the Failure of Absolute PPP." *American Economic Journal: Macroeconomics*, 3(1): 91-127.
- [2] **Anderson, James and Eric van Wincoop.** 2004. "Trade Costs." *Journal of Economic Literature*, 42(3): 691-751.
- [3] **Balassa, Bela.** 1964. "The Purchasing Power Parity Doctrine: A Reappraisal." *Journal of Political Economy*, 72: 584-96.
- [4] **Burstein, Ariel, Joao Neves and Sergio Rebelo.** 2003. "Distribution Costs and Real Exchange Rate Dynamics During Exchange-Rate Based Stabilizations," *Journal of Monetary Economics*, 50(6): 1189-1214.
- [5] **Burstein, Ariel, Martin Eichenbaum and Sergio Rebelo.** 2005. "Large Devaluations and the Real Exchange Rate," *Journal of Political Economy*, 113(4): 742-784.
- [6] **Burstein, Ariel, Martin Eichenbaum and Sergio Rebelo.** 2007. "Modeling Exchange Rate Passthrough After Large Devaluations." *Journal of Monetary Economics*, 54(2): 346-368.
- [7] **Crucini, Mario J. and Anthony Landry.** 2012. "Accounting for Real Exchange Rates Using Micro-data." Vanderbilt University, National Bureau of Economic Research, Working Paper No. 17812.
- [8] **Crucini, Mario J. and Mototsugu Shintani.** 2008. "Persistence in Law-of-One-Price Deviations: Evidence from Micro-Data." *Journal of Monetary Economics*, 55(3): 629-644.
- [9] **Crucini, Mario J., Chris I. Telmer and Marios Zachariadis.** 2005. "Understanding European Real Exchange Rates." *American Economic Review*, 95(3): 724-38.
- [10] **Crucini, Mario J. and Telmer, Chris I.** 2012. "Microeconomic Sources of Real Exchange Rate Variation." National Bureau of Economic Research, Working Paper No., 17978.
- [11] **Crucini, Mario J. and Yilmazkuday, Hakan.** 2009. "A Model of International Cities: Implications for Real Exchange Rates" NBER Working Paper No. 14834.

- [12] **Engel, Charles.** 1999. "Accounting for U.S. Real Exchange Rates." *Journal of Political Economy*, 107 (3): 507-38.
- [13] **Engel, Charles and Rogers, John H.** 1996. "How Wide Is the Border?", *The American Economic Review*, 86(5): 1112-1125.
- [14] **Harrod, Roy F.** 1933. *International Economics*. London: James Nisbet and Cambridge University Press.
- [15] **Herrendorf, Berthold and Akos Valentinyi.** 2006. "Which Sectors Make the Poor Countries so Unproductive?" mimeo, Arizona State University.
- [16] **Hummels, David.** 2001. "Toward a Geography of Trade Costs." Unpublished manuscript.
- [17] **Kennedy, Peter.** 2003. *A Guide to Econometrics*, Cambridge, MA: MIT Press.
- [18] **Krugman, Paul.** 1987. "Pricing to Market when the Exchange Rate Changes." in Sven Arndt and J. David Richardson, eds., *Real Financial Linkages Among Open Economies*, Cambridge, MA: MIT Press, 49-70.
- [19] **Samuelson, Paul.** 1964. "Theoretical Notes on Trade Problems." *Review of Economics and Statistics*, 46: 145-154.
- [20] **Sanyal, Kalyan K. and Jones, Ronald W.** 1982. "The Theory of Trade in Middle Products." *American Economic Review*, 72(1): 16-31.
- [21] **Sposi, Michael.** 2013. "Trade Barriers and the Relative Price of Tradables." mimeo.

TABLE 1 – PRICE DISTRIBUTIONS (SUMMARY STATISTICS)

	Traded goods		Non-traded goods	
	U.S.	World	U.S.	World
	No border	Border	No border	Border
Standard Deviation				
Long-run	0.294	0.681	0.543	1.069
Short-run	0.250	0.412	0.258	0.488
Interquartile Range				
Long-run	0.385	0.796	0.616	1.092
Short-run	0.295	0.417	0.295	0.430

TABLE 2 – DISTRIBUTION SHARES

	Estimated from EIU micro-data	U.S. sectoral NIPA data
Minimum	0.20	0.17
First Quartile	0.36	0.36
Median	0.45	0.41
Mean	0.48	0.53
Third Quartile	0.55	0.75
Maximum	1.00	1.00

TABLE 3 – RETAIL PRODUCTIVITY, TRADE COSTS AND BORDERS

	Retail Productivity (B_j)	Distance Elasticity (δ_i)	Border effect (ρ_i)
Minimum	0.50	0.00	0.02
First Quartile	0.75	0.04	0.21
Median	0.97	0.05	0.33
Mean	1.09	0.09	0.52
Third Quartile	1.24	0.08	0.49
Maximum	3.32	2.15	5.30

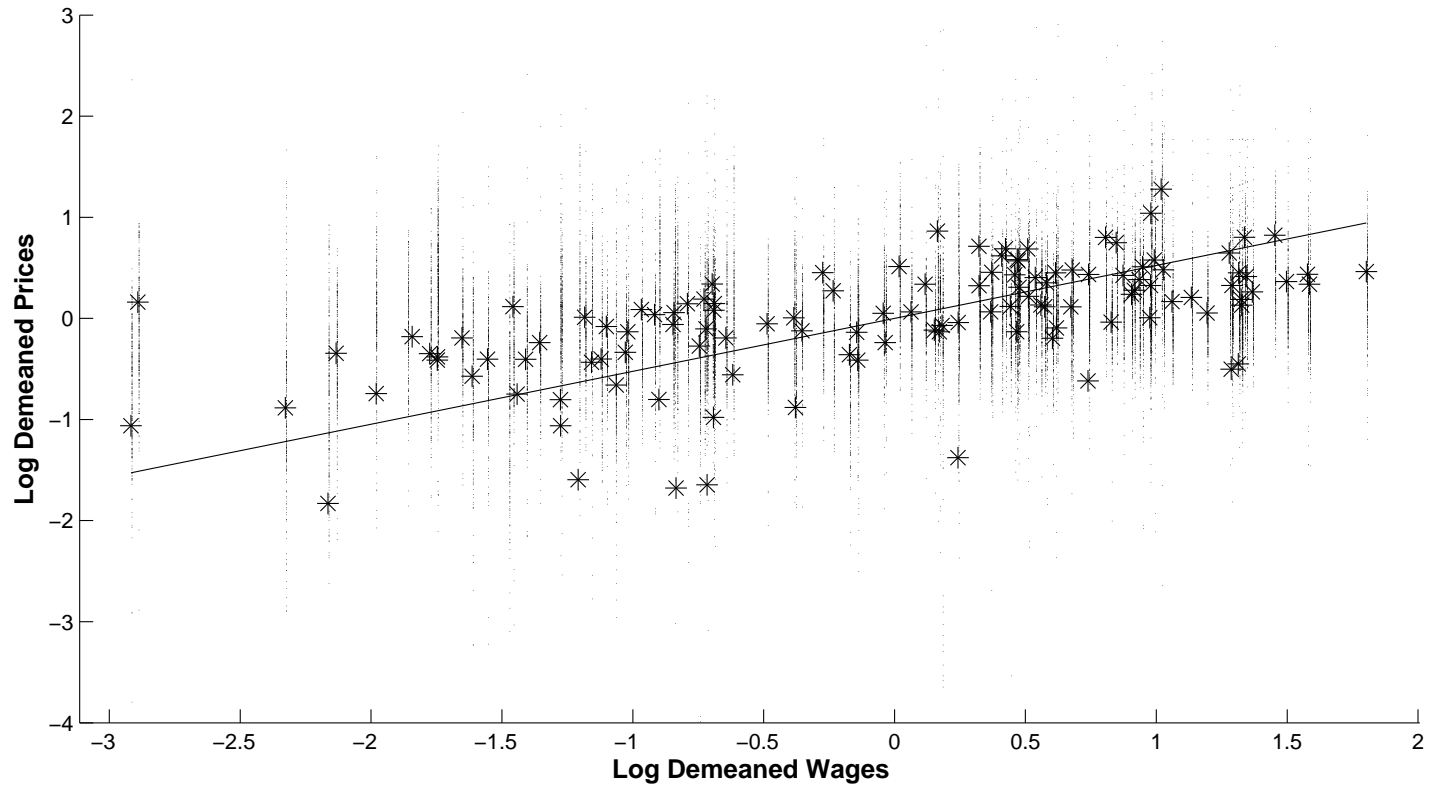
TABLE 4 – ACCOUNTING FOR LONG-RUN DEVIATIONS FROM LOP

	WORLD		LDC		OECD		CAN-US	
	NB	B	NB	B	NB	B	NB	B
Price dispersion	15.02	58.62	24.82	64.84	13.32	31.58	14.20	18.53
Absolute contribution of:								
Border effect	–	10.39	–	11.13	–	7.98	–	7.55
Distance	6.44	13.70	7.57	14.70	6.55	10.42	7.04	8.18
Service wages	0.42	18.70	1.10	10.71	0.19	2.22	0.15	2.77
Retail productivity	2.14	0.04	4.10	7.37	1.76	6.31	2.10	0.63
Residual	6.02	15.79	12.05	20.93	4.82	4.65	4.91	–0.60
Percentage contribution of:								
Border	–	17.7	–	17.2	–	25.3	–	40.7
Distance	42.9	23.4	30.5	22.7	49.2	33.0	49.6	44.1
Service wages	2.8	31.9	4.4	16.5	1.4	7.0	1.1	14.9
Retail productivity	14.2	0.1	16.5	11.4	13.2	20.0	14.8	3.4
Residual	40.1	26.9	48.5	32.3	36.2	14.7	34.6	–3.2
Average miles separating cities	856	4,054	564	3,944	896	3,216	1,083	1,134

TABLE 5 – ACCOUNTING FOR LONG-RUN DEVIATIONS FROM PPP

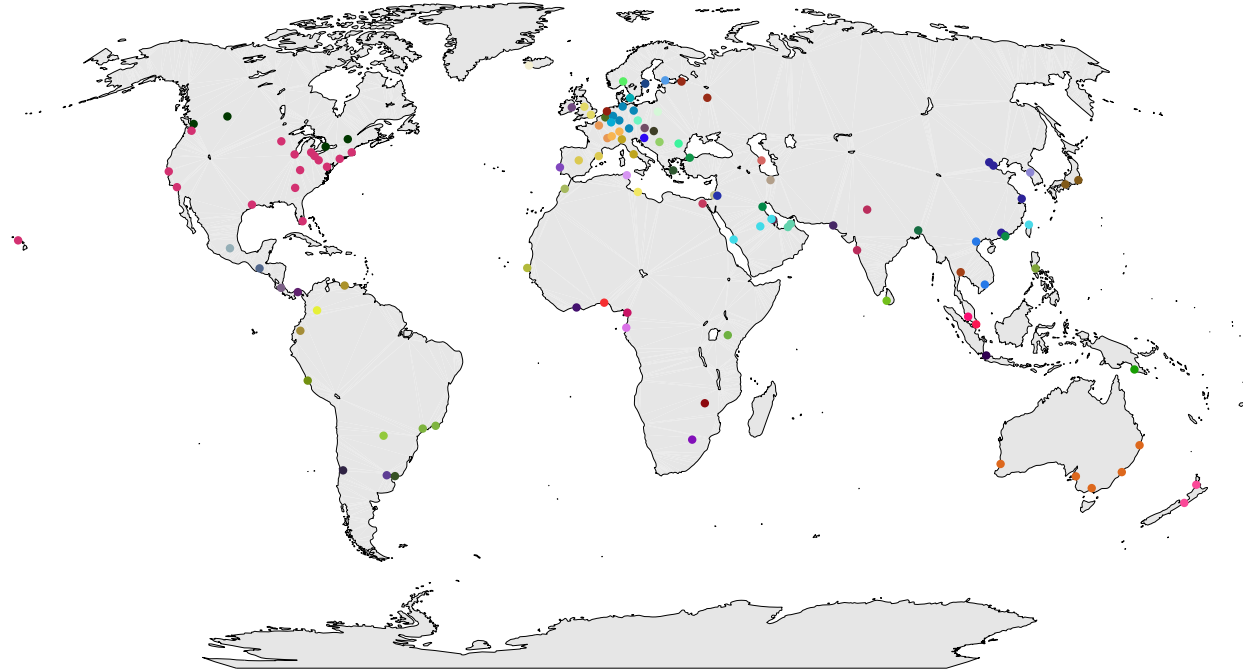
	WORLD		LDC		OECD		CAN-US	
	NB	B	NB	B	NB	B	NB	B
Price dispersion	3.54	15.90	4.53	15.48	3.14	10.07	3.33	5.16
Absolute contribution of:								
Border effect	–	0.12	–	0.18	–	0.23	–	0.18
Distance	0.82	1.01	0.28	1.53	0.85	1.95	0.93	1.26
Service wages	0.28	14.55	0.31	8.78	0.18	1.35	0.16	3.69
Retail productivity	1.82	–0.43	2.49	3.72	1.57	6.13	1.73	–0.25
Percentage contribution of:								
Border effect	–	0.8	–	1.2	–	2.3	–	3.5
Distance	23.2	6.4	6.2	9.9	27.1	19.4	27.9	24.4
Service wages	7.9	91.5	6.8	56.7	5.7	13.4	4.8	71.5
Retail productivity	51.4	–2.7	55.0	24.0	50.0	60.9	52.0	–4.8
Residual	17.5	4.1	32.0	8.2	17.2	4.1	15.3	5.4

Figure 1. Relative price levels and relative wages of domestic help



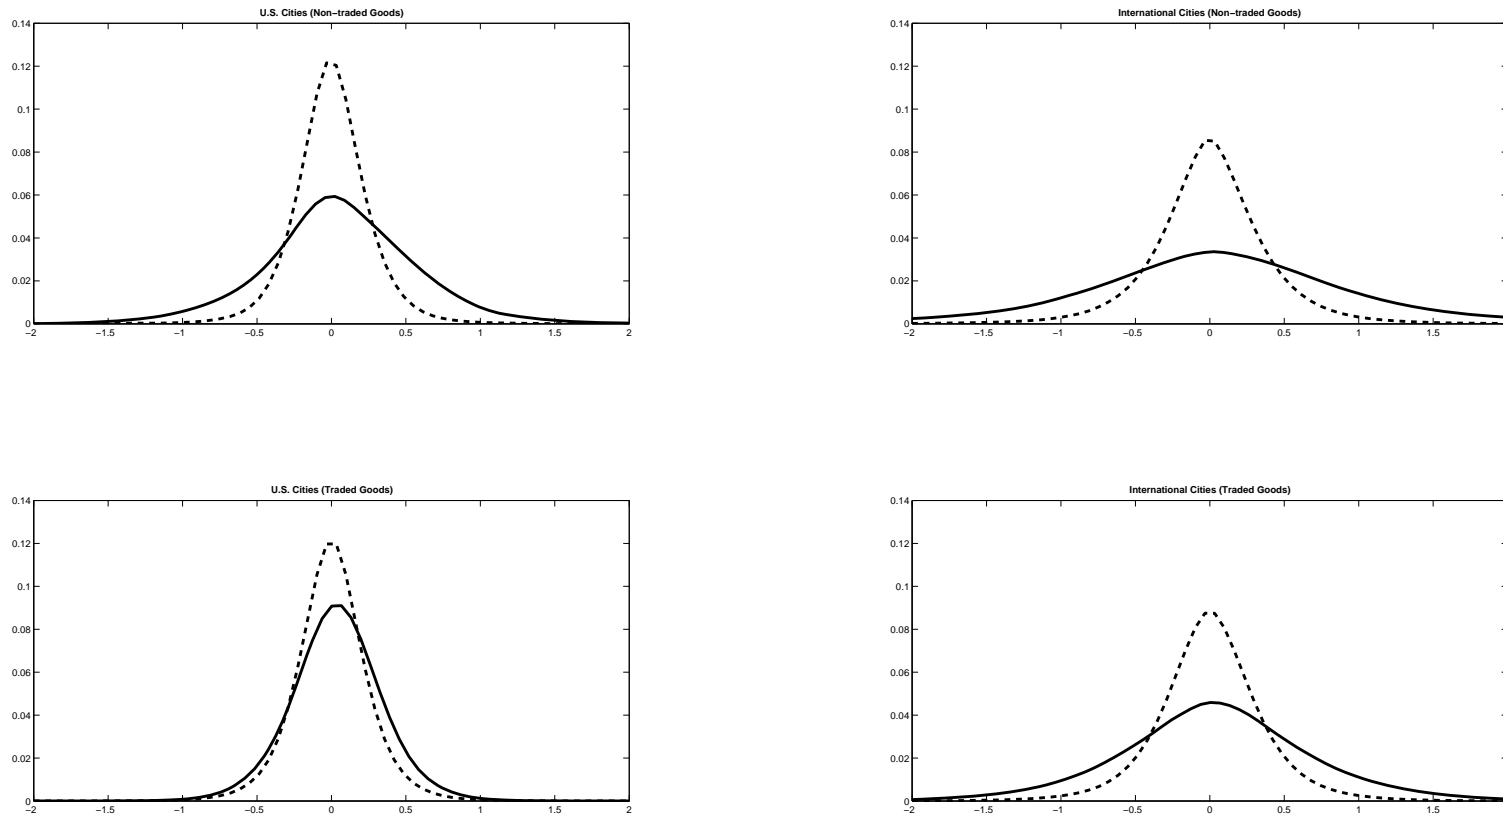
Notes: Each x-coordinate is a city wage relative to the world average. For the large dots, the y-coordinates are the price levels of each city relative to the world price level. The price levels are computed as consumption-expenditure-weighted average of the individual prices. The line through the scatterplot is an ordinary least squares estimate with slope 0.52 and a standard error of (0.05), using the price levels. For the small dots, the y-coordinates are prices of individual goods and services relative to the world average price of those items.

Figure 2. Cities surveyed by the Economist Intelligence Unit



Note: Each symbol marks the location of a city surveyed by the EIU. There are 123 in total. Most of the cities are national capitals located on the coast in the cases of nations not land-locked. Cases in which multiple cities are surveyed in the same country are marked with a common color.


Figure 3. Kernel Density Estimates of Price Distributions



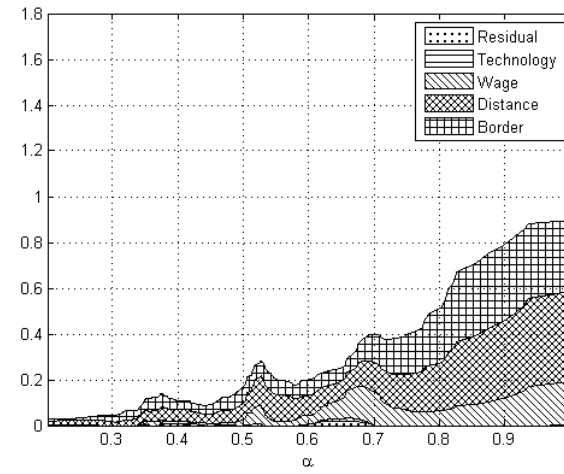
Note: The solid lines are kernel density estimates of the distribution of p_{ijk} , time averaged LOP deviations over the period 1990-2005. The dashed lines are kernel density estimates of the distribution of time series deviations from these long-run values. Each chart contains a different location and commodity grouping as indicated by the headers.

Figure 4. Sources of price dispersion and the distribution share (α_i)

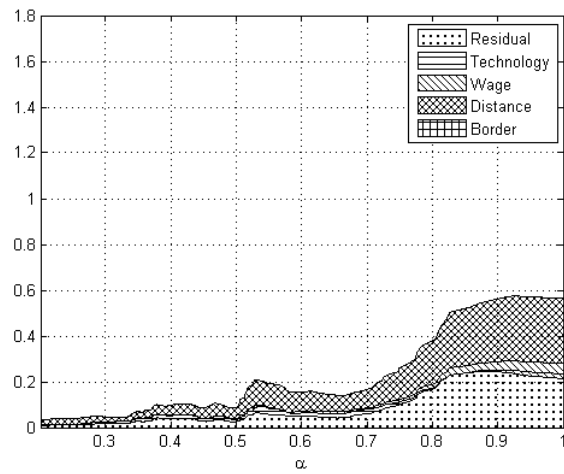
Canada-U.S. Intranational



Canada-U.S. International



OECD Intranational



OECD International

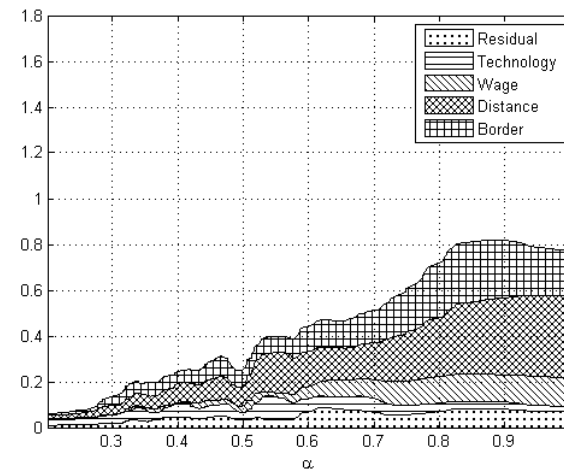




Figure 5. Sources of price dispersion and the distribution share (α_i)


World Intranational



World International



LDC Intranational



LDC International

