A Model of International Cities: Implications for Real Exchange Rates

Mario J. Crucini and Hakan Yilmazkuday[y,z]

October 2008

Abstract

We develop a model of international cities with each city inhabited by two representative agents, one specializing in manufacturing the other in distribution. The manufactured good is exported from the city of production to all destination cities. The distribution sector represents the physical transformation of all internationally traded goods from the factory gate to the final consumer. We allow the distribution share to vary across goods. Using a panel of micro-prices at the city level, we decompose the long-run cross-sectional variance of LOP deviations into the fraction due to distribution costs, trade costs and a residual. For the median good, we find trade costs account for 50 percent of the variance, distribution costs account for about 10 percent. Since the micro-data we use is skewed toward traded goods, we also decompose the variance based on the median good on an expenditure weight based. The tables turn, with distribution accounting for 43 percent and trade costs 36 percent.

JEL Classi...cation: E30, F31, D41

Key Words: Real Exchange Rates; Price Dispersion; Law-of–One-Price; Competitive Equilibrium; Cities

1 Introduction

According to the U.S. National Income and Product Accounts, expenditure by consumers at the retail level is about twice what producers receive for the same goods and services. This difference has come to be called the distribution margin. The distribution margin includes transportation costs from the factory gate to the final point of consumption as well costs and profits at the wholesale and retail stage.

Most existing models abstract from the distribution sector entirely and focus on the fraction of transportation costs attributable to international shipments. Abstracting from the distribution sector is problematic for three reasons. First, the distribution sector may help us to understand the large and persistent deviations from the Law-of-One-Price or Purchasing Power Parity. Second, the general equilibrium interaction of the distribution sector and the production sector is not well understood. Given the prominent role of the dichotomy between traded and non-traded goods in international finance, this is an important omission. Recent evidence also suggests that information technology and scale economies in distribution have altered the efficiency and markup structure of the distribution

Department of Economics, Vanderbilt University, VU Station B #351819, 2301 Vanderbilt Place, Nashville, TN 37235-1819. Phone: (615)-322-7357, Fax: (615) 343-8495. mario.j.crucini@vanderbilt.edu

[y]Department of Economics, Vanderbilt University, VU Station B #351819, 2301 Vanderbilt Place, Nashville, TN 37235-1819. Phone: (615)-343-2472, Fax: (615) 343-8495. hakan.yilmazkuday@vanderbilt.edu

[z]We especially thank Mary Amiti, Eric Bond, Paolo Pesenti and Robert Vigfusson for valuable comments and criticisms. We thank seminar participants at the Federal Reserve Banks of Dallas and New York, and Vanderbilt University; conference participants at the 2008 Midwest International Economics Meeting at The Ohio State University, the 2008 Midwest Annual Meeting participants in Chicago, IL, the Wegman's Conference at the University of Rochester, and the Murray S. Johnson Memorial Conference, The University of Texas at Austin. The authors gratefully acknowledge the financial support of the National Science Foundation (SES-0136979, SES-0524868).

sector. These developments may have fundamentally altered price dispersion and dynamics, both across locations within countries and across countries. Third, the distribution sector includes the medical and education sectors which have grown in importance over time. Much of this market activity is geographically segmented to due the arms-length nature of exchange and public policy decisions. As these sectors tend to rise in expenditure share with the level of development, their economic importance will continue to rise as well.

We have two goals in this paper. The first is to develop a tractable stochastic general equilibrium model of production and distribution at the microeconomic level of individual goods and services across cities. The second is to use the theory to specify a regression model to estimate micro-parameters which in turn are used to conduct a decomposition of the variance of prices across cities into a distribution margin, a trade cost margin and a residual. The distribution margin is further parsed into differences in labor and retail infrastructure costs across cities. The Economist Intelligence Unit retail price data along with supplementary sources for wages at the city level are used in the empirical work. Since the model assumes perfect competition and abstracts from official barriers to trade, the residual in the regression equation is expected to include markups, official barriers to trade and measurement error.

The model assumes that each city is inhabited by two representative agents, a manufacturer and a retailer. The manufacturer produces a single homogeneous good using labor as the only input, the good is shipped to all other cities of the world and deviations in the prices of these traded goods reflect only shipping costs from the factory door to the receiving dock at the retail establishment. The retailer transforms these goods by combining them with her labor and a fixed factor; she may also produce pure services which require no traded inputs at all. The fixed factor is intended to capture fixed retail infrastructure, broadly defined to include land, buildings, equipment and public infrastructure.

The advantage of drilling down to the level of individual goods and services at the city level is that we can learn a great deal about the production structure from the cross-sectional variance in the data. What distinguishes the manufacturing sector from the distribution sector is intimately related to what distinguishes a personal computer from a haircut. Aggregating the data tends to bias results and obscure these differences. For example, if trade costs are symmetric, aggregating across imports and exports has the effect of understating their role. Having cities as the locations allows both greater attention to the spatial dimensions of manufacturing specialization and a more precise measure of the distance between production and consumption locations.

We have two sets of results, one for the sources of LOP variance for the median good, the other for the differences in the sources of variance across goods in the cross-section. For the median good, trade costs account for about 50% of LOP deviations, the distribution margin accounts for about 20% and the remaining 30% is unaccounted for. These fractions are fairly stable across sets of locations that include high and local income countries; they also differ little when a national border is crossed compared to when it is not. The absolute level of cross-sectional variance does rise when a border is crossed or when comparison are made between cities with vastly different wealth levels, but the fractions accounted for by different factors remains more or less the same. One exception is the division of the variance accounted for by the distribution sector into labor and infrastructure components. Variance across low income countries is dominated by differences in the infrastructure component, with labor playing a small role. For other countries, the division of the distribution margin across labor and land is quite similar.

Turning to differences across goods, we find substantial variation at the absolute level of LOP, consistent with Crucini, Telmer and Zachariadis (2005). The structure of the model and methodology allow us to say more about the underlying sources of this heterogeneity. In the international data, the distribution margin accounts for 50% of cross-sectional variance in LOP deviations for the good with the highest distribution share and this fraction falls to a mere 10% as we move to the good with the lowest distribution share. Retail infrastructure accounts for more than 30% of the cross-sectional variance in LOP deviations across Canada and the United States for the good with the highest infrastructure intensity while accounting for virtually none of the variance for the good with the lowest infrastructure intensity.

Our model structure is closest to Giri (2007) who adds a good-specific distribution cost onto the baseline Eaton-

Kortum structure. In his model distribution services are in fixed proportion to the physical units of the base good as in Burstein, Neves and Rebelo (2003) (BNR) with efficiencies drawn from a distribution with a country-specific mean and common world-wide variance. In contrast, we assume that the technological parameter for distribution inputs is good–specific. Given that this margin is measurable in the NIPA, we view this a more plausible way to model the sector than the random efficiency approach. Our model shares with Alvarez and Lucas (2007), Atkeson and Burstein (2007), Eaton and Kortum (2002), and Kanda Naknoi (2008) an interest in the role of traditional trade costs. However, to the extent these papers incorporate a distribution sector, it is a common wedge across all goods in the retail basket, which assumes away any cross-sectional variance in price deviations due to the distribution margin. We find this heterogeneity to be essential in our empirical work.

2 The Model

Each city, indexed by j, is inhabited by two representative agents. One representative agent specializes in the production of a single traded good, also indexed by j. A second representative agent specializes in retail trade and production of non-traded services. Production in the manufacturing sector is proportional to labor input, the factor of proportionality is a random productivity variable. Retail production requires both labor and infrastructure. Productivity varies across cities in both the traded goods sector and the retail sector.

Traded goods are subject to iceberg transportation costs which are good and destination specific. Final goods and local inputs (retail labor and infrastructure) are not traded beyond the city limits. While hours and consumption are both choice variables, the assumptions we make in the model imply constant hours in all sectors in all locations, reminiscent of the Long and Plosser (1983) multi-sector closed economy real business cycle model. Infrastructure, including land, capital and equipment, is in fixed supply (we denote infrastructure as K).

We introduce the good index, i, to distinguish physical objects from the identities of agents and locations only when needed to avoid confusion. When we discuss the flow of goods from one location to another, our convention is that the source is first subscript and the destination is the second subscript. Thus, X_{sd} refers to the shipment of good X, from city s to city d. Given our assumption that individuals at each location specialization, s also indicates the good and the individual to whom the income flows, while d indicates the expenditure side of the equation. $_{sd}$ will be the iceberg shipping cost from the source to the destination. Since there are no durable goods or assets in the model, adding time subscripts is innocuous. We omit them here to focus on the steady-state properties of the model and to simplify notation.

The full solution for quantities and prices is given in the model appendix. In this section we present the overall model and parts of the solution relevant for pricing implications, which is our focus.

2.1 Consumers

Agents preferences are log-additive over consumption and leisure:

$$U\ C_j^A; L_j^A \quad (1 \quad) \log C_j^A + \quad \log L_j^A\ , A = m; s. \tag{1}$$

C_j^A is aggregate consumption and L_j^A is leisure, for an individual working city j. There are two individuals in each city, indexed by $A = m; s$; one is engaged in the manufacture of a single good (m) and the other is engaged in retailing and service activities (s). The consumption aggregate is CES over varieties of manufactured goods produced worldwide:

$$C_j^A = \sum_i^M (\ _i)^{\frac{1}{"}} \ C_{ij}^A{}^{\frac{"\ 1}{"}} \Bigg)^{\frac{"}{"\ 1}}. \tag{2}$$

C_{ij}^A is the consumption of good i in city j by worker type A; $" > 0$ is the elasticity of substitution across goods, $_i$ is a good specific taste parameter and M is the number of manufactured goods in existence. M is also the number of cities given our specialization assumption.

The two agents inhabiting city j, maximize utility (1) subject to their respective budget constraints:

$$\sum_i P_{ij} C_{ij}^A \quad W_j^A N_j^A + \varphi^A H_j K_j \tag{3}$$

where P_{ij} is the price of good i in destination city j. These prices will be the same for all agents in the same location, but differ across locations for reasons described below. For the manufacturing worker the location index, the good index and the individual index are all the same. Each of the two residents of city j earn labor income from their production activities and receive a share of the rental income accruing to the retail infrastructure ($\varphi^m + \varphi^s = 1$), the stock of which is assumed to be fixed at K_j. The price of a unit of retail infrastructure is H_j.

The consumer's problem may be solved in two stages. In the first stage, the consumer chooses aggregate consumption and leisure, subject to a budget and time allocation constraints. In the second stage, the consumer minimizes expenditure across goods. Here we collapse the problem to a single stage for brevity. The key equations from the consumer's problem are:

$$C_{ij}^A = \left(\frac{P_{ij}}{P_j}\right)^{-\theta} C_j^A \tag{4}$$

$$C_j^A = \frac{W_j^A N_j^A + \varphi^A H_j K_j}{P_j} \tag{5}$$

$$N_j^A = 1 - \varphi^A \frac{H_j K_j}{W_j^A} \tag{6}$$

$$L_j^A = \quad + \varphi^A \frac{H_j K_j}{W_j^A} \tag{7}$$

Aggregate consumption is based on the ideal deflator $P_j \quad \sum_i \left(P_{ij}\right)^{1-\theta}\Big)^{\frac{1}{1-\theta}}$, which ensures $\sum_i P_{ij} C_{ij} = P_j C_j$ as well as a mapping from price indices to welfare.

The first equation determines consumption demand for a particular good as a function of the relative price of the good paid by the final consumer and that individual's aggregate consumption level. It is important to note that, P_{ij} is the retail price of good i, in city j; it embodies the cost of local retail services paid to the retailer in addition to the traditional iceberg trade costs of the imported item. The price index, is a weighted average of these retail prices, the empirical counterpart would be the CPI index. The second equation is aggregate consumption of an agent, which is equal to her real income. Real income is the sum of wage income and income from the rental of retail infrastructure, deflated by the price level.

The last two equations determine hours of work and leisure. In the absence of rental income, the two agents would work the same number of hours, independent of their relative wage, due to the offsetting income and substitution effect of wages on effort with Cobb-Douglas preferences. In the presence of rental income the requirement for constant effort in equilibrium is that the ratio of rental income to labor income be constant.

2.2 Manufacturers

The production function for manufactured good, i, is:

$$Y_i = A_i N_i^m \tag{8}$$

where A_i is the productivity level.

Manufacturers choose labor input to maximize profit:

$$\max_{N_i^m} (Q_{ii} Y_i \quad W_i^m N_i^m) \tag{9}$$

the manufacturer receives the factor gate price, Q_{ii}, for every unit produced, no matter where the goods end up being sold. Given our assumptions of constant returns to scale, perfect competition and one factor of production,

the factory gate price equals the manufacturing wage divided by productivity:

$$Q_{ii} = \frac{W_i^m}{A_i} . \tag{10}$$

Given our assumption of specialization, the productivity level in this expression is good and city-specific. The presence of a nation-specific component could easily be treated here as the tendency of productivity levels to be similar across good produced in a particular country.

Retailers in each city purchase the manufactured goods and pay a proportional shipping cost. Thus the retailer purchase price is the factory-gate price marked up by a proportional shipping cost:

$$Q_{ij} = (1 + \quad_{ij}) Q_{ii} \tag{11}$$

where $_{ij} > 0$ is the net transportation cost from i to j. Q_{ij} is the retail price in the destination city and we assume $_{ii} = 0$. Effectively we assume the local manufacturing plant is close enough to the city to ignore local transportation costs. The destination price of the manufactured good depends: i) positively on both the manufacturing wage and the trade cost; and ii) negatively on manufacturing productivity.

2.3 Retailers

The retailer in each city optimally chooses how much of each manufactured good to purchase from various cities of the world. The retailer transforms these goods using a fraction of her time endowment and some amount of the local retail infrastructure. The retailer then sells the resulting composite good in the local retail market. The production function for good i, sold in city j is:

$$R_{ij} = \quad B_j N_{ij}^s \quad {}^{i} (K_{ij})^{1 \quad {}_i} \quad {}^{1 \quad {}_i} (G_{ij}) \quad {}^{i} \tag{12}$$

G_{ij} is the amount of the manufactured good (or input) imported from city i, by a retailer in city j. N_{ij}^s is the fraction of the retailer's time endowment allocated to the transformation of imported good i for local consumption in location j and K_{ij} is the amount of retail infrastructure allocated to retail good i in city j. B_j is labor-augmenting productivity specific to the city (equivalently, the retailer), common to all goods sold there.[1]

We assume the production function is common to all locations, but specific to the good. Our specification is very flexible. It captures pure labor services (e.g., baby-sitting services) with $_i$ equal to zero and $_i$ equal to one; internet purchases (e.g., Amazon.com book purchases), $_i$ equal to one, and points in between.

The retailer in city, j, maximizes profits by choosing the quantity of the traded input, G_{ij}, to import and hours and retail infrastructure to allocate to each activity:

$$\max_{N_{ij}^s ; K_{ij} ; G_{ij}} (P_{ij} R_{ij} \quad W_j^s N_{ij}^s \quad H_j K_{ij} \quad Q_{ij} G_{ij}) \tag{13}$$

At the optimum, the unit price equals marginal cost. Given constant returns to scale and three factors of production, the retail price of good i sold in location j, is a Cobb-Douglas aggregate of the price (inclusive of trade cost) that the retailer paid to acquire the traded input, Q_{ij}, the retailer's market wage, W_j^s, and the price of retail infrastructure, H_j:

$$P_{ij} = MC_{ij}^s = \quad_i \quad W_j^s / B_j \quad {}^{i} (H_j)^{(1 \quad {}_i)} \quad {}^{(1 \quad {}_i)} (Q_{ij}) \quad {}^{i} \tag{14}$$

$$_i \quad {}^{1} \qquad {}_i \quad {}^{i} \quad (1 \quad {}_i)(\quad {}_i) \quad {}^{i} (1 \quad {}_i)^{(1 \quad {}_i)} \quad {}^{(1 \quad {}_i)} \tag{15}$$

The retail price in city j is rising in input prices and falling in retail productivity, B_j.

[1] We could add good-specific productivity of retailers to account for different levels of competency across goods, but we lack productivity data to operationalize this idea.

2.4 Equilibrium

The appendix contains the tedious algebra necessary to arrive at the equilibrium allocations discussed in this section. In the remainder of the paper we specialize the consumption aggregator to be Cobb-Douglas so that we have closed form solutions.

In the global general equilibrium, all the optimality conditions of partial equilibrium must hold for consumers, retailers and manufacturers. In addition, the supply of each good must equal its total demand, including the resources lost to iceberg shipping costs.

$$Y_i = \sum_j G_{ij}(1 + \tau_{ij}) \tag{16}$$

$$= \sum_j G_{ij} + \sum_j G_{ij}\tau_{ij} \tag{17}$$

$$= G_i + T_i \tag{18}$$

In words: the production of good i, Y_i, is exhausted between the global demand for that good by retailers aggregated across destinations indexed by j, G_i, and physical loses due to iceberg costs, T_i.

Each individual has a fixed amount of time to devote to work or leisure, normalized to unity. The constraints for the manufacturer and the retailers are:

$$L_j^m + N_j^m = 1 \tag{19}$$

$$L_j^s + \sum_i N_{ij}^s = 1 \ . \tag{20}$$

The summation in the second time constraint reflects the fact that the retailer must divide her time across the M different retailing activities. The notation implicitly sets the number of goods at the retail level equal to the number of goods in the manufacturing sector. Nothing we derive requires this: we could have some activities that use no traded inputs at all in which case the number of retail goods would exceed the number of manufacturing goods by the number of pure services produced in each city.

The city's retail infrastructure is exhausted across uses:

$$K_j = \sum_i K_{ij} \ .$$

3 The data

Our focus is retail price dispersion across international cities at the microeconomic and macroeconomic level. The data source for prices is the Economist Intelligence Unit (EIU) worldwide retail price survey. The survey spans 123 cities, located in 79 countries. Most of the cities are national capitals. The larger number of cities than countries is due to the fact that the survey also 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. The number of goods and services priced in any given year is 301. The data is annual and the sample period starts in 1990 and ends in 2005. We work with LOP deviations across bilateral city-pairs:

$$q_{ijk} = \ln(S_{jk} P_{ij} / P_{ik}) \ . \tag{21}$$

Figure 1 plots four kernel estimates of q_{ijk}, the distribution we seek to explain. Two of the distributions are confined to city pairs lying within the U.S. We focus on U.S. city pairs because the U.S. is the country with the largest number of cities surveyed. The distributions using a broader set of intranational pairs is not materially different. The other two distributions are confined to city pairs that span a national border. We use all available international data in 2005 in the Figure. We see a clear ranking of price dispersion, with U.S. traded goods being least dispersed, U.S. non-traded goods next, followed by the same ranking of international traded and non-traded goods.

The EIU survey is sparse in terms of wage data. Thus we use supplemental wage data at the country level from the International Labor Organization (ILO) survey of occupational and sectoral wages and at the city level from the Union Bank of Switzerland (UBS) survey. The ILO data are averages for countries. They span 49 sectors, 162 occupations and 137 countries. The sample period is annual from 1983 to 2003. The complete list of these sectors, occupations, and countries is found in Oostendorp (2005). In the raw ILO data, the most common period is the month, followed by the hour, but some countries report weekly pay, others give daily rates for some occupations, and so on. In order to have a comparable wage data across countries, the standardized version of ILO survey by Oostendorp (2005) is used: in cases in which the wage data are reported as hourly or daily, then these wages were made (roughly) comparable with monthly wages by multiplication by 160 and 20 respectively. In order to have the maximum panel size of wage data that are comparable across countries, we use the monthly wages in US dollars that have been obtained by country-specific and uniform calibration in Oostendorp (2005).

Wage data at the city level is more appropriate given the EIU retail price data is city based and the intent of the model. International cities were surveyed by the UBS in 2006. These are hourly wages in US dollars, spanning occupations in 71 international cities, 60 of which are also surveyed by the EIU. Among the 60 EIU cities there two cities from Brazil, Canada, China, France, Italy, Spain and Switzerland; three cities from Germany, and 4 cities from the U.S. The hourly wages have been obtained by dividing the income per year in each occupation by the city level hours of work in a year, where the hours of work have been obtained through surveys.

Bureau of Labor Statistics (BLS) city wage data from the Occupational Employment Statistics (OES) Survey in 2006 are used to complement UBS data. These wage data are hourly wages in US dollars. We draw the same 16 US cities from the BLS survey that are in EIU retail price survey. The combination of UBS and BLS wage data, then, provides wage data for 72 EIU cities. Within these 72 EIU cities, in terms of intranational cities, we have two cities from Brazil, Canada, China, France, Italy, Spain and Switzerland; three cities from Germany, and 16 cities from the US.[2]

In a preliminary part of the analysis, the BLS city wage data are used for broader wage dispersion analysis. These data cover 2 industries, namely production and sales, for 400 cities (on average) within the U.S. in terms of hourly wages from 1999 to 2006.

A number of trade-offs present themselves in terms of the model focus and the available data. Country-level wage data is generally available for more time periods, but fewer locations than city-level wage data. Since the model is explicitly constructed to mimic city level aggregation and steady-state features, ideally one would want long time series at the city level. Unfortunately these are simply not available. We discuss the trade-offs as they arise below.

Land prices and rents are even more difficult to come by than are wages and prices. We use two entries from the EIU survey data: "Typical annual gross rent for top-quality units, 2,000 square meters, suitable for warehousing or factory use" and "Typical annual gross rent for a 1,000 square meter unit in a Class A building in a prime location."

The other two pieces of information are sectoral estimates of the distribution shares, $1 \quad {}_{i}$ which use a combination of U.S. NIPA data and input-output tables. The NIPA data extend to 57 sectors, while the input-output data span 33 sectors. The NIPA shares are computed as the value the producer receives relative to the value consumers pay for the output of a particular sector. This difference includes transportation costs, retail and distribution costs and markups. Sectors involving arms-length transactions, such as medical services are recorded in the NIPA as though the producer and consumer valuation is equal. While this is literally true in some cases, this accounting fails to distinguish local inputs and traded inputs used in the production of services. For these sectors we use the input-output tables to determine the distribution share.

We use all of these sectoral parameter measures to check against good-level parameters that we estimate using our regression framework discussed below. Finally, the greater circle distance between cities in the EIU sample is used to estimate the trade cost component of the LOP deviations at the retail level. In the future, we plan to

[2] In an earlier version of this paper we used PWT per capita annual income data covering the annual period from 1990 to 2004 to proxy for real wages. These data span all 79 EIU countries. The results were qualitatively similar to those reported here.

compare these estimates to measures constructed by Hummels (2000) and others.

4 Microeconomic sources of variation in wages

In the model, wage deviations arise across the retail and manufacturing sectors and across cities. The amount of labor income accruing to the manufacturer relative to the retailer in city j is,

$$\frac{N_j^m W_j^m}{N_j^s W_j^s} = \frac{\sum_i \quad {}_i \quad {}_i \quad {}_i}{\sum_i (1 \quad {}_i) \quad {}_i \quad {}_i} \tag{22}$$

Which is quite intuitive: the numerator is an expenditure share weighted average of labor's share of manufacturing and the denominator is the counterpart in retailing. The appearance of the parameter ${}_i$ in the denominator accounts for the fact that retail production involves some retail infrastructure, unless ${}_i = 1$, in which case we have a labor-only production economy. Note, also, that the ratio is the same in all cities.

As our primary interest is wage variation in order to account for cost variation across cities, we would like to understand the wage ratio and effort ratios separately. The equilibrium relative sectoral wage is given by:

$$\frac{W_j^m}{W_j^s} = \frac{'^m}{'^s} \frac{(1 \quad)}{(1 \quad)} \frac{N_j^s}{N_j^m}$$

Thus, given fixed shares of rental income across agents in the city, relative wages and relative hours move inversely as we would expect. The appendix shows that the equilibrium effort levels are:

$$N_j^m = \frac{(1 \quad {}_0)(1 \quad)}{(1 \quad {}_0) + '^m \quad ({}_0 \quad {}_1)}$$

$$N_j^s = \frac{{}_1 (1 \quad)}{{}_1 + '^s \quad ({}_0 \quad {}_1)}$$

where ${}_0 \quad \sum_i (1 \quad {}_i) \quad {}_i$ and ${}_1 \quad \sum_i (1 \quad {}_i) \quad {}_i \quad {}_i$. Effort in either sector is declining in the share of rental income allocated to the agent, and declining in the preference for leisure ( ), as one would expect.

Substituting these expressions into the wage ratio eventually leads to the following expression for relative wages:

$$\frac{W_j^m}{W_j^s} = \frac{'^s}{'^m} \frac{(1 \quad {}_0) + '^A}{{}_1 + (1 \quad '^A)}$$

$= \quad ({}_0 \quad {}_1)$.

As the retail sector becomes more labor intensive (thus rental income falls), $({}_0 \quad {}_1)$ converges to zero and we revert to the labor-only version of the model with a common fraction of available hours worked by both agents, equal to $(1 \quad)$ and the sectoral wage ratio converges to:

$$\frac{W_j^m}{W_j^s} = \frac{\sum_i \quad {}_i \quad {}_i \quad {}_i}{\sum_i (1 \quad {}_i) \quad {}_i}$$

Turning to wage differences across cities things are much simpler even in the general case:

$$\frac{W_j^s}{W_k^s} = \frac{{}_j \quad {}_j}{{}_k \quad {}_k} \tag{23}$$

$$\frac{W_j^m}{W_k^m} = \frac{{}_j \quad {}_j}{{}_k \quad {}_k} . \tag{24}$$

The cross-city wage differential is the same in both sectors and is determined by the product of the taste and technology parameters in the two locations being compared. The intuition for this result is as follows. Consider, first, the special case in which all goods use traded inputs in the same proportion, ${}_j = \quad$. The bilateral wage

differential then favors the location whose goods are most preferred, $_j$ versus $_k$, a demand-side effect. Next, consider the case in which there is no taste bias across goods $_j =$; then wages are highest for producers of manufactures with the greatest amount of traded input at the distribution stage. Essentially, this determines the productivity of manufacturing effort in delivering a unit of final consumption.

Wage data is available by occupation or sector of employment. Our model focuses on the distinction between goods and services, suggesting the production sector definition is more appropriate. However, we use both data sources as a robustness check.

The more comprehensive of the sources we use is the ILO survey of wage levels across countries. These data span 49 sectors, 162 occupations and 137 countries. The sample period is annual from 1983 to 2003.[3] Because our model is intended to be based on city-level data, our preferred measure is wage data from the UBS that span 14 occupations and 71 international cities for the year 2006.

According to the model, if the retail sector uses only labor and traded goods, the ratio of manufacturing wages to service wages provides on estimate of the overall scale of the distribution sector $W_j^m = W_j^s = \quad = 1 \quad , \quad = \quad P_i \quad i \quad i$. Since we lack consumption expenditure shares at the present time, we associate this with the distribution share alone, using the symmetric taste version of the model, reduces $^{-}$. A direct way to measure the overall size of the distribution sector is to use U.S. NIPA data and input-output data. Crucini and Shintani (2008) do exactly this and find $= 0:57$. The advantage of their calculation is that is it based on expenditure weighting of sectoral 's.

Table 1 reports the sectoral wage ratio averaged across locations as well as the implied value for . It turns out that the direct and indirect (model-based) estimates are exactly equal when we use U.S. wages in production sector relative to the sales sector. The wage ratio in the international data is consistent with a value of 0.52. While this is a modest difference from the U.S. value, the implied manufacturing wage premium is much smaller, a factor of 5 smaller. It could be that relative productivity differences are the cause. Another possibility is that the composition of sectors in responsible for the difference.

As the theory is a two-sector model, any sectoral variation in wages at a particular city is attributable to the ratio of wages in the manufacturing sector to those in the service sector. Variation in wages across sub-sectors are abstracted from entirely. Thus, it is important for our theory of retail price dispersion that wages differ significantly across locations and less so across sectors other than the two we model. And this is what we find.

Table 2 conducts a variance decomposition by sector and country for time-averaged wages in the case of the ILO survey and an analogous decomposition by sector and city for wages in 2006 for the UBS survey data. Since the answer may depend on the set of sectors and locations used, we consider three location groups and allow for many sectors throughout. The location groups are the entire world, the OECD and the LDC.

Based on the ILO wage data: locations account for between 72 and 85 percent of the cross-sectional variation in wages, sectoral differences account for less than 10 percent. The relative importance of location is somewhat lower when the data is organized by occupation, with location effects dropping to between 38 and 65 percent. Most of difference is not attributed to a pure sectoral component, but rather an interaction of location and sector. The UBS tell a similar story to the ILO for location effects, with the occupation effect rising in contribution due to a lower interaction with location compared to the ILO.

We take from this the view that location is a key component of wage dispersion with the precise fractions depending somewhat on the set of locations examined and the precise definition of wage categories.

5 Microeconomic sources of variation in real exchange rates

We turn, now, to our main focus, price dispersion. In the model, prices consumers actually pay may differ from factory gate prices for two reasons. The first is the trade cost to import the good from the foreign production location. The second is the value added by the retailer. To simply the notation, we assume that all international

[3] Useful technical documentation is found in Remco H. Oostendorp (2005).

prices have been converted to common currency units (it does not matter which numeraire is chosen). The ratio of the price of good i in city j relative to k, based on the theory is:

$$\frac{P_{ij}}{P_{ik}} = \frac{W_j^s/B_j}{W_k^s/B_k}^{\ \ {}_i(1 \quad {}_i)} \frac{H_j}{H_k}^{(1 \quad {}_i)(1 \quad {}_i)} \frac{Q_{ij}}{Q_{ik}}^{\ \ {}_i} .$$

Noting that the first term reduces to the ratio of trade costs from the single source of good i to each of the destinations, j and k and taking logs, defines the Law-of-One-Price deviations we seek to explain. The empirical counterpart is: $q_{ijk} = \ln(S_{jk} P^*_{ij} / P^*_{ik})$, where the '$*$' indicates domestic currency values.

$$\begin{aligned} q_{ijk} &= (1 \quad {}_i)[\ {}_i \omega_{jk} + (1 \quad {}_i)\, h_{jk}\,] + \ {}_i \ {}_{ijk} \\ &= \ {}_i \omega_{jk} + \ {}_i h_{jk} + \ {}_i \ {}_{ijk} . \end{aligned} \tag{25}$$

The retail margin is the first term in square braces and is a weighted average of the productivity-adjusted wage and the rental price differential faced by the retailer in each city. The second term is the relative trade cost. The weights attached to the relative input prices in the retail sector depend on ${}_i$. The retail component in its entirely gets weighted by its overall share in the production of the final good, $(1 \quad {}_i)$. The last line, used to specify our regression approach, expresses the relationship in terms of the three key cost ratios, retail wages, rental prices and trade costs.

The aggregate real exchange rate relevant for the theory is immediate since the consumption aggregator is Cobb-Douglas (i.e., $\varepsilon = 1$):

$$\begin{aligned} q_{jk} &= \omega_{jk} + h_{jk} + \sum_i \ {}_i \ {}_i \ {}_{ijk} \\ &\sum_i \ {}_i \ {}_i (1 \quad {}_i), \qquad \sum_i \ {}_i (1 \quad {}_i)(1 \quad {}_i) \end{aligned} \tag{26}$$

The aggregate real exchange rate has a number of interesting features. The distribution components of PPP deviations are driven by the same wage and rental ratios as the LOP deviations, they are just evaluated at the expenditure-weighted means of the LOP deviations. The trade cost component, however is likely to average out across goods since the ${}_{ijk}$ are expected to vary in sign across goods.

5.1 Regression speci…cation

This section conducts a variance decomposition of retail prices into the channels described by the equilibrium model. Adding a measurement error term to the theoretical equation for the LOP deviation, we have:

$$q_{ijk} = \ {}_i \omega_{jk} + \ {}_i h_{jk} + \ {}_i \ {}_{ijk} + \varepsilon_{ijk} \tag{27}$$

We have a wealth of micro-data and yet we face some significant empirical hurdles. We have data on retail prices, wages and rent, but no data on retail productivity or trade costs. We ignore retail productivity differences and use the raw wage ratios in place of ω_{jk}. To isolate the role of trade costs, we employ a two-stage estimation approach.

In the first stage we regress LOP deviations on the wage and rental ratios:

$$q_{ijk} = b_{1i} \omega_{jk} + b_{1i} h_{jk} + b_{ijk} . \tag{28}$$

The residual, according to our model, will be the LOP deviation in the traded component of cost. In practice it will incorporate other sources of deviations as well.

What we do in an attempt to extract the pure trade cost component from these residuals is project them on distance. To accomplish this, we must first define a direction-of-trade indicator function:

$$I_{ijk} = \begin{cases} 1 & \text{if } b_{ijk} > 0 \\ 1 & \text{if } b_{ijk} < 0 \end{cases} \tag{29}$$

where $\hat{\beta}_{ijk} = q_{ijk} \quad \beta_{1i}\omega_{jk} + \beta_{1i}h_{jk}$ from the first-stage regression.

Consider, now, the more elaborate equation for stage two:

$$q_{ijk} = \beta_{2i}\omega_{jk} + \beta_{2i}h_{jk} + \delta_{ijk}I_{ijk}\, d_{jk} + \hat{\beta}_{ijk} \tag{30}$$

$$\beta_{2i} = (1 \quad {}_i)\ {}_i \tag{31}$$

$$\beta_{2i} = (1 \quad {}_i)(1 \quad {}_i) \tag{32}$$

$$\delta_{ijk} = {}_i\ {}_i \tag{33}$$

with the trade cost replaced by $I_{ijk}\ {}_i d_{jk}$. The indicator function ensures the sign of the implied trade cost is consistent with the sign of the residual estimated in stage one. The greatest circle distance between locations j and k is the empirical counterpart to d_{jk} with goods having a different trade cost elasticity with respect to distance, ${}_i$. The benefit of projecting the prices on wages, rents, and the indicator function multiplying distance is that we relegate any sources of variation in retail prices not correlated with wages, rental prices or distance to the error term. This gives us more confidence that the wage, rental, and trade cost components are capturing what the model says they should.

The model seems best suited to describe the long-run properties of real exchange rates since we abstract from nominal exchange rate variation and sticky prices. While we have a long panel of EIU retail price data from which to construct time-averages and target long-run dispersion, we lack city-level panel data on wages. Moreover, the argument could be made for estimating the parameters with a single cross-section. Our benchmark estimation and variance decomposition uses time-average data as available (i.e., for q_{ijk} and h_{jk}) and wage data for a single cross-section in 2006. Wage data from the UBS is used for cities outside of the U.S. and from the BLS in the case of U.S. cities. Preliminary experimentation with alternatives does not seem to alter the main thrust of our results.

We see in Table 3, that the empirical model captures the majority of variation in retail prices across locations for all groupings of the data. The range of variance accounted for is between 70% and 90% for the median good when pooling all international cities or just those in North America. The quality of the fit of the model is excellent over much of the distribution across goods. The lowest quartile mean for the R^2 is a very respectable 0.67 (the OECD cross-border pairs). Thus the empirical model fits well across sub-panels of locations and goods ranging from haircuts to personal computers.

5.2 Variance Decomposition

Interesting and more relevant than explanatory power for the theory is the magnitude and source of variation in the underlying price data. The theoretical model attributes price dispersion to distribution margins, trade costs and potentially their covariance. Using the estimated equations motivated by our theory, we are able to provide a cross-sectional variance decomposition analysis according to the following equation:

$$var_{jk}(q_{ijk}) = \beta_{2i}^2 var_{jk}(\omega_{jk}) + \beta_{2i}^2 var_{jk}(h_{jk}) + (\delta_i \beta_i)^2 var_{jk}\left[\delta_{ijk}\, d_{jk}\right] + var_{jk}[\hat{\beta}_{ijk}] + \text{cov terms}$$

Consider a good which uses no traded inputs at the retail level (${}_i = 0$). The prediction simplifies reduces to:

$$var_{jk}(q_{ijk}) = \beta_{2i}^2 var_{jk}(\omega_{jk}) + (1 \quad \beta_{2i})^2 var_{jk}(h_{jk}) + \text{cov terms}$$

and price dispersion is entirely due to retail costs associated with wage and rental price dispersion, $var_{jk}(\omega_{jk})$ and $var_{jk}(h_{jk})$, respectively. These number naturally depend on the locations pooled in the estimation and we find larger values for geographic dispersion in both wages and prices internationally than internationally. Borders matter.

Consider the opposite end of the continuum, a good with no retail costs at all (e.g., a good available on the internet that trades up to a shipping cost everywhere in the world). Now the expression for the predicted price dispersion reduces to:

$$var_{jk}(q_{ijk}) = \beta_i^2 var_{jk}\left[\delta_{ijk}\, d_{jk}\right]$$

This is an intriguing expression. The coefficient out front is the elasticity of trade cost with respect to distance (we are assuming a log-linear proportional trade cost function as is typical in the gravity literature). The variance of distance is a function of the set of locations under examination (below we contrast locations within North America, the OECD, LDC and all international pairs). As bilateral distance become less symmetric (less equal), trade cost matters more for price deviations.

Beginning with the absolute variances, in the first column of Table 4, we see the lowest retail price dispersion exists within countries, about 0.06. Keep in mind that the U.S. has the largest number of intranational pairs and therefore dominates the intranational samples throughout the table. International price dispersion is highest when we pool the LDC locations (0.42) and lowest when we focus on North America (0.07).

For the median good, distribution costs account for between 5 and 20 percent of overall price dispersion. The wage component tends to account for more of this dispersion than the rental component. An exception is the LDC group where the rental component accounts for 12.6 percent of the dispersion, compared to only 2.5 percent for wages. Trade costs dominate the picture through the table, accounting for as much as 60 percent of the price dispersion for cross-border OECD pairs, to a lower, but significant, 36.1 percent across the Canada-U.S. border. Approximately 30% of the variance is left unaccounted for by the model. This variation could be due to a combination of markup variation, official barriers to trade or measurement error. The covariance across effects is typically less that 5 percent. The bottom line of the analysis of the median good are that trade costs dominate independent of the location or border crossing and that distribution margins are important enough not to ignore.

Variation across goods within the cross-section, is more interesting. Figure 2 shows the variance decomposition at the individual good level as a function of the traded input share, $_{i}$. To make these easier to read we have smoothed the profiles by taking centered moving averages of the variance share across 10 goods. For the good with the lowest traded input share (roughly 0.4), wage dispersion accounts for about 45% of price dispersion. As we move to goods with the highest traded input share (roughly 0.97), wage dispersion accounts for almost none of the price dispersion. Of course if we had a good with no non-traded inputs the contribution would necessarily be exactly zero. The OECD group tells a similar story with about 30 percent of price dispersion accounted for by wage dispersion at one end of the continuum of goods and less than 10 percent contributed for goods embodying mostly traded inputs. The Canada-U.S. pairs have a lower contribution from wage dispersion as we would expect given the similar wage levels of the two countries, the contribution of this component also declines as rises, though not as smoothly as the other groups. In most cases, the falling contribution of wage differences is associated with a rising role for trade costs. The intranational pairs show less heterogeneity in the proportion of variance explained by various components as the trade share of final good production varies. Partly this reflects the lower variance of wages and rent across cities within countries. Nonetheless, the contribution of distribution is not negligible for the intranational pairs either.

Figure 3 displays the same variance decomposition by good plotted against the labor share of total retail cost, $_{i}$. We see the dramatic effect of this parameter on the split of distribution margin variance across labor and rent in the case of Canada and the U.S. as well as the the world. As we move across goods based on this parameter, the contribution of rent goes from zero to about 40% in the Canada-U.S. and from zero to about 20% in the world grouping. The contribution of wage dispersion tends to follow the same pattern in reverse, maintaining the total share of price dispersion due to distribution costs. The OECD is anomalous in the sense that the distribution share contributes about 10% without much variation across goods until we reach very high labor intensities in distribution. However, the variation appears to come at the expense of the residual. Turning to the intranational pairs, the contribution of wage dispersion is rising in its production share as one would expect.

The results for the median good in the EIU cross-section seem to downplay the role of distribution costs relative to trade costs. Given the dramatic differences in how the variance decomposition plays out across goods, the natural question that arises is how representative the EIU sample is of the CPI basket. A second issue is the extent to which the estimated distribution share matches up with the direct measures in the NIPA data.

Regarding the second issue, the average estimated value of the distribution share across goods we are able to use

in the estimation is 0.2. This value is significantly below, 0.5, the average we get when we use merge our micro-data with the U.S. NIPA and use the sectoral values from that source. Moreover, the difference is not due to a few outliers, only 9 out of 160 good are the estimate values above the NIPA counterparts.

To account for this estimation bias and make the results relevant for aggregate consumption, we recompute our variance decomposition using goods with distribution shares in the neighborhood of $= 0:5$. What we do is average the decomposition results across 5 goods on either side of this value. Table 5 reports these findings. We see that the contribution of the distribution margin is much more significant. Wage dispersion alone now accounts for more than one-third of retail price dispersion when all cross-border pairs of cities are pooled (world). The role of wages within the OECD and LDC is more limited suggesting the city pairs that straddle high and low income countries are the reason for the much elevated wage component. It is interesting to note that for the Canada-U.S. pairs, wage dispersion plays a significant role as well. Keep in mind that the absolute dispersion in prices is about 5 time higher for the world group than the Canada-U.S. group.

6 Conclusions

Consumers face prices that are to a varying degree, location-specific. Our model of production and distribution across cities shows how these differences are shaped by the distances separating cities due to trade costs, the good-specific share of retail distribution and its division between local labor and rental costs. While we found trade costs dominated distribution costs of 5 to 1 for the median good in the Economist Intelligence Unit micro-sample, their relative contribution varies greatly across goods. For final goods that involve mostly non-traded inputs, distribution margins dominate trade costs. Given that most of the goods in the EIU have high traded input content, these computations understate the role of distribution margins in the cross-section of retail items typical of the full consumption basket. Using the aggregate distribution share and estimates of the variance decomposition for goods with that share, distribution costs dominated trade costs. In future work we will undertake and analysis of PPP using our model and empirical methodology. We expect the distribution margin will dominate trade costs in this case we well.

7 References

Alvarez, Fernando and Robert E. Lucas, "General Equilibrium of the Eaton-Kortum Model of International Trade, Journal of Monetary Economics, September 2007, 54(6), 1726-1768.

Atkeson, Andrew and Burstein, Ariel. "Pricing-to-Market in a Ricardian Model of International Trade," American Economic Review. 2007, 97(2), 362-367.

Burstein, Ariel, Joao Neves and Sergio Rebelo, "Distribution Costs and Real Exchange Rate Dynamics During Exchange-Rate Based Stabilizations," Journal of Monetary Economics, September 2003.

Crucini, Mario J., Chris I. Telmer and Marios Zachariadis, "Understanding European Real Exchange Rates." American Economic Review. 2005, 95(3), 724-38.

Crucini, Mario J. and Motostugu Shintani, "Persistence in Deviations from the Law-of-One-Price: Evidence from Micro-Data." Journal of Monetary Economics. 2008, Vol. 55, Issue 3, 629-644.

Eaton, Jonathan and Samuel Kortum. "Technology, Geography and Trade." Econometrica. 2002, 70(5), 1741-1779.

Naknoi, K. "Real Exchange Rate Fluctuations, Endogenous Tradability and Exchange Rate Regimes." forthcoming, Journal of Monetary Economics, 2008.

Oostendorp, Remco H. (2003), "The Standardized ILO October Inquiry 1983-2003," mimeo, Free University Amsterdam, Tinbergen Institute, Amsterdam Institute for International development.

Table 1. Mean sectoral wage differentials

	$W_j^m = W_j^s$	Implied
World Manufacture to Sales (ILO)	1.07	0.52
U.S. Production to Sales (BLS)	1.34	0.57

Note: For details on the data sources, see the data appendix.

Table 2. Variance of wage differentials across sectors and locations

	Industry (ILO)			Occupation (ILO)			Occupation (UBS)		
	Location	Sector	Error	Location	Sector	Error	Location	Sector	Error
World									
Proportion of variance	0.85	0.04	0.10	0.65	0.01	0.34	0.66	0.19	0.15
Observations	46	19		136	113		56	14	
OECD									
Proportion of variance	0.84	0.06	0.10	0.64	0.03	0.33	0.51	0.31	0.18
Observations	27	12		26	113		32	14	
LDC									
Proportion of variance	0.72	0.08	0.20	0.38	0.04	0.58	0.48	0.29	0.23
Observations	36	19		109	113		24	14	

Notes: A panel has been selected such that the total number of observations is maximized.

Table 3 - Explanatory power

	First quartile	Median	Third quartile
Panel A: International cities, cross-border pairs			
CANADA-US	0:83	0:90	0:94
OECD	0:67	0:71	0:74
LDC	0:70	0:73	0:75
WORLD	0:69	0:72	0:75
Panel B: Intranational cities, no border			
CANADA-US	0:72	0:77	0:81
OECD	0:71	0:76	0:79
WORLD	0:70	0:75	0:79
LDC	0:70	0:73	0:75

Table 4 - Variance Decomposition (median across goods, = 0:8)

	Total	Fraction of variance account for by:				
		Wages	Land Prices	Trade cost	Error	Covariance
Panel A: International cities, cross-border pairs						
CANADA-US	0.07	10.1	7.8	36.1	12.8	15.5
OECD	0.25	2.3	1.7	60.0	29.4	0.8
LDC	0.42	2.5	12.6	53.9	27.2	0.4
WORLD	0.38	7.7	3.4	50.7	28.1	5.1
Panel B: Intranational cities, no border						
CANADA-US	0.06	4.8	5.9	53.8	24.8	2.3
OECD	0.06	5.1	3.7	51.4	26.5	3.0
LDC	–	–	–	–	–	–
WORLD	0.07	5.2	4.9	55.8	26.1	1.7

Table 5 - Variance Decomposition (aggregate = 0:5)

	Total	Fraction of variance account for by:				
		Wages	Land Prices	Trade cost	Error	Covariance
Panel A: International cities, cross-border pairs						
CANADA-US	0.10	31.8	16.6	38.5	13.1	
OECD	0.36	10.7	6.2	54.6	28.5	
LDC	0.75	15.2	16.7	45.7	22.4	
WORLD	0.66	36.4	6.8	36.0	20.9	
Panel B: Intranational cities, no border						
CANADA-US	0.12	10.0	8.5	57.2	24.3	
OECD	0.13	6.6	11.0	56.0	26.3	
LDC						
WORLD	0.10	8.6	6.3	59.7	25.4	

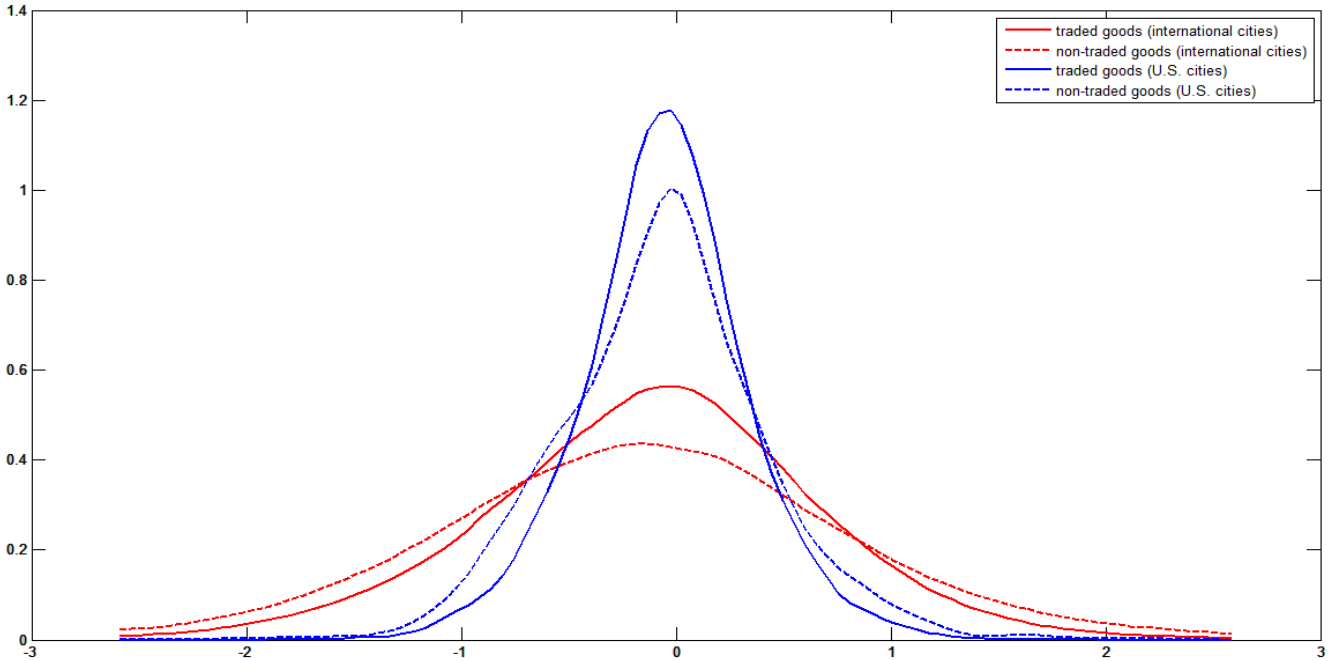

Figure 1. Kernel density estimates of price distributions. Densities ranked from less dispersed to most dispersed: U.S. traded goods, U.S. non-traded goods, international traded goods, international non-traded goods.

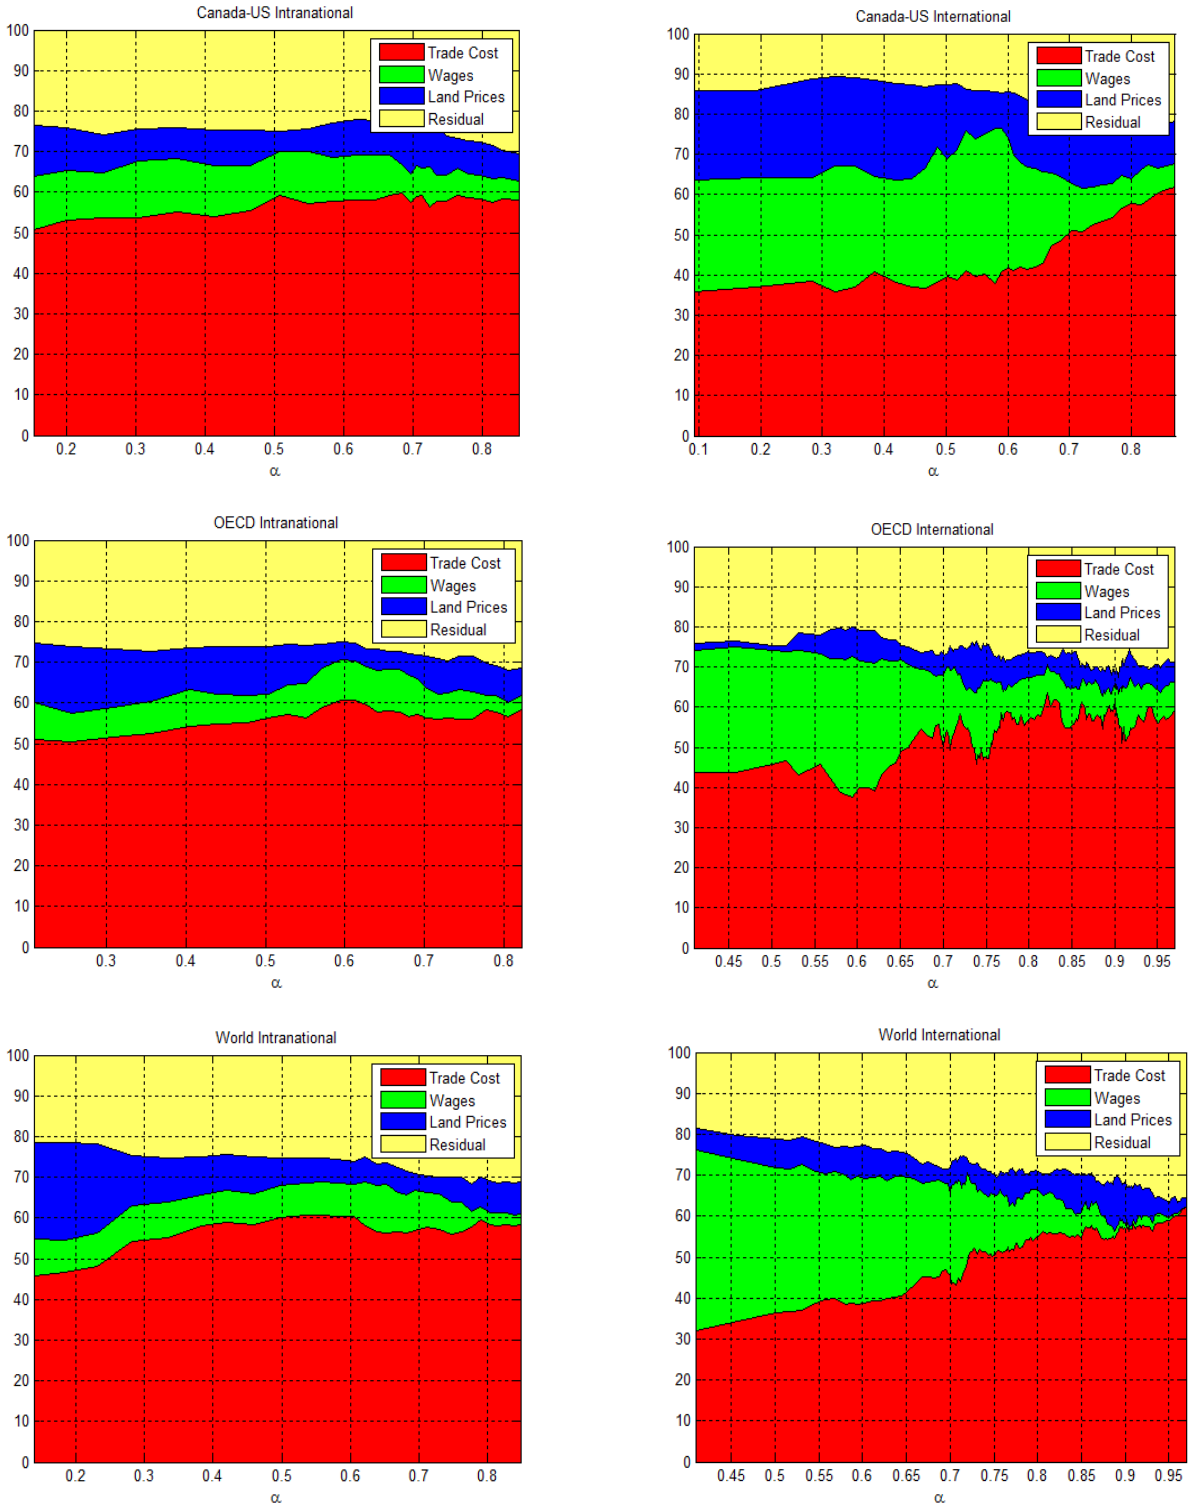

Figure 2. Variance decomposition of lprice dispersion as a function of ¡ the share of traded input costs (x-axis)

Figure 3. Variance decomposition of price dispersion as a function of$_i$ the labor share of non-traded input costs (x-axis)