Carnegie Mellon University
Research Showcase @ CMU

Tepper School of Business

12-2003

Price Dispersion: The Role of Borders, Distance and Location

Mario J. Crucini
Vanderbilt University

Christopher I. Telmer
Carnegie Mellon University, chris.telmer@cmu.edu

Marios Zachariadis
Louisiana State University

Follow this and additional works at: http://repository.cmu.edu/tepper

Part of the Economic Policy Commons, and the Industrial Organization Commons

This Working Paper is brought to you for free and open access by Research Showcase @ CMU. It has been accepted for inclusion in Tepper School of Business by an authorized administrator of Research Showcase @ CMU. For more information, please contact research-showcase@andrew.cmu.edu.

Price Dispersion: The Role of Borders, Distance and Location

Mario J. Crucini*, Chris I. Telmer† and Marios Zachariadis‡

December 2003

Abstract

1. Introduction

We study deviations from the Law-of-One-Price using microeconomic data on the retail prices of approximately 220 individual goods and services across 122 cities located in 79 countries over the period from 1990 to 2000. This paper builds on our earlier work (Crucini, Telmer and Zachariadis (2001)) which focused on how price dispersion within a fixed geography (the European Union) varied by type of good or service. For example, in moving from the least traded good with the largest share of non-traded inputs into retail production (a haircut) to the most traded goods with the lowest share of non-traded inputs into retail production (a desk-top PC), the standard deviation of price across locations dropped from about 32% to 16%. Moreover, the median (and majority) of goods and services fell between these two extremes.

In this companion essay we are interested in the interaction of the characteristics of goods and economic geography in determining relative prices. This boils down to addressing two related questions. At the level of an individual good we

*Department of Economics, Vanderbilt University; mario.j.crucini@vanderbilt.edu.

†Graduate School of Industrial Management, Carnegie Mellon University; chris.telmer@cmu.edu.

‡Department of Economics, Louisana State University; zachariadis@lsu.edu.

ask: How does price dispersion change as we alter the set of locations included in the analysis? We expect the answer to this question to depend whether a border is crossed and the physical distance between the locations. The second question is how these properties differ across goods.

We begin in section 2 with a description of the data and a summary of our main findings with the details for each finding considered in subsequent sections. Section 3 develops a retail price model in which the Law-of-One-Price deviation is a weighted average of a location specific relative price for non-traded inputs (which we associate with relative wages) and a location- and good-specific relative price for traded inputs (which we associate with transport costs). The role of non-traded inputs is to raise prices of all goods in rich countries relative to poor countries, with effect rising with the share of non-traded costs in production. The role of traded inputs is to create a dependence of relative prices on distance (trade costs are assumed to be log-linear in distance), with effect falling with the share of non-traded costs in production.

Section 4 examines the properties of the average (across goods) deviation from the Law-of-One-price, a basic building block of Purchasing Power Parity. We find evidence of a border in the sense that there is greater dispersion in price levels internationally than intranationally. Some, but not all, of this difference is attributed to greater income differences internationally (the Balassa-Samuelson theory). We find no evidence of distance effects in the mean either within or across countries provided we control for income differences internationally. We interpret the absence of distance effects in the means as reflecting an 'averaging-out property' of trade costs in our retail model.

Section 5 studies the variance of Law-of-One-Price deviations across goods for each bilateral pair of cities. We find as we did with the means that borders and income matter. However, we also find that distance matters which we interpret as evidence that the averaging out property of trade costs does not occur in the variances. The distance coefficients for intranational and international bilateral panels are not too different suggesting some robustness to the role of distance in accounting for price dispersion. Taken together, these findings suggest that the border effect is once again relegated to the constant term.

Section 6 moves from cross-section variation to time series variation. Here the unit of observation is the variance of a bilateral relative price over time. This metric is similar to the one used by Engel and Rogers (1996) except that we use the level and they use the difference. Since the level and the difference have similar variability in the micro-data, the transformation used before computing

variability turns out not have a much of an effect on the results. The pattern that emerges here is similar to that found in Section 4 (analysis of means). Distance matters internationally, but not intranationally. Moreover, there is a border in the sense that conditional on distance, international prices are more volatile. While it is trivial to attribute some of this effect to nominal exchange rate variability given the identity linking real and nominal exchange rates, it is far from obvious how to interpret the correlation. On the one hand, domestic nominal prices are not fully responding to nominal exchange rate variation, but on the other hand it is unclear why they should if the absolute deviations are large and reflect real factors that segment markets.

2. Data

The source of our price data is the annual Economist Intelligence Unit retail outlet price survey. Our data begins in 1990 and ends in 2000, effectively creating (ignoring missing data) a balanced 11-year panel of absolute prices for 220 goods and 84 services across 122 major cities across the globe. The total number of countries is 79 countries, with differences between the number of cities and number of countries reflecting the fact that in 58 of the 79 countries the EIU surveys multiple cities. The country with the most intranational observations is U.S., with 16; the next largest number of intranational observations is 5 (Australia, China and Germany)[1] The same basic data source has recently been used by Crucini and Shintani (2002), Parsley and Wei (2002), Rogers (2001) and Engel and Rogers (2003).

Our basic data unit is $P_{ij,t}$, the price, in units of local currency, of good i in location j at time t. For most of our analysis we transform this data into $q_{ijk,t}$, log deviations from the law-of-one-price (LOP) for each bilateral location-pair:

$$q_{ijk,t} = \log(\frac{P_{ij,t}e_{jk,t}}{P_{ik,t}}) \ ,$$

where $e_{jk,t}$ is the nominal exchange rate between locations j and k, in units of location k, and $e_{jk,t} = 1$ if locations j and k are in the same country.

Figure 1 shows estimates of the density function for q_{ijt} for 1990, 1995, 2000; both intranational and international city pairs. Here we normalize the the world

[1] Specifically, the number of intranational cities are (ordered from the most available cities to the least): United States (16), Germany, China and Australia (5), Canada (4), Saudi Arabia (3) and France, Italy, Russia, Spain, Switzerland, UK, India, Japan, Vietnam, New Zealand (2).

average price, rather than use bilateral pairings (i.e. $q_{ijk,t} = \log(\frac{P_{ij,t}e_{jk,t}}{P_{i,t}})$, where $P_{i,t}$ is the average price of good i across locations). We see what is obvious to anyone who has ever traveled between two locations: the LOP is not very useful for describing the properties of $q_{jk,t}$. Moreover, as one might expect, price dispersion is considerably higher across international locations. Much of the existing literature has therefore proceeded by describing a number of interesting empirical regularities inherent in the international relative price distribution. The difficulty with this lies in moving from empirical regularities to economic interpretation. Our approach, therefore, will be to (i) briefly list a set of interesting empirical regularities, and then (ii) flesh out their economic interpretation using a simple production function which relates deviations from LOP in retail prices to deviations from LOP in non-traded inputs (such as labor) and transport costs. At each step we compare and contrast the properties that exist within countries and across countries.

Our data on $q_{ijk,t}$ display the following features. In each case, specific details are deferred until the relevant section of the paper. For properties 1. to 5. we work with q_{ijk} (i.e. $q_{ijk} = \sum_{t=1}^{T} q_{ijk,t}$).

1. *The Balassa-Samuelson Effect for the Mean.* The cross-good average value of q_{ijk} is strongly related to income and labor productivity differences between locations j and k. More precisely, the cross-sectional mean, $E_i(q_{ijk}|\, jk)$ (across goods for each date and location-pair), tends to increase if location j has high income/productivity.

2. *The Balassa-Samuelson Effect for the Variance.* The cross-sectional variance $Var_i(q_{ijk}|\, jk)$ increases in the absolute income/productivity difference between locations j and k.

3. *The Averaging-Out Property.* Once we control for income/productivity differences, $E(q_{ijk}|\, jk)$ is close to zero. Based on our model, we interpret this to imply that once we control for the impact of income/productivity differences on non-traded input costs we are left with the contribution of traded input prices which, according to a standard trade-cost model, tend to average out across goods. That is, for most bilateral location-pairs, there tends to be as many overpriced goods (imports) as underpriced goods (exports).

4. *Distance Does Not Matter for the Mean.* Intranationally, $E(q_{ijk}|\, jk)$ does not depend on distance. Internationally it does, but not once we control

for income/productivity differences. We interpret this latter finding as evidence of a correlation between bilateral income/productivity differences and bilateral distance.

5. *Distance Matters for Cross-Sectional Dispersion.* Defining $Var_i(q^i_{jk,t} \mid jk)$ as the cross-sectional variance — the variance across goods for each date and location-pair — we find that $Var_i(q_{ijk} \mid jk)$ is increasing in the distance between locations j and k. This relationship holds irrespective of whether an international border separates the locations. Moreover, the magnitude of the relationship is not affected by the existence of a border.

6. *Distance Does Not Matter For Intranational Time-Series Dispersion.* We define $Var_t(q_{ijk,t} \mid i, jk)$ as the time-series variance; the variance across time for the relative price of good i between locations j and k. We find that, for intranational location pairs, this variance does not depend on the distance between locations j and k.

7. *Distance Matters For International Time-Series Dispersion.* If a border separates locations j and k, the variance $Var_t(q^i_{jk,t} \mid i, jk)$ depends positively on the distance between locations j and k.

We now demonstrate that each of these empirical regularities is consistent with a retail-good production technology where deviations from LOP are sustained through non-traded intermediate inputs and traded intermediate inputs which are subject to transport costs.

3. Model

This section builds a simple partial equilibrium model retail price determination at the level of individual goods sold in local markets. We assume that trade occurs in intermediate inputs (goods) and retail firms combine local inputs with traded inputs for sale in the local market (allowing for trade in the final products involves a trivial logical extension).

3.1. The Retailer's Problem

In the notation that follows, we drop the time index to conserve notation and reintroduce it when we discuss time series properties. The retail production function

is assumed to take the form:

$$Y_{ij} \equiv (N_{ij})^{\alpha_i}(T_{ij})^{1-\alpha_i}.$$

where N_{ij} is a non-traded (i.e. local) input while T_{ij} is an input which is either exported from or imported into location j.

Two examples may help to fix ideas. Suppose Y_{ij} is a men's haircut in Nashville. The traded input, T_{ij}, might be shampoo. The local inputs are the labor of the barber and the rental cost of the barber shop. If Y_{ij} was a PC sold in Gateway country, the traded input would be the PC itself and the local inputs would be sales personnel and the rental cost of the building housing the sales operation. The value of α is expected to be much closer to one for the haircut than the computer.

The cost function is the solution to the following minimization problem at each date:

$$\min_{\{N_{ij},T_{ij}\}} C_{ij} = P_j^N N_{ij} + P_{ij}^T T_{ij} \tag{3.1}$$

$$\text{s.t. } (N_{ij})^{\alpha_i}(T_{ij})^{1-\alpha_i} \geq Y_{ij} \tag{3.2}$$

where C_{ij} is the cost of producing good i in location j; P_j^N is the cost of a non-traded input, common to all goods but differing across locations; P_{ij}^T is the price of the traded input into production of retail good i in location j.

We have adopted two standard assumptions. The first is that factor mobility is much higher across sectors within a location than across locations – P_j^N is location-specific, not good-specific. The second assumption is that retailers in all locations produce good i using the same production technology – α_i is good-specific, not location-specific.

3.2. Retail Price Determination

Under constant returns to scale, the cost function takes the form: $Y_{ij} \cdot C_i(\mathbf{P}_{ij}, 1)$ where Y_{ij} is the output level, $\mathbf{P}_{ij} = \left(P_j^N, P_{ij}^T\right)$ and $C_i(\mathbf{P}_{ij}, 1)$ is the unit cost function. Under the assumption of perfect competition, the unit cost is also the retail price:

$$P_{ij} = (P_j^N)^{\alpha_i}(P_{ij}^T)^{(1-\alpha_i)} \ . \tag{3.3}$$

To compare this good and location specific price to the price of an identical good sold in another location we work with the bilateral real exchange rate at the level

of an individual good (ie. $q_{ijk} = \ln(E_{jk}P_{ij}/P_{ik})$):

$$q_{ijk} = \alpha_i q^N_{jk} + (1 - \alpha_i) q^T_{ijk}. \tag{3.4}$$

As the equation makes clear, this price deviation is a linear combination of analogous deviations in the non-traded and traded input prices. The weights in the linear combination are the shares of non-traded and traded inputs in production, which add up to unity under our assumption of constant returns to scale.

The retailer solves the same problem in each period so it is valid to add time subscripts to Equation (3.4) to reflect changes are input costs. When we use q_{ijk} it should be understood that we have averaged the good-by-good real exchange rate across time (i.e. $q_{ijk} = T^{-1} \sum_t q_{ijk,t}$) at the outset, thereby eliminating the role of time series variation which is discussed separately.

3.3. The Balassa Samuelson Effect and Trade Costs

Equation (3.4) amounts to little more than an accounting device, relating input prices to output prices. To push the model further in a structural direction we make two assumptions about the properties of the relative prices on the right-hand-side of equation (3.4).

The first assumption is that non-traded relative prices reflect productivity differences across locations. This assumption is based on the logic of the Balassa Samuelson (1964) effect. Because we lack reliable data on productivity we use $z_{jk} \equiv \log(y_j/y_k)$ as a proxy for non-traded productivity difference or q^N_{jk}.

The second assumption is that traded intermediate inputs satisfy the Law-of-One-Price up to a trade cost. Moreover, we assume that trade costs are related to distance as follows: $\log(1+\tau^i_{jk}) = \log(D^{\beta_i}_{\mathrm{jk}})$ when the source is location k and the destination is location j. Thus the log the deviation for the traded relative price is q^T_{ijk} =$\log\{(1+\tau_{ijk})^{I^i_{jk}}\} = I^i_{jk} log(1 + \tau_{ijk}) = I^i_{jk}\beta^i \log D_{\mathrm{jk}}$ where I^i_{jk} is an indicator variable for the direction of the trade flow (equal to 1 if goods travel from k to j, and -1 if good go from j to k.

Combining these two assumptions we arrive at the observable implications of a slightly more structured retail model:

$$q_{ijk} = \alpha_i z_{jk} + (1 - \alpha_i)\beta_i I^i_{jk} \log D_{\mathrm{jk}}. \tag{3.5}$$

where $0 < \alpha_i < 1$ and $\beta_i > 0$.

We see that for give given bilateral pair, the Balassa-Samuelson effect has a common sign, but the magnitude of the impact varies with the share parameter. Greater distance between locations also increases price dispersion, but the magnitude and the sign depend on specifics related to the good.

4. Properties of the Mean

We begin with an analysis of the cross-sectional mean $E_i(q_{ijk}|jk)$, where it is understood that we have averaged the good-by-good real exchange rate across time (i.e. $q_{ijk} = T^{-1}\sum_t q_{ijk}$) at the outset.

Taking a simple average of both sides of the retail pricing equation, we arrive at:

$$q_{jk} = az_{jk} + b_{jk}\log D_{\mathrm{jk}}. \tag{4.1}$$

where $a = \sum_{i=1}^{N}\alpha_i$ and $b_{jk} = N^{-1}\sum_{i=1}^{N}(1-\alpha_i)\beta_i I^i_{jk}$.

4.1. Balassa Samuelson Effects

Ignoring the role of trade costs for the moment, the average deviation reduces to: $q_{jk} = az_{jk}$. Since the parameter a is constant across bilateral pairs by virtue of the common production function, the variance in the mean across jk is simply a scaled version of the variance in income/productivity across locations. Consider a bilateral pair countries with identical (vastly different) levels of income/productivity, we would expect the average Law-of-One-Price deviation to be zero (very large). Another way to visualize this effect is to plot $q_{jk} = az_{jk}$ against income. We expect a strong positive correlation and as Figure 2 amply demonstrates we find one.

4.2. Distance and Trade Costs

Consider, next the implications of the trade cost model and distance. The multiplier on distance is $b_{jk} = N^{-1}\sum_{i=1}^{N}(1-\alpha_i)\beta_i I^i_{jk}$. The dependence on jk is due to the fact that despite the assumption that the production and trade cost parameters are independent of location, trade flows obviously are not.

However, the jk subscripts represent little more than a nuisance because we argue that the coefficient itself is expected to equal zero on a bilateral basis. We refer to this notion as the 'averaging out' property of traded good relative prices. The easiest way to see this is to suppose that the production

and trade cost parameters are common across goods and an equal number of goods are imported and exported. Then the expression is literally equal to zero: $(1-\alpha)\beta N^{-1}\left\{\sum_{i=1}^{N/2}(-1)+\sum_{i=(N/2)+1}^{N}1\right\}=0$. While some interesting asymmetries may give rise to positive or negative coefficient, we expect this property of averaging out of the deviations to prevail for quite general production and trade cost configurations.

4.3. Findings

Figure 3 plots the q_{jk} against bilateral distance. It appears that distance matters internationally, but not internationally. This visual impression is confirmed by a simple pair of regression estimates:

$$|q_{jk}| = \underset{(0.0276)}{0.244} + \underset{(0.0031)}{0.0044} \log D_{\text{jk}} + \varepsilon_{jk} \text{ international} \tag{4.2}$$

$$|q_{jk}| = \underset{(0.0433)}{0.130} - \underset{(0.0059)}{0.0044} \log D_{\text{jk}} + \varepsilon_{jk} \text{ intranational .} \tag{4.3}$$

We use the absolute value of q_{jk} so that the coefficient on distance has a consistent sign across bilateral pairings.

Beginning with the international data we find evidence that distance matters as it should based on the trade cost model. However, the magnitude of the coefficient is small: going from cities that are neighbors to 100 miles apart adds 1% to the price differential, going and additional 2400 miles is needed to add another 0.5%! Controlling for distance, though, the intercept is both highly economically and statistically significant. Thus even after conditioning on distance (a proxy for trade costs), we resoundingly reject the Law-of-One-Price.

The intranational data tells a different story. The distance coefficient is of the wrong sign and statistically insignificant. Using the same interpretation as we did for the international context we reject the Law-of-Price, but unlike the international data distance plays no role at all.

Combining the two results there appears to be a border involving an increase in the unconditional variance in price levels across locations and an increase (from zero) in the impact of geographic distance on price differences.

According to our retail model, though, income/productivity disparities play an independent role. Re-estimating with the absolute value of the log of relative income across bilateral pairs, we have:

$$|q_{jk}| = \underset{(.0272)}{0.2537} - \underset{(.0031)}{0.0022} \log D_{jk} + \underset{(.0031)}{0.0447} |z_{jk}| + \varepsilon_{jk} \text{ international} \tag{4.4}$$

$$|q_{jk}| \quad = \quad \underset{(0.0433)}{0.130} - \underset{(0.0059)}{0.0044} \log D_{jk} + \varepsilon_{jk} \text{ intranational} \tag{4.5}$$

where we have simply repeated our results for the intranational case because we lack data on income levels across regions within countries (assuming that income differences are small across regions within countries, this should not be too problematic, but we will rectify this for the U.S. where we know such data is available).

The absolute deviations of the log of relative income has a positive coefficient as expected. The distance coefficient is no longer not statistically significant as we would expect if the income ratio was effective in picking up the impact of non-traded goods, based on the 'averaging out' property for traded good prices. It is also interesting to note that the reduction in the magnitude of the coefficient on distance coefficient (it actually becomes negative) is consistent with a positive correlation between distance and income differentials.[2]

5. Properties of Cross-Sectional Variance

Next we examine the cross-sectional variance: $Var_i(q_{ijk}|jk)$ which, in the model, is given by:

$$Var_i(q_{ijk}|jk) = a'(z_{jk})^2 + b'(\log D_{\text{jk}})^2. \tag{5.1}$$

$a' = Var(\alpha_i)$ and $b' = var_i\left\{(1-\alpha_i)I^i_{jk}\beta_i\right\}$.

5.1. Balassa Samuelson Effects

Taking the Balassa Samuelson effect in isolation of the trade cost give us:

$$Var_i(q_{ijk}|jk) = a'(z_{jk})^2 \ . \tag{5.2}$$

Thus the variance of Law-of-One-Prices around the mean is increasing in the income/productivity gap. Thus if we select a location pair with very similar incomes there will be very little variation in the size of the deviations across goods. If we take location pairs with vastly different incomes we will find that the Law-of-One_Price deviations are more heterogenous across goods. In other

[2] This classic results is derived as follows. Let the true regression be: $y_i = bx_i + cz_i + \varepsilon_i$. If we estimate: $y_i = bx_i + cz_i + \varepsilon_i$ then the difference between the OLS estimate $\widehat{b}$ and the true one is: $\widehat{b} - b = (x'x)^{-1}x'zc$. In our context $c > 0$ and if $x'z > 0$ the OLS estimate is upwardly biased.

words, the q_{ijk} are expected to be heteroscedastically distributed with a known form of the heteroscedasticity, namely equation (5.2). Plotting the cross-sectional standard deviation of relative prices for each bilateral pair against the absolute value of the logarithm of relative income gives us the positive association we expect to find. Locations with similar income levels have low price dispersion, as income disparities rise, the deviations become more dispersed across goods for that bilateral pair (see Figure 4).

5.2. Distance and Trade Costs

Obviously we will not have an averaging out property in the second moment of the distribution since the summations in the variance expression involve squared terms. The second term in the $Var_i(q_{ijk}|jk)$ expression simplifies somewhat once we impose the condition that the mean is zero (which is plausible given our earlier analysis of the mean):

$$b' = \sum_{i=1}^{N} (1 - \alpha_i)^2 (\beta^i)^2$$

where we have exploited the fact that $(I^i_{jk})^2 = 1$. Unfortunately, the coefficient does not reduce conveniently to a simple function of the average value of either the production parameter or the elasticity of trade cost with respect to distance. If either of the parameters α_i or β_i were constant across goods, we would identify the average value of the other parameter up to a scalar (the scalar being the value of the other parameter). If we had traded good prices on the left-hand-side we would be able to estimate the average β_i. In any case, it must be true that $b' > 0$. Thus distance is expected to matter for the second moment properties of the cross-section. Figure 5 plots $Var_i(q_{ijk}|jk)$ against bilateral distance. We see the positive relationship implied by the above algebra.

5.3. Findings

Turning to our regression results, we see that both income/productivity and distance matter as the theory predicts. We find that the distance coefficient falls in the international case when we add income, but remains highly statistically significant. Recall that it became insignificantly different from zero in the analysis of means, which was predicted by the averaging out property. The second moments should be affected by distance and they are. Moreover, the impact of distance is quite similar across the two panels as long as we control for income disparities in

the international specification. The coefficient on the distance variable is 0.0154 internationally versus 0.011 intranationally. Given the standard error on the intranational distance coefficient we probably cannot reject the hypothesis that the two coefficients are equal (had we not controlled for income we might have been able to claim distance matters more internationally).

$$Var_i(q_{ijk}|jk) = \underset{(0.0070)}{0.5352} + \underset{(0.0001)}{0.0014} (\log D_{\text{jk}})^2 + \varepsilon_{jk} \text{ international} \tag{5.3}$$

$$Var_i(q_{ijk}|jk) = \underset{(0.0062)}{0.5314} + \underset{(0.0001)}{0.0010} (\log D_{\text{jk}})^2 + \underset{(0.0005)}{0.0203} (z_{jk})^2 + \varepsilon_{jk} \text{ international} \tag{5.4}$$

$$Var_i(q_{ijk}|jk) = \underset{(.0187)}{0.3500} + \underset{(.0003)}{0.0008} (\log D_{\text{jk}})^2 + \varepsilon_{jk} \text{ intranational .} \tag{5.5}$$

The economic significance of distance and income are not great relative to the unconditional mean of dispersion in the data. Moving from locations that are 100 miles apart to 1000 miles apart adds about 2.7% to the variance of prices (using the international coefficient). Moving from locations with equal income to a situation where income is twice as great in one location than the other increase the variance from 0.5314 to 0.5456 (1.4%).

6. Properties of Time-Series Variance

Here we employ implications of the stochastic trade cost model of Serci, Uppal and Van Hulle (1995) derived by Lee (2003).

TBA

7. Conclusions

TBA

References

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

[2] Crucini, Mario J. Chris I. Telmer and Marios Zachariadis, 2001, "Understanding european real exchange rates," Vanderbilt University Working Paper No. 01-W20.

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

[4] Engel, Charles and John H. Rogers, 2003, "European product market integration in the wake of EMU," mimeo.

[5] Lee, Inkoo, 2003, "Implications of censoring for absolute and relative deviations from the law-of-one-price," Chapter 2, Vanderbilt University Ph.D. dissertation proposal.

[6] Rogers, John, 2001, "Price level convergence, relative prices, and inflation in Europe," Board of Governors of the Federal Reserve System International Finance Discussion Papers No. 699.

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

[8] Sercu, Piet., Raman Uppal and C. Van Hulle, 1995, "The exchange rate in the presence of transactions costs: implications for tests of purchasing power parity," *Journal of Finance* 50, 1309-1319.