BANK OF CANADA
BANQUE DU CANADA

Staff Working Paper/Document de travail du personnel 2017-12

Accounting for Real Exchange Rates Using Micro-Data

by Mario J. Crucini and Anthony Landry

Bank of Canada Staff Working Paper 2017-12

April 2017

Accounting for Real Exchange Rates Using Micro-Data

by

Mario J. Crucini[1] and Anthony Landry[2]

[1]Vanderbilt University
and
NBER
mario.j.crucini@vanderbilt.edu

[2]Canadian Economic Analysis Department
Bank of Canada
Ottawa, Ontario, Canada K1A 0G9
anthonylandry@bankofcanada.ca

ISSN 1701-9397

Acknowledgements

This paper was written while Mario Crucini was a visiting senior fellow at the Globalization and Monetary Policy Institute at the Federal Research Bank of Dallas. Mario Crucini gratefully acknowledges the financial support of the National Science Foundation. We are grateful to comments from participants at various research conferences and universities. We particularly thank Dean Corbae, Mick Devereux, Rebecca Hellerstein, Karen Lewis, John Rogers and Robert Vigfusson for their detailed comments. The views in this paper are our responsibility and should not be interpreted as reflecting the views of the Bank of Canada, the Federal Reserve Bank of Dallas or the Federal Reserve System.

Abstract

The classical dichotomy predicts that all of the time-series variance in the aggregate real exchange rate is accounted for by non-traded goods in the consumer price index (CPI) basket because traded goods obey the Law of One Price. In stark contrast, Engel (1999) claimed the opposite: that traded goods accounted for all of the variance. Using micro-data and recognizing that final good prices include both the cost of the goods themselves and local, non-traded inputs into retail such as labor and retail space, our work re-establishes the conceptual value of the classical dichotomy. We also carefully show the role of aggregation, consumption expenditure weighting and assignment of covariance terms in the differences between our findings and those of Engel.

Bank topics: Exchange rates; International financial markets; Trade integration
JEL code: F3

Résumé

Selon la dichotomie classique, les biens non échangeables du panier de l'indice des prix à la consommation (IPC) expliquent intégralement la variance au sein des séries chronologiques du taux de change réel agrégé, car les biens échangeables suivent la loi du prix unique. Engel (1999) soutient à l'inverse que ce sont les biens échangeables qui expliquent toute la variance. Nous rétablissons la valeur conceptuelle de la dichotomie classique en nous appuyant sur des microdonnées et en reconnaissant que les prix des biens finaux renferment le coût même des biens et les coûts d'intrants locaux non échangeables, comme ceux du travail et des espaces commerciaux, qui sont liés au commerce de détail. Nous démontrons en outre que l'agrégation, la pondération des dépenses de consommation et la répartition des termes de covariance jouent un rôle dans les différences observées entre nos résultats et ceux d'Engel.

Sujets : Taux de change ; Marchés financiers internationaux ; Intégration des échanges
Codes JEL : F3

Non-Technical Summary

The classical dichotomy predicts that all of the time-series variance in the aggregate real exchange rate is accounted for by non-traded goods in the consumer price index (CPI) basket because traded goods obey the Law of One Price. In stark contrast, Engel (1999) claimed the opposite: that traded goods accounted for all of the variance. Using retail price data at the level of individual goods and services across many countries of the world, we show that the classical dichotomy is a very useful theory of international price determination when applied to intermediate inputs. Specifically, we parse the role of non-traded (e.g., labor and retail space) and traded inputs at the retail level, removing a significant source of compositional bias from the micro-data and highlighting the role of the two inputs in explaining real exchange rate dynamics. We also carefully show the role of aggregation, consumption expenditure weighting and assignment of covariance terms in the differences between our findings and those of Engel. Our results point to the usefulness of microeconomic theories that distinguish traded and local inputs and their composition in final goods.

1 Introduction

One of the most robust facts in international finance is that bilateral real exchange rates are highly variable over time, and typically more so the more variable the bilateral nominal exchange rate. Taking the benchmark of constant real exchange rates, or relative purchasing power parity (PPP), these facts indicate that goods markets are nationally segmented. The practical challenge in addressing this empirical fact with existing quantitative models is knowing where to place the frictions. In contrast to most trade models, consumers obtain the lion's share of their consumption goods through intermediaries—brick-and-mortar retail stores—thus leading to an important role for a "distribution" wedge. We argue that to fully appreciate the role of this form of market segmentation, highly disaggregate data are needed.

The reason is simple: at the level of aggregation of most consumer price index (CPI) sub-indices the underlying retail cost structure of a category of final goods has a substantive mix of the cost of the traded inputs and the local inputs, obscuring their separate role in international relative price dynamics. The food sub-index (a traded good in Engel's classification) is a prime example. Food consists mainly of two sub-categories: groceries and restaurant meals. The U.S. National Income and Product Accounts (NIPA) indicates that for grocery retailers, 36% of the cost is non-traded inputs while for the restaurant industry the share is 75%. Based on inputs, the classical dichotomy would treat groceries as traded goods and restaurant meals as non-traded services. Our approach indexes goods by their sector distribution share dramatically mitigating the impact of these arbitrary classifications.

We have two main findings. First, when we dichotomize the micro-data into traded and non-traded goods and create a two-sub-index version of the aggregate real exchange rate to emulate Engel (1999), the non-traded goods real exchange rate accounts for 66% of real exchange rate variation (for the average OECD country pair)—almost twice that of the traded good real exchange rate. This goes in the direction of the classical dichotomy and contrasts sharply with Engel's assertion that almost all of the variation in the aggregate real exchange is due to the traded real exchange rate. We show that his result is a consequence of his arbitrary rearrangement of terms in the variance decomposition, which has the impact of attributing the covariance between the traded and non-traded real exchange to the traded component while also failing to weight the variance of the traded real exchange rate by its expenditure share.

Second, the classical dichotomy looks much more compelling when we model each final good as a composite of a non-traded input and a traded input using distribution shares from the NIPA and back out the contribution of non-traded and traded inputs to the variance of the aggregate real exchange rate. Non-traded inputs account for 85% of the total variance (for the average OECD country pair). The fractions accounted for vary only modestly across different collections of bilateral city pairs such as the OECD, non-OECD or North America. Simply put, the segmentation of markets in macroeconomic models should be weighted toward non-traded inputs along the lines of Corsetti, Dedola and Leduc (2008), but not with the symmetry across goods they assume. The fact that considerable variation remains even for the traded inputs points to the important quantitative role of traditional trade frictions (official and natural barriers to trade) and markups of imports at the dock that is emphasized in the existing trade and industrial organization literature.

Finally, as the analysis becomes more granular as is necessary to consider cross-sectional differences sector-by-sector or good-by-good, the difference in the roles of traded and non-traded inputs in accounting for wedges in law of one price (LOP) deviations becomes increasingly stark and intuitive. An item with a relatively low cost of distribution (17%), such as a compact car, has only 28% of its LOP variance attributed to non-traded inputs, whereas an item with a very high cost of distribution (85%), a man's haircut with tip included, has 92% of its LOP variance attributed to non-traded inputs. This is not surprising, as a haircut is the archetype non-traded good. This example also illustrates that even services, which are not themselves traded, use a nominal amount of traded inputs, which mitigates the level of segmentation in this large and growing sector.

The rest of the paper proceeds as follow. In Section 2, we present the data. In Section 3, we describe our methodology and compute individual contributions of LOP deviations to aggregate real exchange rate (RER) volatility. In Section 4, we document a striking positive relationship between the magnitude of the contribution of an LOP deviation to aggregate RER volatility and the cost- share of inputs used to produce that good. Then, we develop and estimate a two-factor model, and aggregate these factors to measure the contribution of intermediate traded and local inputs to aggregate RER volatility. In Section 5, we show that our microeconomic decompositions, when aggregated, look very similar to earlier studies using aggregate CPI data, but the economic implications are starkly different. Section 6 concludes.

2 The Data

There are four sources of data used in this study; each is discussed in detail below. The first and primary data source is the Economist Intelligence Unit (EIU) Worldwide Survey of Retail Prices. These are local currency prices collected at the city level (the most commonly known line-item is the Big Mac, as it features regularly in The Economist magazine). The remaining three supplementary data sources are specific to the U.S. and include detailed consumption-expenditure weights from the Bureau of Labor Statistics (BLS), and distribution margins and non-traded input shares computed from Bureau of Economic Analysis sources described below.

The EIU conducts the most comprehensive international price survey by a single agency in a consistent fashion over time and across countries. Begining in 1990, the EIU has collected prices of 301 goods and services across 123 (mostly capital) cities of the world. The panel used in this study is annual, from 1990 to 2015.[1] The 301 line-items include a significant number of cases in which the item is priced in two different types of retail outlets. For example, all food items are priced in both supermarkets and mid-priced stores. Clothing items are priced in chain stores and in a mid-price/branded store. Since the goods and services are not specified down to the level of brands, these different outlet observations of a particular "good" are best thought of as varieties, similar to how goods are differentiated by country of origin in the trade literature (the Armington assumption). The prices are collected from

[1] Spot exchange rates are applied to the city data surveyed by the EIU, and are available along with the price data for each year. The exchange rate reported is the spot rate for the survey date when the data was gathered (usually in September).

the same physical outlet over time, thus the prices are not averages across outlets. These data have been used in a large number of peer-reviewed journal articles, though not for the purpose utilized here, and as far as we know, have not been updated to include the Great Recession (i.e., most studies have used the 1990-2005 window).[2]

The first supplementary data series is consumption-expenditure weights. These data are more aggregated than the EIU prices, leading us to allocate the 301 individual EIU line-items to 73 unique U.S. expenditure categories. We divide the sector level expenditure weights by the number of prices surveyed in each sector, so that each category of goods in the EIU panel has the same expenditure weight as in the U.S. CPI index. Because some sectors are not represented in the EIU retail price surveys, the expenditure shares are inflated by a common factor such that they add up to one.

The second supplementary source is from the Bureau of Economic Analysis Personal Consumption Expenditure Bridge Tables (1992). These tables show the value of consumer expenditures by expenditure category in producers' and purchasers' prices. The macroeconomic literature refers to this as the distribution wedge: the difference between what consumers pay and what producers receive divided by what consumers pay. For example, if final consumption expenditure on bread is \$1.00 and bread producers receive \$0.60, the distribution wedge is 0.40. This wedge includes wholesale and retail services, marketing and advertisement, local transportation services and markups.

For services, however, this is problematic as a meaure of the traded inputs in final consumption. The reason is simple: according to these tables, what consumers pay and what producers receive is the same value. Conceptually, this is inconsistent with the approach used for goods. For example, when a consumer (or that consumer's health insurance provider) receives a medical bill, the charges may include wage compensation for the physician and the cost of goods and non-physician services included in the overall treatment, whether or not it is itemized on the invoice. Since the doctors' services are local inputs while the goods used in production of medical services are traded

[2] Andrade and Zachariadis (2016), Bergin and Glick (2007), Burstein, Eichenbaum and Rebelo (2005), Crucini and Shintani (2008), Crucini, Shintani and Tsuruga (2010), Crucini and Telmer (2012), Rogers (2007).

inputs, we need to separate the two. In these circumstances, we use the 1990 U.S. input-output data to measure non-traded and traded inputs. Each retail item in the EIU panel is reconciled with an input-output sector and assigned a distribution share equal to the share of services that sector purchases.

The median item in the cross-section (including goods and services) has a distribution wedge of 0.41. In other words, the local factor content of a typical item in the consumption basket is about 40%. Weighting these wedges by final expenditure shares provides an assessment of their importance in aggregate consumption expenditure. The expenditure-weighted value of the distribution share varies slightly across sub-groups of locations in our analysis due to modest differences in the availability of micro-data across them: it is 0.48 for U.S. cities, 0.49 for OECD cities and 0.53 for non-OECD cities. Notice that all of these numbers consistently exceed the median across goods for the simple reason that consumption expenditure is skewed toward items with relatively high distribution shares.

Overall, our distribution wedges are similar to those used in Burstein et al. (2003) and Campa and Goldberg (2010). Instead of estimating the size of the distribution sector using aggregate data, Berger et al. (2012) measured the distribution wedges using U.S. retail and import prices of specific items from the U.S. CPI and producer price index (PPI) data. They find that the distribution wedges in these data are distinctively larger than the estimates reported for U.S. consumption goods using aggregate data. Their median U.S. distribution wedge across all items in their cross-section is 0.57 for imports priced on a c.i.f. (cost, insurance, and freight) basis and 0.68 for imports priced on an f.o.b. (freight-on-board) basis. While their dataset allows for a more disaggregated calculation of the distribution wedges, it does not include services, which constitutes a large fraction of consumption expenditure. Importantly for our results, Burstein et al. (2003) and Berger et al. (2012) found that the distribution wedges are stable over time. Therefore, RER variations appear not to be coming from changes in share of cost attributed to distribution wedges, but rather from changes in traded and non-traded input prices that are weighted by these shares in variance decompositions we conduct below.

3 Methodology

Our novel methodology for the variance decomposition of the aggregate real exchange rate involves four essential steps: i) construction of price indices using a common consumption basket (the same items and a common set of expenditure weights); ii) decomposition of the variance of the aggregate real exchange rate into the contribution of each LOP deviation; iii) estimation of a two-factor model to decompose the variation in LOP into the contribution of a non-traded (local) cost component and a traded cost component; iv) macroeconomic aggregation using this two-factor model to assess the role of traded and non-traded inputs (plus a small residual term) in aggregate real exchange rate variation.

3.1 Real Exchange Rates With Common Consumption Baskets

As the cognoscenti of CPI indices know, these indices are explicitly constructed with one goal in mind: to measure inflation of the basket of goods and services representative of consumption patterns of residents of their respective countries.[3] As a practical matter, the contents of the consumption basket vary dramatically across countries. This is formally known in trade theory as home bias: trade costs and taste differences skew the CPI basket away from foreign-produced goods and services, toward domestic ones. The implication of this fact is that real exchange rates constructed from CPI data violate the premise of the basic building block of PPP, the LOP. The LOP is a proposition about the equality of common-currency prices of identical goods.[4] While the EIU goods and services are not identical to the extent of bar-code level

[3] It is also true that in attempting to achieve this goal, national statistical agencies use very different methods to survey retail prices and aggregate them into consumer price indices. The differences include: the geographic scope of the survey; the number of outlets surveyed; the frequency with which consumption expenditure shares are updated; and the treatment of product entry and exit.

[4] Strictly speaking, the LOP is the proposition that, after conversion to a common currency, the price of an identical good or service should differ across countries only to the extent of natural (traditional trade costs) and official (tariffs, quotas, voluntary export restraints) barriers to trade.

data, they have the uncommon virtue of being collected by a single agency in a consistent fashion across time and locations. This allows us to construct a common basket as prescribed by the theory in the sense of including the same line-item list in each country and applying a common set of expenditure weights (common across countries) in the process of aggregation to an aggregate price index.

With these price indices in hand, any change in a bilateral real exchange rate must, by construction, be an expenditure-weighted average of the changes in the LOP deviations of the goods and services included in the price survey. To see this, consider the price index of country j defined as a geometric average of the local currency prices of individual goods and services (P_{ijt}) in city j:

$$P_{jt} \equiv \prod_{i}^{M} (P_{ijt})^{\omega_i} \quad . \tag{1}$$

Here ω_i is the consumption expenditure weight applied to good i.

Our results are not sensitive to the precise expenditure weights used: the crucial part of the construction is that the weights are good-specific and not country-specific. When expenditure shares differ across goods, the change in the aggregate real exchange rate will reflect both changes in relative price of goods and changes in the prices of the same goods, and make the sources of the deviations impossible to isolate.[5]

The bilateral real exchange rate used in this study is the relative cost of this common basket across countries:

$$Q_{jkt} = \frac{S_{jkt} P_{jt}}{P_{kt}}, \tag{2}$$

where S_{jkt} is the spot nominal exchange rate between city j and k, and P_{jt} and P_{kt}, are price levels constructed as shown above. If Q_{jkt} exceeds unity, the price level is higher in city j than city k in year t.

3.2 Microeconomic Variance Decomposition

Because we are able to construct our own price indices and correspondent aggregate real exchange rates, the logarithm of the aggregate RER is exactly

[5] We use U.S. sector-level consumption expenditure weights, assigning each line-item price to a sector and dividing the sector weight equally among the items in the sector.

equal to the expenditure-weighted average of the logarithm of the LOP deviations:

$$q_{jkt} = \sum_{i}^{M} \omega_i q_{ijkt} \ . \tag{3}$$

To see this, substitute the price index (1) into (2) and take logarithms of both sides.

Notice that this is a very high-dimensional object because the sum is taken over about 202 LOP deviations. A standard variance decomposition of the aggregate real exchange rate using these micro-data would entail estimating $M(M-1)/2$ covariances (typically 20,301 covariances for each bilateral city pair in our application). Fortunately, macroeconomic models typically have little to say about, for example, the covariance between the LOP deviations of apples and haircuts.

Instead, our microeconomic variance decomposition exploits the relationship between variance and covariance, reducing the number of moments to compute by a factor of about 100, (i.e., from 20,301 to 202). Taking the covariance of the variables on each side of this expression with respect to the aggregate real exchange rate, q_{jkt}, one arrives at:

$$var(q_{jkt}) = cov(q_{jkt}, \sum_i \omega_i q_{ijkt}) = \sum_i \omega_i cov(q_{ijkt}, q_{jkt}) \ . \tag{4}$$

Dividing all terms on both sides of the equation by the variance of q_{jkt} results in the variance decomposition:

$$1 = \sum_i \omega_i \frac{cov(q_{ijkt}, q_{jkt})}{var(q_{jkt})} = \sum_i \omega_i \beta_{ijk} \ , \tag{5}$$

which has the convenient property of being the same dimension as the vector of goods and services used to construct the aggregate real exchange rate, M.

Simply put: the contribution of the deviation from the LOP of good i to the variance of the aggregate RER is the product of its expenditure share, ω_i, and its good- and location-pair-specific beta, β_{ijk}. This decomposition is analogous to the use of betas to describe the contribution of the return on an individual stock to the variance of the return on the stock portfolio. As in finance applications, what matters is covariance risk. Covariance risk is also a natural metric for macroeconomics as the tendency of prices to exhibit local currency price stickiness will show up as a general tendency for a positive covariance, $\beta_{ijk} > 0$, reflective of exchange rate risk.

The utility of this decomposition is particularly intuitive for stark benchmarks taken from the existing theoretical literature. Suppose all prices are equally "sticky"; that is, fixed in local currency units during the period in which a nominal exchange rate movement occurs and then adjusted to exactly satisfy the LOP at the end of the period. If the nominal exchange rate is the only underlying shock, every single good would contribute exactly the same amount to the variance of the aggregate RER, $\beta_{ijk} = 1$. This characterizes the view that all goods markets are equally segmented, "goods are all alike"; it produces a degenerate distribution of the βs with all the mass at 1. Put differently, a nominal exchange rate change essentially shifts the mean of the price distribution in the home country relative to the foreign country without altering relative prices within either country.

Suppose instead that all traded goods adjusted instantaneously to the nominal exchange rate movement while non-traded goods take one period to adjust. Now non-traded goods account for all of the variance and traded goods for none of the variance. This example characterizes the classical dichotomy of international finance: it produces a degenerate distribution of βs with two mass points. One mass point in the distribution of betas is at the value $\beta^T = 0$, which has mass (probability) equal to the expenditure share for traded goods in the microeconomic sample (0.4). The second mass point in the distribution has probability 0.6 (the fraction of expenditure on non-traded goods) and given the identity solved in equation (5), this mass must occur at the point $\beta^N = (0.6)^{-1} = 1.7$.

Table 1 reports summary statistics for the microeconomic variance decomposition defined by equation (5). The mean beta for non-traded goods does exceed the mean for traded goods in most cases, ranging from a difference of 0.35 (1.09–0.74) for U.S.-Canada city pairs to a low of 0.21 (1.00–0.79) for non-OECD city pairs. The relative standard deviation of the LOP deviations average twice that of the aggregate RER, indicative of considerable idiosyncratic variation in LOP deviations. The mean correlation of LOP deviation and PPP deviation is 0.42 in the pooled sample. The fact that the distribution of betas is not degenerate at unity is consistent with the observation by

Crucini and Telmer (2012)—LOP deviations are not driven by a common factor such as the nominal exchange rate; much of the variation is idiosyncratic to the good.

This begs the question: What does the empirical distribution of β_{ijk} look like in relation to our stark benchmarks? **Figure 1** provides the answer: it presents three kernel density estimates: one density for all goods (grey line), one density for traded goods (red line) and one density for non-traded goods (blue line). The vertical lines denote the mean beta across goods and services within each of the three groupings. These distributions have absolutely no resemblance to either of the two stark views described above. There is far too much heterogenity in the βs to be content with the broad-brushed view that goods markets are equally segmented internationally. Afterall, the support of the distribution extends from -2 to +4! Turning to the classical dichotomy, the density of the βs for the traded goods should be degenerate at 0 (its mean is actually 0.75) and the density of the betas for the non-traded goods should be degenerate at 1.7 (its mean is 1.01). The means of the distributions go in the direction of the classical dichotomy, but the distributions overlap to such a degree as to obscure almost any distinction between them.

In summary, the contribution of individual goods to aggregate RER variability shows a central tendency, but with considerable variation across individual goods. The stark views of local currency sticky prices or the classical dichotomy theories of international price adjustment commonly imposed in the quantitative theoretical DGSE literature are non-starters in describing the underlying micro-data. And yet, things are not so grim. Our variance decomposition approach establishes the empirical validity of the classical dichotomy at the level of inputs into production of final goods, while also showing an important role for incomplete pass-through of nominal exchange rate changes to local currency prices of traded goods (inputs into retail sales), consistent with the existing literature that has focused on prices at the dock.

3.3 Trade in Middle Products

The phrase "trade in middle products" is the title of an often overlooked but very insightful contribution to trade theory by Sanyal and Jones (1982). They model each final good and service as a composite of an internationally traded intermediate input and local factor inputs (largely labor and capital services devoted to retail and distribution). This theory differs in an essential way from the large existing literature emphasizing trade in intermediate inputs as these are theories about the international factor content of goods that are traded across countries. In the Sanyal and Jones model, consumer goods are not traded—which inevitably makes their prices more sensitive to local factor market conditions. Essentially the retail market is a clear point of demarcation that segments final markets while free trade occurs in intermediate traded input (up to official and natural barriers, of course). Essential for our application is the fact that the local factor content varies dramatically across goods.

While Sanyal and Jones did not specify a production function for final goods, here, we follow a large and growing literature (see, for example: Engel and Rogers (1996), Crucini, Telmer and Zachariadis (2005)) and assume that each retail good is a Cobb-Douglas aggregate of local labor and traded inputs, with cost shares α_i and $1 - \alpha_i$, respectively. Two appealing features of the Cobb-Douglas aggegator are the constancy of cost shares across time and the fact that it provides a rationale for the use of geometic averages of prices rather than raw averages employed by many national statistical agencies.

Solving the retail firm's cost minimization problem under perfect competition leads to the following unit price for good i in city j (up to a constant factor of proportionality, ignored here as it is irrelevant in what follows):

$$P_{ijt} = W_{jt}^{\alpha_i} T_{ijt}^{1-\alpha_i}. \tag{6}$$

Taken literally, W_{jt} is the unit input cost of local factors of production, which would include labor and retail space (our notation emphasizes labor inputs), and T_{ijt} the cost of the traded input inclusive of a transportation cost and a markup from the source to the destination. Our empirical methodology requires markups (μ) over marginal need to take one of two forms: i) μ_{it}, a time-varying but common proportional markup specific to the good but common across locations, or ii) μ_{ij}, a good- and destination-specific markup that does not vary over time.

The logarithm of the LOP deviation across bilateral city pair j and k for good i is thus

$$q_{ijkt} = \alpha_i w_{jkt} + (1 - \alpha_i)\tau_{ijkt} \ , \tag{7}$$

where each of the variables is now the logarithm of either the relative local factor input prices, w_{jkt} (variation in international wages of unskilled workers, once wages are converted to common currency and efficiency units), and the LOP deviation of the traded good itself, τ_{ijkt}. Simply put: the LOP deviation for good i, across bilateral city pair, j and k, depends on the deviation of distribution costs and traded input costs across that pair of cities, weighted by their respective cost shares.

This leads to a crisp mapping between the beta for an individual retail good or service and the properties of time series variation of local input costs and traded input costs. To see this, recall the definition of β_{ijk} is:

$$\beta_{ijk} = \frac{cov(q_{ijkt}, q_{jkt})}{var(q_{jkt})}. \tag{8}$$

Substituting the identity $q_{ijkt} = \alpha_i w_{jkt} + (1 - \alpha_i)\tau_{ijkt}$ on the right-hand-side gives:

$$\beta_{ijk} = \frac{cov(q_{ijkt}, q_{jkt})}{var(q_{jkt})} = \alpha_i \frac{cov(q_{jkt}, w_{jkt})}{var(q_{jkt})} + (1 - \alpha_i)\frac{cov(q_{jkt}, \tau_{ijkt})}{var(q_{jkt})}, \tag{9}$$

which can be written more compactly by recognizing the covariance terms on the right-hand-side are also betas,

$$\beta_{ijk} = \alpha_i \beta^w_{jk} + (1 - \alpha_i)\beta^\tau_{ijk} \ . \tag{10}$$

Effectively, this equation is what facilitates the partition of LOP variance into deviations arising from retail and distribution costs (the local inputs) from LOP deviations at the border (the traded inputs).

4 Results

The good-specific betas for the retail prices are directly estimable from the covariance of the LOP deviations and the aggregate real exchange rate (recall these estimates were displayed in **Figure 1**): the factor input betas (β^w_{jk} and β^τ_{ijk}), in contrast, are not. The approach we take is to estimate these as unobserved factors using a simple linear regression model.

The first step in this procedure is acccomplished by expressing the traded input factor as the sum of a component common to all goods and an idiosyncratic component specific to the good: $\beta_{ijk}^{\tau} = \beta_{jk}^{\tau} + \nu_{ijk}$. The contribution of good i to the variation of the bilateral RER across city pair j and k becomes:

$$\beta_{ijk} = \alpha_i \beta_{jk}^{w} + (1 - \alpha_i)\beta_{jk}^{\tau} + \epsilon_{ijk} \ , \tag{11}$$

where $\epsilon_{ijk} = (1 - \alpha_i)\nu_{ijk}$. In the language factor models, α_i and $(1 - \alpha_i)$ are "factor-loadings," on the local and traded factors, β_{jk}^{w} and β_{jk}^{τ}, respectively. In most applications of factor models, the factor loadings are inherently difficult to interpret, but not here. Notice the observables are: i) the estimated betas, β_{ijk}, and ii) the distribution wedges from the U.S. NIPA data, α_i. The unobservables are the two factors of interest, β_{jk}^{w} and β_{jk}^{τ}.

These betas are estimated using a single-variable linear regression:

$$\beta_{ijk} = \beta_{jk}^{\tau} + (\beta_{jk}^{w} - \beta_{jk}^{\tau})\alpha_i + \epsilon_{ijk} \ ,$$

which, written in the more familiar intercept and slope form, is:

$$\beta_{ijk} = a_{jk} + b_{jk}\alpha_i + \epsilon_{ijk} \ . \tag{12}$$

Comparing this equation to the theoretical model, it is apparent that the constant term and the slope parameter identify the two factors of interest:

$$\beta_{jk}^{\tau} = a_{jk} \tag{13}$$

$$\beta_{jk}^{w} = a_{jk} + b_{jk} \ . \tag{14}$$

Figure 2 presents a scatter-plot of the contribution of good i to the variance of the bilateral RER averaged across international city pairs (β_i) against the distribution share for that good (the non-traded input cost, α_i). Since the distribution shares are sector-level while the price data are good-level, the number of y-coordinates is the number of goods assigned to that sector and thus any variation in betas across goods within a particular sector must be attributed to differences in the role of the traded input. Notice the strong positive relationship between a final good's contribution to RER variability and its distribution wedge, the correlation of β_i and α_i, is 0.67. This implies that a larger share of the time series variation in LOP deviation is coming from the local inputs than from traded inputs and this is increasing as we move from the left-side of the figure to the right-side. Conceptually, this is consistent with the view of the classical dichotomy applied at the level of inputs to retail production.

Items near the extremes are the anecdotes we tend to use in classroom discussions of traded and non-traded goods: 1 liter of gasoline (red dot) and a two-bedroom apartment (blue dot), for example. Based on our assignment of the EIU micro-data with the U.S. NIPA data, the cost share of local inputs for gasoline is 0.19 while for a two-bedroom apartment it is 0.93. The betas of these items are 0.59 and 1.12, meaning that a two-bedroom apartment contributes almost twice as much to variation in the aggregate real exchange rate than does retail gasoline. If final goods in the consumption basket could be dichotomized into such stark categories, the classical dichotomy would be much more successful in accounting for real exchange rates at the level of final goods. But, as the figure clearly shows, most consumption goods in the basket have a less stark cost structure. Take two important consumption categories as prime examples: food items have a non-traded input cost share in the neighborhood of 0.35, while the non-traded input share for clothing is about 0.5. The median good in our sample has a non-traded input cost share of 0.41. It is precisely because the cost share of the median good is close to 0.5 that the classic dichotomy applied to final consumption goods performs poorly, muddling the two underlying sources of variation. The cost structure of the typical retail good effectively averages away the differences in the underlying input cost variation. What researchers need to do is apply the theory of middle products where the classical dichotomy is a better description of the underlying stochastic properties—to input prices.

Table 2 reports the estimated factors averaged across city pairs within different country groups.[6] The standard deviations across city pairs are re-

[6] Note that since the distribution shares are more aggregated than the betas, we take simple averages of the betas across i for goods that fall into each sector for which we have distribution wedges. Following this aggregation, equation (12) is estimated by ordinary least squares (OLS) to recover the non-traded and traded factors. We also report results obtained by weighted least squares (WLS) where each observation is weighted by the inverse of the number of goods falling into each distribution-share sector (not shown), they are almost identical to the OLS estimates.

ported in brackets. The differences across groups of locations and individual city pairs is discussed in a subsequent section. The first column pools all city pairs. The traded-factor averages 0.54 while the non-traded factor averages 1.05. This implies that, on average, non-traded inputs contribute twice as much as traded inputs to RER variations.

Recall that the average traded and non-traded goods have betas of 0.75 and 1.01. Notice that the traded input factor is much lower than the average contribution of a traded good to aggregate RER variability while the non-traded factor is about equal to the average contribution of a non-traded good to aggregate RER variability. This reflects two interacting effects. First, the non-traded factor is the dominant source of variation. Second, the average traded good has more of its cost accounted for by non-traded inputs than the average non-traded good has accounted for traded inputs. Thus, most of the bias in attributing non-traded factor content in decompositions of the variance of real exchange rates is due to the so-called traded goods. To see this more clearly, it is productive to examine the cross-sectional variance in the contribution of the non-traded and traded factors at the microeconomic level rather than average across goods as **Table 2** does. We turn to this level of detail next.[7]

Once these two drivers of cost are estimated, it is possible to decompose the variance of each and every item in the basket into these two factors and a small residual term. This is particularly useful in placing a diverse literature intersecting trade, macroeconomics and industrial organization into a common empirical frame. In trade and industrial organization, for example, a researcher typically has very rich information about demand and cost structure about a particular sector and thus the ability to conduct a forensic analyis of the role of local inputs, traded inputs and markups for that sector, but not the ability to plug that implication into the larger picture of aggregate real exchange rate variation. Thus our method allows both cross-industry comparisons and a point of contact with the broader macroeconomic literature on

[7] Figure 3 shows the traded-input factor betas in ascending order (red), alongside the corresponding non-traded factor betas (blue). The red and blue horizontal lines represent the factor means of 0.54 and 1.05. The traded- and non-traded factor medians (0.54 and 1.05) are nearly identical to the means.

retail price adjustment than has been possible in the past. In fact, our variance decomposition may be aggregated to any desired level—both aggregation into final consumption sectors or in terms of the underlying local and traded factor content of those sectors. This flexibility is valued because economic models at the intersection of trade and macroeconomics often use different levels of aggregation depending on the central question of interest.

4.1 Microeconomic Variance Decompositions

Recall that after averaging the estimated equation (12) across jk pairs, we arrive at a decomposition of our original good-level betas:

$$\beta_i = 1.05\alpha_i + 0.54(1-\alpha_i) + \epsilon_i \tag{15}$$

$$= 0.54 + 0.52\alpha_i + \epsilon_i \ . \tag{16}$$

Were a purely traded good to exist at the retail level, $\alpha_i = 0$ and it would be predicted to contribute 0.54 times its expenditure share to aggregate RER variability. At the other end of the continuum, a purely non-traded good involves no traded inputs, $\alpha_i = 1$. If such a good existed, it would be expected to contribute 1.05 times its expenditure share to aggregate RER variability.

Table 3 shows the entire cross-sectional distribution of the good-specific contributions to RER variation, β_i, decomposed in this manner. Goods are ordered from those with the lowest distribution wedge (0.17), an example of which is a "compact car," to goods with the highest distribution wedge (1.00), an example of which is the "hourly rate for domestic cleaning help." Note that each row is an average across goods in a particular sector that shares the same distribution wedge (the second column) and the first column provides one concrete example (good description) from our micro-data.

Since the non-traded input beta, β^w_{jk}, averages 1.05, the contribution of the non-traded input is approximately equal to the distribution share, α_i. By our metric a compact automobile looks a lot like 1 liter of unleaded gasoline, but very distinct from a two-bedroom apartment or the hourly rate for domestic cleaning help. The contribution of LOP variation in each of the former two cases is about 70% traded inputs and 30% non-traded inputs, whereas the latter two are almost entirely driven by the non-traded input factor.

An interesting contrast is fresh fish and a two-course meal at a restaurant. Both are treated as traded goods when CPI data are used to implement the classical dichotomy using existing approaches (such as Engel) because they fall into the same CPI category, food. However, one is food at home (fresh fish) and the other is food away from home (two-course meal at a restaurant). Should they be treated similarly, as food items, or differently based on their underlying factor content? Consistent with the two-factor intermediate input model, **Table 3** provides a definitive answer: treat them differently. Fresh fish is indistinguishable from unleaded gasoline both in terms of the dominating role of traded inputs and the relatively moderate contribution to aggregate RER variation (beta of 0.65). A restaurant meal is dominated by the non-traded factor and contributes 41% more to aggregate RER variability than does fresh fish (here we are assuming the same expenditure share for items of each type).

Consider a good with a median distribution wedge (0.41), such as toothpaste. Despite the fact that the cost of producing this good is skewed moderately toward traded inputs (0.59), non-traded inputs still dominate in accounting for the toothpaste beta, 0.41 versus 0.30 for traded inputs. This reflects the fact that our estimated non-traded factor is twice as important as our estimated traded factor in accounting for variation in the aggregate RER, 1.05 versus 0.54. Stated differently, for the traded input factor to dominate in contribution to variance requires a distribution wedge of less than 0.34 (i.e., a traded input share of more than 0.66). **Figure 4** illustrates this point by showing the traded and non-traded input contributions to real exchange rate variations for the full cross-section of goods and services (i.e., across distribution shares).

4.2 Macroeconomic Variance Decompositions

Macroeconomics is, of course, about aggregate variables. Our thesis is that macroeconomists should aggregate final goods based on their non-traded and traded factor content, where the impact of international trade and nominal exchange rates is more easily distinguished. Our methodology provides that option. Here, we demonstrate its utility.

4.2.1 Aggregation Based on Intermediate Inputs

Recall that the microeconomic variance decomposition of the aggregate RER based on final goods is:

$$1 = \sum_i \omega_i \beta_{ijk}. \tag{17}$$

Substituting our two-factor model for the LOP deviation, $\beta_{ijk} = \alpha_i \beta^w_{jk} + (1 - \alpha_i)\beta^\tau_{jk} + \epsilon_{ijk}$, into this equation provides a theoretically consistent method of aggregating the micro-data based on the theory of trade in middle products:

$$1 = \sum_i \omega_i \left[\alpha_i \beta^w_{jk} + (1 - \alpha_i)\beta^\tau_{jk} + \epsilon_{ijk}\right]. \tag{18}$$

Notice that since the two intermediate factors are assumed to be location-specific, not good-specific, the expression aggregates very simply to provide a two-factor macroeconomic decomposition:

$$1 = \pi \beta^w_{jk} + (1 - \pi)\, \beta^\tau_{jk} + \eta_{jk}, \tag{19}$$

where the weights on the traded and non-traded input factors, π and $(1 - \pi)$, are consumption expenditure-weighted averages of the non-traded and traded input shares into each individual good in the consumption basket. These two weights measure the factor content of aggregate consumption expenditure in terms of traded and non-traded inputs. The residual term, η_{jk}, is an expenditure-share weighted average of the ϵ_{ijk}, which will be a very small number at any reasonable level of aggregation across goods, i.

Recall that the median distribution wedge (α_i) in the micro-data is 0.41. Using U.S. NIPA data and the EIU micro-sample, the weight of distribution in expenditure is estimated to be $\pi = 0.68$. The higher impact at the aggregate level reflects the fact that consumption expenditure is skewed toward services, which are intensive in distribution inputs. The dominant weight on the non-traded input factor, combined with the fact that β^w_{jk} is about twice the magnitude of β^τ_{jk}, is the reason that non-traded inputs dominate the variance decomposition of the aggregate RER by a very large margin. **Table 4** shows just how large.

Table 4 reports the results using ordinary least squares (OLS) estimates (weighted least squares (WLS) results are very similar). Beginning with the averages across the entire world sample, the non-traded factor accounts for about 81% (i.e., 0.71/ (0.71+0.17)) of the variance of the aggregate RER, while traded inputs account for the remaining 19%. The contribution of non-traded and traded inputs is even more dramatically skewed for the U.S.-Canada subsample, with non-traded inputs accounting for 89% and traded inputs accounting for the remaining 11%. These numbers clearly illustrate the main point of our paper: The classical dichotomy is an appealing theory of real exchange rate variations when we model each final good as a composite of a non-traded input and a traded input.

4.2.2 Aggregation Based on Final Goods

To further emphasize the difference between applications of the classical dichotomy at the level of inputs and final goods, we repeat the aggregate exercise applying the classical dichotomy at the level of final goods. To implement this using the micro-data, we must first decide on a definition of a non-traded good. In theory, the micro-data provides an advantage because it allows us to, for example, assign fish to the traded category and restaurant meals to the non-traded category, rather than placing all food in the traded category. The rule we use to be consistent with the intermediate input concept of the classical dichotomy is to categorize a good as a "non-traded good" if it has a distribution wedge exceeding 60%. This cutoff corresponds to a jump in the value of the distribution wedges across sectors from 0.59 to 0.75 (see Table 3 or Figure 2). As it turns out, this categorization matches up very well with the categorical assignments used by Engel (and the very large literature following his approach), who used very aggregated CPI data. The traded-goods category includes: cars, gasoline, magazine and newspapers, and foods. The non-traded goods category includes: rents and utilities, household services (such as dry cleaning and housekeeping), haircuts and restaurant and hotel services.

With the assignments of individual goods and services to these two categories, the aggregate RER is given by:

$$q_{jkt} = \omega q_{jkt}^{N} + (1-\omega) q_{jkt}^{T} \ , \tag{20}$$

where q_{jkt}^{N} and q_{jkt}^{T} are the bilateral RERs for non-traded final goods and traded final goods built from the LOP deviations in the microeconomic data, weighted by their individual expenditure shares.[8]

The variance decomposition of the aggregate RER is conducted using our beta method:[9]

$$1 = \omega\beta_{jk}^{N} + (1-\omega)\beta_{jk}^{T} \ . \tag{21}$$

Table 5 reports the outcome of the variance decomposition arising from this macroeconomic approach. Beginning with the averages across the entire world sample, the non-traded factor accounts for about 64% (i.e., (1-0.46)×1.18) of the variance of the aggregate RER, while traded inputs account for the remaining 36%.

It is instructive to compare Table 5 with Table 1 since they both use final goods as the working definition for traded and non-traded goods. What is the consequence of aggregating the data before conducting the variance decomposition? As it turns out, the betas are very similar across the two approaches. The average beta for non-traded (traded) goods pooling all locations is 1.18 (0.78) using the two-index construct (Table 5) compared with 1.05 (0.54) using the microeconomic decomposition. These are relatively small differences.

The underlying sources of the contribution to variance, however, are different. When using the macroeconomic approach, the non-traded RER contributes more to the variability of the aggregate RER for two reasons. First, the non-traded RER is more highly correlated with the aggregate RER than is the traded RER (0.94 versus 0.86). Second, reinforcing this effect is the fact that the non-traded sub-index of the CPI is more variable than the traded RER (1.25 versus 0.91). In contrast, when the microeconomic approach is used, non-traded and traded goods are not distinguished by the relative volatility of their LOP deviation (at least for the median good). Both types of goods have standard deviations twice that of the aggregate RER. Consistent with the macroeconomic approach, the LOP deviations of the median non-traded

[8] More precisely, the weights used earlier are renormalized to $\frac{\omega_i}{\omega}$ ($\frac{\omega_i}{1-\omega}$) for non-traded (traded) goods so that the weights on the two sub-indices sum to unity.

[9] The relationship between the microeconomic betas of our original decomposition and this two-factor decomposition is straightforward: $\omega_{jk}\beta_{jk}^{N} = \sum\limits_{i \in N} \omega_{ijk}\beta_{ijk}$ and $(1-\omega_{jk})\beta_{jk}^{T} = \sum\limits_{i \in T} \omega_{ijk}\beta_{ijk}$.

good have a higher correlation with the aggregate RER than do the median traded goods (0.51 versus 0.40). Thus, traded goods have more idioynscratic sources of deviations from the LOP than non-traded goods.

5 Related Literature

Our analytic and empirical methods also allow simple and transparent connections to virtually all approaches in the existing liteature.

5.1 Engel (1999)

Engel (1999) conducts a variance decomposition using what we refer to as aggregation based on final goods. His two-sector price index version of the real exchange rate is:

$$q_{jkt} = \omega q^{N}_{jkt} + (1-\omega) q^{T}_{jkt} \ . \tag{22}$$

But Engel does not work with this equation; he rearranges the equation as follows:

$$q_{jkt} = q^{T}_{jkt} + \omega (q^{N}_{jkt} - q^{T}_{jkt}) \ .$$

A few lines of algebra show that the variance decompositions that result, using our beta notation are, respectively,

$$1 = \omega \beta^{N}_{jk} + (1-\omega) \beta^{T}_{jk} \tag{23}$$

and

$$1 = \beta^{T}_{jk} + \omega (\beta^{N}_{jk} - \beta^{T}_{jk}).$$

Under the null hypothesis that the classical dichotomy holds, $\beta^{T}_{jk} = 0$, and the two expressions are mathematically equivalent. But, of course, we know this is not the case. Our estimate of $\beta^{T}_{jk} = 0.78$, so there is no disagreement on the strong rejection of this straw-man null hypothesis that traded final goods obey the LOP.

The title of Engel's paper, however, is "Accounting for U.S. Real Exchange Rate Changes." Accounting refers to the implementation of a variance decomposition, which Engel characterizes as showing: "The outcome is surprising:

relative prices of nontraded goods appear to account for almost none of the movement of U.S. real exchange rates." What is being referred to as relative prices of non-traded goods is actually the relative price of non-traded goods relative to the relative price of traded goods, $q^N_{jkt} - q^T_{jkt}$. The rearrangement of terms in the expression, while innocuous in this application in terms of shooting down a straw-man that the LOP holds for traded goods in the final consumption basket, is highly misleading in the content of a variance decomposition, the stated purpose of the paper. This can be distilled into two very simple points.

First, in any variance decomposition there are covariance terms. Second, in the case of a price index, the components must be weighted by their respective expenditure weights. The second issue is obvious with the replacement of the traded goods expenditure share, $(1 - \omega)$, with a unit coefficient as we move from our decomposition to Engel's, which approximately doubles its influence in the variance decomposition.

To see the role of covariance, it is useful to rewrite the original real exchange rate as follows:

$$q_{jkt} = \omega(s_{jkt} + \widehat{q}^N_{jkt}) + (1 - \omega)(s_{jkt} + \widehat{q}^T_{jkt}), \tag{24}$$

where $q^Z_{jkt} = s_{jkt} + \widehat{q}^Z_{jkt}$, $Z = T, N$, $E(\widehat{q}^N_{jkt}, \widehat{q}^T_{jkt}) = 0$ and s_{jkt} is the nominal bilateral exchange rate, but for the sake of this argument could be any common factor that generates a positive covariance of real exchange rates across individual goods or sectors.

The beta decomposition becomes:

$$1 = \omega(\widehat{\beta}^N_{jk} + \Lambda) + (1 - \omega)(\widehat{\beta}^T_{jk} + \Lambda), \text{ where } \Lambda = cov(s_{jkt}, q_{jkt})/var(q_{jkt}). \tag{25}$$

So our decomposition method partitions any common factor driving international relative prices in a neutral way, according to the expenditure shares.[10]

[10] If the factor loadings on the common factor are different, our approach will still accommodiate this, $1 = \omega(\widehat{\beta}^N_{jk} + \gamma^N\Lambda) + (1 - \omega)(\widehat{\beta}^T_{jk} + \gamma^T\Lambda)$. A natural case is less pass-through of nominal exchange rate changes into local input prices than traded goods, $\gamma^N > \gamma^T$, which would be a natural reason that the overall beta on non-traded goods is larger as we found earlier.

Consider what happens if the terms are rearranged as in Engel (1999). Now the two terms in the variance decomposition are:

$$1 = (\widehat{\beta}_{jk}^{T} + \Lambda) + \omega(\widehat{\beta}_{jk}^{N} - \widehat{\beta}_{jk}^{T}).$$

It is obvious that the common factor cancels out the second term, and gets fully attributed (i.e., not even deflated by the expenditure share of traded goods) to the first term.

Another arbitrary rearrangement of terms would greatly exaggerate the conceptual value of the classical dichotomy. Consider rearranging the real exchange rate to read, $q_{jkt} = q_{jkt}^{N} + (1-\omega)(q_{jkt}^{T} - q_{jkt}^{N})$. The variance decomposition becomes:

$$1 = (\widehat{\beta}_{jk}^{N} + \Lambda) + (1-\omega)(\widehat{\beta}_{jk}^{T} - \widehat{\beta}_{jk}^{N}),$$

yielding the opposite conclusion: now the covariance term is assigned to the non-traded component with a unit weight and the result strongly favors the classical dichotomy. Both conclusions are false.

5.2 Parsley and Popper (2009)

The same issue arises when Engel's method is applied using micro-price data. Parsley and Popper (2009) use two independent retail surveys in the United States and Japan. The U.S. survey is conducted by the American Chamber of Commerce Researchers Association (ACCRA) and the Japanese survey is from the Japanese national statistical agency publication, the *Annual Report on the Retail Price Survey*. Both contain average prices across outlets, at the city level. The Japanese survey is vastly more extensive in coverage of items than the ACCRA survey, since it represents the core micro-data that goes into the Japanese CPI construction. Both data panels are at the city level and thus are quite comparable in many ways to the EIU data used in this paper. Parsley and Popper restrict their sample to items that are as comparable as possible across the two countries. This selection criteria leaves them with a sample of highly traded goods.

To elaborate our method when micro-data are employed, rather than two sub-indices, consider applying item-specific weights, ω_i, to LOP deviations. Starting with the aggregate RER as an expenditure-weighted average of LOP (we abstract from city-pair subscripts here because it does not effect our arguments), we have:

$$q_t = \sum_i \omega_i q_{i,t} \ . \tag{26}$$

Parsley and Popper follow Engel's approach by placing an individual good in the lead position with a unit coefficient as its weight. That is, for each good i, they work with:

$$q_t = q_{i,t} + \left(\sum_g \omega_g q_{g,t} - q_{i,t} \right) \ . \tag{27}$$

That is, the lead term is a single good i, and the remaining terms are all other prices in the panel. Parsley and Popper then compute the variance of the lead term, the LOP variance, and divide it by the total variance of the RER and define this ratio as the contribution of good i to the variance of the aggregate RER.

In terms of our betas, their variance decomposition is equivalent to:

$$\begin{aligned} 1 &= \beta_i + (\sum_g \omega_g \beta_g) - \beta_i && (28) \\ &= \beta_i + (1 - \beta_i) \ . && (29) \end{aligned}$$

The last identity holds because the expenditure weighted average of the betas must equal unity by construction. However, the variance decomposition following our method keeps the weights on all the terms,

$$1 = \omega_i \beta_i + \sum_{g \neq i} \omega_g \beta_g \ . \tag{30}$$

As is evident, the contributions to variance of individual goods in the Parsley and Popper paper are actually equal to our betas. However, in following Engel's approach, they give each good a unit weight. As our decomposition shows, these good-specific betas need to be multiplied by expenditure shares in order to conduct a legitimate variance decomposition. Parsley and Popper end up reconciling 28 items across the U.S. and Japan, 2 of which are services. They compute the contribution to variance at different horizons, including five quarters. At this horizon, the good-specific contributions range from just under 0.5 to about 0.86. Interpreted as betas, these estimates certainly fall within the range we find, which spans negative values to values exceeding

1. However, they are not contributions to aggregate RER variance; to arrive at a legitimate variance decomposition, each beta must be multiplied by its consumption-expenditure weight.

6 Conclusions

Using retail price data at the level of individual goods and services across many countries of the world, we have shown the classical dichotomy is a very useful theory of international price determination when applied to intermediate inputs. Specifically, by parsing the role of non-traded and traded inputs at the retail level, a significant source of compositional bias is removed from the micro-data and differences in the role of the two inputs become patently obvious. Aggregate price indices are not useful in uncovering this source of heterogeneity in LOP deviations for two reasons. First, the dividing line between traded and non-traded goods at the final goods stage is arbitrary and more under the control of officials at statistical agencies whose goal is not to contrast the role of trade across CPI categories of expenditure. Second, even at the lowest level of aggregate possible, most goods and services embody costs of both local inputs and traded inputs. Consequently, the contribution of each LOP deviation to PPP deviations is a linear combination of the two components with the weights on the two components differing substantially in the cross-section.

Our results point to the usefulness of microeconomic theories that distinguish traded and local inputs and their composition in final goods. The method used here also allows for LOP deviations at the level of trade at the "dock." Importantly, the method assigns covariance risk that links microeconomic variables to aggregate variables in a neutral way. The findings that non-traded or local inputs dominate in contribution to the variance of the aggregate real exchange rate points to the need for a hybrid model with a distribution sector and segmentation at the level of traded inputs. Finally, our method provides a mechanism to estimate and calibrate international price risks faced by firms and workers in different cities and countries based on the nature of their specialization.

References

[1] **Andrade, Philippe and Marios Zachariadis**. 2016. Global versus local shocks in micro price dynamics. Journal of International Economics, 98, 78–92.

[2] **Berger, David, Jon Faust, John Rogers and Kai Stevenson**. 2012. Border prices and retail prices. Journal of International Economics, 88(1): 62-73.

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

[4] **Burstein, Ariel, Martin Eichenbaum and Sergio Rebelo.** 2005. Large devaluation and the real exchange rate. Journal of Political Economy, 113(4): 742-784.

[5] **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.

[6] **Campa, José Manuel and Linda S. Goldberg**. 2010. The sensitivity of the CPI to exchange rates: distribution margins, imported inputs, and trade exposure. The Review of Economics and Statistics, 92(2): 392-407.

[7] **Corsetti, Giancarlo, Luca Dedola and Sylvain Leduc**. 2008. International risk sharing and the transmission of productivity shocks. Review of Economic Studies, 75: 443-473.

[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., Mototsugu Shintani and Takayuki Tsuruga**. 2010. Accounting for persistence and volatility of good-level real exchange rates: the role of sticky information. Journal of International Economics, 81, 48-60.

[10] **Crucini, Mario J., Chris I. Telmer and Mario Zachariadis.** 2005. Understanding European real exchange rates. American Economic Review, 95(3): 724-738.

[11] **Crucini, Mario J., Chris I. Telmer**. 2012. Microeconomic Sources of Real Exchange Rate Variation. NBER Working Paper No. 17978.

[12] **Crucini, Mario J., Hakan Yilmazkuday**. 2014. Understanding long-run price dispersion. Journal of Monetary Economics, 66: 226-240.

[13] **Engel, Charles and John Rogers.** 1996. How wide is the border? American Economic Review, 86(5): 1112-1125.

[14] **Engel, Charles**. 1999. Accounting for U.S. real exchange rates. Journal of Political Economy, 107(3), 507-538.

[15] **Parsley, David C. and Helen Popper**. 2006. Understanding real exchange rate movements with trade in intermediate products. Pacific Economic Review, 15(2): 171-188.

[16] **Rogers, John H.** 2007. Monetary union, price level convergence, and inflation: How close is Europe to the USA? Journal of Monetary Economics, 54:3, 785–796.

[17] **Sanyal, Kalyan K., and Ronald W. Jones**. 1982. The theory of trade in middle products. American Economic Review. 72(1): 16-31.

Table 1:

Variance decomposition of real exchange rates

Microeconomic approach, international pairs

	All pairs	OECD	Non-OECD	U.S.-Canada
Std. dev. RER	0.12	0.13	0.10	0.19
Non-traded weight	0.54	0.56	0.52	0.57
Traded weight	0.46	0.44	0.48	0.43
All goods				
Beta	0.81	0.77	0.83	0.81
Correlation	0.42	0.42	0.43	0.49
Rel. std. dev. LOP	2.12	2.08	2.16	1.64
Non-traded goods				
Beta	1.01	1.01	1.00	1.09
Correlation	0.51	0.52	0.50	0.61
Rel. std. dev. LOP	2.14	2.11	2.16	1.79
Traded goods				
Beta	0.75	0.71	0.79	0.74
Correlation	0.40	0.39	0.41	0.46
Rel. std. dev. LOP	2.11	2.07	2.15	1.60
Number of city pairs	4732	993	1295	48

Table 2:

Traded and non-traded input regressions

International pairs

	All pairs	OECD	Non-OECD	U.S.-Canada
beta (traded)	0.54	0.44	0.63	0.35
	(.52)	(.55)	(.45)	(.35)
beta (non-traded)	1.05	1.07	1.03	1.16
	(.33)	(.34)	(.31)	(.32)
slope	0.52	0.63	0.39	0.81
	(.74)	(.76)	(.66)	(.61)
R-squared	0.08	0.10	0.06	0.15
Number of pairs	4732	993	1295	48

Note: Minimum of 4 observations per city pair.

Table 3: Variance decomposition using intermediate input betas

		Contribution			Non-traded
Example	α_i	Non-Traded	Traded	Residual	Cont. (%)
Compact car (1300-1799 cc)	0.17	0.17	0.44	0.02	28%
Unleaded gasoline (1 liter)	0.19	0.19	0.42	-0.02	31%
Fresh fish (1 kg)	0.22	0.21	0.42	0.02	33%
Time (news magazine)	0.32	0.32	0.35	0.03	47%
Toilet tissue (two rolls)	0.34	0.34	0.33	-0.03	51%
Butter (500 g)	0.36	0.36	0.33	-0.01	53%
Aspirin (100 tablets)	0.37	0.36	0.32	0.02	53%
Marlboro cigarettes (pack of 20)	0.37	0.37	0.33	0.03	53%
Electric toaster	0.39	0.39	0.31	0.04	55%
Toothpaste with fluoride (120 g)	0.41	0.41	0.30	-0.02	57%
Compact disc album	0.41	0.41	0.30	-0.08	58%
Insect-killer spray (330g)	0.45	0.45	0.28	0.01	62%
Paperback novel	0.49	0.48	0.26	0.00	65%
Razor blades (5 pieces)	0.49	0.49	0.27	0.00	65%
Batteries (two, size D/LR20)	0.50	0.50	0.26	-0.06	65%
Socks, wool mixture	0.52	0.52	0.24	0.08	68%
Men's shoes, business wear	0.52	0.52	0.24	0.08	68%
Lettuce (one)	0.52	0.52	0.24	-0.01	68%
Frying pan (Teflon)	0.53	0.53	0.24	0.02	68%
Light bulbs (two, 60 watts)	0.57	0.57	0.22	-0.07	72%
Child shoes, sportwear	0.59	0.58	0.21	-0.02	74%
Tennis balls (Dunlop, Wilson or equivalent)	0.59	0.59	0.21	-0.09	74%
Two-course meal at a restaurant (average)	0.75	0.75	0.13	-0.03	86%
Electricity, monthly bill (average)	0.76	0.75	0.12	0.09	86%
Man's haircut (tips included)	0.85	0.85	0.08	-0.03	92%
Taxi, airport to city center (average)	0.86	0.86	0.07	-0.04	92%
Telephone line, monthly bill (average)	0.92	0.89	0.04	0.00	95%
2-bedroom apartment	0.93	0.91	0.04	0.18	96%
Annual premium for car insurance	0.94	0.92	0.03	-0.01	97%
Hourly rate for domestic cleaning help	1.00	1.00	0.00	-0.10	100%

Table 4:

Macroeconomic variance decomposition

Intermediate input approach, international pairs

	All pairs	OECD	Non OECD	U.S.-Canada
Non-traded share (π)	0.68	0.68	0.68	0.68
	(.02)	(.02)	(.03)	(.02)
Contribution of				
Traded inputs	0.17	0.14	0.21	0.11
	(.17)	(.18)	(.15)	(.11)
Non-traded inputs	0.71	0.73	0.69	0.79
	(.22)	(.23)	(.21)	(.22)
Error term	0.11	0.13	0.10	0.10
	(.20)	(.22)	(.19)	(.16)
Number of city pairs	4732	993	1295	48

Table 5:

Variance decomposition of real exchange rates

Macroeconomic approach, international pairs

	All	OECD	Non-OECD	U.S.-Canada
Std. dev. RER	0.12	0.13	0.10	0.19
Non-traded weight	0.54	0.56	0.52	0.57
Traded weight	0.46	0.44	0.48	0.43
All goods				
Beta	0.98	0.97	0.99	0.96
Correlation	0.90	0.90	0.91	0.94
Rel. std. dev. LOP	1.08	1.08	1.09	1.02
Non-traded goods				
Beta	1.18	1.18	1.17	1.17
Correlation	0.94	0.95	0.94	0.97
Rel. std. dev. LOP	1.25	1.25	1.26	1.20
Traded goods				
Beta	0.78	0.76	0.81	0.76
Correlation	0.86	0.84	0.88	0.90
Rel. std. dev. LOP	0.91	0.91	0.93	0.85
Number of city pairs	4732	993	1295	48

Figure 1: Density Distributions of Betas, Microeconomic Decomposition

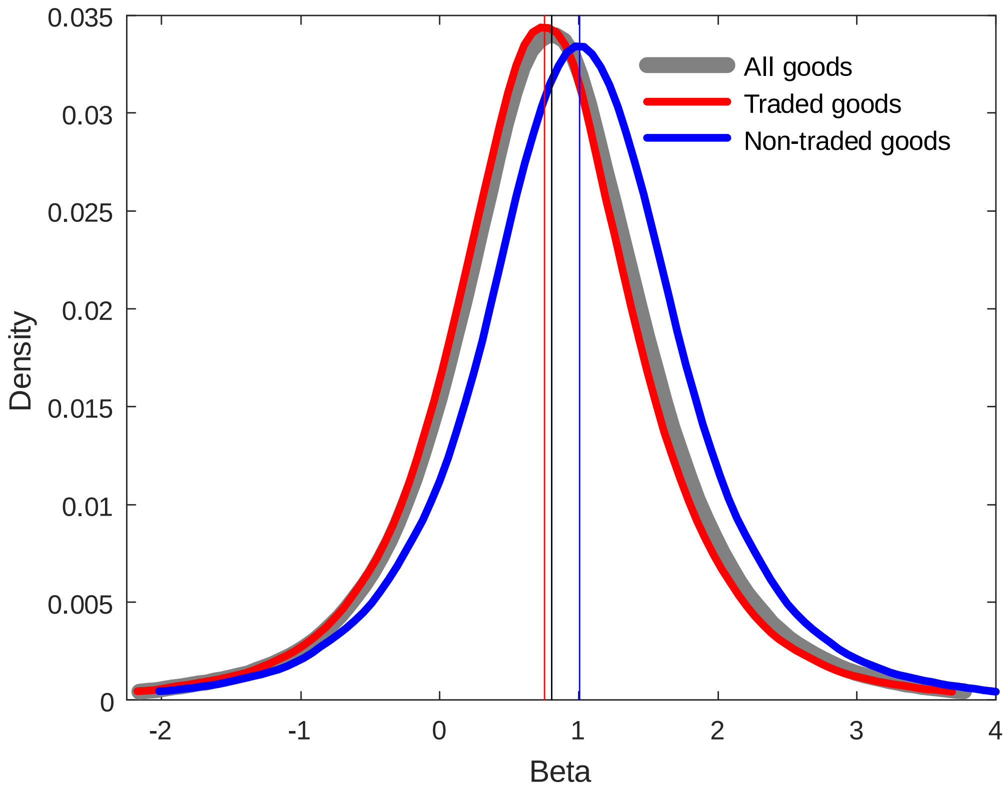

Figure 2: Sectoral Betas and Distribution Shares

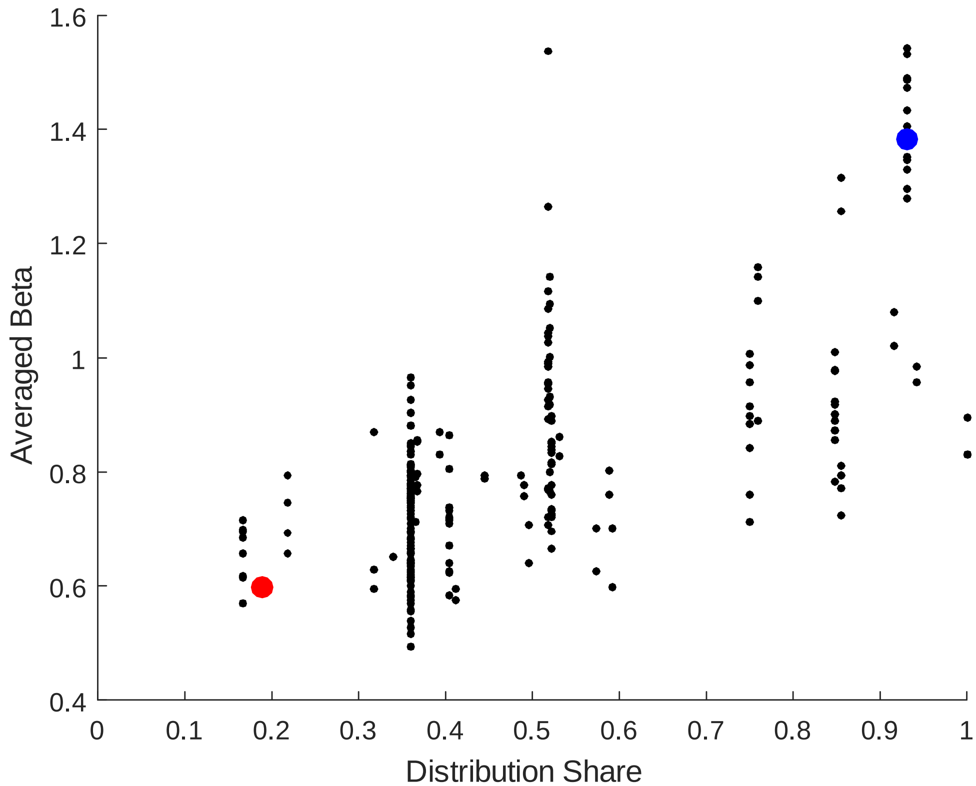

Figure 3: Traded and Non-Traded Inputs Factor Betas

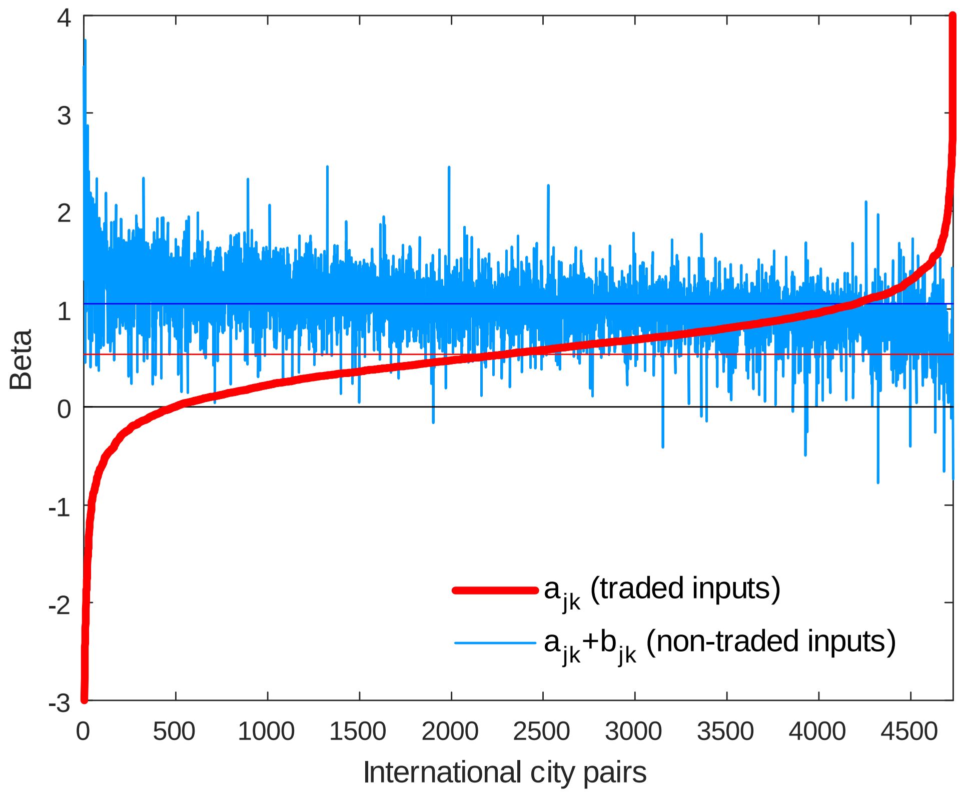

Figure 4: Traded and Non-Traded Inputs Contribution to Real Exchange Rate Variations