Accounting for Real Exchange Rates Using Micro-Data[*]

Mario J. Crucini
Vanderbilt University
NBER

Anthony Landry
Federal Reserve Bank of Dallas
The Wharton School of the University of Pennsylvania

February 2012

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 CPI basket because traded goods obey the Law of One Price. In stark contrast, Engel (1999) found that traded goods had comparable volatility to the aggregate real exchange. Our work reconciles these two views by successfully applying the classical dichotomy at the level of intermediate inputs into the production of final goods using highly disaggregated retail price data. Since the typical good found in the CPI basket is about equal parts traded and non-traded inputs, we conclude that the classical dichotomy applied to intermediate inputs restores its conceptual value.

JEL codes: F4

---

1 Introduction

One of the most stable empirical relationships in international macroeconomics is the comparable volatility of nominal and real exchange rates. This relationship was first documented by Mussa (1986) using a short panel of data following the breakdown of the Bretton Woods system of fixed nominal exchange rates. Subsequent research has shown that this relationship is robust to longer time spans of data and broader cross-sections of countries. Understandably, these observations have been interpreted as evidence that goods markets are nationally segmented.

An enduring explanation for real exchange rate variability is the classical dichotomy of Salter (1959) and Swan (1960). This is the notion that the consumption basket consists of items which are imported, items which are exported and items which are produced only for domestic consumption. The literature refers to items in the first two categories as traded goods and items in the last category as non-traded goods. The use of the term goods is misleading in the sense that some services are traded and some goods are not traded: to avoid confusion with conventional usage, the term goods will refer to all items in the consumption basket. According to the classical dichotomy, traded goods satisfy the Law of One Price (LOP) up to a constant iceberg trade cost and thus the real exchange rate of traded goods is constant over time, leaving the real exchange rate of non-traded goods to account for all of the time series variation in the CPI-based real exchange rate.

In a highly influential paper, Engel (1999) constructs traded and non-traded real exchange rates using sub-indices of the CPIs of the United States and its largest trading partners and conducts a variance decomposition of bilateral aggregate real exchange rates into the contributions of the two sub-indices. In stark contrast to the predictions of the classical dichotomy, he finds that real exchange rates of traded and non-traded goods contribute a comparable amount to aggregate real exchange rate variability. His results have shifted the consensus among economists from the classical dichotomy view to the view that all final goods markets are equally segmented. And yet the sources of market segmentation have been difficult to pin down.

A voluminous literature on exchange rate pass-through shows that move-

ments in nominal exchange rates are not fully reflected in subsequent movements in destination prices (see Campa and Goldberg (2005) for a comprehensive treatment). Thus even the prices of highly traded goods are unlikely to satisfy the LOP at the retail level since they fail to do so at the border. The question then becomes not simply whether one can or should distinguish goods and services in the stark manner suggested by the classical dichotomy, but also whether it is productive to make more subtle distinctions in crafting the architecture of modern international macroeconomic models.

The goal of this paper is to reassess the usefulness of distinguishing items in the consumption basket by the variation in their real exchange rates. Two complementary elements of novelty are introduced. The first is the use of microeconomic retail price data in the variance decomposition of the aggregate real exchange rate. The second is the application of the classical dichotomy to intermediate inputs into the production of individual final goods, rather than dichotomizing the final goods themselves.

To achieve the first step the aggregate real exchange rate for a typical bilateral pair (i.e., suppressing bilateral location indices) is built from the ground up using Law of One Price deviations, q_{it},

$$q_t = \sum_i \omega_i q_{it} \ , \tag{1}$$

where ω_i is an item-level consumption expenditure weight. Due to the very large number of terms on the right-hand-side of this equation a conventional variance decomposition is not feasible. Instead, the covariance of each of the LOP deviations is taken with respect to the aggregate bilateral real exchange rate q_t; dividing all terms on each side of the equation by the variance of q_t gives the desired result:

$$1 = \frac{cov(q_t, q_t)}{var(q_t)} = \sum_i \omega_i \frac{cov(q_{it}, q_t)}{var(q_t)} = \sum_i \omega_i \beta_i \ . \tag{2}$$

The use of the notation β_i is deliberate. It reminds the reader of the betas used in portfolio analysis. In this application, the return on the portfolio is replaced with the relative price of the same basket across two locations; the portfolio weights are the expenditure shares and the LOP deviations take the role of the returns on individual stocks in the portfolio. As a concrete example,

2

values of β greater than unity indicate items in the consumption basket that contribute more than their expenditure share to the variability of the aggregate real exchange rate.

Engel's two-index decomposition is easily constructed by aggregating the LOP deviations into sub-indices for traded goods and non-traded goods:

$$q_t = (1 - \omega)q_t^T + \omega q_t^N \ , \tag{3}$$

where ω and $1 - \omega$ are the aggregate consumption expenditure shares of goods deemed to be either traded or non-traded. The variance decomposition becomes:

$$1 = (1 - \omega)\beta^T + \omega\beta^N \ . \tag{4}$$

These β's are simply expenditure-weighted averages of their microeconomic counterparts β_i. The restrictions of the classical dichotomy applied to these aggregates are: $\beta^T = 0$ and $\beta^N = \omega^{-1}$. The first restriction is the notion that the LOP is assumed to hold for all traded goods and thus also for the aggregate traded real exchange rate index. The second restriction is an implication of the first along with the definition of a variance decomposition. It is important to note that in the aggregate version, the assumption about LOP only needs to hold on average across traded goods. Given that roughly 0.6 of expenditure is attributed to non-traded goods based on the dichotomous classification, $\beta^N = 1.7$ is predicted by the classical dichotomy when applied to final goods.

Using Economist Intelligence Unit (EIU) retail price data and U.S. expenditure weights for the ω_i, the average β for non-traded goods is $\beta^N = 1.03$ and the average for traded goods is $\beta^T = 0.81$. In words: non-traded goods contribute 22% more to the variability of the aggregate real exchange rate than do traded goods. While this difference is substantial and consistent in direction with the classical dichotomy, the fact remains that the contribution of traded goods to the variance is much closer to that of non-traded goods than it is to zero. It is on this basis that a consensus has emerged that most items found in the consumption basket are exchanged in markets that are nationally segmented. And at this level of aggregation, the EIU data broadly support this conclusion and are therefore consistent with Engel's findings using official CPI data.[1]

---

[1]The aggregation used here corresponds closely to Engel's results using CPI indices. He

There are two problems with this interpretation of the evidence. First, the level of aggregation of the CPI typically available to researchers is not designed to achieve a partition of goods into traded and non-traded items. For example, food includes both food away from home and groceries, housing includes both rent and utilities – if forced to make a choice, the food away from home and rent would be placed on the non-traded side of the ledger while groceries and utilities would be placed on the traded side of the ledger. Since much of the data is aggregated to the level of food and housing, the categories becomes less sharp at distinguishing traded and non-traded goods. However, the use of micro-price data in this study completely avoids categorization bias due to aggregation so either the problem is not acute or something else is making the two macroeconomic and microeconomic approaches appear more comparable.

This brings us to the second difficulty, the fact that final goods involve both traded and non-traded inputs in their production. This violates the basic premise of the classical dichotomy as applied to final goods except for carefully chosen anecdotes. In the Economist Intelligence Unit (EIU) retail price survey of 301 goods and services, there are only a handful of items purchased by consumers with no obvious role for traded inputs (e.g., baby-sitting services and monthly salary of a maid) and there are no items that could be reasonably described as purely traded in the sense of the consumer paying shipping costs and purchasing directly from the manufacturer or wholesaler (as would be true of some online purchases, for example). The modal good in virtually all national retail price surveys, including the EIU sample, is a food product sold in a grocery store, which, according to the U.S. NIPA data, has a non-traded input share of 0.41.

Any hope of resuscitating the classical dichotomy therefore seems to require that the distinction between traded and non-traded goods be pushed to the level of intermediate inputs used in the production of goods and services appearing in the consumption basket. The share of these non-traded inputs in marginal cost have been dubbed distribution margins by Burstein, Eichenbaum and Rebelo (2005) and are measured as the difference between

classified services and housing as non-traded and all other sub-indices of the CPI (referrred to as commodities) as traded. We achieve basically the same classification when we label goods with a distribution cost share of 0.6 or higher as non-traded goods.

the price paid by the final consumer and the price received by the manufacturer at the factory gate. Strictly speaking this measure applies only to goods and not to services because the NIPA data treat the market for services as an arms-length transaction between the service provider and the end consumer. That is, a medical bill paid by a consumer or health insurance company would be recorded as having no distribution margin (or non-traded inputs) by this measure. This study uses the input-output tables of the United States to measure the traded and non-traded cost shares of services, following Crucini, Telmer and Zacharadis (2005). It should also be noted that these measures also include any markups over marginal cost at the wholesale and retail levels.

If the cost share of non-traded inputs took on two values, zero and unity, a researcher with measures of the shares could create a dichotomous variable and formally test the validity of the theory using either microeconomic data or aggregate indices. Such a test is infeasible because there are almost no items in the consumption basket at either of these two extremes. According to U.S. NIPA data, the non-traded input share for food products is approximately one-third, the median share in the EIU cross-section is 0.41 and the average is 0.50. The average is particularly telling since such goods should be viewed as equally weighted in traded and non-traded inputs completely obscuring any tendency for the classical dichotomy to hold at the level of intermediate inputs.

By way of example, and to illustrate our variance decomposition methodology, consider the dichotomy of the CPI employed by Engel: commodities and shelter. An example of a commodity in the EIU micro-sample is a gallon of unleaded gasoline and an example of shelter is an unfurnished two-bedroom apartment. We have chosen these two items because they capture extreme values of the non-traded cost share: 0.19 and 0.93, respectively. The betas for these two items are 0.61 and 1.43, respectively. Notice the much greater degree of separation between these two betas than the averages reported above: if non-traded good inputs are proxied by an apartment rental and traded goods by a gallon of unleaded fuel, the difference in the betas is now 82% – twice the value estimated using the traded and non-traded aggregates.

This represents real progress for the classical dichotomy, particularly when one realizes that this difference is likely to be an underestimate of the differences in the underlying input betas because even for this carefully chosen

anecdote the cost shares of non-traded inputs are not 0 and 1, but rather 0.19 and 0.93. Assuming traded inputs and non-traded inputs are different from one another, but the same for a gallon of unleaded gasoline and an unfurnished two-bedroom apartment, they may estimated by solving two equations in two unknowns. Using the betas for unleaded fuel and apartments, the values for the traded and non-traded input betas that solve the two equations are: $\beta^T = 0.40$ and $\beta^N = 1.51$. Note that the estimated non-traded beta is close to the beta for apartments, 1.51 versus 1.43, due to the fact that 0.93 is close to 1, whereas the estimated traded beta is much lower than that of unleaded fuel, 0.40 versus 0.61, due to the fact that 0.19 is a considerable distance from 0. These indirect estimates of the underlying input betas for non-traded and traded inputs differ by 111%.

To summarize, when we use two sub-indices of the CPI to construct a real exchange rate for traded and non-traded goods, non-traded goods contribute about 27% more to aggregate real exchange rates than do traded goods. When we use individual goods and services that most closely resemble traded and non-traded items the gap widens to 82% and when we adjust the estimates to account for the distribution margins for each good, the gap widens further to 111%. These calculations are representative of the broader cross-section and suggest that the classical dichotomy is a very useful description of LOP deviations. Our interpretation of the invalidation of the classical dichotomy at the level of aggregate CPI indices is a combination of not distinguishing intermediate inputs from final goods and of using data too highly aggregated to preserve interesting differences in the traded factor content of final consumption goods. Consistent with the pass-through literature on prices at the dock, significant LOP deviations remain even after controlling for non-traded inputs indicating the existence of market segmentation where the classical dichotomy assumes none exists: a purely traded input. The conclusion drawn from these variance decompositions is that a hybrid model with some market segmentation in traded goods and a good-specific distribution margin is a fruitful avenue for future theoretical and empirical research.

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 RER volatility. In Section 4,

we document a striking positive relationship between the magnitude of the contribution of a 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 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 that the economic implications are very different. Section 6 concludes.

2   The Data

The source of retail price data is the Economist Intelligence Unit Worldwide Survey of Retail Prices. The EIU survey collects prices of 301 comparable goods and services across 123 cities of the world. The number of prices overstates the number of items because many items are 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 mid-price/branded store. The prices are collected from the same physical outlet over time, thus the prices are not averages across outlets. The panel used in this study is annual and spans the years 1990 to 2005. The local currency prices are converted to common currency using the prevailing nominal exchange rates at the time the survey was conducted.

The data are supplemented with two additional sources from the U.S. Bureau of Economic Analysis: the National Income and Product Account and the Industry Economic Accounts Input-Output tables. The first supplementary data series are consumption-expenditure weights. These data are more aggregated than the EIU prices, leading us to allocate about 300 individual retail prices to 73 unique expenditure categories. We divide the sectorial 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. As some sectors are not represented in the EIU retail price surveys, the expenditure weights are adjusted upward to sum to unity.

The second supplementary source are the distribution shares, the fraction of gross-output attributable to non-traded inputs. These include wholesale and

retail services, marketing and advertisement, local transportation services and markups. For goods, this is the difference between what consumer pays and what manufacturer receives divided by that the consumer pays. For example, if final consumption expenditure on bread is \$100 and manufacturers receive \$60, the distribution share is 0.40.

For services, however, what consumers pay and what sellers receive would be the same value by this accounting method. In reality, when a consumer (or that consumer's health insurance provider) receives a medical bill, the charge includes wage compensation for their doctor and the cost of any goods or other services included in the treatment, whether or not it is itemized on the invoice. In these circumstances we use input-output data to measure non-traded and traded inputs. Each retail item in the EIU panel is reconciled with one of these sectors and assigned that sector's distribution share, leading to 30 unique sectorial shares. The median good as a distribution share of 0.41. The economy-wide average distribution share weighted by final expenditure is 0.48 for US cities, 0.49 for OECD cities and 0.53 for non-OECD cities. Overall, our distribution shares are similar to those used in Burstein et al. (2003) and Campa and Goldberg (2010).

3 Microeconomic Decomposition

The theoretical construction of the aggregate real exchange rate appeals to a utility function and the derivation of a corresponding price index. Let C_{jt} denote consumption by individual j at time t consisting of a Cobb-Douglas aggregate of this individual's consumption of various goods (and services), C_{ijt}, with weights ω_i. We have:

$$U(C_{jt}) = \prod_i \left(C_{ijt}\right)^{\omega_i} \ . \tag{5}$$

Note that the j index will also refer to the location where the individual purchases the goods, which given the nature of our data will be a city. Note also that the absence of an individual index on the ω_i means that all individuals have the same preferences.

Solving an expenditure minimization problem produces an ideal price index in the sense that it maps the prices of individual goods and services into a

single consumption deflator with the property that aggregate consumption is consistent with the utility concept defined by the structure of preferences. For the case of Cobb-Douglas preferences, the price index P_{jt}, is a simple geometric average of good-level prices P_{ijt} with the consumption expenditure shares as weights in the average:

$$P_{jt} = \prod_i (P_{ijt})^{\omega_i} \ . \tag{6}$$

This deflator satisfies $P_{jt}C_{jt} = \sum_i P_{ijt}C_{ijt}$, where the quantities of aggregate consumption and consumption of individual goods and services are the optimal levels chosen by consumers in city j, taking prices and income as given.

Converting prices to common currency at the spot nominal exchange rate, leads to the definition of the aggregate real exchange rate (RER) $Q_{jk,t}$ for bilateral city pair j and k as a function of microeconomic relative prices:

$$Q_{jk,t} = \frac{S_{jkt}P_{jt}}{P_{kt}} = \prod_i \left(\frac{S_{jkt}P_{ijt}}{P_{ikt}} \right)^{\omega_i} \ , \tag{7}$$

where S_{jkt} is the spot nominal exchange rate between city j and k. Taking logarithms leads to a relationship in which the RER is a consumption-expenditure weighted average of Law of One Price (LOP) deviations[2]:

$$q_{jk,t} = \sum_i \omega_i q_{ijkt} \ . \tag{8}$$

Our microeconomic variance decomposition is achieved by taking the covariance of the variables on each side of this expression with respect to $q_{jk,t}$ and dividing all terms on each side of the equation by the variance of $q_{jk,t}$:

$$1 = \frac{cov(q_{jkt}, q_{jkt})}{var(q_{jkt})} = \sum_i \omega_i \frac{cov(q_{ijkt}, q_{jkt})}{var(q_{jkt})} = \sum_i \omega_i \beta_{ijk} \ , \tag{9}$$

where

$$\beta_{ijk} = \frac{cov(q_{ijkt}, q_{jkt})}{var(q_{jkt})} = \frac{std\,(q_{ijkt})}{std\,(q_{jkt})} \times corr\,(q_{ijkt}, q_{jk,t}) \ . \tag{10}$$

The contribution of good i to the variance of the aggregate RER is given by $\omega_i \beta_{ijk}$. It is increasing in that good's weight in expenditure, the relative standard deviation of its real exchange rate (relative to the aggregate) and its

---

[2]Note this is also the definition for the aggregate real exchange rate for common CES preferences, up to a first-order approximation.

correlation with the aggregate RER. Quite apart from economizing on degrees of freedom in estimating a variance decomposition, the approach recognizes that we are interested in the covariance of the LOP deviations with the aggregate RER.

To fix ideas, suppose all prices are fixed in local currency units during the period and then adjusted to satisfy the LOP at the end of the period, with a nominal exchange rate change occurring during the period. Every single good in the distribution would contribute exactly the same amount to the variance of the aggregate RER, $\beta_{ijk} = 1$. Suppose instead that all traded goods adjusted instantaneously to the nominal exchange rate movement within the period while non-traded goods took one period to adjust. Now non-traded goods account for all of the variance and traded goods for none, $\beta^N = \omega^{-1}$ and $\beta^T = 0$, where ω is the share of expenditure on non-traded goods.

The first example characterizes the view that all goods markets are equally segmented, that goods are all alike, at least for the issue of understanding real exchange rates. The second example characterizes thrust of the classical dichotomy, there are just two types of goods. The first view produces a degenerate distribution of the microeconomic β's at 1, the second produces two degenerate distributions, one for traded goods at $\beta^T = 0$ and one for non-traded goods at $\beta^N = \omega^{-1} = 1.7$ (using a non-traded expenditure share of 0.6, appropriate for our micro-data).

The obvious question to ask is: what does the distribution of β_{ijk} look like? Figure 1 presents three kernel density estimates: one pools all goods (black line), one pools traded goods (red line) and one pools non-traded goods (blue line). The vertical lines display their averages. This distribution has little resemblance to either of the two views described above. There is far too much variation in the β's to be consistent with the broad-brushed view that goods markets are equally segmented internationally, the support of the distribution extends from -2 to +4. At the same time, the distribution exhibits too much central tendency toward its mean of 0.81 to be consistent with a dichotomous classification of final goods. If the classical dichotomy were to hold in the micro-data, the pooled density should be bimodal with a proportion of the data corresponding to traded goods centered at zero (no deviations) and the remaining proportion centered at 1.7. In fact, traded goods are centered at

0.76 and non-traded goods are centered at 1.03.

Table 1 reports summary statistics for the microeconomic variance decomposition. The mean beta for non-traded goods does exceed the mean for traded goods in most cases, ranging from a difference of 0.27 (1.03-0.76) for all cities pooled together (Figure 1) to a low of 0.04 (0.87-0.83) for U.S.-Canada city pairs. The relative standard deviation of the LOP deviations average twice that of the aggregate real exchange rate, indicative of considerable idiosyncratic variation in LOP deviations. The mean correlation of LOP deviation and PPP deviation is 0.45 in the pooled sample. As Crucini and Telmer (2011) note, LOP deviations are not driven by a common factor such as the nominal exchange rate, much of the variation is idiosyncratic to the good.

In summary the contribution of individual goods to aggregate real exchange rate variability shows a central tendency, but with considerable variation across individual goods. Certainly the distribution is not the stark bimodality expected from the classical dichotomy applied to final goods. Our goal is to maintain the two-factor parsimony of the classical dichotomy, but with the tradability applied at the level of inputs. To accomplish this we first elaborate a simple two factor model that stands in for the two types of inputs. Importantly, the share of non-traded and traded inputs in the cost of the final good is assumed to vary across individual final goods as measured by the distribution margin.

4   The Intermediate Inputs Model

Many researchers have argued that the classical dichotomy is more appropriate to apply at the level of inputs than at the level of final goods. Up until quite recently the data has not been available to conduct a systematic investigation of this hypothesis. We follow Engel and Rogers (1996) and Crucini, Telmer and Zachariadis (2005), and assume that retail prices are Cobb-Douglas aggregates of a non-traded input W_{jt} and a traded input inclusive of a transportation cost from the source to the destination, T_{ijt}:

$$P_{ijt} = W_{jt}^{\alpha_i} T_{ijt}^{1-\alpha_i} \ . \tag{11}$$

The LOP deviation (in logs) becomes,

$$q_{ijkt} = \alpha_i w_{jkt} + (1 - \alpha_i)\tau_{ijk,t} ,$$ (12)

where each of the variables is now the logarithm of a relative price across a bilateral pair of cities. Thus the LOP deviation for good i, across bilateral city pair, j and k, depends on the deviation of non-traded and traded input costs across that pair of cities, weighted by their respective cost shares.

Elaborating on the cost structure of individual goods and services in this way adds an additional layer to the original variance decomposition. The betas for the individual retail prices of final goods may now be expressed as a simple weighted average of the underlying betas for non-traded and traded input prices:

$$\frac{cov(q_{jkt}, q_{ijkt})}{var(q_{jkt})} = \alpha_i \frac{cov(q_{jkt}, w_{jkt})}{var(q_{jkt})} + (1 - \alpha_i)\frac{cov(q_{jkt}, \tau_{ijk,t})}{var(q_{jkt})}$$ (13)

$$\beta_{ijk} = \alpha_i \beta_{jk}^w + (1 - \alpha_i)\beta_{ijk}^\tau$$ (14)

$$= \beta_{ijk}^\tau + \alpha_i(\beta_{jk}^w - \beta_{ijk}^\tau) .$$ (15)

This equation leads to two important insights. First, the β_{ijk} for final goods are predicted to be increasing in the share of non-traded inputs α_i, provided non-traded factor prices contribute more to RER volatility than do traded factor prices, $(\beta_{jk}^w - \beta_{jk}^\tau) > 0$. Note that this is a much weaker condition than the classical dichotomy where the relative prices of traded goods are assumed not vary at all across locations ($\beta_{jk}^\tau = 0$), in which case the model would reduce to $\beta_{ijk} = \alpha_i \beta_{jk}^w$. Second, even if the classical dichotomy holds at the level of traded inputs, it will not hold at the level of final goods since, $\alpha_i \beta_{jk}^w > 0$. Ironically, the anecdotes that are often drawn into the debate are precisely the ones that elucidate the role of traded and non-traded inputs. Namely goods at the extremes of the distribution in terms of high and low values of α_i. Of course anecdotes are misleading unless they help us explain the broader patterns in the data and aggregate RER variability, which is our focus.

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). Two items

toward the extremes of the distribution share are explicitly labelled: 1 liter of gasoline and a 2-bedroom apartment. Based on our reconciliation of the EIU micro-data with the U.S. NIPA data on the distribution share, the distribution share for gasoline is 0.19 while that of a 2-bedroom apartment is 0.93.

Three observations are immediate. First, there is a positive relationship between a final good's contribution to RER variability and its distribution share, the correlation of β_i and α_i is 0.69. Second, goods at the extremes such as fuel and shelter, which are often used to provide anecdotal evidence of traded and non-traded goods, fit the classical dichotomy more closely than goods toward the middle of the distribution, goods with an average distribution share. Third, averaging across goods obscures the role of tradability of intermediate inputs because the median good has a distribution share of 0.41, implying close to equal shares of traded and non-traded inputs in the cost of production. Put differently–through the lens of the intermediate input model–examining the median good is analogous to taking a simple average of fuel and shelter. Doing so averages away the differences in the underlying cost structure of the two goods. In the next section we develop a two-factor model to infer the role of traded and non-traded inputs across the entire distribution of the micro-data.

4.1   Two-Factor Model

The objective of this section is to decompose the good-specific contributions to aggregate RER variation into the role of traded and non-traded inputs used in the production of each good. To accomplish this, we incorporate the fact that the cost of producing final goods involves different shares of non-traded and traded inputs, the α_i parameters measured along the horizontal axis of Figure 2. This answers the question: if the contribution of LOP variation in fuel to the aggregate RER is 5%, how much of this contribution to variance is coming from the traded inputs (gasoline) and how much is coming from the non-traded inputs (the other costs associated with operating a gas station). To achieve this, we estimate a two-factor model of the β_{ijk} for each bilateral city pair. These two factors, one for the non-traded input and one for the traded input will later be aggregated back up to the level of the CPI to determine how much of the variation in the aggregate RER is due to variation in RER

for non-traded and traded input costs.

To reduce the intermediate inputs model to a two-factor structure for each bilateral city pair, the traded factor is assumed to be the sum of a component common to all goods and an idiosyncratic component specific to the good:

$$\beta_{ijk}^{\tau} = \beta_{jk}^{\tau} + \nu_{ijk} \ . \tag{16}$$

The contribution of good i to the variation of the bilateral real exchange rate across city pair j and k is now:

$$\beta_{ijk} = \alpha_i \beta_{jk}^{w} + (1 - \alpha_i)\beta_{jk}^{\tau} + \epsilon_{ijk} \ , \tag{17}$$

where $\epsilon_{ijk} = (1 - \alpha_i)\nu_{ijk}$. In the language factor models, the β_{jk}^{w} and β_{jk}^{τ} are the two factors and α_i and $(1 - \alpha_i)$, their respective factor-loadings.

4.2   Estimation

In the model of the previous section, the observables are the estimated betas, β_{ijk}, and the distribution shares from the NIPA, α_i; the unobservables are the two factors of interest, β_{jk}^{w} and β_{jk}^{τ}. Consider the following linear regression model:

$$\beta_{ijk} = a_{jk} + b_{jk}\alpha_i + \epsilon_{ijk} \ . \tag{18}$$

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

$$\beta_{jk}^{\tau} = a_{jk} \tag{19}$$
$$\beta_{jk}^{w} = a_{jk} + b_{jk} \ . \tag{20}$$

We perform this regression separately for each city pair using the β_{ijk} estimated from the expenditure-weighted version of the aggregate RER to conform with the existing macroeconomic literature. 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 shares. Following this aggregation, equation (18) is estimated by Ordinary Least Squares (OLS) to recover the non-traded and traded factors. We also report results obtained by Weighted Least Squares 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.

Table 2 reports the estimated factors averaged across city pairs within different country groups. The standard deviations across city pairs are reported 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 non-traded factor averages 1.03. 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 0.76 and 1.03. Notice that the traded input factor is much lower than the average contribution of a traded good to aggregate real exchange rate variability while the non-traded factor is coincidentally equal to the average contribution of a non-traded good to aggregate real exchange rate variability. This reflects two interacting effects. First, the non-traded factor is the dominant source of variation. Second, the average traded good has far more non-traded factor input content than the average non-traded good. Thus, most of the bias in attributing non-traded factor content in the decomposition is found in 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 factor at the microeconomic level rather than average across goods as Table 2 does. We turn to this level of detail next.

4.3 The Role of Distribution Margins

Recall that after averaging the estimated equation (18) across jk pairs, we arrive at a decomposition of our original good-level betas:

$$\beta_i \quad = \quad 1.03\alpha_i + 0.54(1 - \alpha_i) + \epsilon_i \tag{21}$$

$$= \quad 0.54 + 0.50\alpha_i + \epsilon_i \ . \tag{22}$$

Simply put, a purely traded good is one which involves no non-traded inputs, $\alpha_i = 0$. If such a good existed in the retail basket, 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 1.04

times its expenditure share to aggregate RER variability.

Table 3 shows the entire cross-sectional distribution of the good-specific contributions to real exchange rate variation, β_i, decomposed in this manner. Goods are ordered from those with the lowest distribution share (0.17), an example of which is a 'compact car,' to goods with the highest distribution share (1.00), an example of which is the 'hourly rate for domestic cleaning help.' Note that each row is an average across goods sharing the same distribution share (the second column) and the first column is just an example of a good found in that sector.

Since the non-traded input beta, β_{jk}^w, average 1.03, 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 are about 70% traded inputs and 30% non-traded inputs whereas the latter two are largely driven by the non-traded input factor. Another interesting comparison is fresh fish and a two-course meal at a restaurant. Both are treated as traded goods when CPI data are used because they fall into the same 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 as food at home and food away from home? 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 dominate role of traded inputs and the relatively moderate contribution to aggregate real exchange rate variation (0.65). A restaurant meal is dominated by the non-traded factor (85 percent) and contributes 35% more to aggregate real exchange rate variability than does fresh fish.

A good with a median distribution share (0.41) is toothpaste. Despite the fact that the cost of producing this good is skewed moderately toward traded inputs (59% traded inputs), non-traded inputs still dominate in accounting for the toothpaste beta, 0.40 versus 0.31 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, 1.04 versus

16

0.54. Stated differently, for the traded input factor to dominate in contribution to variance requires a distribution share of less than 0.34 (i.e. a traded input share of more than 0.66).

4.4 The Role of Location

When focusing on the role of the distribution margin, it was productive to average across bilateral pairs. Similarly, when focusing on the role of location, it is useful to average across goods. Recall, however, that the two estimated factors are location-specific and the group means of Table 2 suggested the presence of variation across location pairs in the two factors. To better visualize the full extent of the variation without presuming a source of the variation across city-pairs, Figure 3 presents kernel density estimates of the non-traded and traded factors.

The figure effectively convey three messages. The first, and central message, is that there is a strong central tendency toward the means initially reported in Table 2 (for both of the factors), supportive of the parsimony imposed by the two factor model. The second message is that the contributions of traded and non-traded inputs to aggregate RER variability are much more easily distinguished than was true of traded and non-traded goods. This is evident in comparing the two distributions in Figure 3 with their counterparts in Figure 1. Third, the dispersion across locations in the estimated factors is significant and greater for the estimated traded factor (red line) than the estimated non-traded factor (blue line), consistent with the impression conveyed by the group-mean coefficients – reported in Table 2.[3]

What is responsible for the variation across location pairs in these distributions? Figure 4 plots the non-traded and traded input betas for each bilateral city pair against the standard deviation of the nominal exchange rate for that bilateral city pair. Two features are notable. First, there is a positive and possibly concave relationship between the variability of the nominal exchange rate and both the non-traded and traded input betas in the variance decom-

---

[3]With regard to the classical dichotomy, we could not reject the hypothesis that the traded-inputs beta is 0 for 56% of the sample. This statistics was computed using a two-tailed t-distribution with 95% confidence intervals.

position. Second, the estimated non-traded input betas lie above the traded input betas.

The positive correlation between the estimated input betas and the volatility of the nominal exchange rate is more subtle. It is important to emphasize that this correlation is not simply a reflection of the positive covariance of nominal and real exchange rates documented in the existing macroeconomics literature. Recall, the betas are components of a variance decomposition and thus have already been normalized by the level of the variance of the bilateral real exchange rate. What is happening as the nominal exchange rate variance increases is that the common source of LOP variation is rising relative to the idiosyncratic sources of variation. This occurs because movements in the nominal exchange rate are almost by definition a common source of variability in LOP deviations.[4] Since the non-traded factor is expected to exhibit less pass-through of nominal exchange rates to local currency prices, at least in the short run, it is also expected that the non-traded factor will lie uniformly above the traded factor. This is evident, the blue dots lie mostly above the red dots at each point along the x-axis (i.e. conditional on a value for the nominal exchange rate variance). This is not to say that changes in nominal exchange rates are causing real exchange rates to vary, to identify the underlying causes of variation would require a richer model. For example, a monetary shock is likely to alter both the distribution of local currency prices and the nominal exchange rate whereas a crop failure in a particular country is unlikely to do either of these things. The thrust of the figure, however, is that real and nominal sources of business cycle variation are likely to play different roles in determining the traded and non-traded input betas we have estimated.

To summarize, we have demonstrated that the classical dichotomy is a very useful theory of the LOP when the theory is applied to inputs. What we do in the next section is show that this is also true at the aggregate level. Moreover,

---

[4]Crucini and Telmer (2011) decompose the variance of the LOP deviations into time series variation and long-run price dispersion using the same data employed here. They find, as Engel did, that the time series variability of real exchange rates of traded goods is comparable to that of non-traded goods whereas the long-run price dispersion goes in the direction of the classical dichotomy with more international price dispersion among non-traded goods. Thus, our paper seeks to resolve the more puzzling feature of the data, its time series properties.

we also show that the EIU data are entirely consistent with the conclusions of the existing literature using aggregative CPI data when the theory is applied to final goods. Our interpretation, however, is very different.

5   Macroeconomic Decompositions

Macroeconomics is, of course, about aggregate variables. Our thesis is that if given the choice, macroeconomists would want to aggregate final goods based on their non-traded and traded factor content. Our methodology attempts to provide that choice. Here, we demonstrate the importance of this choice.

5.1   Aggregation Based on Intermediate Inputs

Recall that the microeconomic variance decomposition of the aggregate real exchange rate based on final goods is:

$$1 = \sum_i \omega_i \beta_{ijk} \tag{23}$$

$$\beta_{ijk} = \frac{cov(q_{ijkt}, q_{jkt})}{var(q_{jkt})} = \frac{std(q_{ijkt})}{std(q_{jkt})} \times corr(q_{ijkt}, q_{jk,t}) \ . \tag{24}$$

Substituting our two-factor model for the LOP deviation, $\beta_{ijk} = \alpha_i \beta_{jk}^w + (1 - \alpha_i)\beta_{jk}^\tau + \epsilon_{ijk}$, into this equation gives the theoretically appropriate method of aggregating the micro-data based on the model of intermediate inputs:

$$1 = \sum_i \omega_i \left[\alpha_i \beta_{jk}^w + (1 - \alpha_i)\beta_{jk}^\tau + \epsilon_{ijk} \right] \ . \tag{25}$$

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

$$1 = \pi \beta_{jk}^w + (1 - \pi) \beta_{jk}^\tau + \eta_{jk} \ , \tag{26}$$

where the weights on the traded and non-traded input factors, π and $(1 - \pi)$ are consumption expenditure-weighted averages of the shares of non-traded and traded inputs into each individual good in the consumption basket. The residual term, η_{jk} is an expenditure-share weighted average of the ϵ_{ijk}.

19

In other words: the variance of the aggregate real exchange rate may be expressed as a weighted average of the two factors estimated from the micro-data. The weight on each factor depends on the relationship between taste parameters and relative prices that determine consumption expenditure shares and production parameters, and relative factor prices that determine distribution shares. Recall that the median distribution share in the micro-data is 0.41. The weight on the non-traded factor, π, turns out to be much greater than this average because consumption tends to be skewed toward services which are intensive in distribution inputs. Using US NIPA data and the EIU micro-sample, $\pi = 0.69$. The dominant weight on the non-traded input factor, combined with the fact that β_{jk}^w is about twice the magnitude of β_{jk}^τ is the reason that non-traded inputs dominate the variance decomposition of the aggregate real exchange rate by a very large margin.

Table 4 shows just how large. The table reports the results using OLS estimates (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 real exchange rate, while traded inputs account for the remaining 19%. The contribution of non-traded and traded inputs is moderately more balanced in the U.S.-Canada sub-sample, with non-traded inputs accounting for 73% and traded inputs accounting for the remaining 27%. Consistent with our earlier microeconomic decompositions, the OECD looks more like the U.S.-Canada sub-sample than does the non-OECD group.

5.2   Aggregation Based on Final Goods

An alternative two-factor macroeconomic model of the real exchange rate is to apply 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, for example, to 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 share exceeding 60 percent. This cutoff corresponds to a jump in the value of the distribution shares across sectors from 0.59 to 0.75 (see Table 3 or Figure 2). Coincidentally, this categorization matches up very well with the categorical assignments used by Engel (1999) who used much more aggregated 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 real exchange rate is:

$$q_{jk,t} = \omega q_{jkt}^N + (1 - \omega) q_{jkt}^T . \tag{27}$$

where q_{jkt}^N and q_{jkt}^T are the bilateral real exchange rates for non-traded final goods and traded final goods built from the LOP deviations in the microeconomic data, weighted by their individual expenditure shares.[5]

The variance decomposition of the aggregate real exchange rate is conducted using our beta method[6]:

$$1 = \omega \beta_{jk}^N + (1 - \omega) \beta_{jk}^T . \tag{28}$$

Table 5 reports the outcome of the variance decomposition arising from this macroeconomic approach. It is instructive to compare Table 5 to 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 location is 1.17 (0.78) using the two index construct (Table 5) compared to 1.03 (0.76) using the microeconomic decomposition. These are relatively small

---

[5]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.

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

differences. The underlying sources of the contribution to variance, however, are different.

When using the macroeconomic approach, the non-traded real exchange rate contributes more to the variability of the aggregate RER for two reasons. First, the non-traded real exchange rate is more highly correlated with the aggregate real exchange rate than is the traded real exchange rate (0.96 versus 0.86). Reinforcing this effect is the fact that the non-traded sub-index of the CPI is more variable than the traded real exchange rate (1.22 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 real exchange rate. Consistent with the macroeconomic approach, the LOP deviations of the median non-traded good has a higher correlation with the aggregate real exchange rate than does the median traded goods (0.55 versus 0.42). Thus traded goods have more idioynscratic sources of deviations from the LOP than do non-traded goods.

5.3   The Role of Location

Location also plays a role in determining the relative importance of traded and non-traded goods in accounting for real exchange rate variability. Figure 5 presents the entire distribution of the β_{jk} for the case in which all locations are averaged (the first column of Table 5). The means for all goods, non-traded goods and traded goods from Table 5 are indicated with vertical lines at 0.97, 1.17 and 0.78, respectively. We see considerable variation in the relative importance of non-traded and traded sub-indices of the CPI across bilateral city-pairs, but the two distributions clearly have a different first moment, which was not at all obvious in the microeconomic distributions of Figure 1.

5.4   The Compositional Bias

To further clarify the difference between the implications of the classical dichotomy applied to intermediate inputs and final goods, this section estimates the compositional bias arising from using final goods to infer the factor content of trade at the level of the two sub-index deconstruction of the aggregate real

exchange rate. Consider the traded and non-traded partition based on final goods and how the variance decompositions relate across the two methods.

The contribution of the non-traded aggregate real exchange rate to aggregate RER variability is:

$$\omega \beta_{jk}^N = \sum_{i \in N} \omega_i \left[\alpha_i \beta_{jk}^w + (1 - \alpha_i) \beta_{jk}^\tau \right] , \tag{29}$$

while the contribution of the traded aggregate real exchange rate to the aggregate RER variability is:

$$(1 - \omega)\beta_{jk}^T = \sum_{i \in T} \omega_i \left[\alpha_i \beta_{jk}^w + (1 - \alpha_i) \beta_{jk}^\tau \right] . \tag{30}$$

Note that each contribution is written on the right-hand-side in terms of the intermediate input model.

Recall, the contribution of non-traded inputs to the variation in the aggregate RER variance according to the intermediate inputs model is actually:

$$\pi \beta_{jk}^w = \left(\sum_i \omega_i \alpha_i \right) \beta_{jk}^w , \tag{31}$$

while the counterpart for the contribution of traded inputs is:

$$(1 - \pi) \beta_{jk}^\tau = \left(\sum_i \omega_i (1 - \alpha_i) \right) \beta_{jk}^\tau . \tag{32}$$

Using these four equations, the bias in the estimate of non-traded inputs to aggregate RER variance arising from using the final goods definition is

$$\pi \beta_{jk}^w - \omega \beta_{jk}^N = \left(\sum_i \omega_i \alpha_i \right) \beta_{jk}^w - \sum_{i \in N} \omega_i \left[\alpha_i \beta_{jk}^w + (1 - \alpha_i) \beta_{jk}^\tau \right] \tag{33}$$

$$= \left(\sum_{i \in T} \omega_i \alpha_i \right) \beta_{jk}^w - \left(\sum_{i \in N} \omega_i (1 - \alpha_i) \right) \beta_{jk}^\tau . \tag{34}$$

and likewise, the bias for traded goods is

$$(1 - \pi) \beta_{jk}^\tau - (1 - \omega_{jk})\beta_{jk}^T = \left(\sum_{i \in N} \omega_{ijk} (1 - \alpha_i) \right) \beta_{jk}^\tau - \left(\sum_{i \in T} \omega_{ijk} \alpha_i \right) \beta_{jk}^w . \tag{35}$$

Table 6 reports the average share of non-traded goods together with the traded and non-traded basket average contributions to RER volatility when all city pairs are used in the comparison. The aggregate contribution of traded goods using aggregation to final goods is about twice the their true underlying contribution based on the intermediate inputs model, 0.34 versus 0.17.

23

5.5 Relation with the Existing Literature

Engel (1999) and other researchers focus on whether RER fluctuations are primarily associated with movement in the relative price of tradable goods across countries or with movements in the relative price of non-tradable to tradable goods. The equation Engel works with is

$$q_t = q_t^T + (1 - \omega)(q_t^N - q_t^T) , \tag{36}$$

while we work with

$$q_t = \omega q_t^T + (1 - \omega)q_t^N . \tag{37}$$

Simple algebra allows us to express Engel's variance decomposition in terms of betas for traded and non-traded baskets:

$$1 = \beta^T + (1 - \omega)(\beta^N - \beta^T) . \tag{38}$$

Engel's variance decomposition split the RER volatility into two components. The first component is the volatility in the traded basket RER. The second component is the variance between non-traded and traded basket prices. The approximate equality in the average betas for non-traded and traded baskets explains why this decomposition attributes almost all of the variance in RER to traded goods: Since the baskets' $\beta^N - \beta^T$ is about 0.39 (using expenditure weights, 0.13 using equal weights) and $1 - \omega$ is less than one, β^T must be close to 1 as Engel's reports. This implies that the remaining volatility must be explained by the traded basket RER. Since β^T is 0.78, 78 percent of RER fluctuations are attributable to movements in the relative price of traded goods.[7]

As we saw in the previous sub-section, an important bias arise when one use traded and non-traded baskets' contributions. Using our two factor model, Engel's decomposition becomes:

$$1 = \beta^\tau + \pi(\beta^w - \beta^\tau) + \varepsilon_{jk} . \tag{39}$$

Since β^τ is 0.54, this implies that 54 percent of RER fluctuations are attributable to movements in the relative price of traded intermediates.[8] From this

---

[7]Using US-CA pairs, 88 percent of RER fluctuations are attributable to movements in the relative price of traded goods compared to 95 percent as reported in Engel's work.

[8]For US-CA pairs, this number rises to 67 percent.

point of view, a considerable amount of volatility is coming from the relative price of non-tradable to tradable intermediates.

A number of researchers modify Engel's decomposition, but largely reinforce his conclusions about traded good prices. Out of concern about the functional form of the Cobb-Douglas aggregates, Betts and Kehoe (2006) work with:

$$q_t = q_t^T + (q_t - q_t^T) , \tag{40}$$

where q_t^T is a producer price index which is not contaminated by non-traded distribution services. This decomposition preserves the identity on the left and right-hand-side of the equality as necessary for a variance decomposition and has the desired attributes of using the producer price index, which is arguably a better proxy for traded goods prices than is an aggregate of consumer prices across highly traded goods. As is also apparent, there is no need to assign expenditure weights in the decomposition. Again, using our variance decomposition the beta-representation of the decomposition is:

$$1 = \beta^\tau + (1 - \beta^\tau) + \varepsilon_{jk} . \tag{41}$$

The variance to be explained is the same as in Engel's original contribution, namely the variance of the aggregate CPI-based real exchange rate. However, the variance of the traded goods prices is different, since producer prices replace the traded-CPI component on the right-hand-side. As one might expect, the variance of producer prices is higher than consumer prices, but the covariance of producer prices with the aggregate CPI may be lower. Since the beta is rising in relative variability and covariance, the implication of replacing a consumer price index for traded goods with a producer price index is expected to be ambiguous. These nuances notwithstanding, Betts and Kehoe find a modest reduction in the contribution of traded goods relative to Engel.

Parsley and Popper (2009) take a microeconomic approach using 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: 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 is quite comparable in many ways to the EIU data. 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. The aggregate real exchange rate becomes:

$$q_t = \sum_i \omega_i q_{i,t} \ , \tag{42}$$

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:

$$q_t = q_{i,t} + \left(\sum_g^M \omega_g q_{g,t} - q_{i,t} \right) \ . \tag{43}$$

Parsley and Popper then compute the variance of the lead term, the LOP variance and divide it by the total variance of the real exchange rate and define this ratio as the contribution of good i to the variance of the aggregate real exchange rate.

In terms of betas, their variance decomposition is:

$$1 = \beta_i + (1 - \beta_i) \ . \tag{44}$$

This is because the expenditure weighted average of the betas must equal unity by construction. However, the variance decomposition following our method is:

$$1 = \omega_i \beta_i + \sum_{g \neq i} \omega_g \beta_g \ , \tag{45}$$

As is evident, the good-specific variance contributions of Parsley and Popper's 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 5 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 real exchange rate 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 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 is evident. 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 as well as an important role of LOP deviations at the level of trade. This points to the need for a hybrid model with a distribution sector and segmentation at the level of traded inputs. In arguing for one stripped down model or another, macroeconomists may unwittingly reject virtually all useful theories of international price determination.

References

[1] **Betts, Caroline and Timothy Kehoe**. 2006. US Real Exchange Rate Fluctuations and Relative Price Fluctuations. Journal of Monetary Economics, 53(7): 1257-1326.

[2] **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.

[3] **Burstein, Ariel, Martin Eichenbaum and Sergio Rebelo.** 2005. Large Devaluation and the Real Exchange Rate. Journal of Political Economy, 113(4): 742-784.

[4] **Campa, José Manuel and Linda S. Goldberg**. 2005. Exchange Rate Pass-through into Import Prices. The Review of Economics and Statistics, 87(4): 679-690.

[5] **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.

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

[7] **Crucini, Mario J., Chris I. Telmer**. 2011. Microeconomic Sources of Real Exchange Rate Variability. Manuscript.

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

[9] **Engel, Charles**. 1999. Accounting for U.S. Real Exchange Rates. Journal of Political Economy, 107(3), 507-538.

[10] **Mussa, Michael.** 1986. Nominal Exchange Regimes and the Behavior of Real Exchange Rates: Evidence and Implications. Carnegie-Rochester Conference on Public Policy, 25: 117-213.

[11] **Popper, Helen and David C. Parsley**. 2006. Understanding Real Exchange Rate Movements with Trade in Intermediate Products. Pacific Economic Review, 15(2): 171-188.

[12] **Salter, W.E.G.** 1959. Internal and External Balance: The Role of Price and Expenditure Effects. Economic Record, 35: 226-38.

[13] **Swan, Trevor W.** 1960. Economic Control in a Dependent Economy, 36: 51-66.

Table 1: Variance decomposition of real exchange rates, microeconomic approach

	All pairs	OECD	Non-OECD	US-Canada
Std. dev. RER	0.14	0.13	0.13	0.10
Number of city pairs	4835	1543	856	52
Non-traded weight	0.57	0.56	0.58	0.55
Traded weight	0.43	0.44	0.42	0.45
All goods				
Beta	0.81	0.84	0.74	0.83
Correlation	0.45	0.43	0.42	0.38
Rel. std. dev. LOP	2.00	2.18	1.98	2.19
Non-traded goods				
Beta	1.03	0.95	1.11	0.87
Correlation	0.55	0.50	0.57	0.46
Rel. std. dev. LOP	2.00	2.10	2.07	1.89
Traded goods				
Beta	0.76	0.81	0.63	0.83
Correlation	0.42	0.42	0.37	0.36
Rel. std. dev. LOP	2.00	2.20	1.95	2.26

**Table 2: Traded and non-traded inputs regressions,
international pairs**

	All pairs	OECD	Non-OECD	US-Canada
beta (traded)	0.54	0.65	0.34	0.67
	(.49)	(.46)	(.52)	(.52)
beta (non-traded)	1.03	0.96	1.12	0.82
	(.34)	(.32)	(.37)	(.26)
slope	0.50	0.30	0.78	0.15
	(.72)	(.64)	(.79)	(.64)
R-squared	0.08	0.05	0.12	0.04
Number of pairs	4835	1543	856	52

Note: Minimum of 4 observations per city pair.

Table 3: Variance decomposition using intermediate inputs betas

| | | Contribution | | | Non-traded |
Example	α_i	Non-Traded	Traded	Residual	Cont. (%)
Compact car (1300-1799 cc)	0.17	0.17	0.44	0.07	28%
Unleaded gasoline (1 liter)	0.19	0.19	0.43	-0.02	30%
Fresh fish (1 kg)	0.22	0.21	0.44	0.07	32%
Time (news magazine)	0.32	0.32	0.36	0.05	47%
Toilet tissue (two rolls)	0.34	0.34	0.34	0.00	50%
Butter (500 g)	0.36	0.36	0.33	-0.01	52%
Aspirin (100 tablets)	0.37	0.36	0.34	0.01	51%
Marlboro cigarettes (pack of 20)	0.37	0.37	0.33	0.06	53%
Electric toaster	0.39	0.38	0.33	0.02	54%
Toothpaste with fluoride (120 g)	0.41	0.40	0.31	-0.05	57%
Compact disc album	0.41	0.41	0.30	-0.05	57%
Insect-killer spray (330g)	0.45	0.45	0.27	-0.03	62%
Paperback novel	0.49	0.47	0.27	-0.03	63%
Razor blades (5 pieces)	0.49	0.49	0.27	-0.01	64%
Batteries (two, size D/LR20)	0.50	0.49	0.27	-0.04	65%
Socks, wool mixture	0.52	0.52	0.25	0.05	68%
Men's shoes, business wear	0.52	0.52	0.25	0.03	67%
Lettuce (one)	0.52	0.52	0.25	0.05	68%
Frying pan (Teflon)	0.53	0.53	0.25	-0.01	68%
Light bulbs (two, 60 watts)	0.57	0.57	0.22	-0.18	72%
Child shoes, sportwear	0.59	0.58	0.21	-0.03	73%
Tennis balls (Dunlop, Wilson or equivalent)	0.59	0.59	0.21	-0.11	74%
Two-course meal at a restaurant (average)	0.75	0.75	0.13	0.00	85%
Electricity, monthly bill (average)	0.76	0.75	0.13	0.04	86%
Man's haircut (tips included)	0.85	0.85	0.08	-0.04	91%
Taxi, airport to city center (average)	0.86	0.85	0.07	-0.01	92%
Telephone line, monthly bill (average)	0.92	0.90	0.04	0.03	96%
2-bedroom apartment	0.93	0.91	0.04	0.24	96%
Annual premium for car insurance	0.94	0.93	0.03	0.01	97%
Hourly rate for domestic cleaning help	1.00	1.00	0.00	-0.08	100%

Table 4: Macroeconomic variance decomposition, intermediate input approach

	All pairs	OECD	Non OECD	US-CA
Non-traded share (π)	0.69	0.69	0.69	0.68
	(.02)	(.20)	(.02)	(.02)
Contribution of				
Traded inputs	0.17	0.20	0.10	0.21
	(.15)	(.14)	(.16)	(.17)
Non-traded inputs	0.71	0.66	0.78	0.56
	(.23)	(.22)	(.25)	(.18)
Error term	0.12	0.14	0.12	0.23
	(.20)	(.22)	(.21)	(.21)
Number of city pairs	4835	1543	856	52

Table 5: Variance decomposition of real exchange rates, macroeconomic approach

	All	OECD	Non-OECD	US-Canada
Std. dev. RER	0.14	0.13	0.13	0.10
Number of city pairs	4835	1543	856	52
Non-traded weight	0.57	0.56	0.58	0.55
Traded weight	0.43	0.44	0.42	0.45
All goods				
Beta	0.97	0.98	0.96	0.99
Correlation	0.91	0.92	0.89	0.94
Rel. std. dev. LOP	1.06	1.06	1.08	1.05
Non-traded goods				
Beta	1.17	1.14	1.23	1.10
Correlation	0.96	0.96	0.96	0.96
Rel. std. dev. LOP	1.22	1.18	1.29	1.15
Traded goods				
Beta	0.78	0.82	0.69	0.88
Correlation	0.86	0.88	0.82	0.92
Rel. std. dev. LOP	0.91	0.94	0.86	0.96

Table 6: Compositional bias

	Aggregate using final goods	Aggregation using Intermediate inputs
Non-traded share	0.57	0.69
Contribution		
Traded	0.34 (34%)	0.17 (19%)
Non-Traded	0.66 (66%)	0.71 (81%)

Figure 1: Density Distributions of Betas, Microeconomic Decomposition

Figure 2: Sectoral Betas and Distribution Shares

Figure 3: Density Distributions of Factor Betas

Figure 4: Factor Betas and Nominal Exchange Rate Volatility

Figure 5: Density Distributions of Betas, Macroeconomic Decomposition