NBER WORKING PAPER SERIES

MICROECONOMIC SOURCES OF REAL EXCHANGE RATE VARIABILITY

Mario J. Crucini
Christopher I. Telmer

Working Paper 17978
http://www.nber.org/papers/w17978

NATIONAL BUREAU OF ECONOMIC RESEARCH
1050 Massachusetts Avenue
Cambridge, MA 02138
April 2012

Financial support of the National Science Foundation, SES #1030164 and SES # 0524868 is gratefully acknowledged. The views expressed herein are those of the authors and do not necessarily reflect the views of the National Bureau of Economic Research.

Microeconomic Sources of Real Exchange Rate Variability
Mario J. Crucini and Christopher I. Telmer
NBER Working Paper No. 17978
April 2012
JEL No. E3,F31,F4

ABSTRACT

We provide three sets of variance decompositions on microeconomic international relative price data. The first shows that the overall distribution of absolute deviations from the Law of One Price (LOP) is dominated by cross-sectional variation in long-term averages, not by time-series variation around the long-term averages. The second shows that time-series variation in changes in LOP deviations is dominated by idiosyncratic, goods-specific variation, not by aggregate variation such as that arising from nominal exchange rates. The third shows that time-series and cross-sectional variance are connected across goods. Goods that exhibit high cross-sectional variance also exhibit high time-series variance. Moreover, when this connection is made conditional on the tradeability of a goods, a two-factor structure for the goods-specific cross-section is revealed. We argue that this factor structure, in addition to our other variance decompositions, is informative for the construction of models that can synthesize the micro and macroeconomic behavior of relative prices.

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

Christopher I. Telmer
Tepper School of Business
Frew and Tech Streets
Carnegie-Mellon University
Pittsburgh, PA 15213
chris.telmer@cmu.edu

1 Introduction

This paper provides some variance decompositions on microeconomic international relative price data that are meant to inform the burgeoning literature bringing together international macroeconomics and trade. International macro typically focuses on the intertemporal dimension of prices. Trade places more weight on the spatial dimension. The intersection between the two (*e.g.*, Burstein and Atkeson (2008), Ghironi and Melitz (2005)) emphasizes both time and space. Our variance decompositions are explicitly designed with this in mind. We parse the total variation in Law-of-One-Price (LOP) deviations into cross-sectional and time series components. The former is the variance in LOP deviations across locations and goods, after averaging away the time series variation. We associate it with the trade literature, which emphasizes imperfect competition and barriers to trade. The latter is the variation around these cross-sectional, long-run averages. We associate it with DGSE models, which emphasize productivity and policy shocks and how they interact with nominal rigidities. Our goal is to better understand each of these dimensions of the data and, in particular, any interactions between them.

Our results boil down to three sets of variance decompositions. What they show is as follows. First, cross-sectional variance in long-term *absolute* deviations from the Law of One Price (LOP) is large relative to time-series variance. Second, time-series variance in *changes* in LOP deviations is dominated by good-specific variation, not country-specific variation such as arises from nominal exchange rates. Third, cross-sectional and time-series variation are *connected* when we look across goods. This connection has several facets. The simplest is *unconditional*: goods that exhibit high cross-sectional variance in long-term, mean LOP deviations also exhibit high time-series variation around these long-term means. A richer pattern arises when we *condition* on the tradeability of goods. Non-traded goods exhibit relatively *high* cross-sectional variation, but relatively *low* time-series variation. We show that this implies the existence of a second source of cross-sectional (across goods) variation, above and beyond tradeability. This second 'factor' generates a positive association between cross-sectional and time-series variation, and it is large in the sense that it clouds the negative relationship associated with tradeability. We argue that these variance decompositions provide valuable information for the construction of models aimed at synthesizing the microeconomic and macroeconomic behavior of international relative prices.

We now elaborate on each set of results in turn. To better understand the first set, consider a specific good: an apple. We first compute the absolute LOP deviation for apples between 123 major cities in the world, for each year between 1990 and 2005. We then compute the time-averaged, long-term LOP deviation — the 'fixed effect' — for each pair of cities. We find that cross-sectional variation in these fixed effects is large relative to time-series variation around them. That is, apples simply tend to be expensive in some cities and cheap in others, this city-specific tendency is stable over time, and the time-variation that does exist is relatively small. Our data include the prices of many other goods and services, not just apples. This relative price behavior is broadly representative of most goods and services in the typical urban consumption basket.

What does this tell us about economic models? Consider sticky-price models, for example. They emphasize frictions in the mechanism through which prices *change*. They often ignore frictions which cause long-term LOP deviations. Our results suggest that what they are ignoring is large. The overall distribution of LOP deviations is dominated by something that sticky-price models ignore, basically by-construction. Is this likely to matter? We think so. It seems likely that whatever are the frictions underlying long-term LOP deviations, these frictions also play a role in governing how the deviations change over time. Much of our paper is dedicated to substantiating this. For now, we simply emphasize an interpretative point. Suppose that one is ultimately interested in the transmission of shocks. We don't deny that at a sufficiently short horizon, goods prices — in particular international goods prices — are sticky, and that understanding the frictions that drive this is important. But nevertheless, a friction that generates a 5 or 10% change in a relative price seems less striking in a world where long-term LOP deviations average 50% than in a world where they average zero.

Our second set of results address real and nominal exchange rate *variability* more directly. To motivate them, consider the Canada-U.S. real exchange rate, depicted in Figure 1. This graph is representative of the common wisdom — often attributed to Mussa (1986) — that real and nominal exchange rates are basically the same thing. This evidence has motivated much economic discussion and model building. It is at the root of the notion that nominal exchange rate variability creates allocative distortions in international consumption and investment decision making. It plays a central role in Rogoff's (1996) 'PPP Puzzle:' the statement that PPP deviations are too large and persistent to be reconciled by some combination of nominal rigidities and real shocks. Finally, it is often used to motivate sticky-price models. A caricature of Figure 1's interpretation in this context is that the world is described by fixed prices in domestic and foreign currency units and that nominal exchange rate variability simply 'shifts around' the entire distribution of individual goods prices.

Figure 1

Canada-U.S. Real and Nominal Exchange Rates

(Annual Data: Indexed 1971=1)

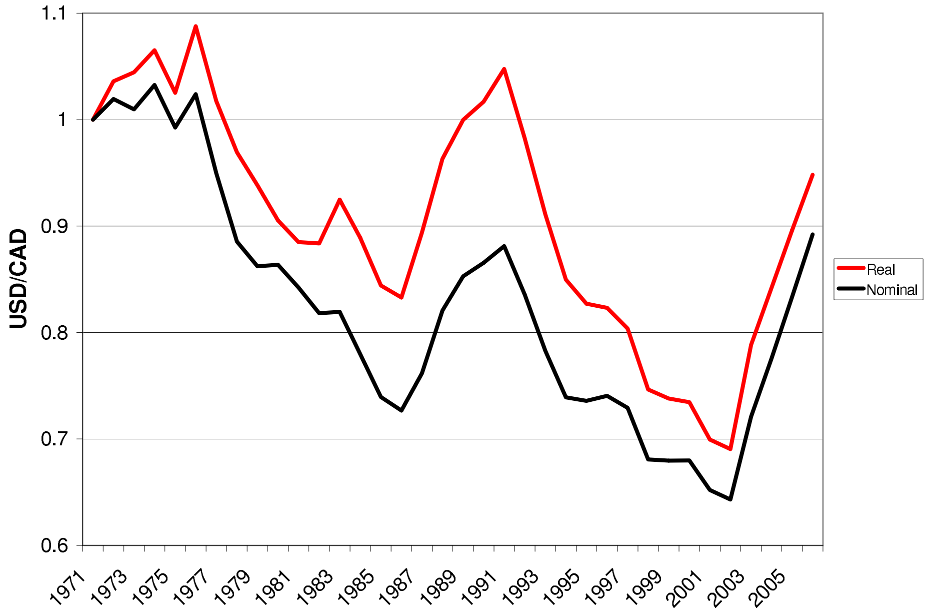

We ask if these types of interpretations are consistent with the behavior of *changes* in microeconomic LOP deviations. Following the above example, we first compute the city-specific relative price of apples, 1990-2005. We do so on a bilateral-pair basis (*e.g.*, the relative price of apples between Pittsburgh and Toronto, 1990-2005). We then compute changes in these relative prices, 1991-2005, and do the same thing for bananas, toaster ovens, haircuts and many other goods and services. Finally, for each bilateral city-pair we decompose the variance of the changes into two orthogonal components: a city-pair-specific component and a good-specific component.

What we find is that the magnitude of the city-specific component is small relative to the good-specific component. This is true for city-pairs which are within a country. It is only slightly less true for for city-pairs that are not. For example, among intra-U.S. city-pairs the average amount of total variation attributable to the city-specific component is 1.8%. The analogous number for U.S.-Canadian city-pairs is 6.9%. One's natural inclination is to attribute the difference to nominal exchange rates. This inclination is correct; the correlation between the U.S.-Canadian city-specific component and changes in the nominal exchange rate is 0.93. The main point, however, is that roughly 93% of the variation in changes in U.S.-Canada LOP deviations has nothing to do with nominal exchange rates. It is specific to the goods themselves.

What does this tell us about economic models? That they place heavy weight on country-specific price shocks (like the nominal exchange rate) at their peril. What about the empirical evidence on aggregate real exchange rates discussed above? Our results *do not* contradict them. When we aggregate across the microeconomic prices our real exchange rate picture is very similar to Figure 1. What our results do contradict, however, are some of the popular interpretations of Figure 1. If you think that real exchange rate variability is driven by nominal exchange rates moving around a distribution of microeconomic sticky prices, you are wrong. There is a great deal of movement *within* the distribution. Microeconomic prices in local currency units move around a lot more than the nominal exchange rate does.

At this point, one might argue that our results are not surprising. Qualitatively, we agree. It is not surprising that long-term deviations from LOP exist between two cities. Retail goods are bundles of traded and non-traded goods and this is exactly what we should expect. Neither is it surprising that changes in good-level relative prices are more variable than changes in aggregate prices. This is an obvious implication of averaging. The importance of our point, then, is quantitative, not qualitative. It is important to know that 2/3 of the total variation in LOP deviations is cross-sectional. It is important to know that changes in nominal exchange rates account for less than 10% of the variation in good-specific LOP deviations. This being said, the final section of our paper goes further, drawing some informative *connections* — both qualitative and quantitative — between the long-term and the short-term, and between the tradeability of goods and whatever else it is that drives price dispersion.

Our third set of variance decompositions articulates these connections. We begin by showing that there exists an unmistakable positive relationship between cross-sectional and time-series variance. That is, if apples are associated with relatively large long-term, mean LOP deviations, then the variation around these means also tends to be relatively large. This is an *unconditional* statement (in the goods dimension). One's first reaction to it, motivated by years of theory and evidence, is to ask whether the tradedness of goods can shed any light. We do so, *conditioning* on tradeability in the cross-section. What we find, at first blush, might seem to contradict the unconditional evidence. Consistent with the standard Balassa-Samuelson (Balassa (1964), Samuelson (1964)) idea, non-traded goods exhibit more variance than traded goods, but this is only the case for cross-sectional variance. In the time-series dimension, traded goods exhibit larger deviations from LOP. Where is the contradiction? If non-traded goods have high cross-sectional variation, and cross-sectional variation is positively associated with time-series variation, then shouldn't non-traded goods also have relatively high time-series variation?

The answer is no. There is no contradiction. Instead, there is a second source of variation across goods, above-and-beyond tradeability. Moreover, this source of variation is not just noise. It must display a systematic pattern. It must (i) move both cross-sectional and time-series dispersion in the same direction, and (ii) be large enough to cloud the negative relationship induced by tradeability. We neither model nor measure (directly) this source of variation. Our objective is simply to provide some broad variance decompositions that can inform structural models as well as further empirical work.

It is worth noting that our results, as they relate to tradeability, evoke those of Engel (1999). Broadly speaking, he showed that, in contrast with the basic Balassa-Samuelson idea, variability in *traded* goods prices seems to play an important, perhaps dominant role in accounting for real exchange behavior. His conclusions were based on index-number data, thus emphasizing time-series variation. Our findings are complementary. We show, while the Balassa-Samuelson, traded/non-traded dichotomy gets it right in the long-run, in the short-run we see the opposite, with time-series variation being larger for traded goods. What is going on? We provide additional empirical evidence suggesting that the *shocks* that affect traded goods prices are larger than those affecting non-traded goods prices. Anecdotally, this just says that the shocks that affect the cost of petroleum products and electronic components are large relative to those affecting haircuts (but, nevertheless, the cross-sectional dispersion in haircut prices is relatively large). Another anecdote is that the Headline CPI Index — which includes the highly tradeable food and energy components — is more variable that the Core CPI Index. In Section 5 go beyond anecdotes and show that our microeconomic data displays such behavior.

The remainder of our paper is organized as follows. Section 2 describes our data and the basic definitions and transformations that we employ. Section 3 presents results based on *absolute* LOP deviations and Section 4 turns to *changes* in LOP deviations. Section 5 provides some interpretation of our basic results and then delves further into the linkages between sources of cross-sectional and time-series variation that we discuss above. Section 6 concludes.

2 Data

We use panel data on individual goods prices that vary across both geographic locations and time. This allows us to study both absolute and differenced LOP deviations. The body of work that exploits such data is by now large and growing. Examples include Alessandria and Kaboski (2011), Bergin and Glick (2007), Broda and Weinstein (2008), Campbell and Lapham (2004), Fischer (2012), Fitzgerald and Haller (2008), Gagnon (2010), Gopinath and Rigobon (2008), Gopinath, Gourinchas, Hsieh, and Li (2011), Gorodnichenko and Tesar (2009), Hellerstein (2008), Hummels (2007), Imbs, Mumtaz, Ravn, and Rey (2005) and Simonovska (2010), as well as a previous paper of our own, Crucini, Telmer, and Zachariadis (2005).

The specific data that we use is as follows. We obtain local-currency retail prices from the *Worldwide Cost of Living Survey* coordinated and compiled by the *Economist Intelligence Unit* (EIU). The target market for this data source are corporations seeking to determine compensation levels for employees residing in different cities around the world. While the goods and services reflect this objective to some extent, the sample is broadly representative of what would appear in the consumption basket of an urban consumer.[1] What makes the data attractive for research purposes is the fact that the prices are in absolute currency units and the survey is conducted by a single agency in a consistent manner over time. It also has a limited *intra*-national dimension, thus providing a useful contrast between domestic and international price dispersion.

[1] Rogers (2002) conducts an extensive comparison between the EIU data and data from national statistical agencies. He finds that the EIU data are broadly representative of what the consumer price index data tell us.

More specifically, the EIU dataset consists of local-currency retail prices, inclusive of sales tax, on as many as 301 goods and services, sampled in 123 cities from 78 different countries. The data are annual, 1990-2005. The country with the most intranational observations is the U.S., with 16, followed by Australia, China and Germany with 5, Canada with 4, Saudi Arabia with 3, and Brazil, France, Italy, Russia, Spain, Switzerland, UK, India, Japan, Vietnam, New Zealand with 2. A number of recent papers have used this data, including Crucini and Shintani (2008), Engel and Rogers (2004), Parsley and Wei (2000) and Rogers (2002).[2].

We denote P_{ijt} as the local-currency price of good i in city j in year t and $S_{jk,t}$ as the date t nominal exchange rate between cities j and k, in units of city k ($S_{jk,t} = 1$ if cities j and k share the same currency). We transform prices into bilateral log deviations from the law-of-one-price (LOP):[3]

$$q_{i,jk,t} = \log(\frac{P_{ijt}S_{jk,t}}{P_{ikt}}) \ , \tag{1}$$

In words, these LOP deviations are the date t (log) prices of good i in city j in units of good i in city k. Appendix A provides additional details, including how we clean the data and deal with missing observations.

Figure 2 shows estimates of the density function for $q_{i,jk,t}$ for 1990, 1995, 2000 and 2005, for both international city-pairs and U.S. city-pairs (the graph is quite similar for intranational pairs more broadly). The graph shows that dispersion in good-by-good LOP deviations is large, and substantially larger once we include a

[2]See http://bertha.tepper.cmu.edu/telmerc/eurostat for a list of all goods-and-services and all cities

[3]A previous version of the paper employed an alternative LOP measure, the log deviation from the cross-city geometric average: $\log(S_{jn,t}P_{ijt})/\sum_{j=1}^{M}\log(S_{jn,t}P_{ijt})$ where M denotes the total number of cities and n denotes the numeraire currency in units of which all prices are expressed (our measures of price dispersion are independent of the choice of the numeraire currency). While this definition results in lower overall LOP variability (by construction), the main message of our paper remains unchanged.

wide array of international location-pairs.

3 Variance in Absolute LOP Deviations

We begin by decomposing the variation in $q_{i,jk,t}$, good-by-good, into a cross-sectional and a time-series component:

$$\begin{aligned} Var_{jk,t}(q_{i,jk,t}\,|\,i) &= Var_{jk}\Big(E_t(q_{i,jk,t}\,|\,i,jk)\Big) + E_{jk}\Big(Var_t(q_{i,jk,t}\,|\,i,jk)\Big) \quad (2) \\ &= T_i + F_i \;\; . \quad (3) \end{aligned}$$

Our notational conventions are slightly non-standard. The conditional mean and variance operators, $E_x\,(\cdot\,|\,y)$ and $Var_x\,(\cdot\,|\,y)$, denote the mean and variance calculated by integrating across the variable(s) x while conditioning on the variable(s) y. So, for instance, $E_t[q_{i,jk,t}\,|\,i,jk]$ is the mean of the time series of relative prices for good i between cities j and k and $Var_{jk}(E_t[q_{i,jk,t}\,|\,i,jk])$ is the cross-sectional variance, across location-pairs, in these time-series means.

To more easily interpret Equation (2), consider its individual pieces. First, $E_t[q_{i,jk,t}\,|\,i,jk]$ is the mean (over time) of the relative cost of good i between cities j and k. If, for example, j = New York and k = Toronto, and if this mean is positive, then good i tends to be more expensive in New York than in Toronto in a long-run sense. The first term in the decomposition, T_i, is the cross-sectional variance — across location-pairs — of these long-run means. It asks "how much of the total variation for good i is due to long-run, city-specific 'fixed effects?' The second term, F_i, captures time-series variation around the long-term means. It is the average (across location-pairs) time-series variance in the absolute LOP deviation for good i. It asks "how much of the total variation for good i is due to *shocks* that die out over time?"

Figure 3 provides an (intentionally stark) illustration. It plots the time series of LOP deviations between North American city-pairs for a typical non-traded good and a typical traded good: haircuts and apples, respectively. The haircut graph is dominated by long-run means. Cross-sectional variance in these means is measured by T_i from Equation (3). Time series variance *around* the long-run means is relatively small and is measured by F_i. Numerically, T_i and F_i are, respectively (for haircuts) 86% and 14% of the total variance. Apples tell quite a different story. Time-series variance plays a much larger role. The analogous numbers are 48% and 52%.

Why is this decomposition an interesting one? Because economic models often make stark assumptions about its components. The archetypical trade model assumes that the differences between home and foreign prices reflect tariff and trade barriers, which vary across goods and locations, $T_i > 0$, but not time, $F_i = 0$. The archetypical business cycle model assumes that unexpected shocks generate transitory fluctuations in international relative prices, $F_i > 0$, away from a steady-state in which the LOP holds, $T_i = 0$. Our notation is chosen with this in mind. The letter T represents 'trade costs and trade theory' and the letter F represents 'frictions, finance, and fluctuations.'

Table 1 moves from the anecdotal examples of Figure 3 to a more systematic examination. It reports the average estimate (averaged across goods i) of T_i and F_i from Equation (3). City pairs separated by an international border are separated from those that are not. Consider first the total variance, $Var_{jk,t}(q_{i,jk,t} \,|\, i)$. Among U.S. cities the estimate is 0.128. Interestingly, the estimate is essentially unchanged for Canada-U.S. city pairs. Once we include all international OECD city-pairs, in contrast, the total variance increases to 0.221. Including all international city pairs (*i.e.*, including non-OECD cities) further increases the variance to 0.275.

What's driving this? One's natural inclination might be to attribute it to variation in nominal exchange rates. The incremental increase in variance going from the U.S. to the OECD to the world is 0.093 and 0.147, respectively. These values are in the same ballpark as that of the variance of changes in nominal exchange rates, averaged across countries, for OECD pairs and world pairs, respectively.[4] So, are nominal exchange rates at the root of increasing LOP variability? The remainder of the table says *no*. The majority of the total variance in LOP deviations — roughly 60% on average — is associated with long-run, good-and-city-specific "fixed effects." Almost by their very nature, these things are unrelated to nominal exchange rate variability combined with sticky prices, the story that motivates this entire line of reasoning.

To summarize, between 60 and 70 percent the total variation in absolute LOP deviations attributable to long-run LOP deviations. The lion's share of what determines the international micro price distribution falls under the realm of trade theory. Relatively little seems (directly) related to 'frictions, finance, and fluctuations' as is discussed above.

3.1 Good versus City Specific Variation

The previous results average across goods. They tell us how important long-run means are for the LOP deviations of the *average good*. A related question decomposes variation in these long-run means into that which is city-pair specific and

[4]Implicit in this comparison is an assumption about the persistence of real exchange rates. In related work we find that the variance of the *absolute level* of the micro LOP deviations is roughly the same as the variance of the *change in the absolute level*. If micro real exchange rates follow an AR(1), this means that the autocorrelation is 0.5. In addition it means that it is coherent to add estimates of the variance of the *change* in nominal exchange rates to estimates of the variance of the *level* of the real exchange rate, as we're doing here.

that which is not. That is, denoting the long-runs means as $\mu_{i,jk} \equiv E_t(q_{i,jk,t} \mid i, jk)$ (*i.e.*, the term in the first large parentheses of Equation (2)), we write

$$\begin{aligned} \mu_{i,jk} &= \mu_{jk} + \varepsilon_{i,jk} \\ \Longrightarrow Var_{i,jk}(\mu_{i,jk}) &= Var_{jk}(\mu_{jk}) + Var_{i,jk}(\varepsilon_{i,jk}) \ . \end{aligned} \tag{4}$$

This decomposition examines the relevance of statements like '*goods are always expensive in New York.*' The previous section showed us that the '*goods are always expensive*' part describes the majority of the overall variation in LOP deviations. Here, we examine the '*in New York*' part. We ask how much of the variation in the long-run means is related to city-wide factors like rent and wages. What's left over is good-specific; '*wine is always cheap in Barcelona, but lots of other goods are not.*'[5]

In Table 2 we see that city effects, μ_{jk}, account for 15 to 20 percent of the total variation in the long-term means. The particular set of city-pairs — ranging from all pairs of U.S. cities to all global pairs — is important for the total variation in $\mu_{i.jk}$, but much less so for the fraction attributable to city effects. The relatively small magnitude of this number seems important to us. It is suggestive of the importance of either good-specific factors — *e.g.*, pricing-to-market being more prevalent for some goods than others — or of good-specific sensitivity to common, city-specific factors. An example of the latter is non-traded input costs (*e.g.*, wages, rent, distribution services) varying across cities, with different goods using different shares of non-traded versus traded inputs. Section 5 elaborates on these themes.

Table 3 conducts a similar decomposition as Equation (4), but with the time-

[5] We do not include a good-specific term, μ_i, in the decomposition (4) because, by construction, the long-term average LOP deviations, $\mu_{i,jk}$, average to zero across jk: $E_{jk}\mu_{i,kj} = 0$

series variances from Equation (2) instead of the time-series means. Denoting $\sigma^2_{i,jk} \equiv Var_t(q_{i,jk,t}\,|\,i,kj)$ (*i.e.*, the term in the second large parentheses of Equation (2)), we decompose the good-city-pair specific time-series variances into three pieces:

$$\sigma^2_{i,jk} = \sigma^2_i + \sigma^2_{jk} + \varepsilon_{i,jk} \quad . \tag{5}$$

The interpretative questions are similar to those above. The city-pair effect, σ^2_{jk} asks '*are there city pairs for which the (absolute) LOP deviations exhibit relatively much (relatively little) time-series volatility?*' If, for example, nominal exchange rate variability is important, this should show up as such a city-pair effect.[6] Similarly, high transport costs for cities on the geographic periphery may translate into wider no-arbitrage regions, more LOP variability, and a larger value for σ^2_{jk}. The second term, σ^2_i, captures good-specific variation. It answers questions like '*are LOP deviations for some goods more volatile than for others?*' A common interpretation of this involves traded versus non-traded goods.

What we see in Table 3 is fairly striking. Good-specific effects are substantially more important than city-specific effects. The unconditional time-series variance of (absolute) LOP deviations seems more affected by things like a good's tradeability than things like nominal exchange rate variability and city-wide productivity shocks. Specifically, for the OECD set of city pairs, the good-specific variation is 17% of the total variation whereas city-specific variation accounts for only 9%. For U.S. cities and Canada/U.S. cities the difference is even larger at 25 or 30%

[6] This exchange-rate interpretation depends critically on the extent to which our data has a good balance of intra versus international city pairs. This varies a lot across the different city-pair sets in Table 3. For 'Canada-U.S.' about 1/3 of all city pairs are separated by a border. This reflects a large number of U.S. cities (16), versus just 4 Canadian cities. For 'OECD' and 'World,' in contrast, borders separate about 90% and 97% of city pairs, respectively. This reflects a large number of countries — 16 out of 78 — for which we only have data on one city per country (see Section 2 for additional details). The effect of nominal exchange rates in Table 3, then, is quite blurred. We rectify this in Section 4 where, instead of decomposing the *variance of the shocks* as in Equation (5), we decompose the *shocks themselves*.

versus 2 or 3%. For all 'World' city pairs, in contrast, things go the other way at 12 versus 15%. In Section 4 we show that the main reason for the latter is nominal exchange rate variability.

4 Variance in Changes in LOP Deviations

We now examine the behavior of *changes* in good-by-good LOP deviations.

$$\Delta q_{i,jk,t} \equiv q_{i,jk,t} - q_{i,jk,t-1} \quad .$$

One advantage of working with first-differences is that our findings do not depend on the precision with which we estimate the long-term LOP means, $\mu_{i,kj}$, that permeate much of the analysis in Section 3 (*i.e.*, the first LHS term in Equation (2)). In addition, much of the existing body of empirical (*e.g.*, Figure 1) work is based on first-differences, something necessitated by using index-number data. What we will see here is that our results are broadly consistent with the existing evidence, while at the same time offering new insights that derive from having absolute prices and a large cross-section.

We begin with a decomposition that is analogous to the time-series versus cross-sectional decomposition from Equations (2) and (3). Conditioning on each city-pair, jk, this is

$$\begin{aligned} Var_{i,t}(\Delta q_{i,jk,t}\,|\,jk) &= Var_t\Big(E_i(\Delta q_{i,jk,t}\,|\,jk,t)\Big) + E_t\Big(Var_i(\Delta q_{i,jk,t}\,|\,jk,t)\Big) \\ &= C_{jk} + I_{jk} \quad . \end{aligned} \tag{6}$$

This decomposition asks a simple question. Consider a city-pair, jk. At each point in time there is distribution of good-by-good LOP deviations. Equation (6) asks "how much of the change (over time) in the LOP deviations represents changes in the *mean* of the distribution versus movement *within* the distribution, around the (time-varying) mean?" More specifically, Equation (6) decomposes the total variance (of the changes) into two pieces: *common across all goods* (C) and *idiosyncratic to each good* (I). The common piece, C_{jk}, is basically real exchange rate variation. That is, $E_i(\Delta q_{i,jk,t} \,|\, jk, t)$ is the date-t cross-good-average for city-pair jk, a close cousin to the change in the CPI-based real exchange rate.[7] Time-series variation in this average, C_{jk}, measures the contribution of *movement in the mean* of the distribution of LOP deviations. The idiosyncratic piece I_{jk}, in contrast, measures the contribution of *movement around the mean* of the distribution. It takes the cross-sectional distribution for jk at date t, asks how *different* the change in good i's LOP deviation is from good l's, and then time-averages the resulting sequence of cross-sectional variances.

A stark example clarifies what Equation (6) is trying to measure. Consider two cities that use different currencies. Suppose that all local-currency goods prices are fixed, so that all of the variation in $\Delta q_{i,jk,t}$ is due to changes in the nominal exchange rate. Then the change in each good's LOP deviation would be the same at each date t, the cross-sectional variance, $Var_i(\Delta q_{i,jk,t} \,|\, jk, t)$ would be zero, and so would its time-series average, I_{jk}. We would say that all of the LOP variation is 'common,' as reflected by C_{jk}. In the less stark case, local currency prices also move. If this movement is idiosyncratic — if it represents cross-sectional variation *around* variation in the real exchange rate — then I_{jk} will pick this up.

Table 4 reports estimates of the moments in Equation (6). The answer to the questions posed above are quite clear. Very little of the variation is common

[7] Crucini, Telmer, and Zachariadis (2005) show that, for European data, equally-weighted, cross-good averages of LOP deviations behave similarly with respect to CPI-weighted averages. Also, Figure 6 below shows that for each date t, $E_i(\Delta q_{i,jk,t} \,|\, jk, t)$ is highly correlated with a NIPA-based measure of the real exchange rate.

across goods. Almost all of the variation is idiosyncratic, representing movement within the distribution, not movement of the distribution itself. This is particularly true for city-pairs not separated by an international border, where the common variation represents only 1 or 2% of the total. For city pairs that are separated by a border, common variation is larger, but the magnitude remains small. The ratio of the common to total variation is 7.3%, 20.0%, and 22.4% for Canada-U.S., OECD and "World" city pairs, respectively.

Figures 4 and 5 add some color, structure and detail to Table 4. Whereas Table 4 reports cross-city *averages*, the figures report the actual *values* of the ratio $C_{jk}/(C_{jk} + I_{jk})$ for each jk. Figure 4 does so for the Canada-U.S. case. Figure 5 does so for all OECD city pairs.[8] We see, especially in Figure 4, that the border effect holds pointwise, not just on average. We also see some structure emerge when we organize city-pairs by nominal exchange rate variability. This is apparent in Figure 5, where it is clear that higher nominal exchange rate variability is strongly associated with a higher ratio of common to total variability.

The basic point, then, is that common sources of variation — most obviously nominal exchange rates — play a minor role in accounting for variation in changes in international, microeconomic relative prices. It is interesting to contrast this with what we know about macroeconomic relative prices. The Mussa (1986) evidence is that real and nominal exchange rates are basically the same thing. Put differently, nominal exchange rates play a *major* role in accounting for variation in international relative prices. Is there a contradiction here?

Figure 6 shows that there is not. It plots, for the Canada-U.S. case, the NIPA-based change in the real exchange rate and the analogous object from Equation (6),

[8]Including all remaining the remaining city-pairs, beyond the OECD, changes neither the picture nor the point.

a time series of the cross-good average $E_i(\Delta q_{i,jk,t} \,|\, jk, t)$, one line for each bilateral city-pair, jk. The Mussa evidence is quite apparent. The average correlation between the NIPA and micro-based lines is 0.89. The reasoning is quite obvious. Most of the variation in good-by-good LOP changes is idiosyncratic. Once we average-away this variation, we are left with the Mussa facts. To a certain extent this is true-by-construction; averages vary less than the things being averaged. But our basic point is not a tautology. It is simply that the *magnitude* of what is being averaged away is large. At the good-specific level, there is much more going on than movement in nominal exchange rates.

5 Discussion

It is not surprising that long-term deviations from LOP exist between two cities. Retail goods are bundles of traded and non-traded goods and this is exactly what we should expect. Neither is it surprising that changes in good-level relative prices are more variable than changes in aggregate prices. This is an obvious implication of averaging. Qualitatively, then, there is nothing surprising about the above results. Their importance is strictly related to the *magnitude* of the effects.

We now report additional results that go beyond 'magnitude' and shed some light on the the economic forces affecting LOP deviations. The basic idea is to (i) establish links between the time-series and cross-sectional evidence presented above, and (ii) partition the data into the classic traded versus non-traded dichotomy. The combination of (i) and (ii), we think, provides valuable information for the construction of economic models that can synthesize the micro and macro economic behavior of exchange rates.[9]

[9]Some recent examples are as follows. Sposi (2012) incorporates asymmetric trade costs into

5.1 Cross-Sectional and Time Series Variation Are Related

Figure 7 reports scatter plots of T_i versus F_i — cross-sectional variance versus time-series variance — from Equation (3) of Section 3. There is an unmistakable positive relationship, in particular for the 'North America' and the 'Within-Country' subsets of the data. Univariate regressions of F_i on T_i yield a positive and strongly significant slope coefficient for all but the 'OECD Cross-Border Pairs' subset (NE panel), where the coefficient is not significantly different from zero. The slope coefficients (standard errors) are 0.26 (0.09), -0.0129 (0.14), 0.35 (0.11) and 0.52 (0.11) for the NW, NE, SW and SE panels, respectively. The R^2 is 5%, .08%, 6% and 9%, respectively. These results become sharper — coefficients become more positive, standard errors become relatively smaller, the R^2 coefficient becomes as large as 25% ('OECD Within-Country Pairs') — if we break the data down by traded versus non-traded goods (anticipating our next set of results).

The message of Figure 7 is simple. Whatever it is that is generating the relatively large long-run fixed effects that were the focal point of Section 3 — be it geographic distance, non-traded input costs, product differentiability, *etc.* — there seems to be a important link with whatever it is that is generating time-series variation *around* these fixed effects. This link is more pronounced for within-country city pairs, suggesting that nominal exchange rate variation is masking the relationship between the time-series and cross-sectional variation in the data. From the perspective of models, our results suggest a potentially informative class of restrictions. Models of the transmission of *shocks* — *e.g.*, New Keynesian models

the Eaton and Kortum (2002) model and, using Penn World Table sectoral data, finds that trade costs can account for roughly half of the cross-sectional price dispersion. See also Giri (2012). Alessandria and Kaboski (2011) persuasively argue that retail search costs are important for absolute price dispersion. Regarding time-series variation, Kehoe and Midrigan (2007) show that a time-dependent pricing model can generate volatility and persistence in LOP deviations that is qualitatively similar to what we find here.

with nominal rigidities — can be informed by the cross-sectional variation in the data. When attempting to identify the frictions that are important for the time-series behavior of relative prices one should look to the frictions that generate long-run LOP deviations for confirmation.

5.2 Tradeability

Our next set of results relate to the tradeability of individual goods. A useful reference point is Engel (1999) and many papers that have followed. Engel found that, in contrast to the seminal Balassa-Samuelson model (Balassa (1964), Samuelson (1964)), variability in *traded* goods prices seems to play an important, perhaps dominant role in accounting for real exchange behavior. Engel's results, however, are based on macroeconomic index-number data and can only speak to time-series variation. Here, we ask if microeconomic absolute-price data tell a similar story and/or offer any new insights.

Tables 5 and 6 report the same variance decomposition as Tables 1 and 4, but with the goods categorized as being either tradeable or non-tradeable.[10] The basic pattern that emerges is quite sharp: (i) cross-sectional variation is larger for non-traded goods, (ii) time-series variation is larger for traded goods. The latter is more pronounced than the former. This is particularly apparent in Figure 8, where we plot the individual data-points that comprise the averages reported in Table 6.[11]

[10] Appendix A describes the procedure with which we categorize every good as being either tradeable or non-tradeable.

[11] Figure 8 reports data on time-series variation. We do not report the analogous graph for cross-sectional variation because the structure of the data makes it hard to interpret. The reason (admittedly complex) derives from how we choose which dimensions of the data to condition on. Our measures of cross-sectional variation, T_i, condition on goods. There are, of course, some traded goods with relatively high T_i and some traded goods with relatively low T_i. Table 5 shows that, *on average*, T_i is relatively low for traded goods. A plot of T_i for traded versus non-traded goods, analogous to Figure 8, would (weakly) demonstrate this average, but not much more. This is because we lack any organizational structure with which to compare *individual* traded goods to individual non-traded goods. Figure 8, on the other hand, derives from our *city-pair-conditioned*

The figure shows that city-pair-specific time-series variance is almost pointwise greater for traded than for non-traded goods, and that this pattern is almost exclusively driven by the good-specific component of the time-series variance.

The traded versus non-traded dimension of our data, then, is both consistent with Engel's (1999) basic finding, and with the main idea of the Balassa-Samuelson model. Non-traded goods are associated with large, long-run deviations from LOP. The difference with respect to traded goods is not large (Table 5), but it is certainly not the case that traded-good LOP deviations dominate in the cross-section. Traded goods, on the other hand, do dominate in the time-series dimension. As an anecdote, it's what one would expect if (differential) shocks to the cost of petroleum products and electronic components were large relative to those of haircuts, but the long-run cost of getting a haircut in one location is always a lot larger than in another location. As we will see below, this anecdote is consistent with the overall behavior observed in our data.

5.3 Implications

We summarize the above findings as three empirical facts. Define x as a continuous measure of tradeability, and recall that T and F denote cross-sectional and time-series variability, respectively (Equation (3), Section 3). The facts are that

1. $Cov(T, F) > 0$
2. $E(T \mid x)$ is decreasing in x.
3. $E(F \mid x)$ is increasing in x.

data, C_{jk} and I_{jk}. The variability within traded and non-traded goods categories averages out, revealing the striking pattern that we see in the figure.

At first blush there might seem to be a contradiction here. If non-traded goods have high cross-sectional variation (Fact 2), and cross-sectional variation is positively associated with time-series variation (Fact 1), then shouldn't non-traded goods also have relatively high time-series variation (thus contradicting Fact 3)?

The answer is no, not necessarily. This logic is flawed because is confuses unconditional covariance with covariance in conditional means. It ignores the possibility that goods differ in a higher-dimensional way than just tradeability. In fact, these three facts indicate that there must necessarily be another source of variation. To see this clearly, consider the following decomposition of the unconditional covariance:

$$Cov(T,F) \equiv E_x Cov\big(T\,,\, F \mid x\big) + Cov_x\big(E(T\,|\,x)\,,\, E(F\mid x)\big) \;\; ,$$

where, as above, the subscript denotes the variable that the moment is integrating over ('unconditional covariance equals average conditional covariance plus covariance of conditional means'). Fact 1 says that the LHS is positive and Facts 2 and 3 say that the second RHS term is negative. So $E_x Cov\big(T\,,\, F \mid x\big)$ must be positive; conditional on a level of tradeability, high T is associated with high F. This must be strong enough to counter the negative term. This means that there must be a second source of variation because, if not, then $Cov\big(T\,,\, F \mid x\big) = 0$, making $Cov(T,F) > 0$ impossible.

So, there must be a second source of variation. Call it y.

$$Cov(T,F) \;\;\equiv\;\; E_y Cov\big(T\,,\, F \mid y\big) + Cov_y\big(E(T\,|\,y)\,,\, E(F\mid y)\big) \qquad (7)$$

Things aren't as clear now because we don't know what y is. However, it's easy to show that, if T and F are linearly related to x and y, then Facts 1-3 imply that

$E_y Cov\big(T\,,\,F \mid y\big) < 0$. This is just the standard 'omitted variable' story. Think of a scatter-plot of T_i against F_i with two groups of points, one close to the origin and one far away. The variable y is associated with distance from the origin. The variable x, tradeability, is associated with a negative relationship between T and F in each group of points (*i.e.*, conditional on y). This tells us that the second term on the right must be positive and must be relatively large, something we return to when summing-up.

To better articulate what we think is going on, we label y as *differentiability* and write (x, y) as the pair of characteristics that define each good, i. LOP deviations can therefore be written $q_{jk,t}(x, y)$. Next, for every city pair jk with a positive long-run LOP deviation, (*i.e.*, $E_t(q_{jk,t}(x, y) \,|\, jk, (x, y)) > 0$), the symmetric counterpart, kj, must have a negative long-run mean. Without loss of generality, we ignore all of the latter and write

$$q_t(x, y) \quad = \quad \alpha(x, y) + \sigma(x, y)\,\varepsilon_t \;\;, \tag{8}$$

where $\alpha, \sigma > 0$, ε_t is a shock that could have some dynamics, and the unnecessary notation jk has been suppressed. Equation (8) says that each LOP deviation depends on a good-specific, long-run 'fixed effect,' $\alpha(x, y)$, and good-specific time-series variance, $\sigma^2(x, y)$.[12] This equation isn't implementable for us because we lack data on y. It is simply a labeling device that helps us exposit some intuition for what Facts 1-3 are saying.

[12]There is good reason to believe that α and σ might also depend on the particular location pair, jk. For example the long-run LOP deviation α, might depend on the geographical distance between cities j and k. We ignore this because our analysis is focused on variation in the goods-dimension, thus averaging-away these sorts of things. That is, F_i asks, for good i, what is the *average* long-term LOP deviation across all city pairs. Similarly, T_i measures the cross-location, *average* time-series variance.

Fact 1 says that both $\alpha(x,y)$ and $\sigma(x,y)$ are increasing in y. The former is natural. It says that high differentiability permits producers to sustain relatively large long-run markups and LOP deviations. It's also consistent with differentiable goods being more susceptible to persistent measurement error such as an inadequate accounting for quality. The latter — $\sigma(x,y)$ increasing in y — might be less natural, but it's certainly plausible. If, for instance, time-variation in LOP deviations is bounded by goods-market arbitrage, then this says that the arbitrage bands are wider for more differentiable goods. *'A larger LOP deviation is needed to initiate arbitrage in the market for cars than in the market for wheat.'* Alternatively, one can just as easily imagine that differentiable goods are produced using a less-diversified production structure and/or are subject to larger demand shocks.

Facts 2 and 3 say that $\alpha(x,y)$ is decreasing in x, but $\sigma(x,y)$ is *increasing* in x. Again, the former is quite natural; non-traded goods exhibit higher long-run LOP deviations. This is one of the pillars of trade theory. The latter, in contrast, is hard to call natural. Taken at face value, it says that traded goods have 'bigger' LOP deviations than non-traded goods! Indeed, this is what made Engel's (1999) findings so provocative. Equation (8) shines some light. It says that, while the classical traded/non-traded dichotomy gets it right in the long-run, in the short-run traded goods are affected by *shocks* that are more volatile than those affecting non-traded goods. This seems plausible to us. It seems plausible that shocks to the relative cost of haircuts are less volatile than shocks to the relative cost of food.

We offer two observations in support of all this 'plausbility.' The first is obvious. 'Core CPI' is widely-viewed as being more reliable than 'Headline CPI.' The difference between the two is that (i) the former is less volatile than the latter, and (ii) the former excludes food and energy, two of the more tradeable goods out there. We are saying the same thing.

Our second observation is that the qualitative behavior of these CPI indeces shows up in our micro data. Recalling that P_{ijt} is the *local-currency* price of good i in city j at date t, we define the price change, net of the city-specific overall inflation rate, as

$$v_{ijt} \equiv \Delta p_{ijt} - E_i\big(\Delta p_{ijt}\,|\,j,t\big) \ ,$$

where p is the logarithm of P and Δ denotes the change from $t-1$ to t. We then compute the good-specific, average price volatility as

$$v_i = E_j\big(Stdev(v_{ijt})\big) \ . \tag{9}$$

Figure 9 plots the estimated cross-sectional density for v_i, where i is dichotomized by traded/non-traded goods. What we see is consistent with our supposition. This simple measure of the volatility of traded goods prices is 0.14, substantially higher than that for non-traded goods prices, 0.10.

To summarize, consider again the scatter plots in Figure 7. The most basic model of how tradeability affects price dispersion — the Balassa-Samuelson — predicts that points associated with goods that are more tradeable should lie closer to the origin (taken literally, it suggests that the traded-goods points should lie *on* the origin). What we've found elaborates on their basic intuition. In the vertical dimension of the graphs, we *do* see the Balassa-Samuelson idea at work. In the long-run, traded goods exhibit less cross-sectional price dispersion. In the horizontal dimension, in contrast, the opposite is true. Traded goods exhibit more time-series variation, something we've attributed to higher-volatility in local-currency prices. Finally, looking back at Equation (7), we've found evidence of a second source of (cross-good) variation, above-and-beyond tradeability. This second source, y, must move both cross-sectional and time-series dispersion in the same direction and its effect must be large enough to cloud the negative relationship between T and F that is induced by tradeability (*i.e.*, $Cov_y(E(T \mid y)\,,\, E(F \mid y))$ is large and positive). We've labeled this source of variation 'differentiability,' but this was only for expository purposes. Theory is rife with other candidates, such as those related to the distribution of final goods. It could also just be *i.i.d.* measurement error. We leave the modeling and measurement of y for future work and hope that our simple variance decompositions can help discriminate between alternative candidates.

6 Conclusions

This paper's objective is to document some facts about microeconomic international relative prices that are informative for economic models and for the interpretation of aggregate data. We find that long-term LOP deviations — good-and-city-specific "fixed effects" — dominate the distribution of international relative prices. Nominal exchange rates play a relatively minor role in determining the relative cost of goods and services between, say, Tokyo and Los Angeles. Some goods are just always expensive in Tokyo and others are just always expensive in LA.

This is a statement about *absolute* deviations from LOP. Much of the existing literature on the behavior of real exchange rates — both aggregate and less-aggregate data — has focused on *changes*. Our second set of findings speaks to this evidence. We find that changes in nominal exchange rates play a relatively minor role in whatever it is that moves microeconomic relative prices across national borders. Other shocks (and probably a good dose of time-varying measurement error) are far more important.

When we consider the long-run dispersion and the short-run dispersion together, we find that there are some connections that are informative for what distinghishes goods from one another. The traditional characteristic, tradeability, does have some explanatory power in the cross-section. But there's much more going on. A second source of cross-sectional variation is more important than tradeability, accounting for what is essentially a 'level effect.' It moves both the cross-sectional and the time-series variance in the same direction. It also shows us that the effect of tradeability is more complex than the classical model suggests. While more tradeability is associated with less long-term, cross-sectional dispersion, it is also associated with *more* time-series dispersion.

These results raise a number of interesting questions for further research. First, what *is* the second source of cross-good dispersion in LOP deviations? In Section 5 we labeled it "differentiability," and we provided some discussion that is consistent with this label, but due to data limitations we cannot *really* say much more than 'there's a second source of variation and here's what it looks like.' Second, *why* do traded goods display more time-series variance than non-traded goods? We provided evidence — both anecdotal and statistical — suggesting a very simple answer; 'the shocks are bigger.' But a richer explanation obviously needs to go deeper. The set of goods that are traded is certainly an endogenous outcome. Perhaps what we are finding here actually motivates trade? If supply and demand shocks to the traded goods sectors of the economy are inherently larger, then perhaps a role played by trade is to facilitate the smoothing of such shocks, so that a region hit by a bad apple harvest doesn't have to dramatically switch to the consumption of pears? Our hope is that our simple variance decompositions may be informative for the quantitative study of questions like this one.

Finally, the broad economic implications of our results — even the simpler ones — seem interesting to us. For example, there is a popular notion that nominal exchange rate 'noise' distorts the international flow of goods and capital. This now seems less convincing. Yes, as Mussa (1986) so provocatively pointed out 20 years ago, nominal exchange rates and aggregate, CPI-based real exchange rates are essentially the same thing. This is true in our data just as it was in Mussa's. Yes, nominal exchange rates seem disconnected in many ways from macroeconomic fundamentals. But does this mean that nominal exchange rates are distorting the allocative role of the international price system? Our results indicate that there's a lot going on *within* the distribution of international prices which is not apparent in the behavior of the mean of the distribution. Exchange rates govern the mean, not variation around the mean. Whatever it is that is driving the variation, it is this that probably plays the dominant role in determining allocations. Firms don't export and import the CPI basket. They import and export goods and services. The consumer price signals which inform these goods flows are subject to many shocks of which the nominal exchange rate is but a relatively small one.

References

Alessandria, George A., and Joseph P. Kaboski, 2011, Pricing to market and the failure absolute ppp, *American Economic Journal: Macroeconomics* 3, 91–127.

Andrade, Phillipe, and Marios Zachariadis, 2012, Global versus local shocks in micro price dynamics, Unpublished manuscript, University of Cyprus.

Balassa, Bela, 1964, The purchasing power parity doctrine: A reappraisal, *Journal of Political Economy* 72, 584–596.

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

Broda, C., and David Weinstein, 2008, Understanding international price differences using barcode data, NBER Working Paper 14017.

Burstein, Ariel, and Andrew Atkeson, 2008, Pricing to market, trade costs, and international relative prices, *American Economic Review* 98, 1998–2031.

Campbell, Jeffrey R., and Beverly Lapham, 2004, Real exchange rate fluctuations and the dynamics of retail trade industries on the u.s.-canada border, *American Economic Review* 94, 1194–1206.

Crucini, Mario J., and Mototsugu Shintani, 2008, Persistence in law-of-one-price deviations: evidence from micro-price data, *Journal of Monetary Economics* 55, 629–644.

Crucini, Mario J., Christopher I. Telmer, and Marios Zachariadis, 2005, Understanding european real exchange rates, *American Economic Review* 95, 724–738.

Eaton, Jonathan, and Samuel Kortum, 2002, Technology, geography and trade, *Econometrica* 70, 1741–1779.

Engel, Charles, and John H. Rogers, 2004, European product market integration after the euro, *Economic Policy* 19, 347–384.

Engel, Charles M., 1999, Accounting for u.s. real exchange rate changes, *Journal of Political Economy* 107, 507–538.

Fischer, Christoph, 2012, Price convergence in the emu? evidence from micro data, Forthcoming, *European Economic Review.*

Fitzgerald, Doireann, and Stephanie Haller, 2008, Exchange rates and producer prices: Evidence from micro-data, Unpublished manuscript, Stanford University.

Gagnon, Etiene, 2010, Price setting during low and high inflation: Evidence from mexico, *Quarterly Journal of Economics* 124, 1221–63.

Ghironi, Fabio, and Marc J. Melitz, 2005, International trade and macroeconomic dynamics with heterogeneous firms, *Quarterly Journal of Economics* 120, 865–915.

Giri, Rahul, 2012, Local costs of distribution, international trade costs and micro-evidence on the law of one price, *Journal of International Economics* 86, 82–100.

Gopinath, Gita, Pierre-Olivier Gourinchas, Chang-Tai Hsieh, and Nicholas Li, 2011, International prices, costs and markup differences, *American Economic Review* 101, 2450–2485.

Gopinath, Gita, and Roberto Rigobon, 2008, Sticky borders, *Quarterly Journal of Economics* 123, 531–75.

Gorodnichenko, Yuriy, and Linda L. Tesar, 2009, Border effect or country effect? seattle may not be so far from vancouver after all, *American Economic Journal: Macroeconomics* 1, 219–241.

Hellerstein, Rebecca, 2008, Who bears the cost of a change in the exchange rate? pass-through accounting for the case of beer, *Journal of International Economics* 76.

Hummels, David, 2007, Transportation costs and international trade in the second era of globalization, *Journal of Economic Perspectives* 21, 131–154.

Imbs, Jean, Haroon Mumtaz, Morten Ravn, and Helene Rey, 2005, Ppp strikes back: Aggregation and the real exchange rate, *Quarterly Journal of Economics* 120, 1–43.

Kehoe, Patrick, and Virgiliu Midrigan, 2007, Sticky prices and real exchange rates in the cross-section, Unpublished manuscript, New York University.

Mussa, Michael, 1986, Nominal exchange rate regimes and the behavior of real exchage rates: Evidence and implications, *Carnegie-Rochester Series on Public Policy* 25 (Autumn): 117-214.

Parsley, David, and Shang-Jin Wei, 2000, Explaining the border effect: the role of exchange rate variability, shipping costs, and geography, NBER Working Paper No. 7836, August, forthcoming *Journal of International Economics*.

Rogers, John H., 2002, Monetary union, price level convergence, and inflation: How close is europe to the united states?, *Federal Reserve Board International Finance Discussion Paper* No. 740.

Rogoff, Kenneth, 1996, The purchasing power parity puzzle, *Journal of Economic Literature* 34 (2), 647–668.

Samuelson, Paul A., 1964, Theoretical notes on trade problems, *Review of Economics and Statistics* 46, 145–154.

Simonovska, Ina, 2010, Income differences and prices of tradables, NBER Working Paper 16233.

Sposi, Michael, 2012, Trade barriers and the relative price tradables, Working Paper, University of Iowa.

Appendix A

Data Description

The EIU dataset consists of local-currency retail prices, inclusive of sales tax, on as many as 301 goods and services, sampled in 123 cities from 78 different countries. A list of all goods-and-services and all cities is available at http://bertha.tepper.cmu.edu/telmerc/eurostat. The data are annual, 1990-2005. The country with the most intranational observations is the U.S., with 16, followed by Australia, China and Germany with 5, Canada with 4, Saudi Arabia with 3, and Brazil, France, Italy, Russia, Spain, Switzerland, UK, India, Japan, Vietnam, New Zealand with 2. Average nominal exchange rates for each calendar year are also provided by the EIU. We performed an extensive set of reliability checks using our own nominal exchange rate data. Aside from the Turkish lira, which we replaced with data from Datastream and IFS, the EIU data proved to be of high quality.

We transform the data as follows. We denote P_{ijt} as the local-currency price of good i in city j in year t and $S_{jk,t}$ as the date t nominal exchange rate between cities j and k, in units of city k ($S_{jk,t} = 1$ if cities j and k share the same currency). We transform prices into bilateral log deviations from the law-of-one-price (LOP):[13]

$$q_{i,jk,t} = \log(\frac{P_{ijt}S_{jk,t}}{P_{ikt}}) \ , \tag{A1}$$

The data contain many missing observations. We clean the data using the following algorithm. First we choose a set of cities (*e.g.*, OECD, North America, etc.).

[13] A previous version of the paper employed an alternative LOP measure, the log deviation from the cross-city geometric average: $\log(S_{jn,t}P_{ijt})/\sum_{j=1}^{M}\log(S_{jn,t}P_{ijt})$ where M denotes the total number of cities and n denotes the numeraire currency in units of which all prices are expressed (our measures of price dispersion are independent of the choice of the numeraire currency). While this definition results in lower overall LOP variability (by construction), the main message of our paper remains unchanged.

Second, for each year, we eliminate any good that has data for less than 50% of the cities. Third, we remove outliers, defined as a relative price greater than 3 or less than 1/3. Fourth, we repeat step 2. This defines our basic data-structure, having done a first-pass removal of outliers and goods with too little cross-location data. Our panel is unbalanced, having (for example) eliminated some goods for earlier years but not later years. Subsequent calculations involve further cleaning. Whenever we compute a time-series moment, we insist that each good-city pair have a complete set of time-series observations. For cross-sectional moments we insist that a given good (or city) have 75% or more of the cities (goods) represented. Finally, we have extensively experimented with different cleaning algorithms (*e.g.*, eliminating goods with less than 75% of observations, using totally balanced panels, etc..) and our results do not change in qualitatively-important ways. Our view is that for the broad variance decompositions that are the focal point of this paper, an unbalanced panel is unlikely to be a problem.

Our traded versus non-traded classification is based on country-specific ratios of imports to GDP, based on the 6-digit Harmonized System concordance between import values available through the TRAINS dataset and the individual goods from the EIU prices database. This measure is averaged across locations for each good to obtain good-specific measure of tradeability. This leaves us with, typically (across the different subsets of cities that we examine) about 50 non-traded goods and 200 traded goods. The fact that traded goods are over-represented in the EIU data is unavoidable. For a more detailed discussion of this, and many more pros and cons of the EIU data, see Andrade and Zachariadis (2012) who provide a very thorough and careful analysis, including tables that list the traded versus non-traded goods that are very similar to our study.

Matlab code is available upon request from the authors. The code will read in the generic form that the EIU distributes the data in (in Excel spreadsheets). The data is, by now, widely available, at many academic libraries for example.

Table 1
Variance in Absolute LOP Deviations

	Total	Cross Sectional	Time Series	Cross Sectional / Total
U.S.	0.128	0.071	0.061	0.559
Canada-U.S.				
Combined	0.138	0.078	0.065	0.564
International	0.124	0.057	0.073	0.462
Intranational	0.127	0.071	0.061	0.554
OECD				
Combined	0.215	0.143	0.077	0.666
International	0.221	0.147	0.079	0.667
Intranational	0.123	0.069	0.058	0.561
World				
Combined	0.270	0.186	0.089	0.690
International	0.275	0.190	0.091	0.691
Intranational	0.118	0.065	0.057	0.551

The table reports cross-good averages of the variables T_i and F_i from Equations (2-3), reproduced here:

$$\begin{aligned} Var_{jk,t}(q_{i,jk,t}\,|\,i) &= Var_{jk}(E_t[q_{i,jk,t}\,|\,i,jk]) + E_{jk}[Var_t(q_{i,jk,t}\,|\,i,jk)] \\ &= T_i + F_i \ . \end{aligned}$$

Section 3 provides definitions and a detailed discussion. The first column of numbers are $\sum_i (T_i + F_i)/M$, where M is the number of goods. The second and third columns report the analogous averages of T_i and F_i, respectively. The last column is simply the ratio of the second to the first. The extent to which the sum of columns 2 and 3 is inconsistent with column 1 is due to missing observations. Figure 7 provides information on the cross-good dispersion in T_i and F_i.

Table 2
City-Specific Variation in Long-Run Average LOP Deviations

	U.S.	Canada-U.S.	OECD	World
Total	0.0781	0.0838	0.1528	0.1821
City-Pair	0.0130	0.0109	0.0231	0.0362
Residual	0.0652	0.0731	0.1308	0.1491
City-Pair/Total	0.1660	0.1297	0.1509	0.1989

The table reports a variance decomposition of the variation in long-run average LOP deviations into a city-specific component and a residual. Specifically, the time-average of the LOP deviation for good i between city-pair jk is denoted $\mu_{i,jk} \equiv E_t(q_{i,jk,t} \,|\, i, jk)$. This is then written

$$\begin{aligned} \mu_{i,jk} &= \mu_{jk} + \varepsilon_{i,jk} \\ \Longrightarrow Var_{i,jk}(\mu_{i,jk}) &= Var_{jk}(\mu_{jk}) + Var_{i,jk}(\varepsilon_{i,jk}) \ . \end{aligned}$$

Rows 1, 2 and 3 report estimates of the total variance, $Var_{i,jk}(\mu_{i,jk})$, the city-pair-specific variance, $Var_{jk}(\mu_{jk})$ and the residual variance, $Var_{i,jk}(\varepsilon_{i,jk})$ for U.S. city pairs, Canada-U.S. city pairs, all OECD pairs and all global pairs, respectively. The last row reports the fraction of the total variance attributable to city-effects. This decomposition is the same as "city-pair-specific dummy variables," or "1-factor ANOVA." Additional details and discussion are provided in Section 3.1.

Table 3

City-Specific and Good-Specific Variation in Time-Series Variances of Absolute LOP Deviations

	U.S.	Canada-U.S.	OECD	World
Total	0.00346	0.00364	0.00432	0.00585
Good	0.00105	0.00090	0.00073	0.00069
City-Pair	0.00007	0.00009	0.00040	0.00090
Residual	0.00236	0.00268	0.00327	0.00438
Good/Total	0.30491	0.24833	0.16920	0.11717
City-Pair/Total	0.02157	0.02590	0.09229	0.15385

The table reports a variance decomposition of the total cross-sectional variance of the time-series variances of absolute LOP deviations into city and good-specific sources of variation, and a residual. More specifically, we denote the time-series variance of the absolute LOP deviation for good i between city-pair jk as $\sigma^2_{i,jk} \equiv Var_t(q_{i,jk,t} \,|\, i, jk)$. We write

$$\begin{aligned} \sigma^2_{i,jk} &= \sigma^2_i + \sigma^2_{jk} + \varepsilon_{i,jk} \\ \Longrightarrow Var_{i,jk}(\sigma^2_{i,jk}) &= Var_i(\sigma^2_i) + Var_{jk}(\sigma^2_{jk}) + Var_{i,jk}(\varepsilon_{i,jk}) \ . \end{aligned}$$

Rows 1, 2, 3 and 4 report estimates of the total variance, $Var_{i,jk}(\sigma^2_{i,jk})$, the good-specific variance, $Var_i(\sigma^2_i)$, the city-pair-specific variance, $Var_{jk}(\sigma^2_{jk})$ and the residual variance, $Var_{i,jk}(\varepsilon_{i,jk})$ for U.S. city pairs, Canada-U.S. city pairs, all OECD pairs and all global pairs, respectively. The last two rows report the fractions of the total variance attributable to good and city-pair effects. This decomposition is the same as "city-pair and good-specific dummy variables," or "2-factor ANOVA." Additional details and discussion are provided in Section 3.1.

Table 4
Variance of Changes in LOP Deviations

	Total	Common (across goods)	Idiosyncratic (good-specific)	$\frac{\text{Common}}{\text{Total}}$
U.S.	0.042	0.001	0.041	0.017
Canada-U.S.				
Combined	0.045	0.002	0.043	0.041
International	0.048	0.004	0.045	0.073
Intranational	0.043	0.001	0.042	0.018
OECD				
Combined	0.050	0.009	0.041	0.185
International	0.051	0.010	0.041	0.200
Intranational	0.041	0.001	0.040	0.017
World				
Combined	0.060	0.013	0.047	0.218
International	0.061	0.014	0.047	0.224
Intranational	0.041	0.001	0.040	0.018

The table reports a decomposition of the variance of *changes* in LOP deviations into common and idiosyncratic components. "Common" means "time-variation, common to all goods for each location-pair," and "idiosyncratic" means "time-variation that is good-specific, for each location pair." More specifically, the table reports cross-location averages of the variables C_{jk} and I_{jk} from Equation (6), Section (4), reproduced here:

$$\begin{aligned} Var_{i,t}(\Delta q_{i,jk,t}\,|\,jk) &= Var_t\Big(E_i(\Delta q_{i,jk,t}\,|\,jk,t)\Big) + E_t\Big(Var_i(\Delta q_{i,jk,t}\,|\,jk,t)\Big) \\ &= C_{jk} + I_{jk} \quad . \end{aligned}$$

C_{jk} measures, for a given city pair jk, variation in LOP changes that are common across goods. I_{jk} measures variation (in changes) that is specific to each good. Section 4 provides a more detailed discussion. The first column of numbers ('Total') contains $\sum_{jk}(C_{jk}+I_{jk})/N$, where N is the number of city-pairs. The second and third columns report the analogous averages of C_{jk} and I_{jk}, respectively. The last column is simply the ratio of the second to the first. The extent to which the sum of columns 2 and 3 is inconsistent with column 1 is due to missing observations. Figure 8 provides information on the cross-location dispersion in C_{jk} and I_{jk}.

Table 5
Variance in Absolute LOP Deviations: The Effect of Tradeability

	Cross Sectional		Time Series	
	Traded	Non-Traded	Traded	Non-Traded
U.S.	0.063	0.113	0.064	0.048
Canada-U.S.				
Combined	0.070	0.138	0.068	0.054
International	0.051	0.099	0.076	0.057
Intranational	0.063	0.142	0.064	0.051
OECD				
Combined	0.135	0.181	0.079	0.067
International	0.140	0.182	0.081	0.069
Intranational	0.060	0.113	0.061	0.048
World				
Combined	0.184	0.198	0.088	0.098
International	0.188	0.200	0.089	0.100
Intranational	0.057	0.106	0.059	0.045

The table reports cross-good averages of the variables T_i and F_i, broken-down by the tradeability of the goods. These variables are defined by Equations (2-3), which we reproduce here:

$$\begin{aligned} Var_{jk,t}(q_{i,jk,t}\,|\,i) &= Var_{jk}(E_t[q_{i,jk,t}\,|\,i,jk]) + E_{jk}[\,Var_t(q_{i,jk,t}\,|\,i,jk)] \\ &= T_i + F_i \;\; . \end{aligned}$$

Section 3 provides further definitions and a detailed discussion. This table is identical to Table 1, except for the traded, non-traded distinction.

Table 6

Variance in *Changes* in LOP Deviations: The Effect of Tradeability

	Cross Sectional		Time Series	
	Traded	Non-Traded	Traded	Non-Traded
U.S.	0.001	0.002	0.045	0.023
Canada-U.S.				
Combined	0.002	0.003	0.047	0.023
International	0.004	0.005	0.049	0.024
Intranational	0.001	0.002	0.046	0.023
OECD				
Combined	0.010	0.010	0.044	0.023
International	0.010	0.010	0.044	0.023
Intranational	0.001	0.001	0.043	0.022
World				
Combined	0.013	0.015	0.050	0.030
International	0.014	0.016	0.050	0.031
Intranational	0.001	0.002	0.043	0.023

The table reports cross-city-pair averages of the variables C_{jk} and I_{jk}, where the goods used in the calculations are categorized as either traded or non-traded. The variables C_{jk} and I_{jk} represent "common time-series variation" and "good-specific time-series variation." They are defined by Equation (6), Section 4, which is reproduced here:

$$\begin{aligned} Var_{i,t}(\Delta q_{i,jk,t}\,|\,jk) &= Var_t\Big(E_i(\Delta q_{i,jk,t}\,|\,jk,t)\Big) + E_t\Big(Var_i(\Delta q_{i,jk,t}\,|\,jk,t)\Big) \\ &= C_{jk} + I_{jk} \ . \end{aligned}$$

Section 4 provides further definitions and a detailed discussion. Figure 8 plots the city-pair specific estimates of C_{jk} and I_{jk} that are the data used in to compute the averages in this table.

Figure 2
Distribution of LOP Deviations: $\log q_{i,jk,t}$

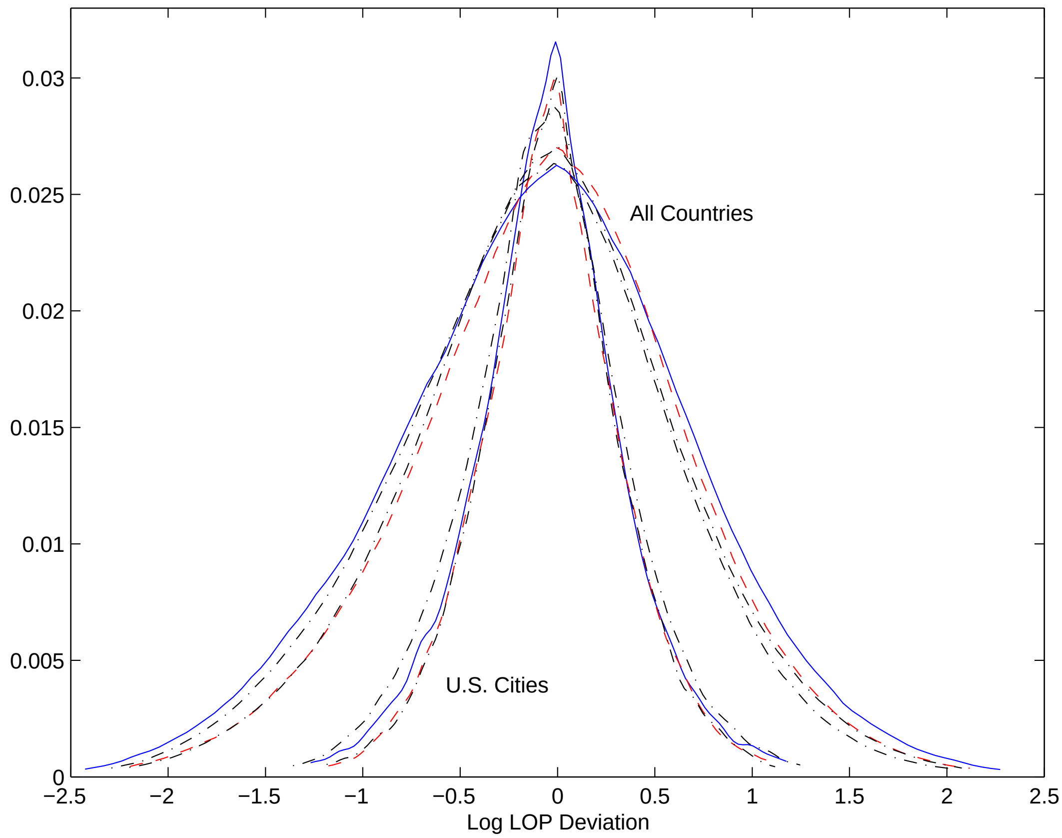

Density of log LOP deviations across (i) U.S. and Canadian location-pairs, and (ii) all location-pairs. The densities describe variation in $q_{i,jk,t}$ across goods i and location-pairs jk, where $q_{i,jk,t}$ is computed as

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

and P and S are local-currency prices and the nominal exchange rate. For each set of cities there are four lines: $t = 1990$, 1995, 2000 and 2005.

Figure 3
Decomposing the Variance in Absolute LOP Deviations

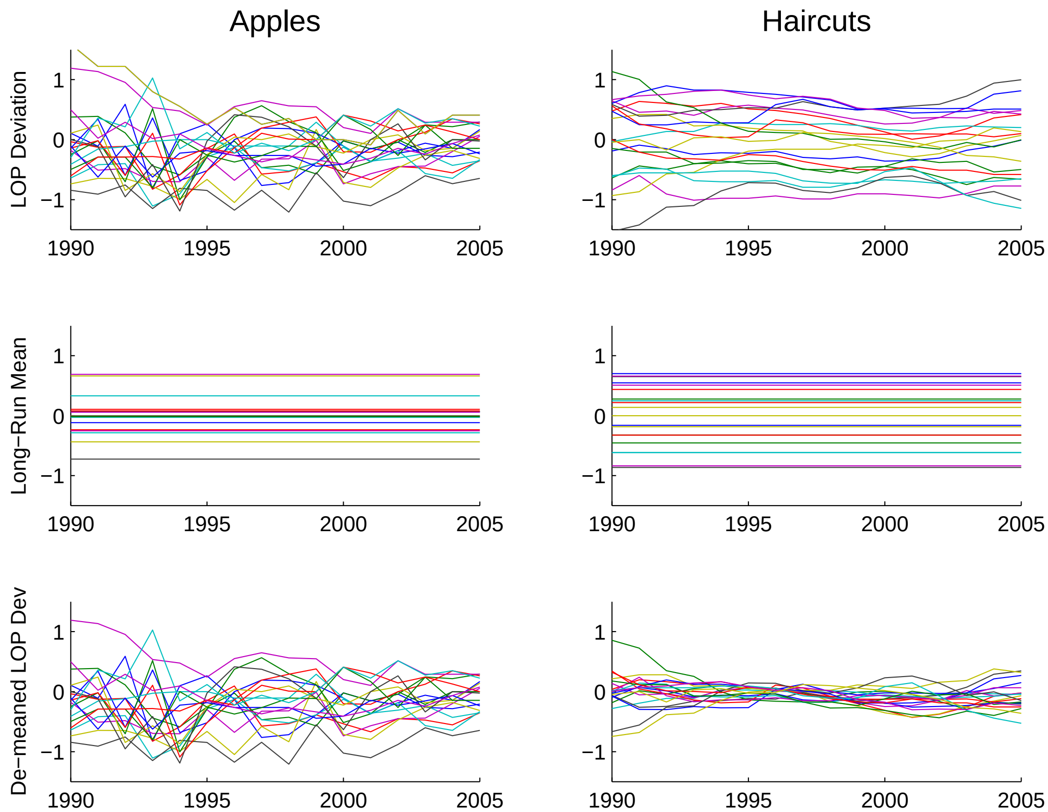

The intent of this figure is to help the reader visualize the variance decomposition in Equation (3) of Section 3. Start with the upper-left graph. Each line represents the time-series of log LOP deviations for a traded good, apples, between two *particular* North American cities (*e.g.*, Toronto/Pittsburgh). Our data represent 16 U.S. cities and 4 Canadian cities, which makes for 190 bilateral pairs. The graph, therefore, could have as many as 190 lines on it. To make it legible, we randomly select and plot just 20 of these pairs (the qualitative content of the graph is not changed as we increase the number of pairs and choose a different random sample). The upper-right graph does the same thing for a non-traded good, haircuts. The middle-two graphs plot the *time-series mean* of each of the lines in the upper-two graphs. The *cross-sectional variance* of these time-series means is our variable T_i, i=(apples, haircuts), from Equation (3). Finally, the lower-two graphs are simply the de-meaned versions of the upper-two graphs. Our variable F_i from Equation (3) is the *cross-sectional mean* of the time-series variance of each of these lines. Further discussion is provided in the text of Section 3.

Figure 4

Common Variation in Changes in Good-by-Good LOP Deviations Canada-U.S. City-Pairs

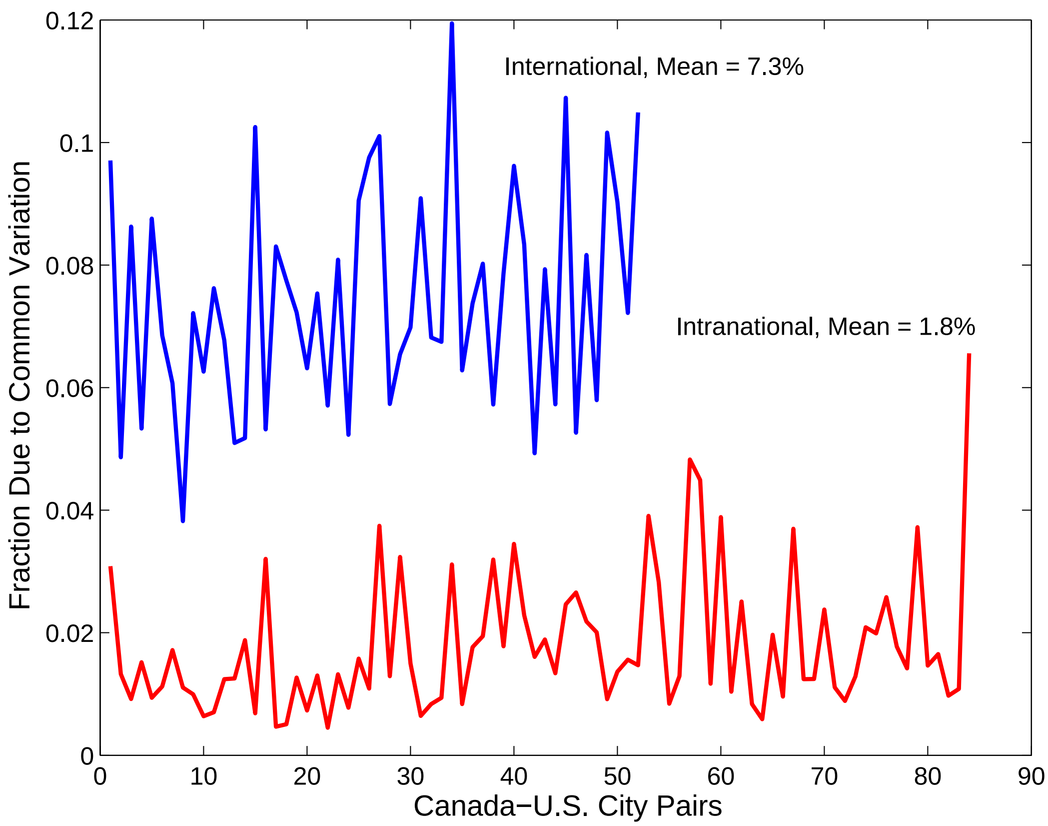

The horizontal axis represents Canada-U.S. bilateral city-pairs. The city-pairs are sorted by geographic distance, from closest apart to farthest apart. There are 4 Canadian cities and 16 U.S. cities, which make for 190 bilateral pairs. Of these, 54 were eliminated due to insufficient data, leaving 84 intranational pairs and 52 international pairs. The upper (blue) graph represents the latter and the lower (red) line represents the former. The vertical axis represents the fraction of the variance in the good-by-good *changes* in LOP deviations attributable to a source of variation which is common across all goods (for each city-pair). This fraction is, for each city-pair jk, $C_{jk}/(C_{jk} + I_{jk})$ from Equation (6) in Section 4 of the text.

Figure 5
Common Variation in Changes in Good-by-Good LOP Deviations OECD City-Pairs

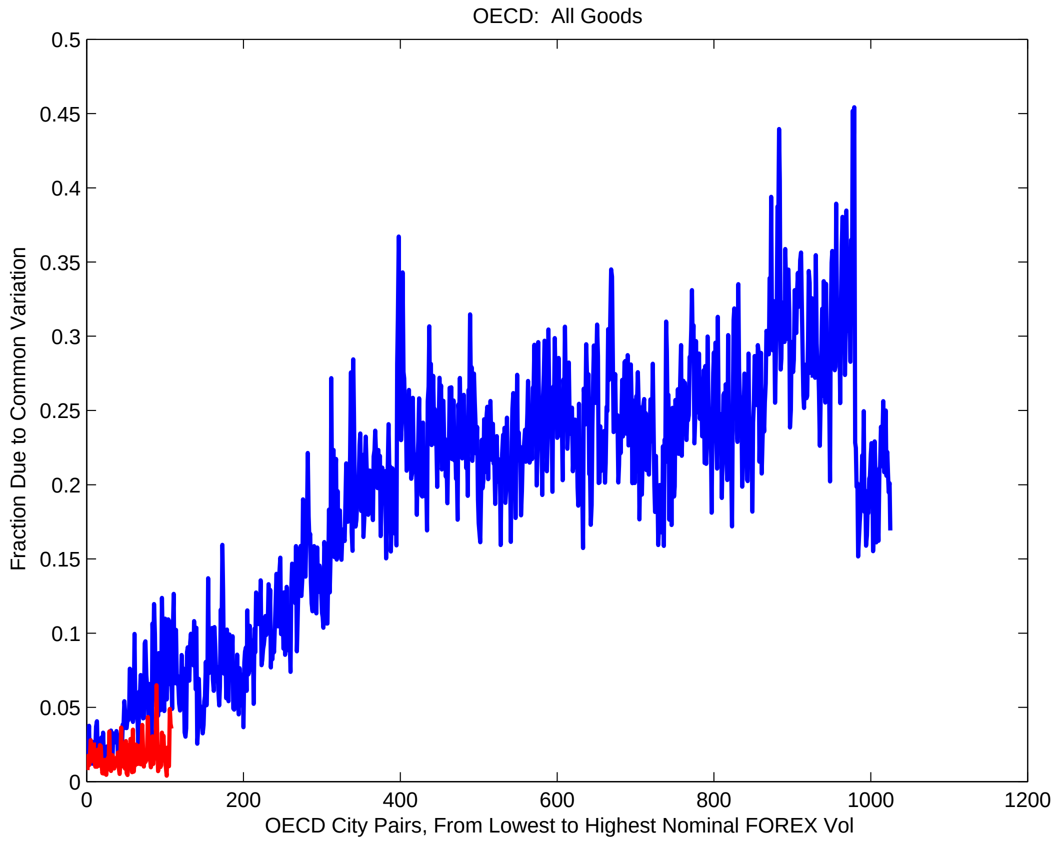

The horizontal axis represents OECD bilateral city-pairs. The city-pairs are sorted by the volatility of the nominal FOREX depreciation rate between each pair of cities, from lowest to highest. The vertical axis represents the fraction of the variance in the good-by-good *changes* in LOP deviations attributable to a source of variation which is common across all goods (for each city-pair). This fraction is, for each city-pair jk, $C_{jk}/(C_{jk}+I_{jk})$ from Equation (6) in Section 4 of the text. The top (blue) line represents international city-pairs and the bottom (red) line represents intranational city-pairs. The last bunch of international city-pairs all include Istanbul. A list of the city-pairs and depreciation volatilities is available at:

`http://bertha.tepper.cmu.edu/eurostat/oecd1.txt`

Figure 6

Common Component of Changes in LOP Deviations and Nominal Depreciation Rate

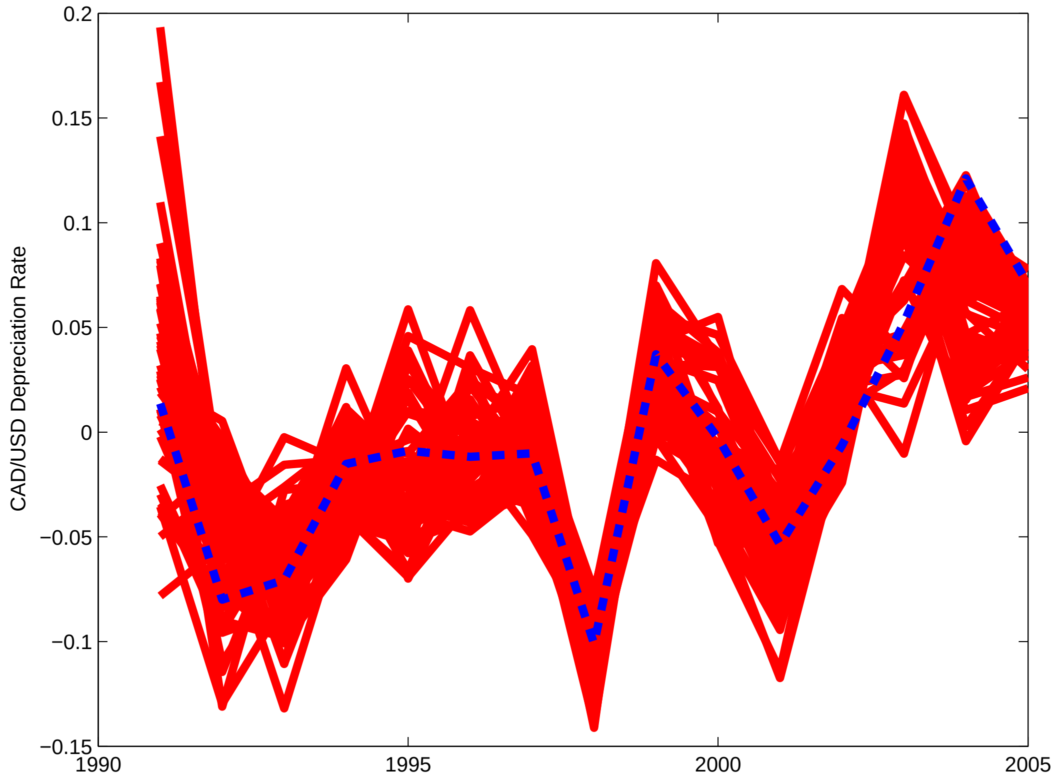

This graph corresponds to the Canada-U.S. city pairs represented in Figure 4. The dashed (blue) line is the annual depreciation rate in the U.S.-Canada nominal exchange rate (USD per CAD). It is computed using the annual, nominal exchange rate data provided by the EIU, which corresponds to the times during the year during which the goods were sampled. However, the graph is not changed much if one computes annual, nominal exchange rates as averages of daily exchanges rates over the calendar year. The solid (red) lines are the 'common factors,' $E_i(\Delta q_{i,jk,t} \mid jk, t)$, extracted from the *changes* in LOP deviations that are the basis of the variance decomposition, Equation (6) in Section 4 of the text. There is one solid (red) line per city-pair, jk. These lines are simply the time t, cross-good averages of changes in LOP deviations. The average correlation between each of the solid (red) lines and the dashed line is 0.89.

Figure 7
Variation in Absolute LOP Deviations: Time-Series Versus Cross-Sectional

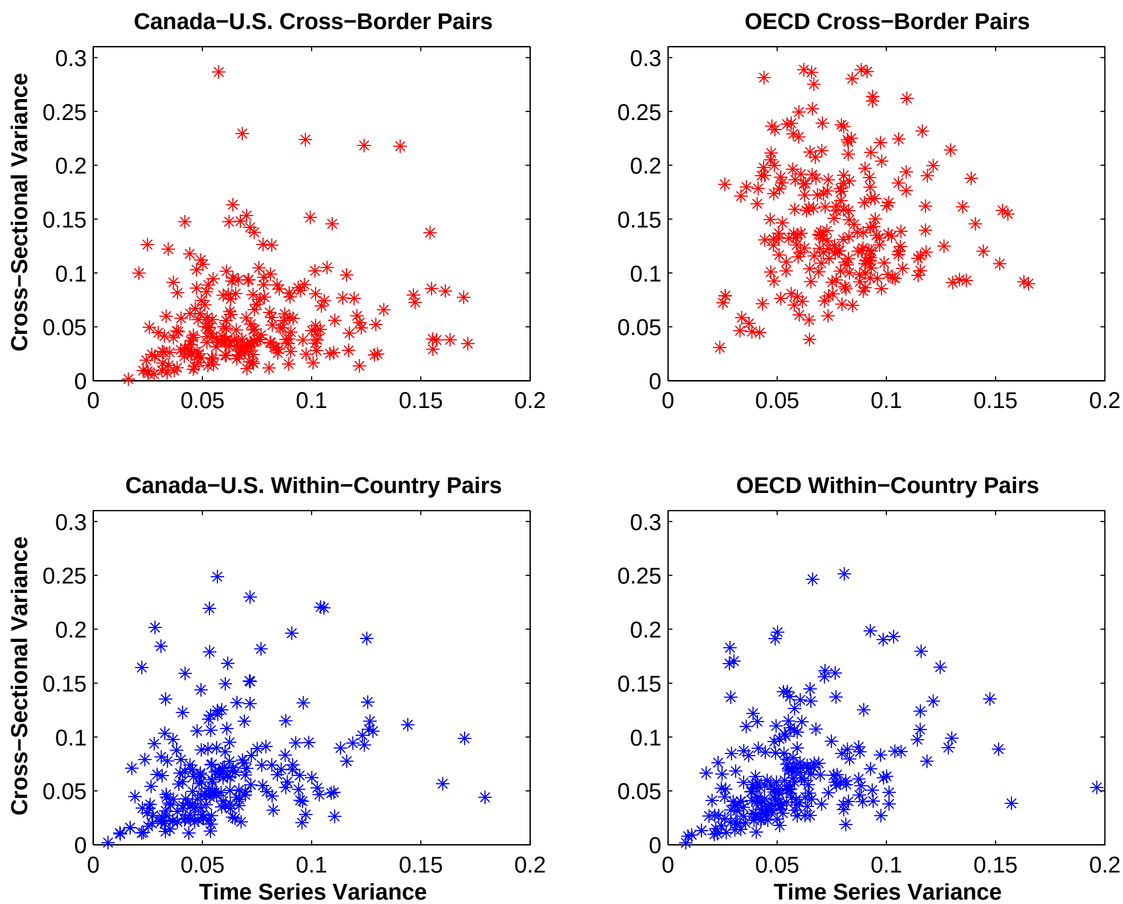

The figure reports scatter plots of time-series variance T_i (horizontal axes) versus cross sectional variance F_i (vertical axes), defined by Equation (3) in Section 3, reproduced here:

$$\begin{aligned} Var_{jk,t}(q_{i,jk,t}\,|\,i) &= Var_{jk}\Big(E_t(q_{i,jk,t}\,|\,i,jk)\Big) + E_{jk}\Big(Var_t(q_{i,jk,t}\,|\,i,jk)\Big) \\ &= T_i + F_i \ . \end{aligned}$$

Regression slope coefficients, (standard errors) are 0.26 (0.09), −0.0129 (0.14), 0.35 (0.11) and 0.52 (0.11) for the NW, NE, SW and SE panels, respectively.

Figure 8

Time-Series Variance in LOP Changes: Common and Good-Specific Variation Distinguished by Tradeability

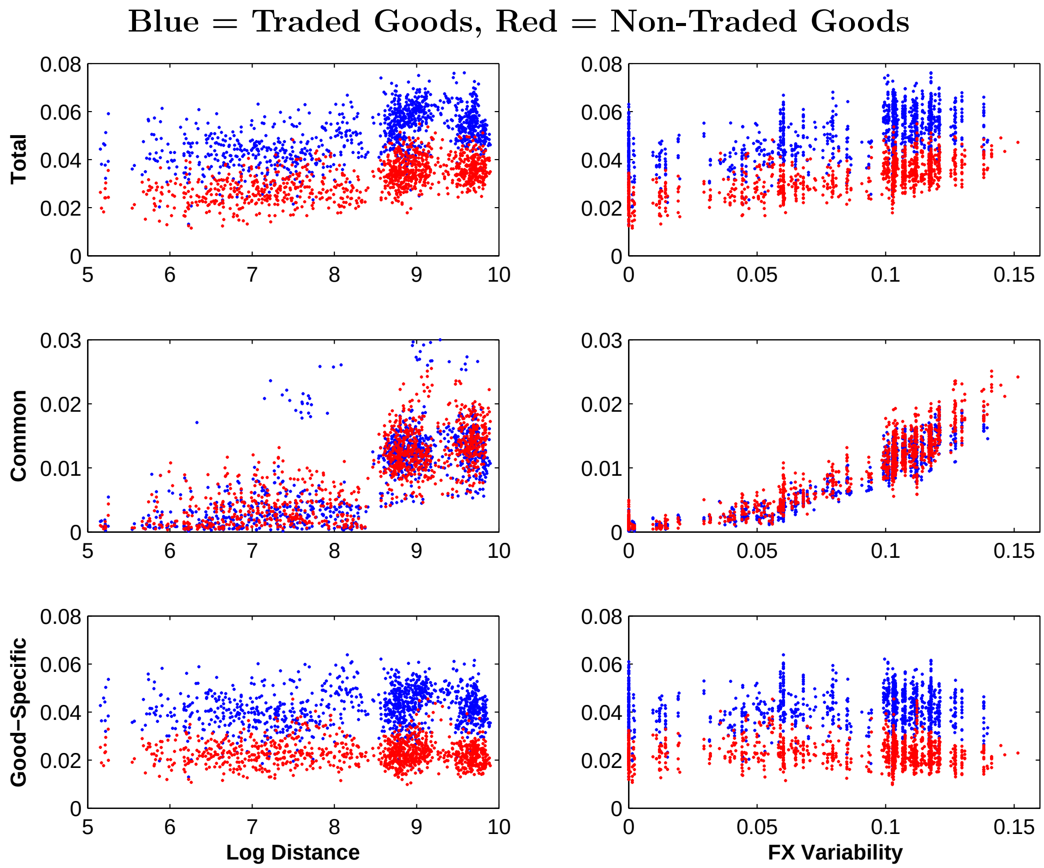

This graph shows that variation in *changes* in LOP deviations — "time-series variation" — is in general larger for traded goods than non-traded goods, and that this effect is almost entirely a feature of good-specific time-series variation as opposed to location-pair specific variation. Each point corresponds to a city-pair-specific object from Equation (6), Section 4, reproduced here:

$$\begin{aligned} Var_{i,t}(\Delta q_{i,jk,t}\,|\,jk) &= Var_t\Big(E_i(\Delta q_{i,jk,t}\,|\,jk,t)\Big) + E_t\Big(Var_i(\Delta q_{i,jk,t}\,|\,jk,t)\Big) \\ &= C_{jk} + I_{jk} \ . \end{aligned}$$

The graph corresponds to all OECD city pairs. The middle two panels report C_{jk}, the common variation across all goods for city-pair jk. The bottom two panels report I_{jk}, the good-specific variation, and the top two panels report the total variation, $C_{jk} + I_{jk}$. The city pairs are organized by either distance from one-another (left-most column of panels), or by nominal exchange rate variability (right-most column of panels).

Figure 9
Distribution of Good-Specific, Local Currency Price Volatility

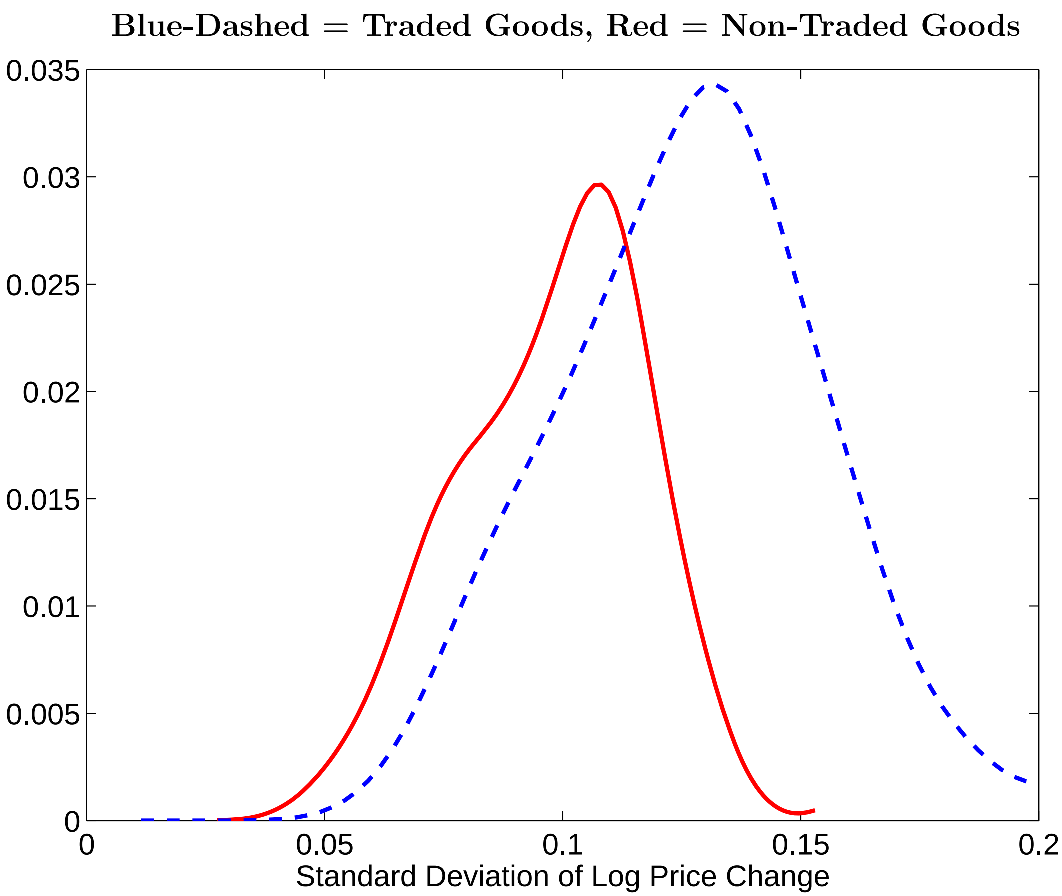

This graph plots kernel density estimates of the distribution of v_i, parsed into traded and non-traded goods, from Equation (9) from Section 5.3. The variables v_i are, for each good i, the averages, across locations, of the time-series volatilities of log changes in local-currency prices. The location-specific inflation rates are removed before computing the volatilities. A high value for v_i simply says that, for this good, the volatility of local-currency price changes is large, on average, across cities. See Section 5.3 for an algebraic derivation.