Preliminary: Please Do Not Circulate

Microeconomic Sources of Real Exchange Rate Variability

Mario J. Crucini* and Chris Telmer,†

April 2008

Abstract

We document two sets of empirical results on the behavior of annual-frequency microeconomic international relative prices. 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 idiosyncratic variation, not country-specific variation such as arises from nominal exchange rates. If you think that annual real exchange rate variability is indicative of 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.

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

†Tepper School, Carnegie Mellon University; chris.telmer@cmu.edu

1 Introduction

We document two sets of empirical results on the behavior of microeconomic international relative prices. 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 idiosyncratic variation, not country-specific variation such as arises from nominal exchange rates.

To understand these results more clearly consider a specific good: an apple. Consider first the cross-sectional results. We begin by computing the LOP deviation for apples between 123 major cities in the world, for each year between 1990 and 2005. We then compute the time-averaged 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 price behavior is broadly representative of most goods and services in the typical urban consumption basket.

What does this tell us about economic models? Sticky-price models 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. Is this likely to matter? We think so. It seems likely that whatever frictions underly long-term LOP deviations are also important for how LOP deviations change over time. At a minimum, our results seem important in an interpretative sense. We don't deny that at a sufficiently short horizon goods prices — in particular international goods prices — are sticky. But 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 following graph of the (indexed) Canada-U.S. real exchange rate.

Figure 1
Canada-U.S. Real and Nominal Exchange Rates
(Annual Data: Indexed 1971=1)

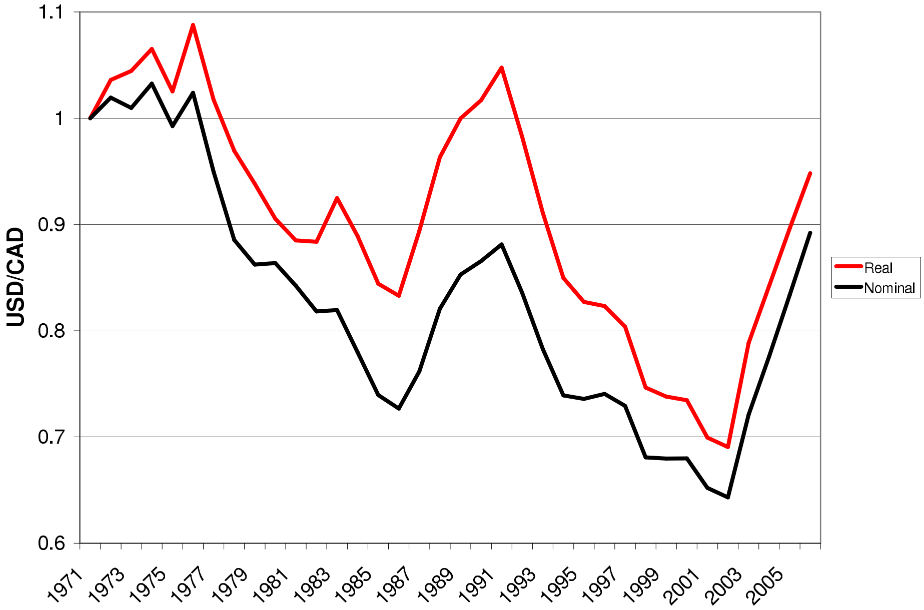

This graph is representative of the common wisdom — often attributed to Mussa (1986) — that real and nominal exchange rates are basically the same thing. This evidence has motivated much economic discussion and model building. It is at the root of the notion that nominal exchange rate variability creates 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.

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.

2 Data

The source of our micro-data on retail prices is 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 (2004), 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 are in the same country). We use two transformations of these prices into log deviations from the law-of-one-price (LOP). The simplest is the good-specific log deviation from the cross-city geometric average:

$$q_{ijt} = \frac{\log(S_{jn,t}P_{ijt})}{\sum_{j=1}^{M}\log(S_{jn,t}P_{ijt}}) \quad , \tag{1}$$

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).

For some questions it will be important to express the LOP deviations in terms of bilateral location-pairs:

$$q_{i,jk,t} = \log(\frac{P_{ij,t}S_{jk,t}}{P_{ik,t}}) \ , \tag{2}$$

These LOP deviations are the date t (log) prices of good i in city j in units of good i in city k.

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 wide array of international location-pairs.

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

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:

$$Var_{jk,t}(q_{i,jk,t} \,|\, i) = Var_{jk}(E_t[q_{i,jk,t} \,|\, i, jk]) + E_{jk}[Var_t(q_{i,jk,t} \,|\, i, jk)] \quad (3)$$

$$= T_i + F_i \;. \quad (4)$$

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 interpret equation (3), 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 LOP deviation for good i between cities j and k.

Figure 3 provides an 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. City-pairs which are separated by the U.S.-Canada border are distinguished from those which are not. The haircut graphs seem to be dominated by long-run means. The variable T_i from equation (4) captures how much of the total variance is attributable to cross-sectional variance in these long-run means. The variable F_i from equation (4) captures what's left: time-series variation around these long-run means. The apples graphs, not surprisingly, seem to indicate that more of the total variation is attributable to the time-series than is the case for the haircut graphs. This seems particularly true for the cross-border pairs.

Why is this decomposition interesting? 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. Consistent with the notation in equation (3), we use the bilateral-pair LOP definition from equation (2).[3] Table 1 reports the average estimate (averaged across goods i) of T_i and F_i from equation (4). 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 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.

The upshot of Table 1 is that, for international city-pairs, roughly 2/3 of the total variation in absolute LOP deviations is attributable to long-run LOP deviations. Our reading of this is that the lion's share of what determines the international micro price distribution falls under the realm of trade theory and has little to do with 'frictions, finance, and fluctuations' as is discussed above.

[3]The appendix shows that, while the alternative definition in equation (1) results in lower overall LOP variability (by construction), our main message remains unchanged.

[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.

4 Variance in Changes in LOP Deviations

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

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

Motivated by Figure 1 from the introduction, we ask to what extent variation in $\Delta q_{i,jk,t}$ is common across all goods for the location-pair jk and to what extent it is idiosyncratic to the individual goods.

Consider a location-pair, jk. We conduct a simple one-way analysis of variance. A useful way to represent this is by defining $f_{jk,t}$ as a common source of variation across all goods i, and $\varepsilon_{i,jk,t}$ as an idiosyncratic source of variation, specific to good i:

$$\Delta q_{i,jk,t} \quad = \quad f_{jk,t} + \varepsilon_{i,jk,t} \quad , \tag{5}$$

where $E(\varepsilon_{i,jk,t} \,|\, f_{jk,t}) = 0$ so that,

$$Var(q_{i,jk,t}) \quad = \quad Var(f_{jk,t}) + Var(\varepsilon_{i,jk,t}) \quad .$$

We can identify $f_{jk,t}$ by averaging across goods: $f_{jk,t} = E_i(\Delta q_{i,jk,t} \,|\, jk, t)$. We then report the variance ratio,

$$\frac{Var(f_{jk,t})}{Var(q_{i,jk,t})} \quad , \tag{6}$$

and we also report the estimates of $f_{jk,t}$ and ask what they might be?

In Figures 4 and 5 we start with the set of Canada-U.S. locations. Figure 4 plots an estimate of the ratio (6). The upper line represents city-pairs separated by a border and the lower line represents city-pairs within either Canada or the U.S.. As one might expect, variation attributable to the common component is larger when a border is involved: 6.9% versus 1.8%, on average. Perhaps less expected is the *magnitude* of the cross-border common variation. On average, 93% of the variation in the LOP changes is good-specific as opposed to country-specific.

What *is* this common component? It is mostly the nominal exchange rate. Figure 5 plots the estimated values of $f_{jk,t}$ alongside the actual depreciation rate on the USD/CAD nominal exchange rate. There is one time-series for $f_{jk,t}$ for each cross-border bilateral city-pair. While there is substantial variation across city-pairs, the message is pretty clear. Most of the common variation in LOP changes is associated with changes in the nominal exchange rate. The average correlation of $f_{jk,t}$ with the nominal USD depreciation rate is 0.89.

It is important to understand that, according to Figure 5, our results *are perfectly consistent* with the existing literature (*e.g.*, Mussa (1986)) which demonstrates a high correlation between nominal exchange rates and *aggregate* real exchange rates. That is, each of the solid (red) lines in Figure 5 is computed as

$$f_{jk,t} = \sum_{i}^{N_{jk}} \Delta q_{i,jk,t}/N_{jk} = \sum_{i}^{N_i} \left(q_{i,jk,t} - q_{i,jk,t-1}\right)/N_{jk} \quad ,$$

where N_i is the number of goods for cities j and k. The variable $f_{jk,t}$ is the change in the equally-weighted average LOP deviation across cities j and k. This is a close cousin of what is typically meant by the aggregate real exchange rate. The main difference is equal weighting versus consumption expenditure weighting. In related work (*e.g.*, Crucini, Telmer, and Zachariadis (2005)) we find that the weighting scheme doesn't matter much for these types of calculations.

Figure 7 expands the analysis to include many more international location-pairs (and a few more intranational pairs). Subject to a missing data criteria, bilateral pairs for all OECD cities in our dataset are included. The horizontal axis is sorted from lowest to highest nominal depreciation rate variability. Compared to Canada-U.S., the common component of the variation is substantially larger for city-pairs with high nominal exchange rate variability. For Amsterdam-Tokyo, for example, the common component accounts for 30% of the variation in the cross-section. Nevertheless, the fraction attributable to factors other than the nominal exchange rate remains quite large.

5 Conclusions

This paper's objective is to document some facts about microeconomic international relative prices which 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.

To us, the economic implications of these results seem important. 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 are essentially the same thing. This is true in our data just as it was in Mussa's. 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. 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

Crucini, Mario J., and Mototsugu Shintani, 2004, Persistence in law-of-one-price deviations: evidence from micro-price data, Unpublished manuscript, Vanderbilt University.

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

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

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.

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	0.146	0.084	0.067	0.573
International	0.126	0.058	0.074	0.462
Intranational	0.133	0.075	0.063	0.564
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

Table 2
Variance in Absolute LOP Deviations: Traded Goods

	Total	Cross Sectional	Time Series	Cross Sectional / Total
U.S.	0.122	0.063	0.064	0.517
Canada-U.S.				
Combined	0.133	0.070	0.068	0.525
International	0.121	0.051	0.076	0.419
Intranational	0.122	0.063	0.064	0.514
OECD				
Combined	0.209	0.135	0.079	0.646
International	0.215	0.140	0.081	0.649
Intranational	0.116	0.060	0.061	0.513
World				
Combined	0.266	0.184	0.088	0.692
International	0.271	0.188	0.089	0.693
Intranational	0.112	0.057	0.059	0.508

Table 3
Variance in Absolute LOP Deviations: Non-Traded Goods

	Total	Cross Sectional	Time Series	Cross Sectional / Total
U.S.	0.156	0.113	0.048	0.721
Canada-U.S.				
Combined	0.162	0.116	0.050	0.717
International	0.138	0.086	0.056	0.628
Intranational	0.152	0.109	0.047	0.717
OECD				
Combined	0.244	0.181	0.067	0.743
International	0.247	0.182	0.069	0.738
Intranational	0.156	0.113	0.048	0.720
World				
Combined	0.290	0.198	0.098	0.682
International	0.294	0.200	0.100	0.681
Intranational	0.148	0.106	0.045	0.718

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

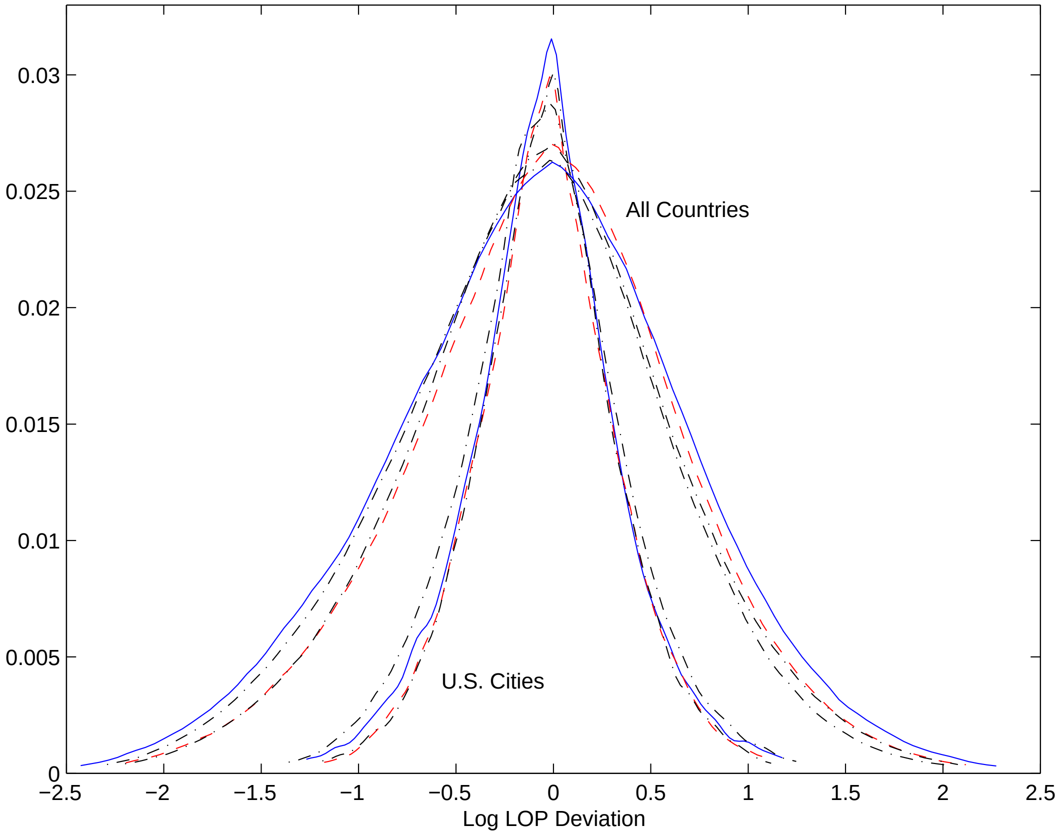

Density of log LOP deviations across (i) U.S. and Canadian location-pairs, and (ii) all location-pairs. The densities describe variation in $q_{i,jk,t}$ across goods i and location-pairs jk, where $q_{i,jk,t}$ is computed as

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

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
Time Series of LOP Deviations: Canada-U.S. Location-Pairs

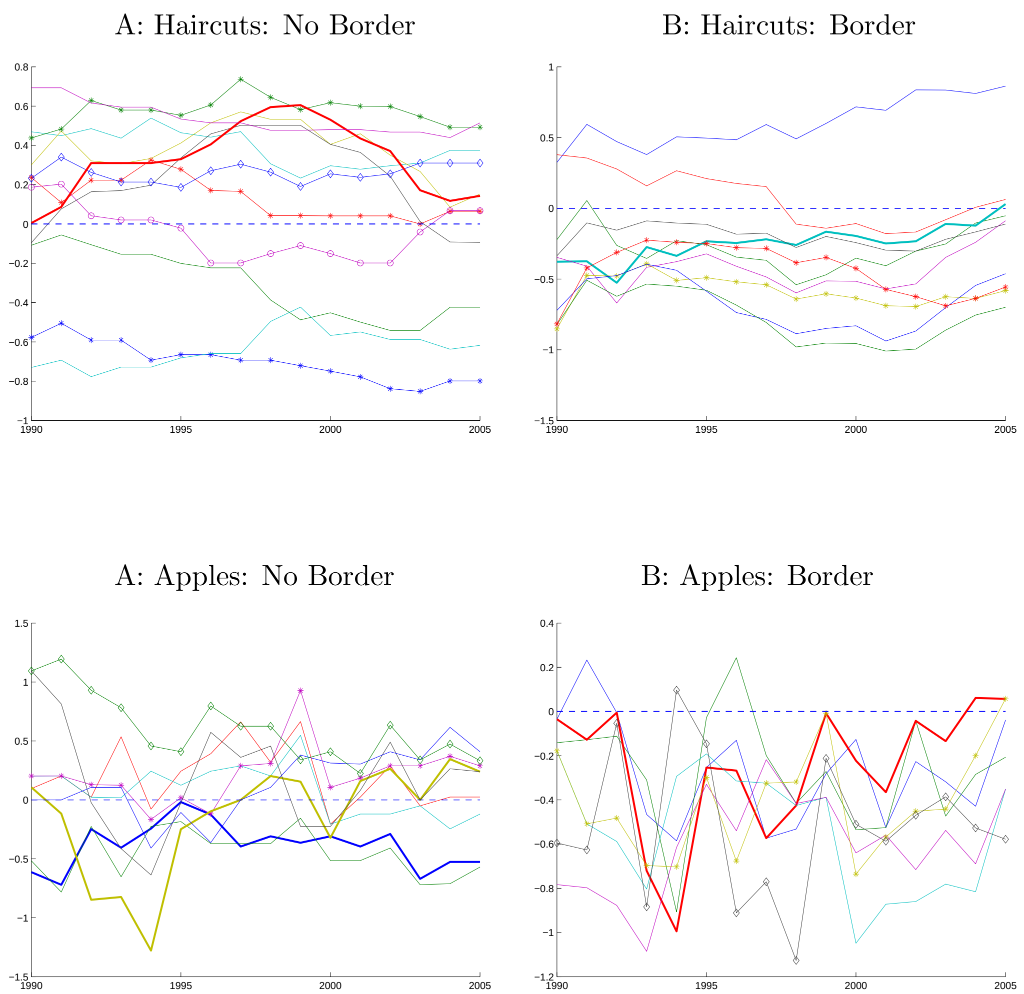

Each graph shows the time-series of log LOP deviations for a particular good, between two North American locations, distinguished by whether or not the cities are separated by the Canada-U.S. border. Our data represent 16 U.S. cities and 4 Canadian locations, which makes for 190 bilateral pairs. To make the graph legible, we randomly select 10 location-pairs for each graph. The qualitative content of the graphs are not changed as we increase the number of pairs and choose a different random sample.

Figure 4
Common Variation in Changes in Good-by-Good LOP Deviations Canada-U.S. City-Pairs

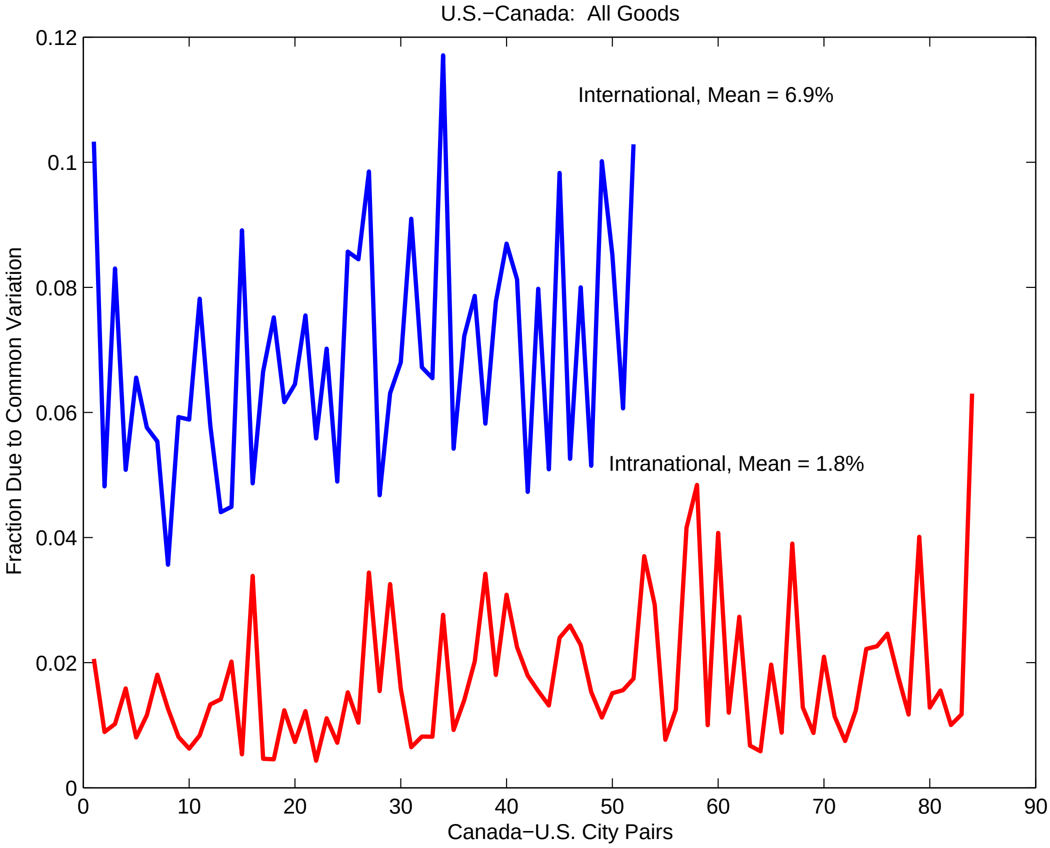

The horizontal axis represents Canada-U.S. bilateral city-pairs. The city-pairs are sorted by geographic distance, from closest apart to farthest apart. There are 4 Canadian cities and 16 U.S. cities, which make for 190 bilateral pairs. Of these, 54 were eliminated due to insufficient data, leaving 84 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). See equation (6) in the text.

Figure 5

Common Component of LOP Deviations and Nominal Exchange Rate

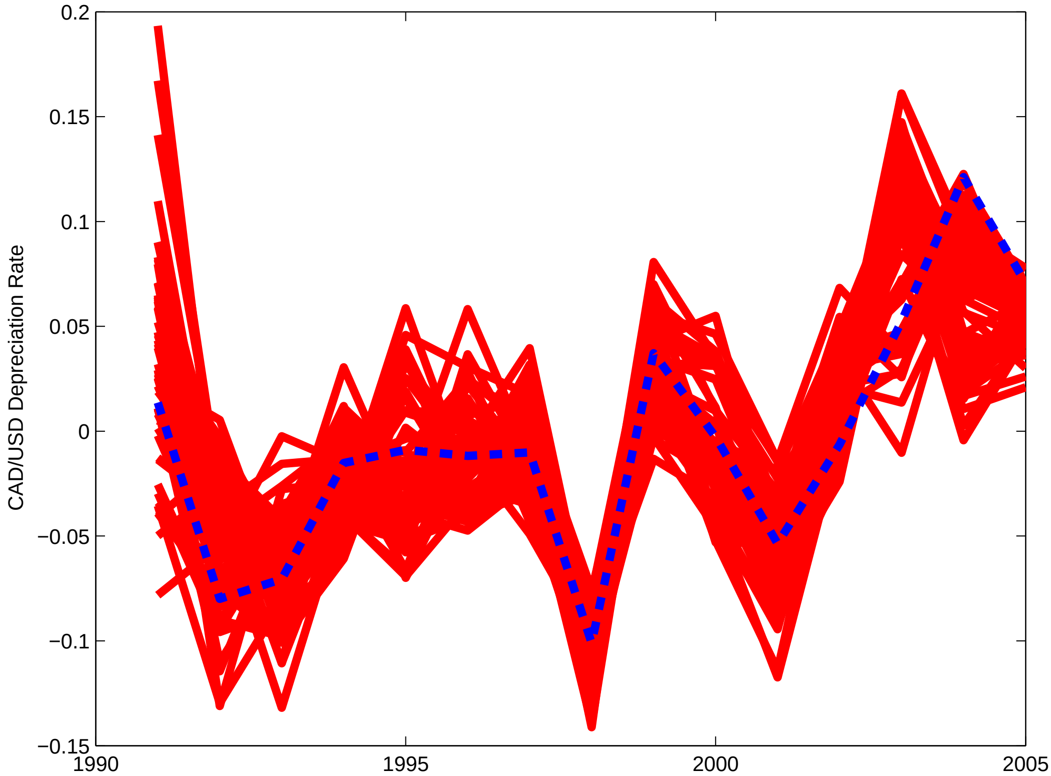

The dashed (blue) line is the annual depreciation rate in the U.S.-Canada nominal exchange rate (USD per CAD). It is computed using the annual, nominal exchange rate data provided by the EIU, which corresponds to the times during the year during which the goods were sampled. However, the graph is not 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,' $f_{jk,t}$, extracted from the *changes* in LOP deviations from equation (5), one for each Canada-U.S. location-pair, jk. The average correlation between each of the solid lines and the dashed line is 0.89.

Figure 6

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

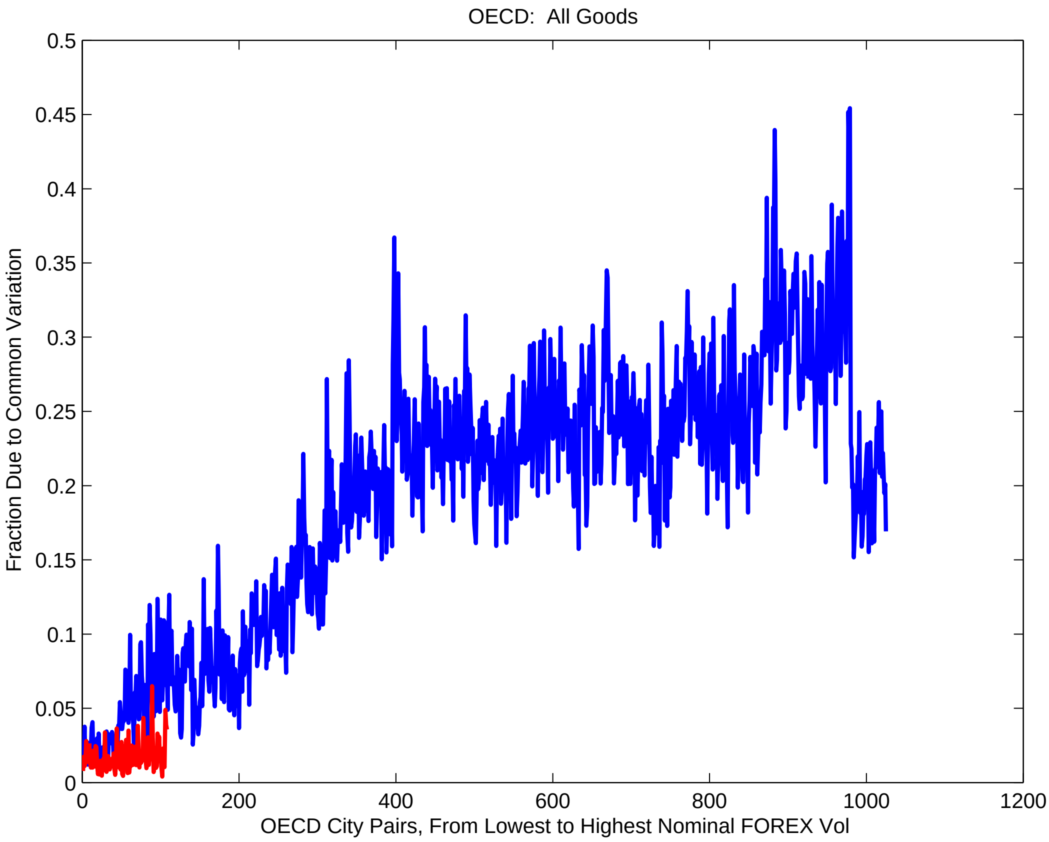

The horizontal axis represents OECD bilateral city-pairs. The city-pairs are sorted by the volatility of the nominal FOREX depreciation rate between each pair of cities, from lowest to highest. The vertical axis represents the fraction of the variance in the good-by-good *changes* in LOP deviations attributable to a source of variation which is common across all goods (for each city-pair). See equation (6) in 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 7
Variation in Absolute LOP Deviations: Time-Series Versus Cross-Sectional