UNDERSTANDING EUROPEAN REAL EXCHANGE RATES

by

Mario J. Crucini, Chris I. Telmer, and Marios Zachariadis

Working Paper No. 01-W20R

September 2001
Revised June 2004

DEPARTMENT OF ECONOMICS
VANDERBILT UNIVERSITY
NASHVILLE, TN 37235

www.vanderbilt.edu/econ

Understanding European Real Exchange Rates*

Mario J. Crucini,[†] Chris I. Telmer,[‡] and Marios Zachariadis[§]

First draft: June 1998
Current draft: June 2004

Abstract

We study good-by-good deviations from the Law-of-One-Price for over 1,800 retail goods and services between all European Union (EU) countries for the years 1975, 1980, 1985 and 1990. We find that for each of these years, after we control for differences in income and value-added tax rates, there are roughly as many overpriced goods as there are underpriced goods between any two EU countries. We also find that good-by-good measures of cross-sectional price dispersion are negatively related to the tradeability of the good, and positively related to the share of non-traded inputs required to produce the good. We argue that these observations are consistent with a model in which retail goods are produced by combining a traded input with a non-traded input.

*The predecessor of this paper circulated as: "What can we learn from deviations from the law of one price?" first presented at the 1998 NBER Summer Institute. In addition to participants at numerous seminars and conferences, we thank Stephen Cosslett, Mick Devereux, Paul Evans, Burton Hollifield, Kjetil Storesletten and Stan Zin for helpful comments. We are particularly indebted to Charles Engel who provided many detailed and constructive comments on earlier versions. We gratefully acknowledge the financial support of the National Science Foundation (SES-0136979) and the Carnegie Bosch Institute.

[†]Department of Economics, Vanderbilt University; mario.j.crucini@vanderbilt.edu

[‡]Carnegie Mellon University and Universitat Pompeu Fabra; chris.telmer@cmu.edu

[§]Department of Economics, Louisiana State University; zachariadis@lsu.edu

1 Introduction

The Law-of-One-Price (LOP) states that identical goods in different countries should have identical prices, once the prices are expressed in common currency units. Purchasing Power Parity (PPP) is the notion that this should hold on average; similar baskets of goods should cost the same once expressed in common units. Each of these propositions is essentially a statement about the *cross-sectional* distribution of international relative prices. Due to data limitations, however, most empirical work has examined the *time-series* distribution of international relative prices. That is, because most data takes the form of index numbers, most of what we know about LOP and PPP deviations involves the volatility and persistence of *changes* in relative prices. We know relatively little about the absolute relative prices themselves. This is particularly troublesome given that economic theory places much starker restrictions on absolute LOP deviations than on their changes.

This paper uses a novel dataset on absolute LOP deviations to bridge this gap. Our data are local-currency retail prices for a broad set of goods and services in all European Union (EU) countries over five-year intervals between 1975 and 1990. The data are quite comprehensive, covering most CPI categories, and are collected with the explicit goal of generating cross-country comparisons of individual goods and services which are as similar as possible.

We focus on the cross-sectional distribution of the LOP deviations from this dataset. For each good, the deviation is defined relative to that good's intra-European average price. We have two sets of results, the first relating to the mean of the distribution and the second relating to the variance. For the mean, we compute, for each country, the average LOP deviation across goods. We find that, after controlling for national income and value-added tax (VAT) differences, most of the means are quite close to zero. In only 4 of 47 cases is a mean greater than 10% in absolute value. This phenomenon is quite stable over the time period we study. In simple words, if we consider most pairs of EU countries with similar incomes and VAT's, at any point in time between 1975 and 1990, there are roughly as many overpriced goods as there are underpriced goods.

Our results about the mean go against the grain of many previous studies which have emphasized large and persistent deviations from PPP. Most of these studies, however, have involved the U.S. dollar and have involved time-series variation over short (*i.e.*, quarterly) horizons. Our results, in contrast, only apply to intra-European prices and are characteristics of absolute price dispersion at four different points in time (1975, 1980, 1985 and 1990). They are also consistent with several independent data sources (documented below). Our main point is that one should be wary of extrapolating the 'consensus' about real-exchange rate behaviour. 'Large and persistent deviations from PPP' is not a good description of Europe between 1975 and 1990.

While the mean is across goods for each country, the variance is across countries for each good. That is, we focus on good-by-good dispersion in the absolute LOP deviations. We find that much of this dispersion can be attributed to two classic characteristics of goods: how tradeable they are and how tradeable the inputs required to produce them are. For example, the average amount of price dispersion in our data is 28%. However, the average is 25% versus 40% for traded versus non-traded goods. Similarly, the average is 25% versus 32% for goods requiring a below-average versus an above-average share of non-traded inputs. Combining these effects in a regression framework, we find that if we consider a non-traded good with the maximum share of non-traded inputs, the predicted price dispersion is 43%. In contrast, the prediction for the most-traded good requiring the lowest share of non-traded inputs is just 12%. The difference of 31% is large, relative to both the unconditional amount of dispersion in the data, 28%, and the range of the good-specific dispersion measures, 2% to 82%. Interestingly, our estimates suggest that the lion's share of the price dispersion is attributable to the tradeability of inputs, not tradeability of the final good.

Our results stand on their own as empirical facts. However, in order to provide a tighter link between facts and theory we begin our paper by outlining a simple model of retail price determination. Retail goods are produced by combining a traded input with a non-traded input. LOP deviations are a convex combination of input cost deviations. This simple theory provides a coherent economic context for the motivation and interpretation of our regressions. It also provides functional form restrictions, which we test. We find that, while the explanatory variables suggested by the theory are supported by the data, the functional form restrictions are not. We discuss reasons for the latter and what it suggests about future research directions.

Our work is most closely related to a large body of empirical work on the international comparison of microeconomic prices. Most data either take the form of index numbers across relatively broad sets of goods, or absolute prices across a very narrow sets of goods. Examples of the former are Engel (1993), Engel and Rogers (1996), Giovannini (1988), Isard (1977), Rogers and Jenkins (1995) and Richardson (1978). Examples of the latter are Cumby (1996), Flam (1992), Rogoff, Froot, and Kim (1995), Ghosh and Wolf (1994), Haskel and Wolf (2000), Knetter (1989, 1993), Lutz (2004) and Parsley and Wei (2000). Our data, in contrast, are distinguished by absolute prices across a very broad set of goods. This allows us to say something akin to absolute PPP and also to relate absolute price dispersion to the characteristics of goods in the cross-section.[1]

[1]Some recent work that has used a broad cross-sectional dataset on absolute prices from The Economist Intelligence Unit includes Crucini and Shintani (2004), Engel and Rogers (2004) Parsley and Wei (2003) and Rogers (2001). This data is quite broad in terms of countries, but

The remainder of the paper is organized as follows. We begin in Section 2 by outlining a conceptual framework, both in terms of data objects and economic theory. Section 3 describes our data. Our analysis is organized around understanding the mean and the variance of the cross-sectional distribution of the LOP deviations from this data. Section 4 examines the mean and Section 5 examines the variance. Section 6 concludes.

2 Conceptual Framework

Before describing our data it is useful to develop a conceptual framework with which to organize it. The price data we seek to understand is denoted P_{ij}: the price of good i in country j, expressed in units of some numeraire currency. None of our analysis will depend on what this numeraire currency is. We transform the price data into log deviations from the geometric-average European price, which we denote q_{ij}:

$$q_{ij} \equiv \log P_{ij} - \sum_{j=1}^{M} \log P_{ij}/M \quad , \tag{1}$$

where M is the number of countries. Note that q_{ij} is independent of the numeraire.

Most of our paper attempts to relate dispersion in q_{ij}, across countries j, to economically meaningful characteristics of the goods i. Following Balassa (1964), Baumol and Bowen (1966), Samuelson (1964), Ethier (1979), Kravis and Lipsey (1983) and Stockman and Tesar (1995), we focus on two characteristics: the international tradeability of the good and the amount of non-traded inputs required to produce the good. To be concrete, consider a typical 'non-traded good,' a taxi ride. As is often asserted, all retail goods involve significant amounts of non-traded inputs such as labor. A taxi ride is no exception. However, less well appreciated is the fact that all non-traded goods involve traded inputs, in this case the automobile and the gasoline. We use the following simple framework to organize this view of what distinguishes goods in different locations.

We view retail goods as being produced by combining a non-traded input with a traded input. With perfect competition, Cobb-Douglas technology and constant returns to scale, we have:

$$P_{ij} = W_j^{\alpha_i} T_{ij}^{(1-\alpha_i)} \quad , \tag{2}$$

lacks the detail and, to some extent, the comparability, of the goods in our Eurostat source. It also covers a more recent time period, 1990-present. In addition Crownover, Pippenger, and Steigerwald (1996) use data on price levels from Internationaler Vergleich der Preise fur die Lebenshaltung published by the German Statistical Office. The data is annual from 1927 to 1992 and covers 6 major industrialized countries.

where W_j is the cost of the non-traded input in country j (*e.g.*, the wage rate), α_i is the share of the non-traded good required to produce good i, and T_{ij} is the cost of the traded input for good i in country j. Inherent in the notation are two standard assumptions. The first is that factor mobility is much higher across sectors within a country than across countries; W_j is country-specific, not good-specific. The second is that retailers in all countries produce good i using the same production technology; α_i is good-specific, not country-specific.

Taking logs of equation (2) and subtracting the geometric average (across j) gives:

$$q_{ij} \quad = \quad \alpha_i w_j + (1-\alpha_i)t_{ij} \ , \tag{3}$$

where, following equation (1), w_j and t_{ij} denote log differences from the cross-country geometric average. Equation (3) says that deviations from LOP should be related to cross-country differences in non-traded and traded factor input costs, as well as differences in the production share attributable to each. Differences in non-traded input costs are the crux of the classic Balassa-Samuelson hypothesis, with w_j being positive for countries with higher productivity in the traded relative to the non-traded sector. Differences in traded costs are often thought of as deriving from transport costs, with t_{ij} being positive for an importer of good i and negative for an exporter.

Our empirical work is organized around equation (3). In Section 4 we begin by examining the country-specific cross-sectional means, which we denote $E(q_{ij}\,|\,j)$. These variables are close cousins of real exchange rates. We find that w_j is important in the sense that the mean is negative for relatively poor countries. We also find that, having controlled for income differences, $E(q_{ij}\,|\,j)$, is often quite close to zero for most European countries. The interpretation offered by equation (3) is that t_{ij} changes sign a lot across goods. In other words, countries import some goods, export others, and on average the effects on LOP deviations cancel out.

In Section 5 we examine the cross-sectional variances, denoted $Var(q_{ij}\,|\,i)$. That is, we examine how cross-country dispersion in LOP deviations varies across goods. According to equation (3),

$$Var(q_{ij}\,|\,i) \quad = \quad \alpha_i^2\, Var(w_j) + (1-\alpha_i)^2\, Var(t_{ij}\,|\,i) + 2\alpha_i(1-\alpha_i)\, Cov(w_j, t_{ij}\,|\,i) \tag{4}$$

We estimate a cross-sectional regression (across i) corresponding to this equation. We use industry-level data on the share of non-traded inputs required for production as a proxy for α_i. We use industry-level data on the tradeability of the final good — as measured by international trade flows divided by total output — to proxy for $Var(t_{ij}\,|\,i)$. We find that an economically important portion of the price dispersion across goods can be accounted for by tradeability and non-traded input cost shares.

Having laid out our data requirements, we now describe our data.

3 The Data

Corresponding to equation (1), we begin by describing our data on local currency prices, P_{ij}, of goods i in countries j.

Our price data are from Eurostat (1975–1990), a series of publications by Eurostat, the statistical agency of the European Community.[2] The publications contain the results of four surveys of retail prices in the capital cities of European Union countries for each of the years 1975, 1980, 1985 and 1990.[3] For our purposes it is sufficient to treat the surveys as four separate cross-sections. However it is important to note that the goods in each cross-section maintain a high degree of comparability, across both locations and time. Table 1 presents basic information about what these cross-sections entail. The 1975 survey covers nine EU countries. Greece, Portugal and Spain are added in 1980. Austria is added in 1985. The number of goods also grows over time, from 658 items in 1975 to 1,896 items in 1990.[4] Finally, a substantial fraction of goods are labeled as 'branded goods.' The importance for our purposes is in terms of comparability of goods. In many cases we are literally talking about the same automobile, portable radio or type of cheese.

As Table 1 indicates, there are a great number of missing observations in the price surveys: 13% in 1975, an abrupt increase in 1980 to 36%, and a similar level in 1985 and 1990. This increase does not reflect a discrete change in the data collection procedure, but instead (we hypothesize) the inclusion of lower-income countries which tend to consume fewer items of the survey's basket of goods. The number of missing observations also differs systematically across countries from a low of about 9% to a high of about 55%. Belgium is consistently the country with the fewest missing observations while Ireland is consistently the country with the most missing observations.

Since our main focus is on explaining price dispersion across countries, we eliminate any good which has an insufficient number of cross-country observations, which we define as 4 in 1975, 5 in 1980, and 6 in 1985 and 1990. The increments reflect the fact that the number of countries in the sample increases over time. We also control for gross measurement error by eliminating goods for which the common-currency price differs from the good-specific median by a factor of 5 or more. These filters reduce our sample of goods from a total of 5,449 to 3,545 with the details for individual years provided in Panel B of Table 1. Of these remaining

[2]The data were not made available to us by Eurostat. We obtained hard-copies of the publications from various libraries and had the data entered into an electronic format by a professional data-entry firm.

[3]Exceptions are the survey data for Germany and the Netherlands in 1980, 1985 and 1990 where the prices are averages across cities within each country.

[4]A comprehensive list of the goods is available at http://bertha.tepper.cmu.edu/eurostat. The raw data will also be made available at this location in the future.

data, the proportion of missing observations never exceeds 25%. Our survey data also contain a large number of brand-name goods, typically accounting for about one-third of the goods that we utilize.

Table 2 reports a number of individual records from the 1985 survey with the goods chosen to be representative of the various categories of goods and services contained in our dataset. All the surveys have a similar structure involving a Eurostat code, a detailed description of the particular good, the units of measure and columns of price data. The retail prices are cash prices paid by final consumers and therefore include taxes, such as VAT (we control for VAT below). The prices are themselves averages of the surveyed prices across different city-wide sales points.[5]

As is evident from the sample of goods reported in Table 2, the surveys are as comprehensive as those used to construct national consumer price indices. We see food items, clothing, major appliances, automobiles, services, and so on. Although Eurostat reports the prices in local currency units, Table 2 presents prices in Belgian francs to facilitate comparisons. The deviations we see from the Law-of-One-Price are suggestive of what is to come. The rental cost of a television, for instance, varies widely across countries whereas the dispersion in the cost of rice is considerably smaller.

The goods-descriptions published by Eurostat are abbreviated versions of those used by the statistical agency to compile the data. The level of detail in the published version also varies across goods. In particular, goods can be placed into two categories: those indicated as selected brands (s.b. in Table 2) and those without such a designation. The reason provided by Eurostat for the selected brand designation is the need for confidentiality. While we might like to know which automobiles are Mercedes and which are Volvos, the record does not provide us with the necessary details. However, it is important to note that the survey is explicitly designed to assure comparability of goods across locations.

One last issue is the exact timing of the surveys. While we find it convenient to refer to our cross-sections by year, in reality the price data for each cross-section is collected in a sequence of surveys. In what we call '1985,' for instance, the prices of most services were collected in September-October, 1985, while prices of most clothing items were collected in December of 1984. The nominal exchange rate data with which we convert prices into a common currency takes explicit account of this timing, taking the form of averages of daily data over the relevant time intervals.

[5]The procedure for selecting sales points follows the practice used to construct national consumer price indices. Sales points are selected by the national statistical offices so that the sample is representative of the distribution of prices in the city with more observations collected for goods having greater price dispersion within the city.

3.1 Supplemental Data

Following the discussion in Section 2, we supplement our retail price data with data on tradedness and production structure. Because these variables are unavailable at the level of the individual good, we assign each good to an industry and use the industry-level measure in place of the good-specific measure.

The variable α_i in equation (4) is the non-traded input share for good i. We measure this as the ratio of non-traded input costs to total cost where both numbers are computed from the 1988 input-output tables of the United Kingdom.[6] Table 3 contains the cost shares of non-traded intermediate inputs by industry. The values range from a low of 4.6% for tobacco to a high of 31.8% for forestry and fisheries.

The other variable in equation (4), $Var(t_{ij} \,|\, i)$, is the cross-sectional variance (across locations j for each good i) in the traded input cost. We proxy this with the tradeability of the final good, measured as the ratio of the total trade among the countries in our sample in a particular industry divided by total output of that industry across the same countries.[7] The (admittedly coarse) idea is that, at the retail level, the intermediate input is often very similar to the final good itself. A seller of computers combines labor with computers and sells computers. The same applies to the seller of furniture. Since computers are more tradeable than furniture, however, the cross-sectional dispersion in the 'intermediate input cost' for computers should be lower than that for furniture. Table 4 reports the trade shares. Among traded goods, average shares for the period 1974 to 1990 range from a low of 13% for tobacco products and 15% for printing, publishing and allied industries, to highs of 122% for professional goods and 121% for office and computing machinery. The average trade share across goods for this period is substantial: 53%.

Comparing some of the numbers in Tables 3 and 4 suggests that the distinction between tradeability and trade in 'middle products' is important. Leather products, for instance, are highly tradeable in the sense that their trade share is high and their non-traded input share is low. The exact opposite is true of most service-related goods. Forestry products (telecommunications), in contrast, are quite tradeable (non-tradeable), yet they require a large (small) amount of non-traded inputs. From an economic standpoint it is unclear whether a 'non-traded' good with substantial traded inputs will exhibit more or less price dispersion than

[6]Non-traded inputs are assumed to include: utilities, construction, distribution, hotels, catering, railways, road transport, sea transport, air transport, transport services, telecommunications, banking, finance, insurance, business services, education, health and other services.

[7]We use the actual trade share whenever trade data is available and assign an index of zero otherwise. The industries assigned zero trade shares are: restaurants and hotels, transport, storage and communication, inland transport, maritime transport, communication, financing, insurance, real estate and community, social and personal services.

a ‘traded good’ with substantial non-traded inputs. Our empirical analysis is designed to separately identify these two economic effects.

The remaining supplemental data we use is data on European nominal exchange rates, national income and VAT rates. Details are provided in the data appendix.

4 LOP Deviations: The Mean of the Distribution

Figure 1 summarizes the empirical distribution of deviations from the LOP, q_{ij} defined in equation (1). The figure contains one chart for each country. Each chart reports one kernel estimate of the density of q_{ij} for each available year. We see four striking features. First, and most obviously, LOP deviations can be large, with the support of most of the densities being on the order of plus or minus 100%. Second, there is a clear income effect. The densities of the relatively poor countries — Greece, Portugal and Spain in particular — seem to be located farther to the left. Third, for the other (relatively wealthy) countries, most of the densities seem to be located near zero. That is, there seem to be roughly as many overpriced goods as underpriced goods. Finally, this latter phenomenon seems quite stable over time. In spite of many relatively large nominal exchange rate movements, the location and shape of the densities don’t seem to move around much.[8]

To be more precise, Table 5 reports the sample means, $E(q_{ij}\,|\,j)$, of the density estimates from Figure 1. The countries are organized from richest to poorest. The values in the left-most columns of Panel A confirm what we see in Figure 1; many of the means for the relatively rich countries are within the interval $\pm 10\%$. There are, however, some exceptions, most notably Denmark. With this in mind, the remainder of the table makes the following adjustments.

1. First, we use industry/year-specific data on VAT rates to reduce each price in our dataset to be a before-VAT rate. The VAT data are described and documented in Appendix A. Note that Denmark is by far the highest VAT country. We then recompute the values q_{ij} and their cross-good averages. The results are reported in Table 5, in the left-most columns of Panel B.

2. Second, we adjust the means for relative income differences via the regression,

$$E(q_{ij,t}\,|\,j,t) = a + bz_{j,t} + \text{residuals} \quad , \tag{5}$$

where $z_{j,t}$ denotes the (log) per-capita income difference between country j and the EU average in date t, and (as above), $E(q_{ij,t}\,|\,j)$ denotes the (cross

[8]There are several glaring exceptions but, interestingly, they go against the notion that movements in the distribution are driven by movements in nominal exchange rates. Between 1975 and 1980, for instance, average prices in Ireland and the U.K. increased substantially relative to Europe (see also Table 5), yet their currencies actually *devalued* relative to Europe.

goods i) average LOP deviation for country j relative to the EU average at date t.[9] The results are reported in the rightmost columns of Table 5.

Net of these adjustments, the average deviations from LOP are substantially reduced and are quite small in magnitude. Out of 47 average LOP deviations, only 4 are outside the interval $\pm 10\%$. The standard deviations of the averages (across countries), shown in the final row of each panel, are reduced by almost half. Figure 2 provides emphasis by plotting the average LOP deviations before and after income/VAT adjustment.[10]

What does the simple model of Section 2 have to say about these results? First, the interpretation of the income adjustment is obvious; it represents cross-country variation in w_j ala Balassa-Samuelson. Second, the fact that (net of VAT) the remaining average price differences are small suggests that t_{ij} changes sign a lot across goods. Countries import some goods, export others, and on average the effects on LOP deviations cancel out.

Further evidence is provided in Table 6, where we distinguish the averages $E(q_{ij} \,|\, j)$ by the tradedness of the goods across which the average is taken. We see that average prices for traded goods tend to be much closer to zero that for non-traded goods. While this is interesting in and of itself, our model suggests two possibilities. First, cross-sectional variation in the α_i's for non-traded goods could bias our income adjustment, which (in Table 6) is done based on the *average* traded/non-traded LOP deviation. Second, it could be the case that, among the set of non-traded goods for a particular country, there is less 'averaging out' associated with some traded inputs being exported and others imported. Our supplementary data are not yet rich enough to distinguish between these possibilities.

5 LOP Deviations: The Variance of the Distribution

We now examine good-by-good price dispersion which, as in equation (4), we measure as the variance of q_{ij} across countries j: $\sigma_i^2 \equiv Var(q_{ij} \,|\, i)$. Figure 3 plots kernel estimates of the density of σ_i for each year. We see that different goods can

[9]The time subscript is included here to indicate that, in our relative income adjustments, we incorporate the important changes which occurred during 1975-1990 (*e.g.*, Ireland).

[10]In a previous version of the paper, available at http://bertha.tepper.cmu.edu/eurostat, we show that CPI expenditure share weighted averages tell very much the same story as the equally-weighted averages reported here. In addition, we perform a number of consistency checks on our data using different data sources. We show that (i) before income/VAT adjustment, our estimates of (absolute) average LOP deviations are quite similar to the absolute PPP data available from the OECD and from the Penn World Tables (PWT), and (ii) when we first-difference our data (necessarily over 5-year intervals), the implications for correlations between changes in real and nominal exchange rates (*e.g.*, Mussa (1986)) are consistent with those from the PWT, the OECD and the national CPI data contained in the International Financial Statistics (IFS).

have substantially different amounts of price dispersion, with the sample standard deviations ranging from 2% to 80%. While it might seem that price dispersion increased after 1975, this is instead an artifact of the sample of countries expanding in 1980 to include lower income countries.

Table 7 reports means of the distributions from Figure 3 broken down by year, tradeability and the fraction of non-traded inputs required for production. We see strong evidence that these variables are important determinants of good-by-good price dispersion. Average price dispersion across traded goods is 25% versus 40% for non-traded goods. The average for goods with above average share of non-traded inputs is 32% versus 25% for goods with a below average share.

Next we consolidate this information in a regression framework.

5.1 Regressions

We begin by slightly rewriting our model of price dispersion, equation (4),

$$\sigma_i^2 \quad = \quad \alpha_i^2\, Var(w_j) + (1-\alpha_i)^2 x_i \tag{6}$$

where, for notational simplicity, we define $x_i \equiv Var(t_{ij}\,|\,i)$. Recall that, as is discussed in Section 3.1, we do not have direct data on t_{ij}. Instead, we proxy x_i itself with a measure of the tradeability of the final good. We also assume that the conditional covariance from equation (4) is zero.[11] Based on this, Table 8 reports estimates of the parameters of the following regression:

$$\sigma_i^2 \quad = \quad a + b\alpha_i^2 + c(1-\alpha_i)^2 x_i + \text{residuals}\ . \tag{7}$$

We see that the coefficients are all highly significant and that the regression function explains a significant portion of the cross-sectional variation in our data on price dispersion. Moreover, the coefficients are economically significant. For example, if we consider the good with the smallest share of non-traded inputs α_i, and the smallest x_i (the latter meaning that the good is highly traded), the predicted price dispersion (based on the pooled regression) is 0.12 (in standard deviation). The opposite case — the good requiring the most non-traded inputs with the highest x_i (lowest tradeability) — gives a value of 0.43. The increase of 0.31 is substantial, both relative to the unconditional (average) dispersion in the data (standard deviation of roughly 0.28), as well as the range of σ_i which is 0.02 to 0.82.

[11] We do not have data which allows us to identify variation in $Cov(w_j, t_{ij}\,|\,i)$ across goods i. Given this, the only alternative is to assume a constant instead of zero. The difference, however, is minor, resulting in only a slightly more complex quadratic function via the inclusion of the term $\alpha_i(1-\alpha_i)$. Moreover, when we do this, we find that (i) the hypothesis that the coefficient on the additional variable is zero cannot be rejected, and (ii) the economic implications of the remaining parameter estimates are almost identical to those of equation (7).

What does theory add here? There are two main differences between the regression (7) and the (obvious) reduced-form regression of σ_i^2 on α_i and x_i. The first is the quadratic functional form. This is a relatively uninteresting artifact of our choice of price-dispersion measure, the variance. Empirically, it turns out to be of little importance. The second difference relates to economics. The essence of the interaction term, $(1-\alpha_i)^2 x_i$, is that the explanatory power of the tradeability variable, x_i, should diminish for goods requiring a high share of non-traded inputs, α_i. If most of the cost of a taxi ride is labor, then the relative price of gasoline shouldn't matter much.

We examined the importance of the interaction term in several ways. All of them indicated that this restriction is inconsistent with our data. For example, if one writes-out equation (4) as a quadratic function of five variables, two of those will be α_i and x_i. In no case were we able to reject the hypothesis that the coefficients on the remaining variables were equal to zero. Similarly, if we add the interaction term to the simple linear specification involving α_i and x_i, we cannot reject the hypothesis that its coefficient is zero. Finally, we took a subsample of goods with relatively high non-traded input shares and regressed the price dispersion measures from this subsample on our tradeability measure. We found strong evidence that the opposite is true. That is, the tradeability measure x_i has more explanatory power conditional on high non-tradedness, not low non-tradedness.

Overall, the message is that α_i and x_i have strong explanatory power for cross-sectional price dispersion, but that the functional-form restrictions from our theory are not an important part of this. This could represent a limitation of the theory. It could also represent a limitation of our data on the characteristics of individual goods. Aside from the fact that we are forced to use industry-level data to try to explain good-specific price dispersion, both of our explanatory variables are proxies for what the theory ultimately calls for. Most serious is $Var(t_{ij}\,|\,i)$ — the dispersion in traded-input costs — which we proxy with the tradedness of the final good. This makes sense if we are talking about a computer retailer putting together a wholesale computer with labor and a shop in order to sell a computer, but it makes less sense if we are talking about a taxi-driver putting together a taxi and some gasoline with labor in order to sell a cab ride. Future work should concentrate on developing better microeconomic data on the characteristics of individual goods.

6 Concluding Comments

A considerable body of empirical and theoretical work in international economics attempts to answer two basic questions. The first question is: What determines whether a deviation from the LOP will be large or small? The second question is: What determines whether a deviation from the LOP will be enduring or short-lived?

The empirical literature on the first question is very limited due to a paucity of absolute price data on comparable goods across international locations. The theoretical literature, in contrast, is well-developed and offers a number of potential answers. The models of Dumas (1992) and Sercu, Uppal, and Hulle (1995) emphasize the magnitude of shipping costs in consumption goods and physical capital, respectively. The models of Krugman (1987) emphasize imperfect competition. The models of Balassa (1964) and Samuelson (1964) emphasize productivity differences across traded and non- traded goods. The models of Ethier (1979) and Jones and Sanyal (1982) emphasize that much of international trade takes place in intermediate inputs, not final goods and services. In each case, the first-order restrictions from theory are on absolute LOP deviations. Our paper represents the development of better data to test these restrictions. What we find is encouraging from the perspective of theory. An economically-important part of the heterogeneity across goods in LOP deviations is related to the tradeability of the good and the share of non-traded inputs required to produce the good.

The empirical literature on the second question is vast, in part because persistence may be studied using relative versions of the LOP and PPP, thus eliminating the need for absolute price data. The consensus is that the half-lives of deviations from parity last between 3 and 5 years.[12] The theoretical literature, however, has focused almost exclusively on the adjustment (or lack of adjustment) of domestic-currency prices to changes in the nominal exchange rate. In particular, the sticky-price models which have followed Obstfeld and Rogoff (1995) are distinguished by assumptions of when firms can change their prices and, therefore, impose first-order restrictions on changes in LOP deviations. Given the persistence of the deviations, most economists believe that sticky-prices in its various forms cannot account for the observed persistence in real exchange rates.

Is there an important link between these questions? We think so, in spite of the fact that the two literatures addressing them have evolved more or less independently. To understand this, consider the "PPP Puzzle" posed by Rogoff (1996); the persistence and conditional volatility in PPP and LOP deviations are too large to be consistent with sticky-price models in conjunction with data on price-setting behaviour. Models of goods-market arbitrage and non-traded inputs, in contrast, restrict the unconditional volatility of PPP and LOP deviations. This is where the link lies. If the unconditional volatility restrictions are supported by the data — as our study suggests — then we are led to rethink our interpretation of the time-series evidence. Rogoff (1996), for instance, interprets the evidence as suggesting that "International goods markets, though becoming more integrated all the time, remain quite segmented, with large trading frictions across a broad

[12]See, for example, Frankel and Rose (1996), Murray and Papell (2003), Rogoff (1996).

range of goods.” Our view is that this is too broad a brush. While it is certainly true for Big Macs, the set of goods for which it is suspect goes beyond gold bullion. Moreover, theory is informative for which goods are in the set and which are not. A promising avenue for future work involves further integration of arbitrage-based models — which impose absolute bounds on LOP deviations — with models which describe behaviour within the bounds, such as sticky-price models and models of traded and non-traded goods.

References

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

Baumol, William J., and William G. Bowen, 1966, *Performing arts: The economic dilemma* (Twentieth Century Fund: New York).

Crownover, Collin, John Pippenger, and Douglas G. Steigerwald, 1996, Testing for absolute purchasing power parity, *Journal of International Money and Finance* 15, 783–796.

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

Cumby, Robert, 1996, Forecasting exchange rates and relative prices: Is what you want what you get with mcparity?, NBER Working Paper No. 5675, July.

Dumas, Bernard, 1992, Dynamic equilibrium and the real exchange rate in a spatially separated world, *Review of Financial Studies* 5, 153–180.

Engel, Charles, 1993, Real exchange rates and relative prices: an empirical investigation, *Journal of Monetary Economics* 32, 35–50.

———, and John H. Rogers, 1996, How wide is the border?, *American Economic Review* 86, 1112–1125.

———, 2004, European product market integration after the euro, *Economic Policy* 19, 347–384.

Ethier, Wilfre J., 1979, Internationally decreasing costs and world trade, *Journal of International Economics* 9, 1–24.

Eurostat, 1975–1990, Price structure of the community countries, Luxembourg: Offices of Official Publications of the European Community, issues from 1975, 1980, 1985 and 1990.

Flam, Harry, 1992, Product markets and 1992: full integration, large gains?, *Journal of Economic Perspectives* 6, 7–30.

Frankel, Jeffrey A., and Andrew K. Rose, 1996, A panel project on purchasing power parity: Mean reversion within and between countries, *Journal of International Economics* 40, 209–224.

Ghosh, Atish R., and Holger C. Wolf, 1994, Pricing in international markets: lessons from the economist, NBER WP No. 4806.

Giovannini, Alberto, 1988, Exchange rates and traded goods prices, *Journal of International Economics* 24, 45–68.

Haskel, Jonathan, and Holger C. Wolf, 2000, Why does the law-of-one-price fail?: A case study, Unpublished manuscript, George Washington University.

Isard, Peter, 1977, How far can we push the law of one price, *American Economic Review* 67 (3), 942–948.

Jones, Ronald W., and Kalyan K. Sanyal, 1982, The theory of trade in middle products, *American Economic Review* 72, 16–31.

Knetter, Michael, 1989, Price discrimination by us and german exporters, *American Economic Review* 79, 198–210.

——— , 1993, International comparisons of pricing-to-market behavior, *American Economic Review* 79, 473–486.

Kravis, Irving B., and Robert E. Lipsey, 1983, Toward an explanation of naional price levels, Princeton Studies in International Finance, No. 52, Princeton: Princeton University.

Krugman, Paul R., 1987, *Pricing to market when the exchange rate changes.* In *Real-financial linkages among open economies* (MIT Press, Cambridge), pp. 49-70.

Lutz, Mattias, 2004, Pricing in segmented markets, arbitrage barriers and the law of one price: evidence from the european car market, *Review of International Economics* 12-3, 456–75.

Maskus, Keith, and Allan Webster, 1995, Factor specialization in U.S. and U.K. trade: Simple departures from the factor-content theory, *Swiss Journal of Economics and Statistics* 1 (131).

Murray, Christian J., and David H. Papell, 2003, Do panels help solve the purchasing power parity puzzle?, Unpublished manuscript, University of Houston.

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.

Obstfeld, Maurice, and Kenneth Rogoff, 1995, Exchange rate redux, *Journal of Political Economy* 103, 624–660.

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.*

——— , 2003, How big and heterogeneous are the effects of currency arrangements on market integration? a price-based approach, Unpublished manuscript, Vanderbilt University.

Richardson, David J., 1978, Some empirical evidence on commodity arbitrage and the law of one price, *Journal of International Economics* 8, 341–351.

Rogers, John H., 2001, Price level convergence, relative prices, and inflation in europe, *Federal Reserve Board International Finance Discussion Paper* No. 699.

——— , and Michael Jenkins, 1995, Haircuts or hysteresis? sources of movements in real exchange rates, *Journal of Internationa Economics* 38, 339–360.

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

Rogoff, Kenneth S., Kenneth A. Froot, and Michael Kim, 1995, The law-of-one-price over 700 years, NBER WP No. 1532.

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

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

Stockman, Alan C., and Linda L. Tesar, 1995, Tastes and technology in a two-country model of the business cycle, *American Economic Review* 85, 168–185.

A Data Appendix

National retail price data. The retail price data were compiled and published by Eurostat, the statistical agency of the European Community, in cooperation with the national statistical agencies of the countries that participated. While all of the price data we utilize is published we are unaware of available electronic copies. All prices refer to cash prices paid by final consumers, including taxes, both VAT and any others paid by the purchaser. Sales points are selected is such a way that the sample selected is representative of the distribution in the capital city. Prices are collected at different locations so that the average price is representative of the distribution within the city. These data were not available electronically so we had a private firm key-punch the commodity codes, commodity descriptions and prices into an electronic format.

Data Reconciliation. In order to explain the price dispersion across goods that exists in our data we constructed measures of tradeability and the costs of non–traded inputs into production. The constructed variables for tradeability, and non–traded inputs from input-output tables are available at different levels of detail. For this reason, and in order to make the most of the information available for each of these factors, we matched the retail price data with each of the variables using two-digit, three-digit, and four-digit classifications depending on the level of detail available for each of the variables rather than attempting to match all variables using the same level of detail. The input-output data are also available at a three-digit level of detail that extends to four-digits for some industry groups. In order to reconcile the data as accurately as possible, we used the ISIC codes and descriptions available in the User Guide of the OECD International Sectoral Database. A description of each of our variables is as follows.

- *Tradeability.* We obtained data on imports, exports, and gross output for the period 1973 to 1990 from the 1994 edition of the OECD STAN Database. We use thirty-two non-overlapping subdivisions of manufacturing for which sufficient data are available and to the extent that they are relevant to the commodities in our price dataset. We also constructed additional tradeability indices for agriculture (sector 1) and electricity, gas, and water (sector 4) using the 1994 edition of the OECD Sectoral Database, which provides value-added instead of gross output data. Unfortunately, we only have this data for six countries: Belgium, Denmark, France, Germany, Italy, and the United Kingdom. We assigned a zero trade share to the following industries: restaurants and hotels, transport, storage and communication, inland transport, maritime and air transport, communication, financing, insurance, real estate and business services, community, social and personal services. Otherwise, our measure of tradeability is calculated as:

$$\theta_k = \frac{\sum_{j=1}^{m_k} (X_{kj} + M_{kj})}{\sum_{j=1}^{m_k} Y_{kj}} \quad ,$$

  where for each sector k we sum over all countries j which have data for that sector. X_{kj} (M_{kj}) stands for exports (imports) of sector k from country j and Y_{kj} stands for the gross output of sector k by country j.

- *Input–Output Data.* We use the input-output matrix for the United Kingdom in 1988. These data were compiled by Keith Maskus and Allan Webster (1995). We thank Tom Prusa for suggesting this data source, which is available at the National Bureau of Economic Research home page. Non–traded inputs are assumed to include: utilities, construction, distribution, hotels, catering, railways, road transport, sea transport, air transport, transport services, telecommunications, banking, finance, insurance, business services, education, health and other services. We compute the cost share of non–traded intermediate inputs computed as,

$$\Phi_k = \sum_{s=1}^{S} \phi_{ks} \quad ,$$

where ϕ_{ks} is the share of non–traded intermediate input s in the total cost of the output of sector k.

Nominal exchange rates. We obtained daily data on nominal exchange rates from the New York Federal Reserve Bank's web page. The Eurostat survey provides a 2-3 month window during which the survey was conducted for each good (for each of our 4 cross-sections). Accordingly, the nominal exchange rate we use is an average of the daily values over the respective months for each good. Data and programs are available at http://bertha.tepper.cmu.edu/eurostat.

VAT data. We obtained country-specific data on VAT rates for 23 different categories of goods and services for each of the years 1975, 1980, 1985 and 1990. We then categorized each of our goods into one of these categories and then divided each individual price by it's associated country/year-specific VAT rate to arrive at before-VAT prices. All of our analysis in Tables 6-8 and Figure 3 use these before-VAT prices. The VAT rates acros EU countries are taken primarily from the European Commission publication "VAT rates applied in the member states of the European Community" (2002), the OECD publication "Taxing Consumption", and the Ernst and Young publication "Vat and Sales Taxes Worldwide: A Guide to Practice and Procedures in 61 Countries" (1996). Secondary sources on VAT rates include the Tax Executives Institute publication "Value-Added Taxes: A Comparative Analysis" (1992), Alan Tait's "Value Added Tax: International Practice and Problems" (1988), and George Carlin's "Value-Added Tax: European Experience and Lessons for the United States" (1980). The data are available at http://bertha.tepper.cmu.edu/eurostat.

Real GDP Per-Capita was obtained from Penn World Tables 6.1 for each of the years corresponding to our cross-sections.

A statistical appendix and more extensive data appendix are available from the authors upon request.

Table 1
Scope of the Price Surveys

	1975	1980	1985	1990
Panel A: Raw survey data				
Number of countries	9	12	13	13
Number of goods	658	1090	1805	1896
Proportion missing	13%	36%	38%	44%
Least missing[a]	9%	23%	25%	32%
Most missing[a]	27%	47%	53%	55%
Proportion of branded goods	31%	42%	48%	54%
Panel B: After eliminating goods with insufficient data and outliers				
Number of countries	9	12	13	13
Number of goods	594	686	1164	1101
Proportion missing	10%	17%	19%	23%
Least missing[a]	4%	3%	7%	9%
Most missing[a]	22%	28%	37%	34%
Proportion of branded goods	31%	28%	33%	38%

The countries are Belgium, Denmark, France, Germany, Ireland, Italy, Luxembourg, Netherlands, and United Kingdom (all years), Greece, Portugal, Spain (added in 1980) and Austria (added in 1985). In Panel B, we eliminate any observation which has a relative price more than or less than 5 times that of the EU median. Then we eliminate goods for which there are more than 4 missing observations in 1975, 5 in 1980, and 6 in 1985 and 1990 (this ensures that, for each good, we have data for at least half of the countries). The category 'Proportion of branded goods' refers to the proportion of goods for which two or more brands exist in the price survey. Explicit descriptions of the goods are available at http://bertha.tepper.cmu.edu/eurostat/.

[a]Belgium is the country with the least number of missing observations in every year and Ireland is the country with the greatest number of missing observations in each year except for 1990 in which case the United Kingdom has the most missing.

Table 2: Sample Records from Price Survey

Code	Good Description	Units	AUS	BLG	DEN	FRF	GER	GRC	IRL	ITL	LUX	NET	POR	ESP	UK
11111	Long grn rice, carton	500g	26	50	48	52	52	42	51	35	52	36		39	32
11251	Chicken-fresh, 70% pres	1.2kg	107	204	207	203		115	159	163			88		148
11421	Cond. milk 9-10 % b.f.	410g		48			56	31				54	56	42	42
11621	Dried almonds	100g	28	54		61	34	28	43	35	54	49	37	49	50
11911	Ground blended coffee	1kg	234	439	425	350				586		355	451	437	446
13112	Liqueur - s.b.	0.7L		579	1294	542	599	589	980	414	502	599	1152	491	910
21112	Man's k-way jacket	1nb	1584	544	1596	1022	824	1373	1331	872	996	1196	904	713	
22121	Ladies boots, box caf	1pr	3720	4027	4095	3733	3692	2645	2362		4422	2615	2745	2951	3176
41112	Chest of drawers	1nb	2834	6988	2192	11400	4889	2904	5078	7186	3990	7742	3529	4630	4065
42111	Spring mattress : s.b.	1nb	4070		3923	3253	2674					3803			
43121	Dishwasher: s.b.	1nb		23480						23548			22705	22823	
43161	Iron : steam, s.b.	1nb	2436		2201	1925	2089	3018		1916	1745	1810	2495	2450	1485
45111	Washing powder : s.b.	700g	414	299		337	316		310		313				
52211	Hearing aid : s.b.	1nb				29528	25976					23343	17734	19134	26054
61113	Car: 1200-1700cc, s.b.	1nb	357298		524411	341213						346786	356516		
61212	Bicycle : racing, s.b.	1nb				15147	19314				18000	17788			15238
63111	Bus fare, 6km, no transf	1tik	52	30	50	58	43	9	41	12	25	39	17	16	33
71131	Record player, s.b.	1nb	7091	5190		5158	4172	10955		4719	4848	4861	4194	7828	
71311	Cassette for game : s.b.	1nb	873	850		1299			948	1072	1066	400			764
72221	Rental of television	1mo		2296	1340	1631	1191	2962	1382	1285	3136	1490	1670	1795	1487

Sample of 20 records from the original 1,805 records of the 1985 survey. Currency units are Belgian Francs. Countries are Austria (AUS), Belgium (BLG), Denmark (DEN), France (FRF), Germany (GER), Greece (GRC), Ireland (IRL), Italy (ITL), Luxembourg (LUX), The Netherlands (NET), Portugal (POR), Spain (ESP) and the United Kingdom (UK). Branded goods are denoted 's.b.' ('selected brand'). Prices are rounded to the nearest franc, although the raw data (and our analysis) are in units of francs and centimes. Blank entries denote missing data.

Table 3
Non-Traded Input Cost Shares

Category	Input share	Category	Input Share
Agriculture	0.144	Glass	0.195
Forestry and fishing	0.318	Clay refractories	0.201
Milk and milk products	0.080	Cement, concrete	0.259
Meat, fruit, veg, fish processing	0.096	Iron and steel; plant	0.111
Oils and fats, grain products	0.124	Non-ferrous metals	0.092
Bread, biscuits etc	0.104	Metal castings etc	0.119
Sugar, confectionery	0.139	Office machinery, computers	0.103
Foods, nes	0.174	Other machinery	0.116
Alcoholic drink	0.097	Telecomms equipment	0.120
Soft drink	0.160	Domestic electric appliances	0.135
Tobacco	0.046	Electronic consumer goods	0.113
Woven textiles	0.095	Electronic components	0.147
Hosiery, other knitted goods	0.097	Electric lighting equipment	0.106
Carpets etc	0.112	Shipbuilding and repairing	0.134
Clothing and furs	0.096	Motor Vehicles and parts	0.090
Leather and leather goods	0.073	Other vehicles	0.110
Footwear	0.086	Instrument engineering	0.149
Wood furniture	0.107	Other manufacturing	0.122
Paper and board products	0.135	Utilities	0.157
Printing and publishing	0.211	Hotels, catering etc	0.144
Synthetic resins, man-made fibers	0.134	Railways	0.272
Paints, dyes etc	0.154	Road transport etc	0.151
Soap and toiletries	0.226	Air transport	0.253
Chemicals nes	0.133	Transport services	0.204
Mineral oil processing	0.054	Telecomms and postal	0.124
Rubber and plastic products	0.127	Business services etc	0.134
		Other services	0.304

Values are based on the author's calculations, deriving from the 1988 input-output matrix for the U.K., compiled by Maskus and Webster (1995) and available at http://www.nber.org. Further details are available in the data appendix.

Table 4
Trade Share

Industry	Trade Share			
	1975	1980	1985	1990
Agriculture, hunting, forestry and fishing	0.26	0.41	0.38	0.56
Food	0.27	0.26	0.29	0.31
Beverages	0.20	0.25	0.28	0.33
Tobacco	0.06	0.17	0.21	0.21
Textiles	0.44	0.53	0.64	0.70
Wearing apparel	0.48	0.47	0.58	0.71
Leather products	0.47	0.63	0.77	0.82
Footwear	0.42	0.58	0.73	0.77
Furniture and fixtures	0.17	0.23	0.29	0.32
Paper and paper products	0.40	0.45	0.54	0.61
Printing and publishing	0.15	0.15	0.16	0.15
Industrial chemicals	0.58	0.75	0.90	0.98
Other chemicals	0.29	0.44	0.54	0.57
Chemical products, n.e.c.	0.25	0.46	0.59	0.65
Misc. products of petroleum and coal	0.62	0.46	0.42	0.36
Rubber products	0.40	0.50	0.58	0.62
Plastic products, n.e.c.	0.21	0.27	0.30	0.28
Pottery, china etc.	0.23	0.28	0.26	0.25
Non-metal products, n.e.c.	0.19	0.19	0.21	0.22
Iron and steel	0.37	0.43	0.48	0.52
Non-ferrous metals	0.69	0.66	0.71	0.73
Fabricated metal products, except machinery and equipment	0.28	0.34	0.39	0.40
Office and computing machinery	1.08	1.24	1.35	1.40
Machinery and equipment, n.e.c.	0.56	0.56	0.60	0.62
Electrical machinery	0.38	0.44	0.52	0.61
Radio, television and communication equipment	0.50	0.49	0.62	0.79
Electrical apparatus, n.e.c.	0.39	0.47	0.49	0.54
Shipbuilding and repairing	0.53	0.40	0.42	0.44
Motor vehicles	0.51	0.58	0.66	0.72
Motorcycles and bicycles	0.55	0.50	0.62	1.00
Professional goods	1.00	1.13	1.52	1.49
Other manufacturing n.e.c.	0.94	1.20	1.39	1.47
Electricity, gas and steam and water	0.63	1.16	0.84	0.68

Source: OECD Sectoral Database and OECD STAN Database. The following industries (not shown) have been assigned a zero trade share: restaurants and hotels, transport, storage and communication, inland transport, maritime and air transport, communication, financing, insurance, real estate and business services, community, social and personal services.

Table 5
Average (Across Goods) LOP Deviation

Panel A: Without VAT Adjustment

Country	z_j	Without Income Adjustment 1975	1980	1985	1990	With Income Adjustment 1975	1980	1985	1990
Luxembourg	0.299	–0.009	0.030	–0.043	–0.027	–0.097	–0.045	–0.137	–0.165
Denmark	0.260	0.177	0.215	0.222	0.219	0.081	0.132	0.125	0.155
Netherlands	0.126	0.023	0.018	–0.055	–0.007	–0.031	–0.022	–0.088	–0.030
France	0.126	0.119	0.112	0.075	0.072	0.084	0.072	0.034	0.039
Belgium	0.122	0.042	0.044	0.067	0.053	0.006	0.002	0.032	0.023
Germany	0.105	0.066	0.089	0.013	0.053	0.046	0.055	–0.028	0.028
Austria	0.096			0.096	0.078			0.065	0.049
Italy	0.061	–0.055	–0.087	–0.007	0.034	–0.050	–0.105	–0.030	0.014
U.K.	0.035	–0.200	0.050	–0.001	–0.037	–0.213	0.054	–0.011	–0.038
Spain	–0.192		–0.100	–0.082	–0.067		–0.018	0.013	0.015
Greece	–0.257		–0.170	–0.170	–0.173		–0.098	–0.067	–0.023
Ireland	–0.323	–0.195	–0.029	0.048	0.012	–0.052	0.104	0.170	0.102
Portugal	–0.457		–0.213	–0.174	–0.254		–0.046	0.008	–0.114
Std Dev	0.232	0.129	0.123	0.109	0.118	0.094	0.077	0.083	0.083

Panel B: With VAT Adjustment

Country	z_j	Without Income Adjustment 1975	1980	1985	1990	With Income Adjustment 1975	1980	1985	1990
Luxembourg	0.299	0.029	0.056	–0.020	0.018	–0.035	0.002	–0.089	–0.082
Denmark	0.260	0.143	0.108	0.124	0.144	0.074	0.048	0.053	0.098
Netherlands	0.126	0.024	–0.001	–0.073	–0.012	–0.015	–0.031	–0.097	–0.029
France	0.126	0.074	0.065	0.041	0.056	0.049	0.036	0.011	0.032
Belgium	0.122	0.016	0.015	0.036	0.046	–0.011	–0.016	0.010	0.024
Germany	0.105	0.078	0.077	–0.002	0.060	0.063	0.053	–0.031	0.042
Austria	0.096			0.044	0.050			0.021	0.028
Italy	0.061	–0.042	–0.097	–0.028	0.032	–0.039	–0.110	–0.045	0.017
U.K.	0.035	–0.152	0.056	0.016	0.007	–0.161	0.059	0.009	0.005
Spain	–0.192		–0.018	0.003	–0.048		0.041	0.072	0.012
Greece	–0.257		–0.088	–0.083	–0.169		–0.036	–0.009	–0.061
Ireland	–0.323	–0.200	–0.060	0.035	0.029	–0.096	0.037	0.123	0.094
Portugal	–0.457		–0.134	–0.087	–0.251		–0.013	0.044	–0.150
Std Dev	0.232	0.111	0.078	0.060	0.103	0.076	0.050	0.062	0.069

The variable z_j denotes average log income differences, 1975-1990, relative to the EU average. The remaining data are equally-weighted averages of LOP deviations, across goods i, for each country j: $E(q_{ij} \mid j)$. Standard errors were reported in the previous version of the paper (available at http://bertha.tepper.cmu.edu) but are omitted here for space considerations. All standard errors are between 0.01 and 0.02 with the exception of those from Greece, Portugal and Spain which are between 0.015 and 0.025. Income and VAT adjustment are described in the text and the appendix. For these calculations we eliminate multiple brands of the same good.

Table 6: Average (Across Goods) LOP Deviation by Tradedness

	1975			1980			1985			1990		
	All	Traded	Non-Traded	All	Traded	Non-Traded	All	Traded	Non-Traded	All	Traded	Non-Traded
Luxembourg	-0.035	0.006	-0.273	0.002	0.009	-0.031	-0.089	-0.064	-0.239	-0.082	-0.043	-0.318
Denmark	0.074	0.082	0.020	0.048	0.067	-0.064	0.053	0.063	0.001	0.098	0.101	0.083
Netherlands	-0.015	-0.023	0.016	-0.031	-0.055	0.109	-0.097	-0.108	-0.027	-0.029	-0.049	0.094
France	0.049	0.045	0.063	0.036	0.026	0.098	0.011	0.001	0.068	0.032	0.038	-0.005
Belgium	-0.011	-0.012	-0.003	-0.016	-0.034	0.094	0.010	0.008	0.026	0.024	0.018	0.067
Germany	0.063	0.049	0.137	0.053	0.040	0.128	-0.031	-0.046	0.051	0.042	0.022	0.149
Austria							0.021	0.010	0.083	0.028	0.015	0.124
Italy	-0.039	0.000	-0.247	-0.110	-0.097	-0.186	-0.045	-0.026	-0.165	0.017	0.039	-0.108
U.K.	-0.161	-0.151	-0.222	0.059	0.036	0.184	0.009	-0.026	0.217	0.005	-0.056	0.326
Spain				0.041	0.044	0.028	0.072	0.079	0.017	0.012	0.015	-0.020
Greece				-0.036	-0.010	-0.160	-0.009	0.007	-0.111	-0.061	-0.035	-0.278
Ireland	-0.096	-0.115	0.017	0.037	-0.002	0.261	0.123	0.085	0.351	0.094	0.059	0.298
Portugal				-0.013	0.026	-0.222	0.044	0.076	-0.133	-0.150	-0.117	-0.301
Std Dev	0.076	0.076	0.151	0.050	0.047	0.153	0.062	0.059	0.158	0.069	0.058	0.211

See Notes to Table 5. Data here are identical (income/VAT adjusted) except for the distinction between traded and non-traded goods. A ‘non-traded good’ is defined as having a trade share equal to zero according to Table 4.

Table 7
Averages of Good-Specific Measures of Price Dispersion

	1975	1980	1985	1990	Average
All Goods	0.2290	0.2941	0.3024	0.2855	0.2777
Non-Traded Goods	0.3138	0.4146	0.4252	0.4537	0.4018
Traded Goods	0.2164	0.2750	0.2846	0.2596	0.2589
Above Avg Share of Services	0.2619	0.3372	0.3464	0.3378	0.3208
Below Avg Share of Services	0.2116	0.2703	0.2779	0.2551	0.2537

Values are averages of good-specific measures of price dispersion. Each value is the average (across goods i) of $Var(q_{ij} \,|\, i)^{1/2}$, the good-by-good sample standard deviation, where the standard deviation is across countries j.

Table 8
Regression Estimates

$$\sigma_i^2 \;=\; a + b\,\alpha_i^2 + c\,(1-\alpha_i)^2 x_i + \text{residuals}$$

Year	a	b	c	R_a^2	R_d^2
1975	0.083	0.299	0.090	0.433	0.120
	(0.006)	(0.115)	(0.013)		
1980	0.099	0.921	0.051	0.291	0.107
	(0.007)	(0.143)	(0.012)		
1985	0.115	0.791	0.059	0.319	0.106
	(0.006)	(0.118)	(0.008)		
1990	0.122	1.033	0.100	0.537	0.291
	(0.005)	(0.106)	(0.008)		
Pooled	0.102	0.868	0.062	0.417	0.129
	(0.003)	(0.064)	(0.005)		

Estimates of the parameters of the regression at the top of the page, based on equation (7) in the text. σ_i^2 denotes our good-specific measure of price dispersion: the sample variance of q_{ij} from equation (1) in the text. α_i denotes the non-traded input share for good i and x_i is (the negative of) the trade share of good i (see Section 3.1 for details). Standard errors in parentheses. Because our explanatory variables are averaged across different numbers of goods within industry groups, the residuals will be heteroskedastic. We therefore use a feasible GLS estimator (details are provided in the appendix). R_a^2 denotes the regression R^2 that results in averaging the dependent variable in the same manner as the explanatory variables. R_a^2 denotes the regression R^2 from the 'raw regression' where, necessarily we are trying to account for variation in σ_i^2 using variables which have (potentially) had some of their explanatory variation averaged-away.

Figure 1
Empirical Distributions of LOP Deviations

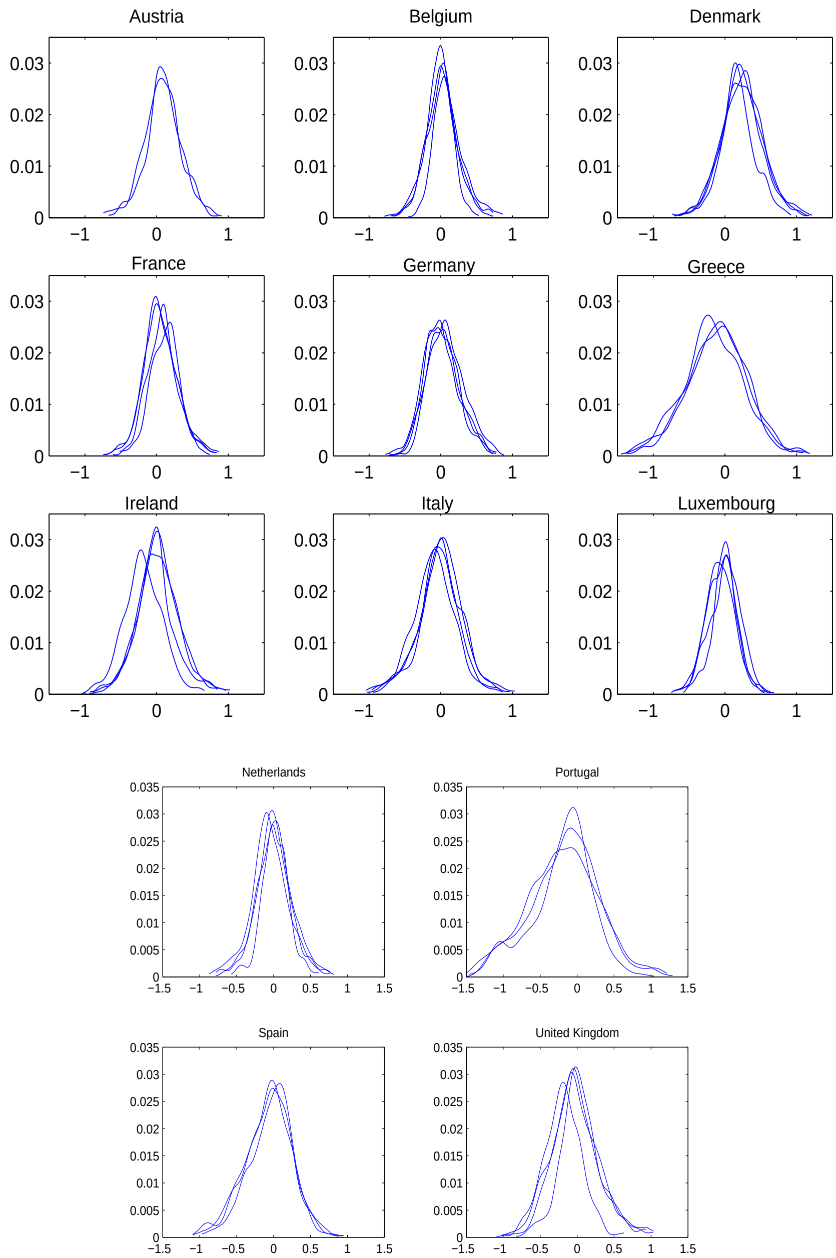

Each line represents an estimate of the density of good-by-good deviations from the Law-of-One-Price, relative to the European average price, for each of the years 1975, 1980, 1985 and 1990. The exceptions are Austria, where we do not have data for 1975 and 1980, and Greece, Spain and Portugal, where 1975 is missing.

Figure 2
Average LOP Deviations

(raw data (circles) and income/VAT adjusted (asterisks))

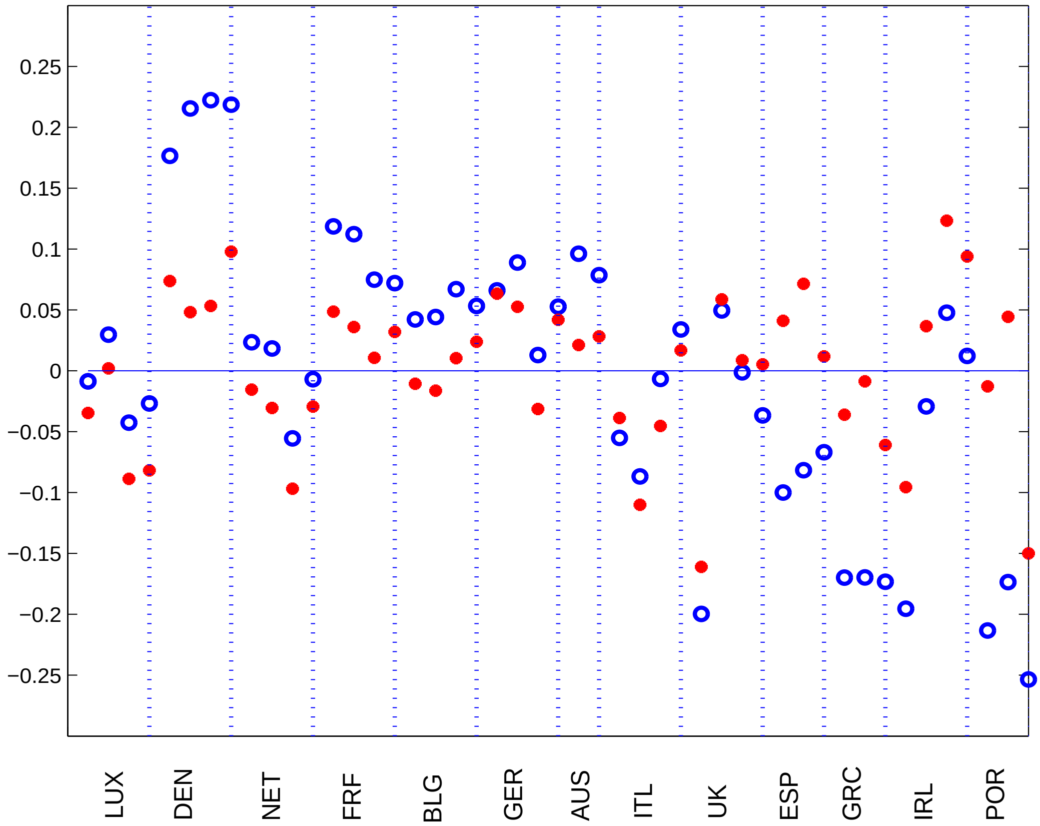

Each point is the average LOP deviation for a particular country/year. Vertical dashed lines delineate countries. Each point between these dashed lines is one particular year for that country. Thus, the first 4 points represent Luxembourg for 1975, 1980, 1985 and 1990, the next 4 represent Denmark for the same years, and so on. The countries are organized from lowest to highest income, just as in Table 5. Finally, open circles represent unadjusted data (*i.e.*, the means from Table 5 in the leftmost columns of Panel A), whereas asterisks represent means which have been adjusted for income and VAT differences (*i.e.*, from the rightmost columns of Panel B in Table 5).

Figure 3
Empirical Distributions of $\boldsymbol{Var}(q_{ij}\,|\,j)^{1/2}$

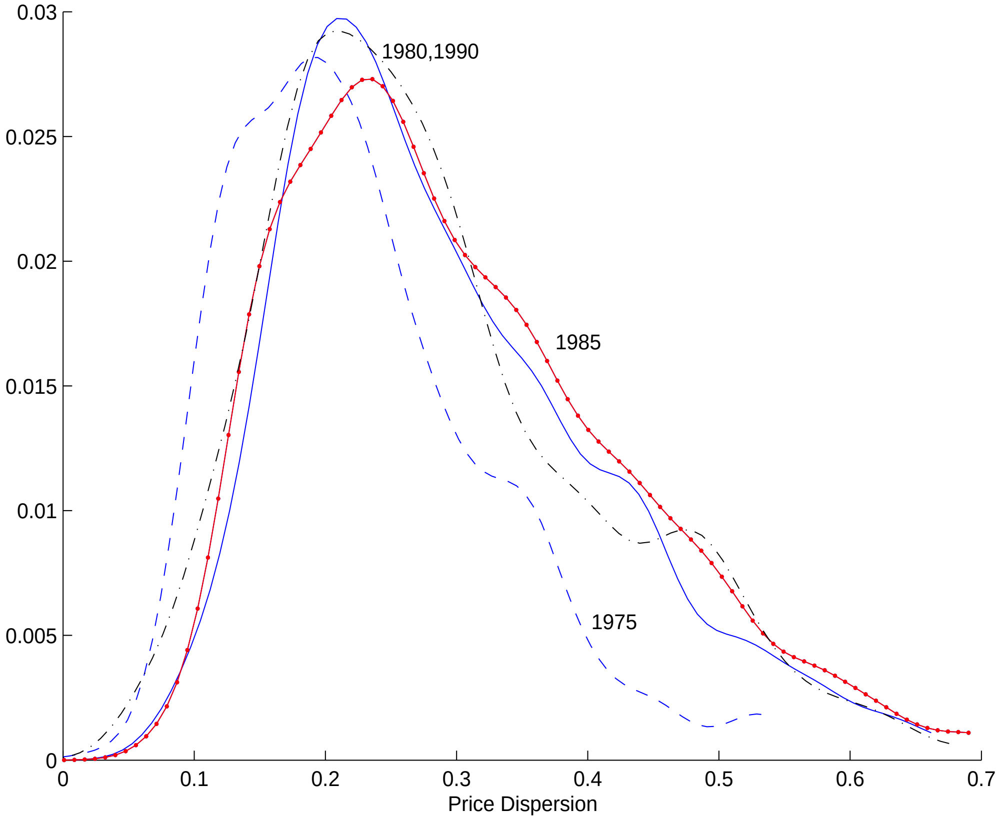

Each line represents an estimate (for each year) of the density of $Var(q_{ij}\,|\,i)^{1/2}$, the standard deviation of the LOP deviation for good i across countries j.