DISPERSION IN REAL EXCHANGE RATES

by

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

Working Paper No. 00-W13

June 1998
Revised May 2000

DEPARTMENT OF ECONOMICS
VANDERBILT UNIVERSITY
NASHVILLE, TN 37235

www.vanderbilt.edu/econ

Dispersion in Real Exchange Rates*

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

First draft: June 1998
Current draft: May 2000

Abstract

Using cross-sectional data on local currency prices of over 1,800 retail goods and services across 13 European countries in the mid 1980's, we characterize the behavior of average relative prices — 'real exchange rates' — as well as dispersion around these averages. We find that the averages are surprisingly close to what purchasing power parity would suggest. In other words, in the mid 1980's, averages of ratios of foreign to domestic prices (across goods for a particular pair of countries) provide surprisingly accurate predictions of most nominal cross-rates. Variation around the averages, however, is large but is found to be related to economically meaningful characteristics of goods such as measures of international tradeability, the importance of non-traded inputs into production and the geographical distance between product markets. Using data on product brands, we find that product heterogeneity is at least as important as geography in explaining relative price dispersion.

*A 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 seminars and conferences, we thank Stephen Cosslett, Charles Engel, Paul Evans and Stan Zin for helpful comments and suggestions. We gratefully acknowledge the financial support of the Carnegie Bosch Institute.

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

[‡]Graduate School of Industrial Administration, Carnegie Mellon University; chris.telmer@cmu.edu.

[§]Department of Economics, The Ohio State University; zachariadis.1@osu.edu.

As long as anything like free movement of merchandise and a somewhat comprehensive trade between the two countries takes place, the actual rate of exchange cannot deviate very much from this purchasing power parity. Cassel (1918, p. 413)

1 Introduction

A time immemorial question in international economics asks what determines the relative price of a good, or a basket of goods, in one country vis-a-vis another. The answer has important implications for research areas ranging from international industrial organization to international finance. The 'law of one price' (LOP) suggests that the answer is trivial: common-currency prices of the same good should be the same. Anyone who has traveled knows that this should not be taken literally. The notion of 'purchasing power parity' (PPP) suggests that something akin to the law of one price should hold at a more aggregate level, involving broad baskets of goods and services. The evidence suggests otherwise, where the common finding is that deviations from PPP can be large and long lasting. The goal of this paper is to bridge the gap between what we know about disaggregate relative price behavior as embodied in deviations from the law of one price and aggregate relative price behavior measured as deviations from PPP. We attempt to do so by characterizing the mean and the variance of the distribution of relative prices for a cross-sectional dataset on the retail prices of roughly 1,800 different goods and services across 13 European countries in the mid-1980's.

What we know about PPP and how it relates to individual goods and services has increased substantially since Gustav Cassel wrote the above quote some 80 years ago. At the aggregate level, an extensive literature has shown that deviations from PPP based upon consumer price indices can be large and persistent with half lives (presuming that real exchange rates are stationary processes) in the neighborhood of three to five years. Using less aggregative data, but data which is still represented by index numbers, papers by Engel (1993), Engel and Rogers (1996), and Rogers and Jenkins (1995) show that economically large deviations from PPP are not a mere artifact of examining CPI baskets of questionable comparability. An important implication of these papers is that international borders represent something special for price determination; patterns of price dispersion across countries seem to be substantially different than those across locations within a country. Finally, a large and growing literature, including papers by Froot, Kim and Rogoff (1995), Giovannini (1988), Gosh and Wolf (1994), Haskel and Wolf (1998), Isard (1977), Parsley and Wei (1999) and Richardson (1978), shows that at even lower levels of aggregation — sometimes down to the individual goods — the prices of similar goods across borders can be vastly different and these differences can persist for long periods of time. Again, this body of work demonstrates that the well-known results on PPP violations at the aggregate level are not simple manifestations of comparing apples and oranges.

Our work continues along the lines of many of the above papers in that we examine relative price behavior for a highly disaggregated collection of goods. What distinguishes our study is the breadth of our dataset and, as a result, the extent to which we are able to relate price dispersion to economic characteristics of particular subsets of goods. A useful framework in which to consider the questions we ask involves the extent to which the definition of a good should have a location-specific component. For instance, it seems natural that a bottle of beer served on Las Ramblas, in Barcelona, is a different commodity than an identical bottle of beer served in the Squirrel Hill Cafe, in Pittsburgh. On the other hand, it seems just as natural that there exists some link between the prices at which these goods sell which goes beyond location and distance. Indeed, a rich literature, including (among many others) papers by Alessandria (1999), Betts and Kehoe (1999), Balassa (1964), Baumol-Bowen (1966), Harrod (1933), Knetter (1989, 1993), Krugman (1987), Samuelson (1964), Ethier (1979), Stockman and Dellas (1989) and Stockman and Tesar (1995), argue persuasively that a lower-dimensional relationship (relative to location) should characterize international relative prices. Among other things, this literature focuses on measures of tradeability and various aspects of industrial organization as being important attributes which should help explain both disaggregate and aggregate price dispersion.

The role played by our study in this context is one of identifying how important the location aspect of a good is relative to its other, economically relevant, attributes. For example, the literature on CPI-based measures of PPP, and to a certain extent the contributions made by Charles Engel and his co-authors, suggests that location plays a paramount role (*i.e.*, national goods markets are strongly segmented) and that price dispersion is not strongly related to a small number of economic factors. The work of Michael Knetter, on the other hand, suggests that one should find an increasing amount of relative price dispersion as one considers products which are more highly differentiated. Our goal is to ask to what extent these viewpoints are supported by the data.

Our findings are as follows. Our data display large deviations from the law of one price; of the 8,400 relative price comparisons which are possible, roughly 4,800 feature deviations of 20% or more. On average, however, there are roughly as many goods which are overpriced as there are goods which are underpriced. That is to say, equally-weighted averages of the ratios of foreign to domestic prices give a surprisingly (in light of previous work) accurate prediction of the nominal exchange rate. Specifically, the PPP predicted nominal exchange rate is within 10% of the actual market rate for Austria, France, Germany, Ireland, Italy, Luxembourg, the Netherlands, and the United Kingdom, using Belgium as numeraire. Denmark is somewhat anomalous in this context in that the kroner is overvalued by 16%. The nominal exchange rates are predicted to be between 16% and 25% above market rates in Greece, Portugal, and Spain, consistent with existing work that finds prices tend to be lower in poorer countries. We also find that a value-weighted average, analogous to the CPI, provides a less accurate prediction of the nominal exchange rate, in particular if expensive items such as automobiles are included. In this sense, our results are consistent with the literature on aggregate, real exchange rate variability. At the very least, our results are suggestive of the importance of using a very broad basket of goods to assess the purchasing power of a country's currency.

Our findings regarding average relative prices should certainly be interpreted with caution. The weakness of our dataset, obviously, is that it represents only a single cross section, taken at a particular point in time. In regard to characterizing the dispersion in relative prices, however, the size of our cross-sectional sample leaves us on stronger ground. Focusing on three economic factors — the tradeability of the good in question, the tradeability of the inputs required to produce the good, and a measure of how differentiated the good is – we find convincing evidence that an important component of price dispersion is attributable to what mainstream economic theory suggests. Price dispersion falls by about 17% as we move from non-traded goods to goods with trade shares at the sample average. Goods requiring a large share of non-traded inputs have greater dispersion; if the fraction of non-traded inputs doubles from 10% to 20% of total cost, price dispersion rises by 8%, going from roughly 23% to 25%. Finally, price dispersion is positively related to the physical distance between the capital cities from which our goods prices arise. The overall message we take from our dataset is that simple economic forces play an important role in international price setting. Just as the law of one price should not be taken literally, neither should the completely segmented goods markets view of the world.

One final issue we examine is how price dispersion and product heterogeneity are related. Our cross section contains many price observations for different brands of the same good. We use this subset of the data to shed light on how market structure relates to price dispersion. Models featuring monopolistic competition, for instance, predict that firms can sustain price differences across differentiated products within countries, while models of geographic price discrimination predict that firms can sustain price differences across countries for identical products. In order to quantify the relative importance of these features of market structure we ask the following question: "How does the dispersion of prices across brands of the same good within a country compare to the dispersion of prices across countries of a particular brand?" We find very little tendency for the dispersion of prices of identical goods across countries to exceed the dispersion of different brands of the same good within countries. One important exception is automobiles where the price dispersion across countries for a given make and model is comparable or higher to the dispersion across makes and models in a given location.

The remainder of the paper is organized as follows. We begin in Section 2 by describing our data. While international price indices are of questionable comparability across countries, the structure of the Eurostat Survey that we utilize is designed to provide exact comparability of individual goods across countries. Thus we can be confident that when price deviations arise they reflect differences in prices not differences in the characteristics of the goods. After describing the data we present some basic descriptive statistics to illustrate both the tendency toward the LOP and the large differences that exist across individual goods used in the comparison. Section 3 contains the main results of the paper. Here we compute a measure of price dispersion for each good in our sample and relate this dispersion measure to features of the good. Section 4 examines the geographic patterns of overall price dispersion and the implications of those patterns for consumer expenditure. Section 5 concludes with remarks about the implications of our findings for the large and growing theoretical literature on the dynamics of international relative prices.

2 The Data

We begin with an overall description of our data, emphasizing the design of the price survey. More specific details are relegated to Appendix A. The majority of our data is comprised of local currency prices on 1,805 different goods and services across 13 European countries. The source is the survey, "Price Structure of the Community Countries in 1985" published by Eurostat, the Statistical Office of the European Communities.[1]

With the help of the individual national statistical agencies, Eurostat has in effect produced an international price survey comparable to what the U.S. Bureau of Labor Statistics (BLS) compiles as an input into the construction of urban consumer price indices. Unlike the BLS survey, however, the Eurostat survey has a location specificity which proves very useful for our questions. That is, the BLS refers to 'typical purchases,' so that the cost of soda pop might refer to Pepsi in New York City and Coke in Los Angeles. The design of the Eurostat survey, in contrast, would result in the prices of both goods in both locations being reported. The end result is that the European price survey is probably the best absolute price data available for purposes of location-to-location price comparisons, arguably better than that available for price comparisons within the U.S.[2]

While the survey is explicitly designed to give exact comparability of goods across locations, the completeness of the descriptions provided varies to a certain degree. Eurostat provides two explanations for this shortcoming of the published record: 1) space constraints (some descriptions would require a paragraph or more) and 2) confidentiality. One good, for instance, is identified as 500 grams of long-grain rice, packaged in a plastic bag. Another is a 'selected brand' of automobile with an engine size between 1,200 and 1,700 cubic centimeters. In this latter case there are actually multiple brands (92 in total), but none of the brands are explicitly identified. Goods in our survey, then, are rarely identified up to the manufacturer and/or product brand, but are defined in such a way as to make one reasonably confident that comparisons are being drawn across goods which are quite similar. The label 'selected brand,' for instance, suggests that we are unlikely to end up comparing the cost of a Toyota with that of a Mercedes Benz.

The price data itself was actually collected in a sequence of surveys spanning the period 1984-1986. Clothing, footwear and household textiles were surveyed in Autumn 1984; durable household goods in Spring

[1] The 1985 Eurostat price survey was an important part of the basic price data used in the International Comparison Project as described in Summer and Heston (1991).

[2] We thank Alan Heston for pointing this out.

1985; services in Summer 1985; glassware and other household articles in Autumn 1985; food, beverages and tobacco in Spring 1986 and health services in Summer 1986. 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 season in question.

The retail prices are cash prices paid by final consumers, inclusive of all taxes such as VAT. For 12 of 13 countries these surveys were conducted exclusively in their respective capital cities. The exception is Germany, where the survey was country wide. The prices we use in our analysis are averages of the surveyed prices across different city-wide sales points, with more observations collected for goods having greater price dispersion. Sales points were selected by the national statistical offices so that the sample is representative of the distribution of prices in the capital city.

For illustration purposes, Table 1 reports a number of individual records on goods and services which are chosen to be representative of the various overall categories contained in our dataset. Two goods from the food and beverages category, for instance, are long-grained rice and a particular brand of liqueur. An example of a consumer durable is a selected brand of dishwasher with 5 cycle settings. Examples of services are the cost of bus fare for travel of about 6 km and a change of bus and the rental of a television for one month. In each case, Table 1 illustrates what our data are comprised of: a Eurostat code, a detailed description of the particular good, the units of measure and 13 columns of price data. Although Eurostat reports the latter in local currency units, Table 1 represents prices in Belgian francs. 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 much smaller.

One final issue involves how we deal with missing observations, of which there are many (indicated with a '-1' in Table 1). Our main focus is on explaining price dispersion across countries. Consequently, we eliminate any good which has an insufficient number of cross-country observations, which we define as 6 or less (out of a possible 13). We also control for gross measurement error by eliminating any particular observation in 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 1,805 to 1,164. Of these remaining data, 19.3% of the observations are missing, something which will be important to keep in mind when interpreting our statistical analysis. For calculations which require a numeraire to be defined, we use Belgium for the simple reason that it has the fewest missing observations (6.8% of the 1,164 goods). The country with the most missing data is Ireland, with 37% missing data.

3 Purchasing Power Parity

The use of PPP for determining the value of currency has a long and colorful history. One of the more well known applications of the PPP doctrine occurred in the debate surrounding Britain's return to the gold standard in April 1925. Keynes (1932) argued aggressively against returning to the prewar parity which was close to that implied by movements in the ratio of wholesale price indices across Britain and the United States. He favored the use of retail prices which suggested the pound would be more than 10% overvalued at the prewar parity. A natural place to begin our analysis, then, is with an investigation of the link between retail prices and nominal exchange rates.

Let N be the number of goods, M the number of countries and p_{ij} the local currency price of good i in country j. The real exchange rate for good i in country j (relative to the numeraire country, n) is then $q_{ij} = e_j p_{ij}/p_{in}$, where e_j is the spot exchange rate for country j in units of the numeraire country and p_{in} is the local currency price of good i in the numeraire country. Thus when q_{ij} exceeds unity, good i is more expensive in country j compared to Belgium (our numeraire).

We summarize these good-by-good real exchange rates in Figure 1 which presents their empirical distribution. Each line is a smoothed kernel estimate of the distribution of log real exchange rates for a particular country (compared to Belgium, the numeraire) except for the thicker line which pools all countries together.

Viewing these real exchange rate distributions for the first time, we were struck both by their symmetry about zero and the remarkable dispersion across goods (evident in the support of the distribution extending to plus and minus 150%). The observation that the mean is close to zero is evidence in favor of the PPP doctrine: purchasing power is at approximate parity across countries, at least when goods prices are averaged on an equally weighted basis. Given the choice however, one would want to shop around because the wide dispersion across goods implies that expenditure savings are large when goods are purchased in their lowest price location.

Of course the dispersion of prices across goods is likely to make the weighting scheme of some importance for the estimate of the mean real exchange rate of an individual country. To assess the role of weighting in tests of PPP we construct an aggregate real exchange rate by averaging the underlying log real exchange rates good-by-good,

$$q_j = \sum_{i=1}^{N} \gamma_{ij} \log q_{ij} \ , \tag{1}$$

where γ_{ij} are weights which sum to unity. We consider both equal weighting, $\gamma_{ij} = 1/N$ (the implications of which could be gleaned from the location of the distribution in Figure 1), and value weighting. We define value weights as follows,

$$\gamma_{ij} = \frac{1}{2}\left(\frac{p_{ij}}{\sum_{i=1}^{N} p_{ij}} + \frac{p_{in}}{\sum_{i=1}^{N} p_{in}}\right) . \tag{2}$$

The value weights can differ dramatically across goods as is evident in comparing the most expensive item with the least expensive item in our panel: a luxury car in Denmark is worth 3.7 million Belgian Francs which is about 2 million times more expensive than a small roll in Portugal, worth 1.5 Belgian Francs. From the point of view of PPP we should weight goods by something like their shares in national consumption (which we do not have at the level of aggregation of our data) which will of course differ from either of our weighting schemes. In value weighting the prices, we are effectively assuming that one of each item in the survey is purchased in a given interval of time so that total expenditure is just the sum of the prices. Large ticket items (e.g. automobiles) get a much larger weight in the value–weighted scheme (compared to equal weighting) as they should, but this weight is perhaps too large. Returning to the extreme example of small rolls and luxury cars, if the average household purchased ten rolls per day and a luxury car every ten years, the consumption expenditure weights would be on the order of 70 to 1, while value weighting gives a ratio of 2 million to 1. Equal weighting does the reverse, assigning too much weight to small ticket items and too little weight to large ticket items. Thus we expect that the two schemes give a reasonable sense of the bounds of different weighting methods on the PPP prediction.

We implement these weighting methods as follows. For each bilateral calculation we include all goods prices for which we have an observation for both countries. However, to avoid double counting similar items, we eliminate multiple brands. That is, given some good for which we have multiple brands, we keep the good which has the maximum number of observations across the 13 countries. In the event of a tie, we randomly choose which good to keep. The resulting samples contain differing numbers of observations for each bilateral comparison, typically exceeding 700 good-specific relative prices.

Table 2 reports the results. Examining the mean real exchange rate computed using an equally-weighted average of the good-by-good log real exchange rates across all goods and countries equals -6.5%: on average Belgium is moderately more expensive than the other countries in the sample. In fact, only Austria, Denmark, and France are more expensive than Belgium by this metric. In the poorer countries – Greece, Ireland, Portugal and Spain, though not Ireland – prices are much lower than in Belgium as we would expect given the well known positive correlation between price levels and per capita income documented, for example, by Summers and Heston (1991) for a much larger sample of countries. Eliminating these countries gives an overall average price difference of just -1.1%. Thus after accounting for the potential impact of wealth on price levels, it would be difficult to imagine stronger evidence in favor of PPP.

In contrast, value weighting gives rise to larger differences and reverses the correlation between price levels and per capita income. The third column shows that this is almost entirely due to a small number of goods that are very influential in a value-weighted sense: automobiles. Omitting automobiles from the calculation moves the mean real exchange rate back towards the prediction under equal weighting. However, value weighting does tend to mitigate the correlation between average price levels and per capita income even with automobiles omitted.

The case of Luxembourg is of particular interest because Belgium and Luxembourg have maintained a fixed exchange rate for some time (beginning well before the price survey and continuing to the present day). In Table 2, we see that the mean difference in prices between Luxembourg and Belgium, while modest, is not significantly different from what is found in comparisons involving other wealthy countries. What we do see, however, is considerably less dispersion across goods; the standard deviation of real exchange rates is 0.3 for Luxembourg, compared to values at or above 0.4 for other countries. Note also the sharp peak in Luxembourg's density at zero in Figure 1; about 40% of Luxembourg's prices are within 10% of those in Belgium compared to half that many in the cases of the remaining 11 countries. The finding of less dispersion is consistent with the view coming out of the time series literature that a fixed exchange rate regime mitigates relative price variation (here measured across goods rather than over time). The fact that substantial variation remains suggests that more is going on, however, than local currency pricing.

The centrality of the good-by-good real exchange rates around the LOP is consistent with the time series evidence that finds real exchanges rate are stationary stochastic processes. If the real exchange rate of each good was an independent random walk we would expect the distribution in Figure 1 to look like a uniform distribution, not the strongly peaked one we find in the data. Stationarity, though, is assumed in most theoretical models and is thus not informative in choosing among them. Our goal in the remainder of the paper is to explain the dispersion in real exchange rates across goods using insights from several branches of the theoretical literature.

4 Understanding Dispersion

Measures of central tendency aside, a striking feature of our price data is the amount of dispersion in relative prices. This is evident in both Figure 1 as well as Table 3, where we provide country-specific empirical probability values for deviations from the law of one price of varying magnitudes. We see that roughly one third of the prices in our sample differ from Belgian prices by between 20% and 50%. Roughly 20% of the prices are even more dispersed, deviating from Belgian prices by more than 50%. Even greater price disparities exist for Greece, Portugal and Spain, where more than one third of the relative prices deviate by more than 50%. Table 3 also exhibits negative skewness, indicating that, relative to the overall cross-section, goods and services in Belgium are expensive. This 'country effect' — something emphasized by Cumby (1996) and Rogoff (1996) in a time series context for Big Mac prices — is one of the first 'economic' characteristics we deal with below.

Dispersion in relative prices, then, is large. Our goal in the remainder of this section is to ask to what extent this dispersion is systematically related to economically meaningful characteristics of the individual goods. For instance, are the relative prices of services more dispersed than those of manufactured goods? If two goods are produced by industries with very different industrial organizations, do we see different patterns in price dispersion? Or, in contrast, are the locations of the markets in which goods are sold of primary importance? If so, we would expect to see price dispersion distributed somewhat uniformly across goods with different economic characteristics.

Prior to answering these questions, we perform some simple manipulations of our data and then provide summary statistics regarding price dispersion. Recalling that p_{ij} are the local currency prices and e_j are the nominal exchange rates, we define z_{ij} as the percentage by which p_{ij} deviates from the cross-country, good-specific mean,

$$z_{ij} = \frac{e_j p_{ij}}{\sum_{j=1}^{M_i} e_j p_{ij}/M_i} - 1,$$

where M_i is the number of cities for which we have price observations on good i. Note that, unlike our analysis in the previous section, these measures of price dispersion are independent of the numeraire, for which $e = 1$. In addition, we do not remove data on goods for which we have multiple brands, thereby making the number of goods which pass our selection criteria equal to 1,164.

A natural place to begin is to ask how much of the variation in the z_{ij}'s is attributable to a 'country effect:' variation across j which is common for all i. Such effects may arise as a result of nominal exchange rate movements, variation across countries in VAT's, differing levels of income and wealth, and so on. In Table 4 we report the country specific means of the percentage price deviations,

$$z_j = \sum_{i=1}^{N_j} z_{ij}/N_j,$$

where N_j is the number of price observations we have on country j. Note that these numbers, reported in Table 4, differ somewhat from those in Table 2 because we normalize by the mean price for each good so that dispersion is computed relative to the cross-country mean for each good not relative to the price of each good in Belgium.

We see that there are some important country-specific effects. Goods are really expensive in Denmark and relatively inexpensive in Portugal, for instance, but the key implication is that somewhere between 89 and 93 percent of the variation in our data goes beyond country effects. We therefore proceed by removing the country means and analyze the data,

$$\hat{z}_{ij} = z_{ij} - z_j.$$

The implication is that all subsequent analysis is net of cross-country variation in things like VAT, so long as these effects apply uniformly across goods.

We now transform our data into a single vector of good-specific measures of relative price dispersion. Defining $mad(\cdot)$ as the mean absolute deviation, we work with,

$$y_i = mad(\hat{z}_{ij}),$$

which is less sensitive to large outliers than is the standard deviation. The vector, y, of 1,164 mean absolute deviations constitutes the fundamental data we seek to explain. Figure 2 plots these data in the order in which they appear in the Eurostat price survey. If the law of one price held for each good, this graph would be the zero line. In contrast, we see that price dispersion is large and highly variable across goods. The mean and standard deviation of y_i are 23 percent and 10 percent, respectively. Figure 2 also contains a limited amount of information on how price dispersion relates to characteristics of goods. This is because the 'good number' — the variable on the horizontal axis — is not determined randomly, but by the Eurostat classification system. Groceries and other food products, for example, appear at the beginning of our cross section whereas automobiles appear towards the end. On the other hand, further down in the ordering, the costs of various types of telephone calls fall between the costs of international airplane tickets and those of portable radios. The ordering of the goods, therefore, does have some economic content, but it is far from precise. This leads us to cautiously view the apparent pattern in price dispersion — the correlation between adjacent data points is 0.41 — as our first hint of some underlying economic content.

There is a more precise way in which we can use the Eurostat classification system as a benchmark. Our data come organized into 190 categories of goods and services, ranging from dried bread products to automotive repair. There are, on average, 6.12 goods per category (the maximum, minimum and median, respectively, are 31, 1 and 5). A measure of how much of the variation in Figure 2 can be explained by 190 category-specific means, therefore, provides a measure of how much of the variation we could hope to explain based on the Eurostat classification. If we had 1,164 categories, for example, this 'upper bound' would be meaningless, at 100 percent. Given many less categories than goods, however, the question has content. The answer is 48 percent. Loosely speaking, once we remove country effects, there is about as much price variation within Eurostat categories as there is across them.

It is important to interpret this number — the 'upper bound' of 48 percent — as a loose benchmark and not much more. One reason is that there remains a great deal of heterogeneity within Eurostat categories. Pianos and outboard boat motors belong to the same category, for instance, as do frozen and fresh foods. It seems plausible that there exists an alternative partition, one with more economic content, which can account for more of the cross-sectional variation. In addition, there are certainly a number of important economic interactions across Eurostat categories. The industrial organizations in which pianos and outboard motors are produced are likely to be related to those of wooden furniture and automobiles, respectively. The differences in shipping costs across frozen and fresh foods is likely to be related to those which distinguish dried from fresh flowers. Incorporating factors such as these may yield a number substantially higher than 48 percent. In spite of this, we find the latter to be an informative benchmark. It indicates that a relatively high dimensional classification system — a system which is therefore unable to offer much economic insight — can account for only half of the variation we seek to explain.

One final set of summary statistics are reported in Table 5. We group the data into 111 different categories and then group these categories into 9 broader categories, ranging from groceries to services. The spirit in which these data are reported is very much that of 'summary statistics,' intended to provide the reader some interpretational latitude of their own. What we take from Table 5 is analogous to what we take from Figure 2. That is, the data are suggestive of a meaningful relationship between dispersion and the characteristics of goods, but a more structured quantitative analysis is needed. Specifically, we see more price dispersion in services and 'vices' (alcohol and tobacco), suggesting that tradeability and excise taxes are important. The relatively small degree of dispersion associated with petroleum products, electronic goods and transportation products — all tradeable goods — supports this. There are many examples of goods which are highly differentiable — wine, liquor, bottled water, flowers, large automobiles — which exhibit a relatively large amount of price dispersion. One might, on the other hand, suspect that these goods are more likely to involve comparisons of what are essentially different commodities (although the explicit spirit of the Eurostat survey is to avoid this). In any case, it is clear that there are many different interpretations of this data. We now turn to a much more parametric analysis to see if these common anecdotal interpretations prove robust.

4.1 Explanatory Variables

We now seek to characterize price dispersion, y, in terms of a relatively small number of characteristics of goods. To do so, we focus on three issues emphasized in the international trade and industrial organization literature: the extent to which final goods are internationally tradeable, the extent to which inputs required to make these goods are tradeable, and the competitive structure of the markets in which final goods are sold. Tradeability is certainly the most obvious candidate. A large literature in international economics dichotomizes goods as traded or not traded, with the law of one price holding for the former but not the latter. A provocative example is provided by Rogoff (1996), who shows that deviations from the law of one price are insubstantial for gold bullion but enormous for MacDonalds Big Mac hamburgers. Related evidence is provided by Cumby (1996). Tradeability in inputs has also seen attention, for example in the theory of trade in 'middle products' (e.g. Sanyal and Jones (1982)). In this context, we would expect international price dispersion to be present for all retail prices but that it would be larger for goods requiring larger amounts of domestic non–traded inputs. Finally, the industrial organization of the markets in which final goods are sold has received growing attention in the international literature (see Krugman (1995) and Feenstra (1995) for comprehensive reviews of this literature). These papers introduce imperfect competition in final goods markets such that price differences are sustained across national boundaries, perhaps indefinitely. Based on insights from this literature we expect the magnitude of the differences in international prices to be related to the extent of competition at the industry level and the elasticity of product demand.

Beginning with our measure of tradeability, we restrict ourselves to a good–specific rather than a country–specific measure. By this we mean the index should depend on the characteristic of the good and not the specific tastes or endowments of a particular country. Toward this end, we aggregate over countries and compute the trade share as,

$$\theta_k = \frac{\sum_{j=1}^{m_k} (X_{kj} + M_{kj})}{\sum_{j=1}^{m_k} Y_{kj}}, \tag{3}$$

where X_{kj} (M_{kj}) denotes exports (imports) of sector k in country j expressed in U.S. dollars and Y_{kj} is gross output of sector k in country j. The index m_k accounts for the fact that we do not have the same amount of sectoral data for all countries (see Appendix A). We employ this as our tradeability proxy though in cases where trade data is not available and the degree of tradeability is obvious, an index of zero is assigned. The sectors 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. The other trade shares range from a low of 20% for tobacco and manufactures to a high of 140% for office, computing and accounting machinery.

The second characteristic of goods that we measure has to do with the fact that all retail goods involve some type of transformation in the process of moving from the port of importation to the destination of final sale. For example, long grain rice might have similar prices at ports in Europe once transport costs and insurance are accounted for but we expect differences to arise at the retail level due to differences in rental costs, transportation from the port to final destination of sale and marketing costs, to name just a few. Our second measure of the characteristic of a good is the cost share of non–traded intermediate inputs computed as,

$$\Phi_k = \sum_{s=1}^{S} \phi_{ks} , \tag{4}$$

where ϕ_{ks} is the share of non–traded intermediate input s in the total cost of the output of sector k. 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. The cost share of non-traded intermediate inputs ranges from a low of 0.11 for motor vehicles to a high of 0.32 for transport, storage, and communication.

4.2 Regression Framework

Given our measures of price dispersion, with N observations on y, our goal is to characterize the variation in y in terms of some vector of explanatory variables, x, on which we have N observations. Assuming that the regression of y on x is linear, we have,

$$y_i = \alpha + x_i \cdot \beta + u_i \ , \tag{5}$$

where u_i is *i.i.d.*. The main problem we have — and one that any study of highly disaggregate data is likely to have — is that our observations on some elements of x are aggregated to a larger extent than those on y. Take for example, our measure of international tradeability. While we have data on price dispersion of, say, many different types of electronic goods, our measure of tradeability is limited to one aggregative value for electronic goods in general. In a nutshell, the variable we seek to characterize — good-specific price dispersion — is observable at a much 'finer' level than the variables we seek to characterize it with.

This type of aggregation has important consequences for statistical inference. In general, it will generate a heteroskedastic pattern in the variance of regression error terms (especially in finite samples) and, just as importantly, make goodness-of-fit measures difficult to interpret. In Appendix B we formulate a statistical framework which we use to incorporate these effects, thereby allowing us to obtain consistent, efficient estimates of β, it's standard errors, and meaningful goodness-of-fit measures.

Briefly, we define our data as being partitioned into G distinct 'groups,' $g \in \{1, 2, \ldots, G\}$, each of which yields N_g sample observations. Examples of groups are textiles, automobiles and personal care products. We modify the population regression, (5) as follows,

$$y_{ig} = \alpha + x_{ig} \cdot \beta + w_{ig} \cdot \gamma + u_{ig} \ , \tag{6}$$

where the subscript i, g denotes the ith observation on a good from group g, and w denotes those elements of x (from equation (5)) on which we *do* have completely disaggregate data. Define the within-group sample mean for x as $\bar{x}_g$,

$$\bar{x}_g = \frac{1}{N_g} \sum_{i=1}^{N_g} x_{ig} \ . \tag{7}$$

The regression equation, (6), can be written,

$$y_{ig} = \alpha + \bar{x}_g \cdot \beta + w_{ig} \cdot \gamma + (x_{ig} - \bar{x}_g) \cdot \beta + u_{ig} \ , \tag{8}$$

which is a statistical regression of y onto $\bar{x}_g$ and w, as long as x and w are uncorrelated, something we assume. Sampling variation in $\bar{x}$, however, generates heteroskedasticity in the error term, $(x_{ig} - \bar{x}_g) \cdot \beta + u_{ig}$, a feature which is particularly important in our dataset, where there is a great deal of variation in the within-group sample size (N_g ranges from a minimum of 1 to a maximum of 280). We take two approaches in estimating the parameters of (8), each of which turn out to yield qualitatively similar results. First, we estimate the regression (8) using generalized least squares (GLS), having characterized the exact form of the heteroskedastic covariance matrix for the errors (details are provided in Appendix B).

Second, we average our disaggregate data, y and w, within groups, g, dictated by our aggregative data, x. Specifically, we average across equation (8) and estimate the following,

$$\begin{aligned} \frac{1}{N_g}\sum_{i=1}^{N_g} y_{ig} &= \alpha + \bar{x}_g \cdot \beta + \frac{1}{N_g}\sum_{i=1}^{N_g} w_{ig} \cdot \gamma + \frac{1}{N_g}\sum_{i=1}^{N_g} u_{ig} \\ \Rightarrow \bar{y}_g &= \alpha + \bar{x}_g \cdot \beta + \bar{w}_g \cdot \gamma + \bar{u}_g \ , \end{aligned} \tag{9}$$

where the sample averages, $\bar{y}_g$, $\bar{w}_g$ and $\bar{u}_g$ are defined in the obvious way and are understood to depend explicitly on the values N_g. Estimates based on equation (9) are also obtained using feasible GLS, given our knowledge of the covariance matrix of $\overline{u}$, which is a simple function of the values N_g. The main disadvantage associated with equation (9) is that it averages away potentially informative variation in y and w. The advantages are a simpler form of the covariance matrix and a more easily interpretable goodness-of-fit, due to the fact that we are not trying to explain variation in y with variation in x, after having removed a great deal of the latter (something inherent in equation (8)).

It is worth noting that, should the variables w be omitted, GLS estimates of β based on equation (9) will be numerically identical to OLS estimates based on equation (8). Standard errors, however, may be quite different owing in large part to a much smaller number of observations associated with (9).

Turning to specifics, we use the following variables. First, the aggregative data we use — which correspond to x — are trade shares and cost shares of non-tradeable inputs. Both these measures are more aggregative than our underlying price data so we grouped our price observations into ISIC revision 3 categories which could then be reconciled with the trade and input-output measures: details are provided in the data appendix. The data corresponding to w are a set of dummy variables.

Prior to reporting regression results, it is informative to view bivariate scatter plots of our measures of price dispersion, y, against our two chief candidates for characterizing the variation in x: tradeability and the cost share of non-traded inputs. Figure 3 plots dispersion against trade share, while Figure 4 presents the same data after having aggregated price dispersion measures across goods with the same trade share. Figure 3 clearly illustrates how our data limitation — not observing good-specific measures of tradeability — can cloud our view of what determines price dispersion. In general, by aggregating across groups of the variable on the horizontal axis, we 'pile up' observations in a vertical line originating at our estimate of the group-specific conditional mean. This will be true *even* if the population regression errors are zero. Figure 3 clearly exhibits this type of behavior and, thus, demonstrates how a strong relationship between dispersion and trade share may be masked by the more aggregated form of the trade data.

Figure 4 provides the scatter plot of dispersion versus trade share, but where we have averaged across dispersion according to the grouping dictated by aggregation within trade share groups. While this graph tends to provide a more accurate picture (relative to the top panel) of the variation in dispersion not accounted for by trade share, it still must be interpreted with caution. The reason is that some data points are based on many more observations on price dispersion than others. In one case (the highest vertical point) we have only one dispersion observation, whereas for others we have as many as 280. With this in mind, the lower panel also plots the bivariate, linear regression function of price dispersion on the trade share, where the coefficients are obtained by GLS (see equation (9)), using the 39 data points in the graph. The form of the covariance matrix will result in points for which we have more observations getting more weight (hence the often used terminology 'weighted least squares'). In addition, we highlight four points for which we have the most observations with asterisks and four points for which we have the least observations with squares.

The overall message we take from Figure 4 foreshadows a main message of our paper. Figure 4, indicates that tradeability is an important determinant of international price dispersion. It is, however, obviously far from the whole story. What we mean by 'important' has a number of facets. First, the slope coefficient is certainly significant, and of the right sign as suggested by theory. Its magnitude suggests that if we consider two goods, one non-tradeable and one with an average level of tradeability (0.75 using our metric), price dispersion will fall (on average) by roughly 22 percent, going from 27 percent to 21 percent. Second, while it may not be obvious from Figure 4, the R^2 from the GLS regression of dispersion on trade share is 38 percent. Again, because of the importance of the weighting of the points in Figure 4, one must interpret the visual evidence with caution. Finally, one must interpret the amount of variation in price dispersion which can be accounted for by tradedness in the context of Figure 2, where we saw that at least half of the overall variation in our price dispersion measures is probably unexplainable in some low dimensional sense.

Figures 5 and 6 have similar structure to Figures 3 and 4, except that we graph price dispersion against the cost share of non-tradeable inputs. All the same caveats apply in regard to disaggregate versus aggregate data. Again, Figure 6 suggests an economically important role played by the explanatory variable. The graph suggests that if the cost share of non-tradeable inputs used to produce a good doubles, from 10 to 20 percent, price dispersion increases by 18 percent, going from roughly 20 to 23.6 percent. This effect is, again, statistically significant and will only become more so once we use disaggregate price dispersion data which uses all the good-by-good dispersion measures.

Table 6 reports the results of estimating equation (8) using both the tradeability and cost share of non-tradeable inputs as regressors. We also include dummy variables for services, electronic goods, vice goods, and large cars. The categories for these dummy variables follow the groupings of Table 5. We incorporated these fixed effects to account for obvious patterns that emerged from the descriptive statistics. Recall that both vice goods and large automobiles had unusually large price dispersion which may be attributable to national differences in excise taxes and significant price discrimination, respectively. A dummy for services is included to distinguish the role that we ascribe to trade in price dispersion and the standard practice of treating all services as non–traded and all goods as traded. Lastly, we include a dummy variable for electronic goods because they have unusually low price dispersion among manufactured goods. The coefficients on these dummy variables all have the anticipated signs, the fixed effects for large cars and vice goods are very large economically, services less so, while for electronic goods the coefficient is tiny and not statistically significant.

The regression results support the hypothesis that both the share of trade and the cost share of non-tradeable inputs matter for international price dispersion. The coefficients are of the anticipated sign: more trade reduces price dispersion and more non–tradeable inputs increases retail price dispersion. The coefficient on the trade share suggests that if we consider two goods, both with the minimum amount non–traded inputs (0.10), one non-tradeable and one with an average level of tradeability (0.75 using our metric), price dispersion will fall (on average) by roughly 17 percent, going from 23 percent to 19 percent. The coefficient on non-traded inputs implies that, for a non–traded good, if the cost share of non-tradeable inputs used to produce a good doubles, from 10 to 20 percent, price dispersion increases by about 8 percent, going from roughly 23 to 25 percent.

We explain a large fraction of the dispersion in the sectoral data, with the R^2 from the GLS regression equalling 0.33. We also report the fraction of variance explained based on the raw data (0.126), which is not as informative since it primarily reflects the different levels of aggregation of the regressors and regressand.

4.3 Pricing to Market

Arthur Pigou (1920) defined price discrimination as being present when different groups of consumers pay different prices for identical goods. In the "pricing to market" literature (e.g. Krugman (1987)) price discriminating oligopolist suppliers use their market power to sustain price differences across national boundaries. Identical goods, then, could sell at different prices across countries even when converted to a common currency. Alternatively, the goods might not actually be identical in which case monopolistically competitive firms could charge different prices depending on the elasticity of substitution between them. Assuming that international goods are homogenous when in fact they are different varieties of the same good would, under monopolistic competition, lead to unfounded rejections of the law of one price.

Our panel data is sufficiently rich that we can shed light on these two alternative views of the microeconomic structure of goods markets. The procedure comes down to a two-way analysis of variance. The first dimension of the variance captures the price differences across brands of the same good within a country. We refer to these as brand effects: the price differences domestic consumers pay for differentiated brands of the same good. The second dimension of the variance captures the differences in price of the same brand across countries. We refer to these as country effects: the price differences international consumers pay for identical brands of a particular good. We gauge the relative importance of product heterogeneity and geographic price discrimination by comparing their contributions to total price variance.

The sub-sample we use consists of those goods in our dataset for which prices are collected for multiple brands of otherwise homogeneous products. Using the entire collection of goods for which brands are indicated would allow us to include 950 different brands of goods. Due to the sparseness in available data for some individual brands or entire categories of goods we have adopted the following criteria for selecting data into this part of our analysis.

First, we exclude a good if the price survey contains less than four different brands since this would limit out ability to infer the variance of price across brands for that good. Second, we exclude a brand when price observations are available for less than four countries since including it would limit our ability to infer variance of its price across countries.

Table 7 presents some detail on the goods included in our subsequent analysis. Our selection criteria resulted in a sample of 35 goods and total of 287 brands. The number of brands per good ranges from a high of 28 brands for automobiles with engines between 1,200 cubic centimeters and 1,700 cubic centimeters, to a low of 4 brands for whisky. The average number of brands per good is 8 and the average number of price observations per good is 82.

We convert each national currency price in our sub-sample to a common numeraire and then take the log of these prices to obtain, $y_{ij}^h = \log(e_j p_{ij}^h)$, the log price for brand h of good i, in country j. Next we estimate three linear regressions for each good separately. The first regression includes both a brand and country dummy:

$$y_{ij}^h = \alpha_i + d_j\beta_j + d_h\delta_h + \varepsilon_{ij}^h \tag{10}$$

where i denotes the good, α_i is a constant scalar. Defining H_i as the number of brands of good i, y_{ij}^h is the $H_iN\times1$ vector of log prices, d_j is an $H_iN \times N$ matrix of country-specific dummy variables for country j capturing country effects common across brands, d_h is an $H_iN \times H_i$ matrix of brand-specific dummy variables capturing effects specific to brand h but common to all countries, β_j and δ_h are the respective $N \times 1$ and $H_i \times 1$ coefficient vectors for each of the two effects, and ε_{ij}^h is the $H_iN \times 1$ vector of residuals.

In order to enable estimation for each good i, we set β_2, the country-specific parameter for Belgium, and δ_1, the parameter for the first brand of each good, equal to zero. Thus we estimate deviations from the first brand of each good in Belgium. The decomposition results were not sensitive to the choice of country.

Next we regress the price observations for each good i on country-specific dummy variables,

$$y_{ij}^h = \gamma_i + d_j\theta_j + \upsilon_{ij}^h \tag{11}$$

and then on brand-specific variables,

$$y_{ij}^h = \mu_i + d_h\phi_h + u_{ij}^h \ . \tag{12}$$

In the absence of missing observations these last two regressions are unnecessary because the explanatory variables would be orthogonal and estimating the effect of brand and country separately leads to the same answer as estimating them jointly (using equation (10)). In this case, the sum of squares explained by the full model is equal to the sum of squares explained by the other two.

We refer to the sum of squares for the regression using country dummy variables as the primary measure of the sum of squares explained by country-specific effects alone, and similarly the sum of squares from the second univariate regression is referred to as the primary measure of the sum of squares explained by brand-specific effects alone. We refer to secondary measures of the sum of squares explained by country or brand-specific effects as the difference between the sum of squares explained by each of the univariate models and that explained by the full model. For example the secondary measure of the sum of squares explained by country-specific effects is obtained as the increment to the explained sum of squares achieved by running the full model with both country-specific and brand-specific dummy variables on the residuals of the model with only brand-specific effects. We report the average of the primary and secondary measures of brand and country effects in the results.

Table 8 reports the total variance of prices across brands and countries for each of the goods separately and the decomposition of the variance into brand effects, country effects, and a residual term. The last column gives a sense of the importance of brand effects versus country effects by recording brand when the proportion of variance explained by brand effects exceeds the proportion of variance explained by country effects by more than the variance left unexplained. The term country records the opposite result while the term ambiguous is left for the remaining cases.

We see that brand and country effects combine to capture a considerable amount of the variation with the fraction left unexplained typically below 10%. In terms of relative importance, the brand effect dominates the country effect in 17 of the 29 cases that are not ambiguous. Taking a simple average of the variance ratios across all goods gives a similar result: brand effects are moderately more important than country effects.

The last row reports the results of pooling all observations into a single panel. We accomplish this by first removing the average price of each good from each brand and then treat all the price observations as if they were *brands* of a single good. We perform the same regressions as before to decompose the variance into brand effects, country effects, and a residual. We find that brand effects explain about 55% of the variance while country effects capture less than 8%.

There are some interesting differences across the individual goods. Goods with particularly strong country effects include whiskey, cigarettes, various services related to automobile maintenance (replacement of tires, brake lining, or clutch linings) and automobiles themselves. In the case of cigarettes, 84% of the price dispersion is accounted for by country effects compared to only 6% for brand effects. The numbers are almost identical for whiskey. Our suspicion is that the large country effects for these commodities reflect differing levels of excise tax on these goods across countries. Similarly automobiles exhibit substantial geographic price dispersion, with about two-thirds of the dispersion accounted for by country effects and one-third accounted for by brand effects. These findings are broadly consistent with work by Knetter that emphasizes the role of market power in accounting for international price dispersion in the automobile industry. Finally, the dispersion in the price of services associated with automotive repair are exactly what one would expect given the hypothesis that non-tradeables (i.e. labor services) will exhibit a relatively large degree of price dispersion. We find that more than 80% of the dispersion in this case is attributable to variation across countries. Most of the remaining goods have as much, or more, dispersion in prices across brands within countries as they do across countries.

Table 9 repeats this exercise using the eight original EC member countries. Recall that these countries are on average more similar in wealth levels and closer geographically than the larger group of countries. As one might expect, the main impact on our results is that the country effects are now less important. In this case 23 of 32 cases that are unambiguous indicate that brand effects are more important. The pooled regression indicates that brand differences explain three-fourths of the relative price variation whereas country effects account for a mere 3%.

The overall picture emerging from these comparisons is that the ability of manufacturers to price discriminate across national boundaries is comparable to their ability to differentiate their product lines within countries. We find these results surprising in the sense that we would have expected national borders to matter more.

5 The Geography of Price Dispersion

We found, in the previous section, that countries in the geographic periphery of Europe are evidently more subject to price discrimination than countries that are close neighbors. To evaluate the role of economic geography more directly, we follow the empirical trade literature that relates the volume of trade between a pair of countries and the geographic distance that separates them.

To do so, we use the same measure of price dispersion across locations as we did in characterizing the dispersion in Section 3, but now we average across goods to get a sense of the overall differences in prices, labelling the result $\bar{z}_{kl}$:

$$\bar{z}_{kl} = N^{-1} \sum_{i=1}^{N} |z_{ik} - z_{il}|, \tag{13}$$

where k and l are two different cities and i indexes the good. The value of this sum is zero only when purchasing power holds exactly for every good in the summation. We explore the proposition that price differences depend positively on geographic distance using the following regression specification:

$$\bar{z}_{kl} = \alpha + \beta_1 d_{kl} + \beta_2 d_{kl}^2 + \nu_{kl}, \tag{14}$$

where d_{kl} is the distance between city k and city l.

We measure distance as the number of kilometers (in thousands) between the cities from which the price data are collected.[3] The thirteen European cities range in distance from as little as 146 km (Luxembourg and Dusseldorf) to as much as 2,866 km (Athens and Dublin). The average distance between a particular city and all the others is a measure of remoteness and indicates that the cities in the periphery of our geographic region are: Athens, Lisbon, and Madrid while Brussels and Luxembourg are the most central (though Paris, Dusseldorf, and Amsterdam are close behind).

Table 10 presents the regression results. Beginning with the specification with both a distance and squared–distance term, the constant term of 0.23 implies a mean absolute deviation of prices across location of about 23% when distance is zero. Ignoring the squared distance terms, the coefficient 0.21 on distance means that 1,000 kilometers of distance approximately doubles the dispersion of prices. Thus dispersion is predicted to be considerable even after controlling for distance which is itself quite economically important.[4]

Thinking of the coefficient on distance as a measure of transportation cost would place such costs at the high end of what is found using export and import price data. Typically, shipping costs are on the order of 10% when averaged across traded goods. The much larger cost estimates obtained here suggests that more is needed to explain price dispersion at the retail level which is perhaps not surprising in light of our earlier findings of economically significant roles for trade shares, non–traded inputs, and imperfect competition in accounting for price dispersion.

To place the geographic price dispersion implied by the distance regressions in perspective consider the following simple experiment. Send one member of each household shopping for a common basket of goods and have them purchase each item in the lowest cost location. We view the difference between this minimized expenditure and expenditure when all shopping is constrained to the home country as a comprehensive measure of the costs of arbitrage.

The results of this exercise are startling. When shopping across all 13 countries we fill our consumption basket with 212 commodities. The minimum cost of this basket is 730,267 Belgian Francs which represents a saving of slightly less than 15% for Belgium, Germany and Luxembourg and between 20% and 30% for Austria, France, Italy, the Netherlands, and the United Kingdom. The savings are much larger in the remaining countries. For example, a resident of Athens is predicted to save 75% of expenditure! As the last column of the table indicates, automobiles are again important in the comparisons. Eliminating them from the calculation generally reduces the savings for the poorer countries and increases them for the richer

[3] We obtained this data from Steve Mitchell's homepage at Fresno State University: http://www.atinet.org/steve/cs150/. The measure is the greatest circle distance between the airports in the respective cities. The cities are: Amsterdam, Athens, Brussels, Copenhagen, Dublin, Dusseldorf, Lisbon, London, Luxembourg, Madrid, Paris, Rome and Vienna.

[4] We also partitioned the data into trade and non-traded goods and very large or very small purchases (in terms of purchase price). The coefficients on distance were similar across cuts at the data but the constant term differed: so the type of good operated like a fixed effect that was not correlated with distance. Non-traded goods were found to have significantly higher price dispersion across locations even after controlling for distance as one might expect.

ones. Thus savings are much more evenly distributed across countries when the comparison is restricted to exclude automobiles.

Using a smaller subset of countries with more complete price observations yields similar results in terms of percentage of savings on the shopping trip. The difference is that the sample expands to include many more items since more price observations are available for the richer countries. The minimum expenditure is now 3,037,917 Belgian Francs. The fact that the percentage savings does not change much with such a dramatic shift in the sample suggests that sample selection is not an important source of bias in the comparison of savings across rich and poor countries in the upper panel of the table.

We view these measures as merely illustrative of the potential expenditure savings since they fail to take into account the substitution individuals make when faced with different prices. Details aside, however, the point is that the mean real exchange rate is not very informative about deviations from the law of one price at the level of individual goods. Comparing France and Belgian, the average real exchange rate is almost identically one (see Table 2) yet the expenditure saving is about 14% in Belgium compared to 30% in France. To say that these two markets are integrated based on the mean real exchange rate is absurd without supporting evidence that explains the deviations from the law of one price existing at the micro-level.

6 Conclusion

This paper is the first to study the patterns of real exchange rates across European Community countries at the level of individual goods. The richness of our cross-sectional data allows us to venture where studies based on more aggregative data have not: a characterization of what determines the good-by-good dispersion in absolute deviations from the law of one price. We find that EC currencies had comparable purchasing power in the mid-1980's, at least when we restrict the comparison to countries with comparable wealth levels. This is to say, average real exchange rates are surprisingly close to unity. In contrast, dispersion around these averages is large, implying, for example, that someone living in Germany faces a distinctly different set of relative prices vis-a-vis someone living in France. The bulk of our paper is dedicated to characterizing this dispersion in terms of factors emphasized by economic theory: tradeability, non-traded inputs into production, product heterogeneity, and the distance between the markets. Taken as a whole, our evidence suggests that a substantial fraction of what determines dispersion in real exchange rates is attributable to these types of factors. In a nutshell, while there is certainly some degree of location specificity to what defines 'a good,' there is also an important degree of good specificity which links prices across national markets in a manner which is consistent with basic microeconomic principles.

Our findings have a number of implications for theoretical work in international economics and finance. They suggest, for instance, that the voluminous literature on exchange rate determination and the time series behavior of real exchange rates may be sensitive to both the weighting scheme used to construct a price index and the breadth of the basket of goods underlying the index. Our findings also suggest that the notion that all prices are equally sticky when denominated in units of domestic currency is implausible. A model of real exchange rates more in line with the patterns of dispersion we find is likely to involve a hybrid which incorporates both imperfect competition and other sources of goods market segmentation. The relative importance of nominal and real sources of segmentation remains an open question.

Recent contributions that incorporate some of these elements are found in Betts and Devereux (2000) and Betts and Kehoe (1999). These models like many others in the literature seem to do reasonably well at mimicking the time series properties of aggregate real exchange rates (e.g. persistence and volatility). The type of evidence we present provides a more powerful test of the various microfoundations upon which models of this class are based.

Beyond this, however, our single cross section leaves us unable to be more definitive, something we defer to future work in which we plan to augment our cross section with data from 1975, 1980, and 1990. Aside from interest in the implications of our data for theoretical work on real exchange rates, we will be very interested in what such data show us regarding the various stages of European unification.

Appendix A: Data Sources and Constructs

National retail price data. The retail price data are found in: "Price Structure of the Community Countries in 1985," compiled and published by Eurostat, Brussels, Luxembourg, 1988. The surveys were carried out by Eurostat between the end of 1984 and the beginning of 1986 as follows: Autumn 1984, clothing, footware, household textiles; Spring 1985, durable household goods; Summer 1985, services; Autumn 1985, glassware, other household articles; Spring 1986, food, beverages, tobacco; and Summer 1986, health services. 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. Few prices vary by less the +/- 15% within cities and there is commonly a spread of +/- 50%, even for very precisely defined articles, in some cases with make and model. Prices are collected at different locations so that the average price is representative of the distribution within the city. The original project involved a sample of nearly 2,800 prices but only 1,805 are in the printed source.

The panel data were not available electronically so we had a private firm key-punch the data and commodity descriptions into a spreadsheet. There are 1,805 individual retail prices for the twelve European Community price plus Austria. Entries involve a Eurostat code, a detailed description of the commodity (e.g. long grained rice), the units purchased (in cartons of 500 grams) and the local currency prices of each country. Missing observations are indicated with a -1.

Data Reconciliation. In order to explain the price dispersion across countries that exists in our detailed dataset of 1,805 commodity prices in thirteen European countries as of 1985, we constructed variables that measure tradeability, and measures of costs of non–traded inputs into production. To accomplish this we assigned each one of the 1,805 commodities to a unique International Standard Industrial Classification (ISIC) Revision 2 sector. 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.

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 commodity 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.

Tradeability. We obtained data on imports, exports, gross output, and exchange rates for the period 1980 to 1987 from the OECD STAN Database 1994. This contains data on forty-nine overlapping subdivisions of manufacturing (sector 3) mostly at a three-digit level of detail that extends to four-digits for some industry groups. 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 construct additional tradeability indices for Agriculture (sector 1) and Electricity, Gas, and Water (sector 4) using the OECD Sectoral Database of 1994 which provides value-added instead of gross output data but we have this data only for six countries: Germany, France, Italy, Belgium, the UK, and Denmark.

The sector Electricity, Gas, and Water requires special attention. Using value-added data from the OECD Sectoral Database, the 1985 value of the tradeability index is 0.97 and the average for the period 1980 to 1987 is 1.03. The average for 1970 to 1990 equals 0.72 and reflects more accurately the tradeability of this sector's product since looking at the tradeability series for the period 1970 to 1990 it is evident that the values for the period 1980 to 1985 are outliers.

We do not need information on Mining (sector 2), and Construction (sector 5) since there are no commodities from these sectors in our price data. For Services (sectors 6, 7, 8, and 9) there are no exports and imports data to the best of our knowledge. We assume that commodity services provided by these sectors are not traded and set the tradeability indices equal to zero.

In order to obtain tradeability indices for as many industries as possible we limit the time period to 1980-1987. For 24 of the 32 manufacturing sectors, the number of countries we aggregate over is $m_j = 8$ (the countries are: Austria, Belgium, Denmark, France, Germany, Italy, Spain and the United Kingdom); for 5 sectors the number of countries is 7 and for the remaining 3 manufacturing sectors the number of countries is 6.[5] For agriculture, and electricity, gas, and water we aggregate over 6 countries as data are missing for Spain and Austria.

We construct a variable that captures the degree of tradeability for each industry's commodity as follows:

$$\theta_k = \frac{\sum_{j=1}^{m_k}(X_{kj} + M_{kj})}{\sum_{j=1}^{m_k} Y_{kj}} \tag{15}$$

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–Ouput Data. We use the input-output matrix for the United States in 1987. 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 thank Tom Prusa for pointing us to this data which is available at the National Bureau of Economic Research home page.

Appendix B: Statistical Appendix

We think of the overall commodity space as consisting of G distinct groups of goods, elements of each group having some economically meaningful attributes in common. Our categorization based on ISIC codes, for instance, contains groups such as textiles, automobiles, personal services, and so on. We denote y_{ig} as the measure of price dispersion (defined in the text) for some good, i, in group $g \in \{1, 2, \ldots, G\}$. Similarly, we denote x_{ig} and w_{ig} as vectors of attributes associated with the ith good of group g. The distinction between x and w will involve aggregation within a group, g, for the former.

[5] For industry 3540 we are missing France, for industries 3832, 3839, and 3842 Belgium, and for 3900 Austria. Moreover, for industries 3825, and 3844 we are missing both Belgium and Austria, and for 3843 both Belgium and Denmark.

We assume that the joint distribution of y, x and w is such that the regression of y onto x and w is linear, so that we can write,

$$y_{ig} = \alpha + x_{ig} \cdot \beta + w_{ig} \cdot \gamma + u_{ig} \ , \tag{16}$$

where u_{ig} is *i.i.d.* with mean zero and variance σ^2, for all i and g. What distinguishes groups of goods is the conditional distribution. For x and w we assume variation across groups in the conditional mean but not the conditional variance. Denoting the conditional means, $\mu_g = E(x_{ig} \,|\, g)$ and $\delta_g = E(w_{ig} \,|\, g)$, we assume that,

$$x_{ig} \sim F(\mu_g, \Sigma) \ \ , \ \ w_{ig} \sim F(\delta_g, \Gamma) \ ,$$

for some distribution function, F.

Turning to the basic issue — x being aggregated — suppose that every element of x were only observable up to the within-group mean. Further, suppose that this is not the case for w, where we do observe individual observations. Should we average both y and w in order to estimate β and γ and, just as importantly, their standard errors? If the answer is yes, this might be problematic, depending on the specifics of what w is. For example, suppose that w_{ig} is the cost of the i,gth good and that the costs uniformly distributed throughout each group, g. Then, by averaging within a group, we lose information on how within-group costs are related to price dispersion (another way to say this is that γ won't be identified under these conditions). We now turn to a discussion of the merits of each approach. We first discuss the merits of using the data we have, as-is, and then go on to discuss the advantages of averaging away the intra-group variation in both y and w.

Estimation Based on Raw Data

When intra-group variation in x is averaged away — a data restriction, not a choice — but the intra-group variation in y and w remains, the variation ends up in the error term. To see this, note that the population regression (16) can be written as,

$$y_{ig} = \alpha + \mu_g \cdot \beta + w_{ig} \cdot \gamma + (x_{ig} - \mu_g) \cdot \beta + u_{ig} \ . \tag{17}$$

The (population) error term from the regression of y onto μ_g and w is therefore $(x_{ig} - \mu_g) \cdot \beta + u_{ig}$. This object is cross-sectionally uncorrelated (*i.e.*, $\text{cov}(x_{ig} - \mu_g, x_{jg} - \mu_g) = 0$ for all i and g) and, so long as x and w are orthogonal, is uncorrelated with the regression function, $\alpha + \mu_g \cdot \beta + w_{ig} \cdot \gamma$.

Given that this is the case — the distributional assumptions above imply it — we can write the following variance decomposition:

$$var(y_{ig}) = var(\mu_g \cdot \beta + w_{ig} \cdot \gamma) + var((x_{ig} - \mu_g) \cdot \beta + u_{ig}) \ .$$

The error term is homoskedastic as long as the conditional covariance matrix does not depend on g. That is,

$$var((x_{ig} - \mu_g) \cdot \beta + u_{ig}) = \beta^\top \Sigma \beta + \sigma^2 \ .$$

Estimates of β and γ based on (17) will therefore be consistent and efficient, conditional on the restriction that we lack observations on x_{ig}. The fit of the regression, on the other hand, will understate the fit of the unrestricted regression, equation (16), and must be interpreted accordingly.

Finite sample considerations change matters in an important way. Suppose that we have N_g observations on y, x, and w, from each group g. The sample analog of equation (16) is,

$$y_{ig} = \alpha + \bar{x}_g \cdot \beta + w_{ig} \cdot \gamma + (x_{ig} - \bar{x}_g) \cdot \beta + u_{ig} \ , \tag{18}$$

where,

$$\bar{x}_g = \frac{1}{N_g} \sum_{i=1}^{N_g} x_{ig} \ .$$

Equation (18) is still a regression (*i.e.*, $\text{cov}(\bar{x}_g, x_{ig} - \bar{x}_g) = 0$ for all i and g), but, because of the sampling variance in $\bar{x}_g$, the covariance matrix of the residuals will have a particular, heteroskedastic structure. Given $N = \sum_{g=1}^{G} N_g$ total observations, the covariance matrix is block-diagonal, with each block defined in terms of observations from a given group, g. Each block has off-diagonal terms equal to,

$$cov((x_{ig} - \bar{x}_g) \cdot \beta + u_{ig}, (x_{jg} - \bar{x}_g) \cdot \beta + u_{jg}) = -\frac{1}{N_g} \beta^\top \Sigma \beta^\top \ ,$$

for a given g and $i \neq j$, and diagonal terms equal to

$$var((x_{ig} - \bar{x}_g) \cdot \beta + u_{ig}) = \frac{N_g - 1}{N_g} \beta^\top \Sigma \beta + \sigma^2 \ .$$

What's likely to be most important, therefore, are the off-diagonal terms in each block, where variation in N_g will have a much larger effect.

Finally, should we choose to correct the regression based on equation (18) for heteroskedasticity, we need estimates of Σ and σ. The latter can be obtained via the GLS regressions outlined below. The former is more problematic. In general, we cannot estimate Σ without observing individual observations of x_{ig}. That is, since,

$$\begin{aligned} var(x) &= E[var(x \mid g)] + var(E[x \mid g]) \\ &= \Sigma + var(\mu_g) \ , \end{aligned}$$

we can estimate $\mathrm{var}(\mu_g)$ but we cannot estimate Σ without some information on x itself. So, efficient estimation based on equation (18), in which we do not average away any variation in either x or w, is not possible without further assumptions regarding the conditional covariance matrix, Σ. In what follows we experiment with a number of arbitrary, but sensible, values for Σ and examine the implications.

Estimation Based on Averaged Data

We now consider the merits of estimating β and γ by averaging away the within-group variation in both y and w. In this case, the distinction between x and w is not relevant so, for notational simplicity, we subsume w into x. Averaging equation (16) within groups, we have

$$\begin{aligned} \frac{1}{N_g}\sum_{i=1}^{N_g} y_{ig} &= \alpha + \frac{1}{N_g}\sum_{i=1}^{N_g} x_{ig}\cdot\beta + \frac{1}{N_g}\sum_{i=1}^{N_g} u_{ig} \\ \Rightarrow \bar{x}_g &= \alpha + \bar{x}_g \cdot \beta + \bar{u}_g \ , \end{aligned} \tag{19}$$

where the sample averages, $\bar{y}_g$, $\bar{x}_g$ and $\bar{u}_g$ are defined in the obvious way and are understood to depend explicitly on the values N_g.

Residuals based on equation (19) will also be heteroskedastic, but in a simpler way than those based on equation (18). The covariance matrix is diagonal with the gth diagonal element equal to σ^2/N_g. A consistent, efficient estimator of β is therefore the GLS estimator $\tilde{\beta} = (X^\top\Omega^{-1}X)^{-1}X^\top\Omega^{-1}Y$, where Y and X denote observations on $\bar{y}_g$ and $\bar{x}_g$, respectively, and Ω is the covariance matrix of the (averaged) error terms. Note that, as is straightforward to show, if we omit w from equations (18) and (19), OLS estimates based on equation (18) are numerically identical to the GLS estimates, $\tilde{\beta}$. What turns out to differ in an important way are the standard errors and the measures of fit of the respective regressions.

The main advantage to estimation based on (19) is that we don't need an estimate of Σ, the conditional variance of the averaged regressors, in order to obtain efficient estimates. Goodness of fit measures based on (19) are also easier to interpret, since we are not trying to explain unaveraged variation in x using averaged variation in y. The disadvantages are mainly related to averaging things we do not have to, in particular the variables in w.

Finally, there are a number of well known issues associated with computing goodness-of-fit measures based on a regression where estimates are obtained by GLS. The basic issue is whether one uses residuals based on (19) or residuals based on the standard 'transformed' GLS sample regression,

$$R\bar{y} = R(\alpha\iota + \bar{x}\beta) + R\bar{u} \ ,$$

where $R^\top R = \Omega^{-1}$, $\bar{y}$ and $\bar{x}$ are vectors of sample observations (where, again, w is subsumed into x), and ι is a vector of ones.

Our approach is simple. The quantity we are ultimately interested in is,

$$\frac{var(x_{ig} \cdot \beta)}{var(y_{ig})} \ .$$

A consistent estimator is of this is,

$$R^2 = \frac{\beta' v\hat{a}r(\bar{x}_g \sqrt{N}_g)\beta}{v\hat{a}r(\bar{y}_g \sqrt{N}_g)} \ ,$$

where $\widehat{var}(\cdot)$ denotes sample variance. While this quantity is not guaranteed to lie between zero and unity, our experience is that it gives sensible answers which incorporate the large amount of variation in N_g exhibited by our dataset.

Table A-1. Trade Shares by ISIC Category

Description	ISIC Code	Input Share	Trade Share
Agriculture, hunting, forestry and fishing	1000	0.17	0.40
Food (3110+3120)	3115	0.16	0.28
Beverage industries	3130	0.17	0.27
Tobacco manufactures	3140	0.15	0.20
Manufacture of textiles	3210	0.14	0.60
Manufacture of wearing apparel except footwear	3220	0.12	0.54
Manufacture products except footwear and apparel	3230	0.17	0.70
Manufacture of footwear except rubber or plastic	3240	0.12	0.68
Manufacture of furniture and fixtures except primarily metal	3320	0.21	0.27
Manufacture of paper and paper products	3410	0.15	0.51
Printing, publishing and allied industries	3420	0.24	0.15
Manufacture of industrial chemicals (3511-3513)	3510	0.22	0.85
Manufacture of other chemical products	3520	0.30	0.51
Drugs and medicines	3522	0.30	0.44
Chemical products, n.e.c.	3529	0.21	0.56
Misc. products of petroleum and coal	3540	0.14	0.40
Rubber products	3550	0.16	0.56
Plastic products, n.e.c.	3560	0.16	0.29
Pottery, china and earthware	3610	0.22	0.26
Glass and glass products	3620	0.22	0.47
Other non-metallic mineral products	3690	0.25	0.20
Iron and steel basic industries	3710	0.31	0.46
Non-ferrous metal basic industries	3720	0.20	0.70
Fabricated metal products except machinery and equipment, n.e.c.	3810	0.17	0.38
Manufacture of machinery except electrical (3820-24 and 3829)	3820	0.15	0.71
Office, computing and accounting machinery	3825	0.14	1.41
Other machinery and ordinance	3829	0.16	0.59
Electrical machinery	3830	0.19	0.49
Radio, TV., and communications equipment and apparatus	3832	0.14	0.54
Electrical apparatus and supplies, n.e.c.	3839	0.17	0.43
Shipbuilding	3841	0.14	0.34
Motor Vehicles	3843	0.11	0.64
Motorcycles and bicycles	3844	0.14	0.52
Professional, scientific, measuring and control equipment	3850	0.18	1.39
Other manufacturing industries	3900	0.23	1.35
Electricity, gas and water	4150	0.26	0.72
Restaurants and hotels	6300	0.26	0.00
Transport, storage and communication	7000	0.32	0.00
Inland Transport	7110	0.28	0.00
Maritime Transport	7120	0.27	0.00
Communication	7200	0.20	0.00
Financing, insurance, real estate and business services	8000	0.30	0.00
Community, social and personal services	9000	0.31	0.00

References

[1] Alessandria, George, 1999, "Deviations from the law of one price: preserving market share under the threat of search," mimeograph.

[2] Balassa, Bela, 1964, "The purchasing power parity doctrine: A reappraisal," *Journal of Political Economy* 72, 584-596.

[3] Baumol, William J. and William G. Bowen, 1966, Performing arts: The economic dilemma. New York: Twentieth Century Fund.

[4] Betts, Caroline and Michael B. Devereux, 2000, "Exchange rate dynamics in a model of pricing-to-market," Journal of International Economics 50, 215-244.

[5] Betts, Caroline and Timothy J. Kehoe, 1999, "tradeability of goods and real exchange rate fluctuations," mimeograph.

[6] Cassel, Gustav, 1918, "Abnormal deviations from international exchanges," *Economic Journal* 28, 413-415.

[7] Cumby, Robert, 1996, "Forecasting exchange rates and relative prices: Is what you want what you get with McParity?" NBER Working Paper 5675, July.

[8] Engel, Charles, 1993, "Real exchange rates and relative prices: an empirical investigation," *Journal of Monetary Economics* 32, 35-50.

[9] Engel, Charles and John Rogers, 1996, "How wide is the border," *American Economic Review* 86, 1112-1125.

[10] Ethier, Wilfred, 1979, "Internationally decreasing costs and world trade," *Journal of International Economics* 9, 1-24.

[11] Eurostat, 1988, "Price structure of the Community countries in 1985," Luxembourg: Offices of Official Publications of the European Community.

[12] Feenstra, Robert, 1995, "Estimating the effects of trade policy," in Gene M. Grossman and Kenneth Rogoff, eds., *Handbook of International Economics*, vol. 3. Amsterdam: North Holland.

[13] Froot, K., M. Kim, and K. Rogoff, 1995, "The law of one price over seven hundred years," manuscript, Princeton University.

[14] Giovannini, Alberto, 1988, "Exchange rates and traded goods prices," *Journal of International Economics* 24, 45-68.

[15] Ghosh, Atish R. and Holger Wolf, 1994, "Pricing in international markets: Lessons from the Economist," Working Paper No. 4806, National Bureau of Economic Research.

[16] Harrod, Roy F., 1933, International Economics. London: James Nisbet and Cambridge University Press.

[17] Haskel, Jonathan and Holger Wolf, 1998, "Why does the law of one price fail: A case study," mimeograph.

[18] Isard, Peter, 1977, "How far can we push the law of one price?" *American Economic Review* 67, 942-948.

[19] Keynes, John Maynard, 1932, "Essays in persuasion," New York: Harcourt Brace.

[20] Knetter, Michael, 1989, "Price discrimination by US and German exporters," *American Economic Review* 79, 198-210.

[21] Knetter, Michael, 1993, "International comparisons of pricing-to-market behavior," *American Economic Review* 83, 473-486.

[22] Krugman, Paul, 1987, "Pricing to market when the exchange rate changes," in *Real-financial linkages among open economies* Cambridge:MIT press, 49-70.

[23] Krugman, Paul, 1995, "Increasing returns, imperfect competition and the positive theory of international trade," in Gene M. Grossman and Kenneth Rogoff, eds., *Handbook of International Economics*, vol. 3. Amsterdam: North Holland.

[24] Organization for Economic Co-operation and Development, 1995, *STAN Database for Industrial Analysis*, computer diskettes, Paris: OECD.

[25] Parsley, David and Shang-Jin Wei, 1999, "Border, border, wide and far, how we wonder what you are," mimeograph.

[26] Pigou, Arthur, 1920, "Some problems of foreign exchange," *Economic Journal* 30, 460-472.

[27] Richardson, 1978,"Some empirical evidence on commodity arbitrage and the law of on price," *Journal of International Economics* 8, 341-351.

[28] Rogers, John and Michael Jenkins, 1995, "Haircuts or hysteresis? Source of movements in real exchange rates," *Journal of International Economics* 38, 339-360.

[29] Rogoff, Kenneth, 1996, "The purchasing power parity puzzle," *Journal of Economic Literature* 34, 647–668.

[30] Samuelson, Paul A., 1964, "Theoretical notes on trade problems," *Review of Economics and Statistics* 46, 145-154.

[31] Sanyal, Kalyan K. and Ronald W. Jones, 1982, "The theory of trade in middle products," American Economic Review 72, 16-31.

[32] Stockman, Alan, and Harris Dellas, 1989, "International portfolio nondiversification and exchange rate variability," *Journal of International Economics* 26, 271-289.

[33] Stockman, Alan, and Linda Tesar, 1995, "Tastes and technology in a two-country model of the business cycle," *American Economic Review* 85, 168-185.

[34] Summers, Robert and Alan Heston, 1991, "The Penn World Table (Mark 5): An Expanded Set of International Comparisons: 1950-1988," *Quarterly Journal of Economics* 106, 327-368.

Table 1. Sample records from price survey

Code	Good description	Units	Austria	Belgium	Denmark	France	Germany	Greece	Ireland	Italy	Lux.	Neth.	Portugal	Spain	U.K.
11111	Long grained rice - in carton	500 g	26	50	48	52	52	42	51	35	52	36	-1	39	32
11251	Chicken - fresh, presentation 70 %	1.2 kg	107	204	208	206	-1	116	159	165	-1	-1	89	-1	149
11421	Condensed unskimmed milk 9 - 10 % butterfat	410 g	-1	48	-1	-1	56	32	-1	-1	-1	54	56	42	43
11621	Dried almonds	100 g	28	54	-1	62	34	28	43	35	54	49	37	50	50
11911	Ground blended coffee	1 kg	234	439	426	355	-1	-1	-1	590	-1	355	456	438	447
13112	Liqueur - s.b.	0.7 L	801	645	1190	609	611	468	975	532	498	609	-1	330	936
21112	Mans jacket type k-way	1 nb	1582	544	1598	1021	823	1368	1332	874	996	1195	900	714	-1
22121	Ladies boots, box caf	1 pr	3716	4027	4100	3730	3687	2635	2362	-1	4422	2611	2731	2958	3153
41112	Chest of drawers	1 nb	2832	6988	-1	-1	4880	2996	5058	7297	3990	7724	3577	4659	3934
42111	Spring mattress: s.b.	1 nb	-1	-1	-1	-1	-1	-1	5383	-1	-1	-1	-1	-1	4680
43121	Dishwasher: 5 programmes, s.b.	1 nb	-1	23480	-1	-1	-1	-1	-1	23912	-1	-1	23014	22966	-1
43161	Iron: steam, s.b.	1 nb	1678	2580	-1	2200	1247	3194	1647	1554	-1	1190	-1	1336	1441
45111	Washing powder : s.b.	700 g	100	75	98	68	88	44	59	77	82	85	-1	-1	-1
52211	Hearing aid: s.b.	1 nb	-1	-1	-1	32715	28864	-1	24426	35348	19400	24354	-1	-1	26287
61113	Car: engine between 1200 and 1700 cc, s.b.	1 nb	467864	387740	770809	481011	403928	-1	-1	533097	369000	464874	-1	-1	514101
61212	Bicycle : racing, s.b.	1 nb	-1	12206	-1	-1	-1	16716	-1	-1	-1	11351	-1	-1	-1
62111	Bus, single ticket, about 6 km, with change of bus	1 ticket	52	30	50	-1	43	-1	-1	-1	25	39	-1	-1	51
71131	Record player, stereo, s.b.	1 nb	7085	5190	-1	5144	4164	-1	-1	4792	4848	4849	4252	7877	-1
71311	Cassette for game : s.b.	1 nb	873	850	-1	1295	-1	-1	944	1089	1066	-1	-1	-1	740
72221	Rental of television	1 month	-1	2296	1339	1630	1190	2981	1382	1294	-1	1488	1678	1804	1481

Note: The table presents a sample of 20 records from the original 1,805 records. Records are exactly as they appear in the original source except the currencies have been converted to Belgian francs to facilitate comparisions. The commodities from rows (1,90,180,...1,800) were selected into the table based on the orginal order of goods in the Eurostat survey. Missing data are indicated as '-1.'

Table 2
Log Real Exchange Rates

Country	Equally Weighted	Value Weighted	Value[a] Weighted	Number of Observations
Austria	0.029	0.151	0.078	717
	(0.414)			
Denmark	0.167	0.445	0.250	663
	(0.458)			
France	0.004	0.102	0.014	750
	(0.388)			
Germany	−0.048	−0.056	−0.109	762
	(0.364)			
Greece	−0.247	0.740	0.087	691
	(0.579)			
Ireland	−0.026	0.197	−0.009	579
	(0.462)			
Italy	−0.071	0.113	0.045	745
	(0.507)			
Luxembourg	−0.092	−0.064	−0.110	712
	(0.303)			
Netherlands	−0.110	0.046	−0.110	771
	(0.375)			
Portugal	−0.259	0.374	−0.005	641
	(0.618)			
Spain	−0.162	0.138	−0.181	696
	(0.501)			
United Kingdom	−0.069	0.062	−0.157	755
	(0.446)			

[a]Excludes automobiles.

Note: Entries are q_j, the weighted average, for country j, of good-specific, log relative prices. Denoting q_{ij} as the real exchange rate for good i in country j, and γ_{ij} as the weighting scheme (described explicitly in the text), we compute q_j as $\sum_{i=1}^{N} \gamma_{ij} \log q_{ij}$. For these calculations we eliminate multiple brands of the same good.

Table 3

Deviations From the Law of One Price

Country	Percentage of Log Real Exchange Rates "x" Percent From Zero: 50>x>20	20>x>10	10>x>0	0<x<-10	-10<x<-20	-20<x<-50	|x|>50
Austria	22	12	12	10	9	15	20
Denmark	31	10	9	9	6	9	26
France	18	12	11	14	12	18	15
Germany	14	11	12	14	14	20	15
Greece	13	5	7	8	8	23	36
Ireland	17	7	11	12	9	21	23
Italy	17	10	8	11	9	18	27
Luxembourg	8	9	16	23	15	19	10
Netherlands	10	11	11	17	11	23	17
Portugal	10	6	7	7	6	22	42
Spain	13	7	9	9	8	23	31
U.K.	15	8	10	10	9	25	23
Average	14	8	9	11	9	18	22

Note: Each entry in the table is the proportion of goods for which the log real exchange rate differs from zero by the amount specified. Belgium is the numeraire country. The row "average" is a simple average of the entries in each column.

Table 4

Country Effects

Country	Mean	Standard Error
Austria	0.075	0.008
Belgium	0.023	0.009
Denmark	0.252	0.012
France	0.046	0.009
Germany	−0.026	0.008
Greece	−0.076	0.014
Ireland	0.041	0.013
Italy	0.002	0.011
Luxembourg	−0.086	0.007
Netherlands	−0.068	0.008
Portugal	−0.098	0.016
Spain	−0.064	0.011
United Kingdom	−0.009	0.011
Percent Explained	8.95	
Percent Unexplained	93.23	
Residual	−2.17	

Note: Entries represent the average, across goods for a particular country, of the percent by which each price deviates from its good-specific mean. Specifically, given the price of the ith good in the jth country, and the nominal exchange rate, e_j, the means are calculated as

$$\text{Mean} = \frac{1}{N_j}\sum_{i}^{N_j}\left(\frac{e_j p_{ij}}{\sum_{j=1}^{M_i} e_j p_{ij}/M_i} - 1\right).$$

The residual associated with the variance decomposition is due to missing observations or, equivalently, an unequal number of observations, N_j, for each country, j.

Table 5

Price Dispersion: Homegrown Aggregates

Category Number	Good	Dispersion	Dispersion Rank	Number of Goods
	Groceries	**24.65**	**7**	**356**
1	Dry Groceries	21.96	46	54
2	Frozen Meat	16.32	19	36
3	Storeable Meat	22.31	49	3
4	Frozen Fish	18.72	35	11
5	Storeable Fish	28.48	83	11
6	Storeable Milk	23.19	58	9
7	Cheese, Eggs, Butter	16.95	24	12
8	Vegetable Oil	29.59	88	7
9	Dried Fruits, Nuts	29.71	89	7
10	Canned Fruit Products	27.92	77	29
11	Frozen Vegetables	23.52	61	7
12	Canned Vegetable Products	24.55	65	12
13	Sugar	17.51	29	14
14	Coffee, Tea	22.88	55	9
15	Jam, Honey	17.98	32	11
16	Chocolate, Candy	23.43	60	30
17	Seasoning, Spices	31.55	94	10
18	Bottled Water	39.89	109	19
19	Fresh Fish	28.56	84	3
20	Soda	16.43	21	14
21	Fresh Meat	20.91	42	5
22	Fresh Milk	22.10	47	18
23	Fresh Non-European Fruit	27.83	76	11
24	Fresh European Fruit	31.10	92	3
25	Fresh Vegetables	32.79	98	11
	Vices	**34.16**	**9**	**49**
26	Liquor	28.87	85	20
27	Wine	39.66	107	9
28	Beer	39.74	108	4
29	Tobacco	28.37	80	16

Table 5 (continued)
Price Dispersion: Homegrown Aggregates

Category Number	Good	Dispersion	Dispersion Rank	Number of Goods
	Clothing	**18.74**	**4**	**94**
30	Mens clothes, formal	15.12	14	13
31	Mens clothes, not formal	18.41	33	21
32	Womens clothes, formal	17.21	26	10
33	Womens clothes, non-formal	22.47	51	8
34	Childrens clothes	25.24	67	13
35	Clothing materials (fabric, etc.)	22.93	57	13
36	Mens shoes	16.97	25	8
37	Womens shoes	12.51	3	4
38	Childrens shoes	17.80	31	4
	Residential Home Products	**21.91**	**5**	**278**
39	Floor covering	28.31	78	6
40	Fabric for furniture	24.06	63	28
41	Linen	22.45	50	7
42	Curtains	13.40	7	3
43	Dishes	23.40	59	9
44	Utensils, Pots and Pans, Kitchen Tools	24.48	64	4
45	Detergents, soaps	22.89	56	10
46	Kitchen Aids (Saran Wrap)	32.90	99	17
47	First Aid Supplies	32.61	97	10
48	Eyeglasses, Eyecare products	26.80	73	7
49	Toiletry items, cheap	17.80	30	16
50	Hearing Aids	15.40	16	12
51	Toiletry items, expensive	14.79	12	17
52	Batteries, Light Bulbs	26.59	71	10
53	Home building materials (cement, paint, etc.)	28.37	79	6
54	Hardware (tools)	29.75	90	21
55	Electric Tools	13.70	9	6
56	Wheel Chair, crutches	36.64	104	6
57	Residential furniture	25.54	69	10
58	Fridges	13.38	6	2
59	Washers, Dryers, Dishwashers, Microwaves, Radiator	12.89	5	10
60	Piano, organ	13.78	10	7
61	Typewriters	15.72	18	12
62	Vacuums	20.22	39	17
63	Sewing Machines, Irons	18.47	34	10
64	Small Kitchen Appliances	15.38	15	15

Table 5 (continued)
Price Dispersion: Homegrown Aggregates

Category Number	Good	Dispersion	Dispersion Rank	Number of Goods
	Transportation Products	**16.80**	**1**	**73**
65	Diesel Car	16.33	20	4
66	Small Car	17.24	27	8
67	Medium Car	21.76	44	8
68	Big Car	31.13	93	6
69	Motorcycles	12.55	4	3
70	Outboard motors	14.98	13	9
71	Bicycles	10.44	2	7
72	Tires	9.97	1	14
73	Car parts (small items)	16.78	23	14
	Petroleum Products	**17.53**	**2**	**5**
74	Gasoline	13.52	8	3
75	Motor Oil	21.54	43	2
	Electronics Goods	**18.21**	**3**	**73**
76	Stereo equipment, portable	17.30	28	4
77	TVs, VCRs	15.51	17	4
78	Stereo Equipment, not portable	19.11	38	16
79	Games (electronic, board)	20.27	40	17
80	Camera equipment	22.49	53	15
81	Camera Film	14.00	11	12
82	Audio Supplies (tapes, etc.)	18.80	36	5
	Miscellaneous	**24.43**	**6**	**72**
83	Sporting equipment	16.73	22	8
84	Flowers	32.24	96	8
85	Novels and Magazines	26.88	74	33
86	Jewelry	22.61	54	6
87	Luggage	18.86	37	3
88	Stationary store stuff	29.24	86	14

Table 5 (continued)
Price Dispersion: Homegrown Aggregates

Category Number	Good	Dispersion	Dispersion Rank	Number of Goods
	Services	**30.24**	**8**	**164**
89	Domestic Servant	22.17	48	7
90	Clothing related	25.18	66	19
91	Home maintenance	28.40	81	22
92	Key Cutting	38.31	105	2
93	Film Developing	25.43	68	1
94	Car Repair	31.06	91	31
95	Car Rental	32.91	100	2
96	Driving Lesson, Language Course	33.70	101	2
97	Entertainment (Films, etc.)	34.74	103	3
98	Haircuts	21.88	45	7
99	Meals at Restaurants	23.58	62	6
100	Hotel Lodging	26.53	70	1
101	Camping	20.68	41	6
102	Parking	38.61	106	7
103	Urban Bus Ride	31.64	95	3
104	Urban Tube Ride	26.77	72	4
105	Taxi	28.40	82	3
106	Coach Ride	34.46	102	11
107	Train Ride	42.91	110	3
108	Plane Ticket	27.15	75	5
109	Residential Utilities	22.48	52	14
110	Payphone Calls	49.02	111	3
111	Postage	29.43	87	2

Note: Based on author's computations. Each row presents the good-by-good geographic price dispersion (y_i in the text) averaged across goods in a particular category. For example, in category 90, the good-by-good dispersion measures for all goods assigned to the category "clothing related" are averaged together to get 25.18. The fourth column gives the rank of this dispersion measure relative to other categories, 1 being the lowest amount of dispersion and 111 being the highest amount of dispersion (clothing is somewhat below the median ranking 66 out of 111) while the fifth column indicates the number of goods included in each category (19 individual goods make up the homegrown aggregate clothing.

Table 6

Price Dispersion and Characteristics of the Goods

Variable	Coefficient	Standard Error
Intercept	0.2112	0.0128
Tradeshare	−0.0508	0.0096
Non-tradeable inputs	0.1854	0.0694
Dummy Variables:		
Large car	0.1103	0.0309
Vice goods	0.0857	0.0142
Services	0.0323	0.0123
Electronic goods	−0.0019	0.0154
R^2 (aggregated data)	0.326	
R^2 (dissaggregate data)	0.126	

Note: The table presents the coefficients of estimating regression equation (8) in the text. R^2 (aggregate data) is the R^2 measure of fit from the GLS regression as defined in Appendix B. R^2 (aggregate data) is the fraction of variance explained based on the raw data which is necessarily lower because we cannot explain intrasectoral dispersion in prices given the more aggregate nature of our explanatory variables.

Table 7

Availability of Brand Data

Good	Brands	Observations
Beer in bottle	7	35
Car battery	4	16
Cigarettes light, with filter	7	77
Coffee maker	5	30
Colour film	4	32
Dishwasher	8	96
Electric razor	5	40
Fridge-freezer	8	64
Garden chair	4	16
Garden table	5	25
Hair dryer	5	35
Hearing aid	14	168
Hi-fi cassette, unrecorded	4	32
Motor car diesel	8	72
Auto (less than 1200 cc)	23	299
Auto (1200-1700 cc)	28	364
Auto (more than 1700 cc)	13	143
Motorcycle	25	300
Portable typewriter	5	50
Reflex camera	4	24
First main service, auto	5	60
Replacement of 4 tires, auto	5	60
Replacement of brake linings, auto	9	117
Replacement of clutch linings, auto	10	120
Slide film	4	28
Spark plug	5	35
Television color	7	42
Tennis racket	5	35
Toilet soap	5	40
Tumble dryer	7	49
Tire	9	117
Vacuum cleaner, cylinder	10	120
Video recorder	8	48
Washing machine	8	64
Whisky	4	32
Total	287	2885
Average per good	8	82

Note: The first column reports a general description of the goods for which prices of multiple brands are reported in the price survey. The second column reports the maximum number of brands available for each of these goods. The third column reports the total number of observations available. The last two rows report the total number of brands and available observations and the average number of brands and available observations per good, respectively.

Table 8

Variance Decomposition (13 countries)

Good	Total Variance	Variance Decomposition Brand	Country	Residual	Dominant Factor
Beer in bottle	5.22	0.07	0.45	0.48	Ambiguous
Car battery	1.44	0.72	0.10	0.18	Brand
Cigarettes light, with filter	7.37	0.06	0.84	0.10	Country
Coffee maker	19.16	0.91	0.08	0.01	Brand
Color film	1.26	0.62	0.27	0.11	Brand
Dishwasher	7.56	0.78	0.15	0.08	Brand
Electric razor	5.68	0.77	0.16	0.06	Brand
Fridge-freezer	2.43	0.70	0.20	0.11	Brand
Garden chair	5.78	0.68	0.21	0.10	Brand
Garden table	3.29	0.88	0.02	0.09	Brand
Hair dryer	2.48	0.84	0.14	0.03	Brand
Hearing aid	7.44	0.24	0.61	0.15	Country
Hi-fi cassette, unrecorded	5.17	0.79	0.17	0.04	Brand
Motor car diesel	5.36	0.72	0.24	0.03	Brand
Auto (less than 1200 cc)	12.94	0.31	0.66	0.03	Country
Auto (1200-1700 cc)	19.95	0.30	0.66	0.04	Country
Auto (more than 1700 cc)	30.78	0.39	0.59	0.02	Country
Motorcycle	105.29	0.90	0.09	0.02	Brand
Portable typewriter	20.71	0.85	0.07	0.08	Brand
Reflex camera	1.12	0.64	0.17	0.19	Brand
First main service, car	19.29	0.09	0.49	0.41	Ambiguous
Replacement of 4 tires, auto	12.65	0.01	0.98	0.02	Country
Replacement of brake linings, auto	27.32	0.06	0.79	0.15	Country
Replacement of clutch linings, auto	25.78	0.04	0.93	0.03	Country
Slide film	1.81	0.13	0.63	0.23	Country
Spark plug	2.19	0.36	0.29	0.36	Ambiguous
Television color	2.92	0.58	0.38	0.04	Brand
Tennis racket	1.57	0.38	0.28	0.35	Ambiguous
Toilet soap	1.69	0.00	0.77	0.23	Country
Tumble dryer	8.97	0.90	0.06	0.04	Brand
Tire	4.04	0.37	0.44	0.18	Ambiguous
Vacuum cleaner, cylinder	11.41	0.41	0.26	0.33	Ambiguous
Video recorder	5.88	0.26	0.70	0.04	Country
Washing machine	7.82	0.83	0.14	0.03	Brand
Whisky	3.9	0.05	0.88	0.07	Country
Mean		0.48	0.40	0.13	Brand
Pooled (N=2422)	456.18	0.55	0.08	0.37	Brand

Note: The figures in the table summarize the results from estimating equations (10)-(12). The first column identifies each of the goods for which multiple brand observations are available. The second column is the total variance of price across brands and countries. The next three columns report the fraction (in percent) of this variance attributable to brand effects, country effects, and a residual term. The final column indicates the relative importance of brand or country effects unless the implications of the variance decomposition are ambiguous. Ambiguous here is defined as cases in which the difference in the proportions of the variance explained by the two effects is less than the proportion of the variance unexplained (the column labelled residual).

Table 9
Variance Decomposition (8 Countries)

Good	Total Variance	Decomposition of Variance Brand	Country	Residual	Dominant Factor
Beer in bottle	NA	NA	NA	NA	NA
Car battery	1.44	0.72	0.10	0.18	Brand
Cigarettes light, with filter	2.56	0.04	0.90	0.06	Country
Coffee maker	12.14	0.98	0.02	0.01	Brand
Color film	1.11	0.66	0.23	0.11	Brand
Dishwasher	3.89	0.81	0.11	0.08	Brand
Electric razor	4.65	0.80	0.13	0.07	Brand
Fridge-freezer	1.22	0.81	0.14	0.05	Brand
Garden chair	5.78	0.68	0.21	0.10	Brand
Garden table	3.29	0.88	0.02	0.09	Brand
Hair dryer	1.89	0.82	0.16	0.02	Brand
Hearing aid	2.52	0.33	0.51	0.16	Country
Hi-fi cassette, unrecorded	4.10	0.78	0.18	0.04	Brand
Motor car diesel	3.94	0.77	0.20	0.02	Brand
Auto (less than 1200 cc)	4.67	0.59	0.36	0.04	Brand
Auto (1200-1700 cc)	6.69	0.55	0.41	0.03	Brand
Auto (more than 1700 cc)	10.16	0.79	0.16	0.05	Brand
Motorcycle	73.59	0.96	0.03	0.01	Brand
Portable typewriter	13.8	0.84	0.05	0.11	Brand
Reflex camera	1.12	0.64	0.17	0.19	Brand
First main service, car	6.48	0.17	0.33	0.50	Ambiguous
Replacement of 4 tires, car	3.75	0.00	1.00	0.00	Country
Replacement of brake linings, car	8.07	0.09	0.53	0.38	Country
Replacement of clutch linings, car	5.37	0.10	0.84	0.06	Country
Slide film	1.04	0.10	0.62	0.28	Country
Spark plug	1.81	0.59	0.16	0.26	Brand
Television color	1.78	0.73	0.24	0.03	Brand
Tennis racket	0.97	0.34	0.34	0.32	Ambiguous
Toilet soap	1.22	0.02	0.84	0.13	Country
Tumble dryer	6.5	0.92	0.04	0.03	Brand
Tire	2.64	0.43	0.43	0.14	Ambiguous
Vacuum cleaner, cylinder	5.21	0.72	0.12	0.15	Brand
Video recorder	1.58	0.53	0.42	0.05	Brand
Washing machine	5.24	0.93	0.05	0.02	Brand
Whisky	0.64	0.15	0.78	0.06	Country
Mean		0.57	0.32	0.11	Brand
Pooled (N=1673)	228.96	0.75	0.03	0.22	Brand

Notes: See the notes to Table 8.

Table 10

Regression Results

	Constant	Distance	Squared	R^2
Specification 1	0.23	0.21	-0.04	0.53
	(0.02)	(0.04)	(0.01)	
Specification 2	0.29	0.09		0.48
	(0.02)	(0.01)		

Note: The table reports coefficient estimates, standard errors (in parentheses) and R^2 for the estimation of: $\overline{z}_{kl} = \alpha + \beta_1 d_{kl} + \beta_2 d_{kl}^2 + \nu_{kl}$. Specification 2 imposes the restriction $\beta_2 = 0$.

Table 11
A Pan European Shopping Trip

Panel A: All Countries

Country	All goods: Country-Specific Expenditure	All goods: Savings in Percent	Excluding automobiles: Country-Specific Expenditure	Excluding automobiles: Savings in Percent
Austria	996,707	26.73	312,976	38.06
Belgium	847,890	13.87	311,490	37.76
Denmark	1,494,882	51.15	365,520	46.96
France	1,050,702	30.50	327,107	40.73
Germany	843,672	13.44	278,762	30.45
Greece	2,905,666	74.87	374,355	48.21
Ireland	1,195,108	38.90	292,332	33.68
Italy	925,466	21.09	301,524	35.70
Luxembourg	843,432	13.42	280,432	30.87
Netherlands	964,216	24.26	273,274	29.05
Portugal	1,727,433	57.73	300,229	35.43
Spain	1,427,634	48.85	318,352	39.10
United Kingdom	1,022,277	28.56	254,497	23.82
Minimized	730,267		193867	
Number of goods	212	212	211	211

Panel B: Eight Countries

Country	All Goods: Country Specific Expenditure	All Goods: Savings in Percent	Excluding Automobiles: Country Specific Expenditure	Excluding Automobiles: Savings in Percent
Austria	4,309,557	29.51	1,612,851	36.32
Belgium	3,655,630	16.90	1,441,450	28.74
France	4,248,811	28.50	1,499,887	31.52
Germany	3,481,140	12.73	1,325,208	22.49
Italy	4,427,646	31.39	1,601,270	35.86
Luxembourg	3,322,702	8.57	1,311,902	21.71
Netherlands	3,934,868	22.79	1,268,177	19.01
United Kingdom	3,938,714	22.87	1,204,469	14.72
Minimized	3,037,917		1,027,117	
Number of goods	433	433	429	429

Note: The table reports the total expenditure necessary to purchase a common basket of goods in each country ("country-specific expenditure") and the expenditure required if each item is purchased from the least-cost location ("minimized"). The first panel utilizes a basket of goods for which all 13 countries have prices available while the second does the same exercise but restricted to 8 countries. For these calculations, there are no multiple brands, brands were eliminated in the same way as is discussed in the PPP section.

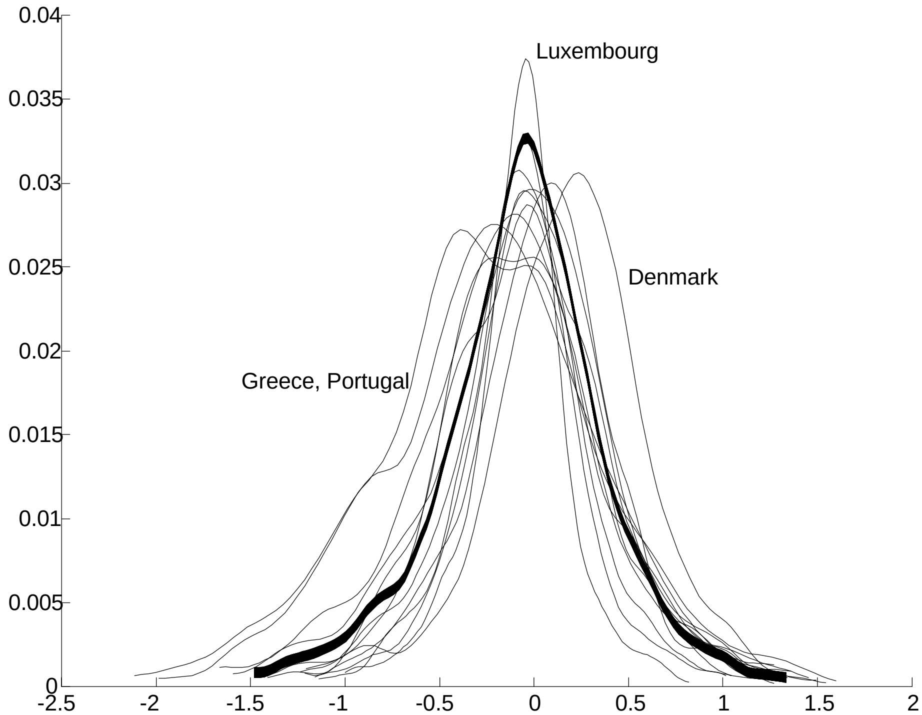

Figure 1: Empirical Distributions of Log Real Exchange Rates

Each thin line represents an estimate of the density of the log relative price, good-by-good, between a Belgium (the numeraire) and some other EC country. More specifically, each line is an estimate of the density, for country j, of $\log q_{ij} = \log(e_j p_{ij}/p_{in})$, where p_{ij} is the price of good i in country j, denominated in domestic currency and e_j is the spot exchange rate between country j and Belgium. The thick line represents the estimated density, pooled across all 12 countries.

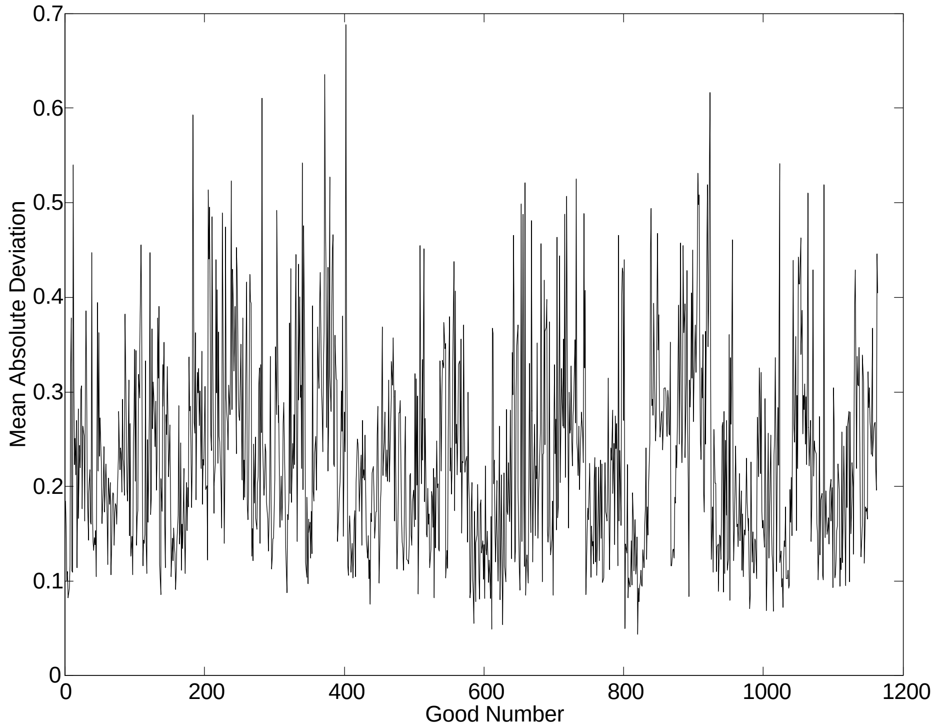

Figure 2: Price Dispersion Versus Eurostat Categories

Data points are the mean absolute deviation, across 13 countries, for each of our 1,164 goods. Prior to computing these statistics we convert common-currency prices into percent deviations from the good-specific mean and then remove the country effects reported in Table 4.

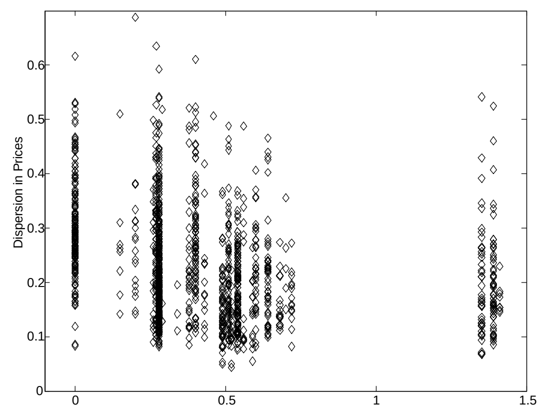

Figure 3: Good-by-Good Price Dispersion Versus Tradedness

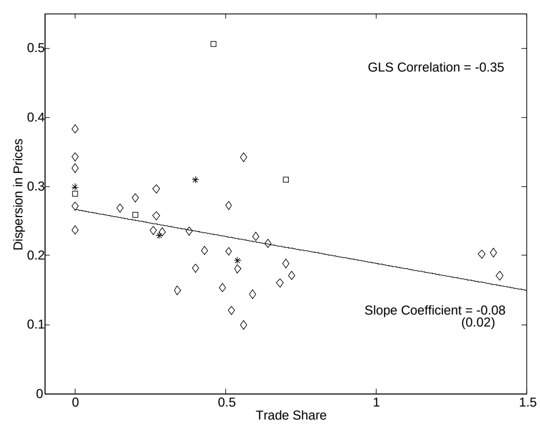

Figure 4: Average Price Dispersion Versus Tradedness

Each data point in the top panel represents the mean-absolute-deviation (across 13 countries) of the price of a particular good, plotted against the trade-share for the ISIC category to which that good belongs (see the text for data definitions). Each point in the bottom panel plots the *average* measure of price dispersion for goods within a particular ISIC category (*i.e.*, in both panels the tradeshare data are 'averaged,' whereas only in the lower panel is the price dispersion data averaged). The downward sloping line in the lower panel is the GLS-estimated regression function, which incorporates the heteroskedasticty associated with unequal numbers of goods in each aggregative group. In the lower panel, we mark the four data points to which our GLS estimator will give the *most* weight with asterisks, and the four points which receive the the *least* weight with squares.

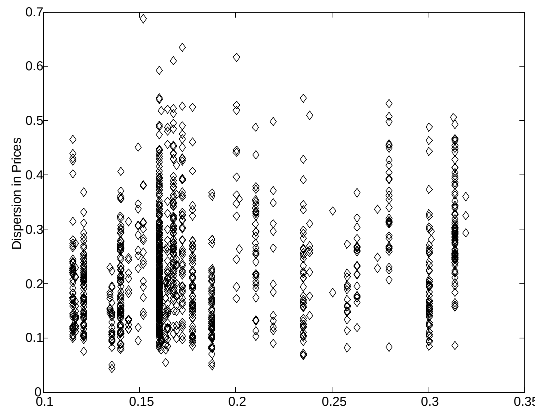

Figure 5: Good-by-Good Price Dispersion Versus Non-traded Input Share

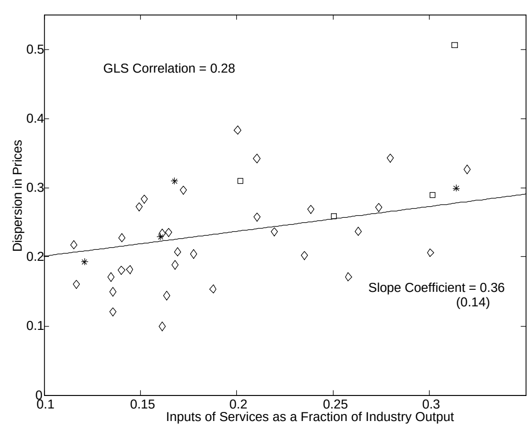

Figure 6: Average Price Dispersion Versus Non-traded Input Share

Each data point in the top panel represents the mean-absolute-deviation (across 13 countries) of the price of a particular good, plotted against a proxy for the fraction of non-tradeable inputs used by the industry to which that good belongs. The proxy for non-tradeable inputs is the ratio of services used as an input to total industrial output. Each point in the bottom panel plots the *average* measure of price dispersion for goods within a particular ISIC category (*i.e.*, in both panels the tradeshare data are 'averaged,' whereas only in the lower panel is the price dispersion data averaged). The upward sloping line in the lower panel is the GLS-estimated regression function, which incorporates the heteroskedasticty associated with unequal numbers of goods in each aggregative group. In the lower panel, we mark the four data points to which our GLS estimator will give the *most* weight with asterisks, and the four points which receive the the *least* weight with squares.