

Understanding International Price Differences Using Barcode Data*

Christian Broda (University of Chicago, GSB and NBER) and

David E. Weinstein (Columbia University and NBER)

November, 2007

Abstract

The empirical literature in international finance has produced three key results about international price deviations: borders give rise to flagrant violations of the law of one price, distance matters enormously for understanding these deviations, and convergence rates back to purchasing power parity are inconsistent with the evidence of micro studies on nominal price stickiness. The data underlying these results are mostly comprised of price indexes and price surveys of goods that may not be identical internationally. In this paper we revisit these three stylized facts using massive amounts of US and Canadian barcode data. We find that none of these three main stylized facts survive. We use our barcode level data to replicate prior work and explain what assumptions caused researchers to find different results from those we find in this paper. Overall, our work is supportive of simple pricing models where the degree of market segmentation across the border is similar to that within borders.

*The authors wish to thank the NSF (grant #0214378) and the Federal Reserve Bank of New York for providing financial support for this project. We would like to thank the Director of Research at the FRBNY, Joseph Tracy, for his early support of this project. We would also like to thank ACNielsen's vice-president of Pricing Research Frank Piotrowski, Ivan Rocabado, and Maura Elhbretch for their careful explanation of the data. Nonetheless, the views expressed here, do not reflect those of the Federal Reserve System, and all errors remain our own. Jesse Handbury provided us with outstanding research assistance. In addition, we would like to thank the Global Financial Markets Initiative at the University of Chicago GSB and the Center for Japanese Economy and Business at the Columbia Business School for research support.

I. Introduction

Data limitations have prevented researchers from comparing the prices of identical goods systematically within and across borders. This restriction has led researchers to infer the extent of market segmentation from the behavior of price indexes, aggregate prices of goods that may not be identical internationally, and a non-random selection of particular goods (e.g. Big Macs). This research has produced three key results about international price deviations: borders give rise to flagrant violations of the law of one price (LOP), distance matters enormously for understanding these deviations, and convergence rates back to purchasing power parity (PPP) are inconsistent with the evidence of micro studies on nominal price stickiness. In this paper we revisit these three stylized facts using massive amounts of US and Canadian barcode data. The results are shocking. None of these three main stylized facts survive. The law of one price in its absolute form holds as well across the border as it does within countries, distance coefficients are three times larger in aggregate data than in micro data, and rates of price convergence within and across borders are fast and completely in line with micro studies. In short, the data is supportive of simple pricing models where the degree of market segmentation across the border is similar to that within borders.

A major difference between our study and that of previous researchers is the data. We bring several databases to bear on the questions we examine. These datasets contain a vast number of products with barcodes in the US and Canada: covering approximately 40 percent of all expenditures on goods in consumption. Since the US and Canada share a common barcode classification system for a large set of consumer goods, we can compare exactly the identical good in each country. Moreover, we have this information within and across 10 cities in the US and 6 regions in Canada. Our data is also vastly richer at the micro level than that used in national statistics. For example, our data contains 700,000 price quotes for the US in a typical year. By contrast the CPI sample is only 5 percent as large. Moreover, unlike all prior work, we have both price and quantity data, which lets us form theory based, as opposed to *ad hoc*, indexes of PPP.

One important feature of the data is that it lets us compare the extent of international market segmentation with segmentation within countries and even within cities. We confirm the early finding by Isard (1977) – and more recently shown by Crucini et al. (2005) – that the LOP

is flagrantly violated in international data. However, we can show that that the LOP is *also* flagrantly violated across cities in the same country. Thus, the observation that an identical can of soda sells at different prices in different countries is not very informative about border barriers because prices vary substantially across space.

Obviously the more interesting question is how much larger are international violations than domestic ones. Here we find the answer to be – not much. In their seminal work, Engel and Rogers (1996) compare border barriers with regional ones by expressing the “width” of the border in terms of distance equivalents. Using barcode data and the same methodology as they do, we find the distance-equivalent border effect to be 3 miles – roughly what one might expect if trucks crossing to border had to stop briefly to fill out some paperwork. In other specifications the “width” of the border rises to a few hundred miles, but never anything close to the tens of thousands of miles found in the original paper and in subsequent work (e.g., Parsley and Wei (2001)).

Our second contribution is to explain why micro data reveals small border effects but aggregate data reveals much larger impacts. We begin by demonstrating that if we form price indexes using our barcode data and then replicate Engel and Rogers (1996), our results are quite similar to theirs. Clearly, something about aggregating micro data causes the border effect to appear larger. We argue that a vast amount of information about market segmentation across space is lost when one uses price indexes. In particular, because aggregate indexes collapse the large within-country idiosyncratic variation of relative goods prices while preserving the variation due to exchange rate movements, they make the cross-country variation appear much larger than the within country variation. Thus, aggregation of individual goods’ prices mechanically serves to amplify the measured impact of the borders on prices. In our data, this unintended consequence of aggregating individual prices into disaggregate product categories is entirely responsible for the large size of the border when using price index data.

We also find a tiny border effect when we look at deviations in the LOP. Here we compare international LOP deviations with those within the US after controlling for distance. Our finding is particularly surprising given that the impact of distance on the price deviations of identical goods is only about one tenth as high as that obtained using price index data. This underscores the role that compositional effects can have in explaining the relationship between price dispersion and distance previously found in the literature. We document that the set of

common goods across cities varies systematically across space and borders and therefore unless all individual prices within the index move together, price indexes will appear to deviate across space and borders simply due to the fact that the underlying weights and goods are different. We next document that the underlying prices within indexes vary enormously across time even for narrowly defined product categories, e.g. “fresh eggs.” This implies that the majority of the increased dispersion in aggregate prices that we observe as the distance between cities rises is not the result of actual deviations from the LOP but rather from compositional effects in the set of goods used to compute city-specific price indexes.

Finally, we turn our attention to understanding what Rogoff (1996) has termed “The PPP Puzzle”: the fact that international price adjustment occurs at much slower rates than what one would expect from micro data. Our first contribution is to find that when we use barcode data, prices converge rapidly to *absolute* PPP. This is in stark contrast to the wealth of literature that has often found slow convergence to *relative* PPP. Our rates of convergence are similar to those found in Parsley and Wei (1996) for convergence within the US and implied by micro studies such as Gopinath and Rigobon (2007). However, when we examine price convergence using price aggregates formed from our barcode data, the PPP puzzle reemerges: convergence is slow to non-existent.

Once again the question arises of why the aggregate results differ so much from those using micro data. We show that this dichotomy is not explained by the type of aggregation “bias” suggested by Imbs et al (2007). The answer in our data is the result of the interaction of two forces. First, we document that convergence rates are highly non-linear. Large relative price deviations disappear very rapidly but small ones are quite persistent. Second, since price indexes are formed by aggregating many prices together, information about relatively large idiosyncratic fluctuations is lost and regressions using aggregate data place more emphasis on the relatively slowly adjusting small price changes. This results in the much lower estimates of convergence using aggregate data.

In sum, we not only show that borders matter little for the LOP, that the role of distance is similarly overemphasized, and that prices of the goods in our sample converge in absolute terms to PPP, but we can also explain why the literature has failed to uncover these facts.

The structure of the paper is as follows. In Section II we provide a review of the theory and the empirical literature on international pricing. In Section III we describe the data and

preview some of the main results and in Section IV we examine the width of the border at the aggregate and micro level and explain the sources of the different results. In Section V we examine the issue of convergence rates to PPP within and across the border both at the aggregate and micro level. In Section VI we provide an explanation for the difference in convergence rates between different levels of aggregation.

II. Theory and Literature Review

The empirical literature on international pricing is vast, and it is useful to have an organizing framework for understanding the prior work. We find it useful to write down the simple prediction of the theory of the LOP in its “exact” form and then contrast these equations with their “approximate” forms, i.e. the equations that are estimated in the literature. The difference between both forms will be instructive in understanding where the problems in the existing tests of this theory lurk. Unfortunately, the empirical literature has not been consistent in its usage of terms like LOP and PPP, especially when narrow aggregates of products are compared. In order to avoid any confusion, we will use the terms LOP and PPP in the same way as in Rogoff (1996) – i.e. if the prices of two goods in different locations are compared, we will refer to that as a test of LOP, and if two price aggregates are compared, we will refer to that as a test of PPP. One of the drawbacks of this approach is that we will refer to some papers as tests of PPP even though the authors refer to their work as tests of LOP. This is regrettable, but because many of the results in this paper turn crucially on what exactly is being tested, we feel it necessary to be precise about our terminology.

The LOP, or what the literature now refers to as “Absolute LOP,” states that the price of an identical good should be the same across locations when denominated in a common currency. Formally, this suggests that P_{uct} (i.e. the price of good u in city or region c in time t) can be written as

$$(1) \quad P_{uct} = E_{cc't} P_{uc't}$$

where $P_{uc't}$ is the price of the good in a different region or country and $E_{cc't}$ is the exchange rate which equals unity if the two cities or regions are in the same country.

Tests of equation (1) have been extremely limited. Previous studies have found that commodities that are traded on organized exchanges, e.g. gold, tend not to have large deviations in the LOP. For the handful of goods that have also been studied, authors have typically found large deviations from the LOP. Examples include the work on Big Macs by Cumby (1996), IKEA sales by Haskel and Wolf (2001), and *The Economist* magazine by Ghosh and Wolf (1994).

A second class of studies has sought to test what might be called “Approximate Absolute LOP:”

$$(2) \quad P_{uct} = E_{cc't} P_{u'c't},$$

where typically goods u and u' belong to a similar product category but are not identical goods. Equation (2) differs from equation (1) in that one is not necessarily comparing the same goods, and hence one cannot distinguish violations in the LOP from violations of the assumption that good u and good u' enter into consumer utility identically. For example, very interesting recent work based on the Eurostat database (c.f., Crucini, Telmer, and Zachariadis (2005) and Crucini and Shintani (2006)) test this form of the LOP. However, it is difficult to know how much of an observed violation in the LOP is due to the fact that borders prevent arbitrage from eliminating price differentials for goods like “lady’s boots” and how much is due to the quality of the sample of lady’s boots varying across countries.

Concern over this unobserved heterogeneity has motivated researchers to examine “Relative LOP,” which we define as follows:

$$(3) \quad \Delta p_{uct} = \Delta e_{cc't} + \Delta p_{u'c't},$$

where lower case letters refer to natural logarithms of the upper case letters, and the Δ ’s refer to time differences. Tests of equation (3) relax the assumption that prices must converge to the same level (perhaps due to a constant trade barrier), and only test whether prices tend to remain a constant level apart.

The micro studies in the literature have typically worked with an equation that might be termed “Approximate Relative LOP:”

$$(4) \quad \Delta p_{uct} = \Delta e_{cc't} + \Delta p_{u'c't}.$$

The major advantage of using equation (4) relative to equation (3) is that it corrects for any unobserved heterogeneity that causes good u and good u' to enter into consumer utility differently. This is what motivated Parsley and Wei (1996) to use this form of the LOP in their

pioneering study of urban prices in the US. Differencing the data does not come without a cost. One can easily imagine that the heterogeneity between two different goods contains a constant component and a time varying component. To the extent that the time varying component is small, estimating equation (4) will be similar to estimating equation (3), but if different goods experience very different shocks across time, it is easy to see how equation (3) might hold closely but equation (4) might be violated severely.

Much of our theory only requires average prices to equilibrate; hence we turn our attention to PPP. In this paper, we derive Absolute PPP by weighting equation (1), summing and then taking logs to produce:

$$(5) \quad \ln \left(\sum_{u \in I_c} w_{uc} P_{uct} \right) = \ln E_{cc't} + \ln \left(\sum_{u \in I_c} w_{uc} P_{uc't} \right)$$

Alternatively, one can first take logs of equation (1) and then sum to produce

$$(6) \quad \sum_{u \in I_c} w_{uc} \ln(P_{uct}) = \ln E_{cc't} + \sum_{u \in I_c} w_{uc} \ln(P_{uc't})$$

There are two important features of equation (5) and (6). First, there is no intellectual content to equations (5) and (6) that is not captured in equation (1). If equations (5) and (6) hold but equation (1) does not, this simply is a statement that there is a weighting scheme that can cause the deviations in equation (1) to cancel. Second, assuming the Absolute LOP holds, Absolute PPP will hold only if one uses the same weights in both locations. Goldberg and Verboven (1995, 2005) and Lutz (2004) have examined variants of equation (6) in which the prices are aggregated together using hedonically adjusted price indexes. However, in these studies, the weights for the goods in the two locations vary, and thus the price indexes are not strictly comparable.

Given the data limitations to find price levels across countries, the literature has in general tended to focus more on Relative PPP. The theoretical version of relative PPP can be written down by first differencing equation (5):

$$(7) \quad \Delta \ln \left(\sum_{u \in I_c} w_{uc} P_{uct} \right) = \Delta \ln E_{cc't} + \Delta \ln \left(\sum_{u \in I_c} w_{uc} P_{uc't} \right)$$

However, all previous work on PPP has focused on what might be termed “Approximate Relative PPP:”

$$(8) \quad \Delta \ln \left(\sum_{u \in I_c} w_{uc} P_{uct} \right) = \Delta \ln E_{cc't} + \Delta \ln \left(\sum_{u \in I_c} w_{uc'} P_{uc't} \right)$$

Prominent studies include Isard (1977) Giovannini (1988), and Knetter (1989, 1993) on average import prices, and Engel (1993), Froot, Kim, and Rogoff (1995), and Rogers and Jenkins (1995) on price indexes.

There are three important differences between equation (8) and equation (7). First, equation (7) may hold but equation (8) will not if the price changes of goods u and u' are different because of idiosyncratic shocks. Second, equation (7) may hold but equation (8) may not if the log price changes of goods u and u' do not equal the simple price changes. Third, equation (8) may not hold because the weights and/or the set of goods on the left hand side do not equal those on the right. This last critique is particularly important because statistical agencies make no effort to insure that international or even urban price indexes use the same weights and/or goods.

Finally, Engel and Rogers (1996) pioneering work deserves special mention. Working around the limitations of existing price data they have instrumented a useful test based on “variance ratio” of price changes. In the simplest form, one can imagine taking the variance of equation (3) and seeing if the variance is larger when c and c' are in different countries relative to when they are in the same country. However, instead of taking the variance of equation (3), Engel and Rogers are forced to work with the variance of equation (8). In section IV we explore the unintended consequences of their tests of (3) based on the relative volatility of the terms in (8).

The foregoing analysis provides a simple roadmap for understanding the way this paper is structured. First we will examine the LOP and PPP in their absolute and relative “exact” forms using thousands of barcode products both within and across borders. Next, every time we find a difference between our results and those of other papers that have examined these relationships in their “approximate” forms we will investigate whether we can replicate the results and pinpoint to the assumption that gives rise to the failure or anomaly. This enables us to not only do precise testing but also understand the previous literature.

III. Data Description

III. A. Overview

A major difference between this paper and prior work is that we bring barcode data to bear on the question of international price differences. We use three datasets that are extracts of ACNielsen's Homescan database. The Homescan database is collected by ACNielsen in the United States and ACNielsen Canada in Canada. In each country Universal Product Code (UPC) scanners are given to a demographically representative sample of households. In the US, approximately 60,000 households in 23 cities receive these scanners and approximately 15,000 households in 6 regions receive them in Canada. Households then scan in every purchase they make. If the purchase is made from a store with ScanTrak technology, the prices of each good are downloaded directly from the store's database. If the good is purchased elsewhere, e.g. on the internet, the household directly enters the price. As such, the database provides us with a vast array of goods with barcodes. The majority of these goods are in the groceries, drugs, and mass merchandise sectors.

Because the full dataset is extremely expensive, we purchased three extracts that we will make use of in this study. The first one is the database that we will refer to as the "US National Database" and was used in Broda and Weinstein (2007). In this extract, we had ACNielsen collapse the city and household dimension of the database, and thus we have price and quantity data on every UPC purchased by the US sample of households for every quarter between 2001:Q1 and 2003:Q4 at the national level. This database contains information on approximately 700,000 goods each year.

The second database, we refer to as the "US Cross-Sectional Database," is new. In this database, we have household level data on every purchase in the fourth quarter of 2003 by a subsample of 3,000 households evenly divided across 10 US cities. In each city, the households were randomly selected from the full sample so that their demographic characteristics match those of the city as a whole.

Finally, the third database, which we shall call the "Canadian Regional Database", is also new. ACNielsen Canada provided us with average price and quantity data by region in Canada for every quarter between 2001:Q1 and 2004:Q4. Table 1 describes the basic statistics of each of these three different databases. As one can see from the table, our data provides a much richer breakdown of prices for this sample of goods than is available in national statistics.

These databases have three key features that lend themselves to the study of pricing in different markets. The first is that we identify different goods using barcodes. Since companies

only use one barcode per good, when we compare goods internationally, we can be confident that we are comparing precisely the same goods. As we will see shortly, previous studies that focused on product categories like eggs, butter, and cheese can be seriously misleading because there is enormous price variation within these categories. Examples of the level of detail in our database are given in Table A1 in the appendix. Second, we can also compare variation of prices across cities within and across borders. This lets us precisely examine the border effect in levels; something no one has done before. Third, because we have both price and quantity data, we know exactly how to weight the goods when building price indexes, which allows us to examine the role that compositional effects play in studies that use national statistic data.

III.B. Data Preview

Before plunging into the econometrics, it is useful to examine the raw data to obtain some intuition for how prices vary across regions and time. The first point that is important to contemplate is the vastness of barcode information that is included in our database. In the US National and Canadian Regional Database there are 700,000 and 490,000 UPCs available, respectively. Even within narrow product categories, consumers have access to an enormous number of different goods. We made use of the US National Database to examine how many UPCs were sold in each of the 123 “Product Groups” in the US. In the ACNielsen classification system, a product group is a highly disaggregated subset of the total database. For example, fresh eggs, ice, and milk are all different product groups. We plot a histogram of the count of the number of UPCs per product group in Figure 1. The first thing that is immediately apparent from the figure is the vast number of UPCs per product group. With the exception of a few product groups – yeast, meal starters, road salt, canning supplies, and contraceptives – all products in the US are comprised of over 200 different UPCs. The typical product group has 2700 different UPCs. Even relatively homogeneous goods like fresh eggs are comprised of 2275 different varieties.

The simple fact that there are many UPCs per product group would be an intellectual curiosity if it weren’t for the fact that the degree of sample overlap varies systematically with variables of interest. In Figure 2, we plot the share of UPCs that are common between cities in the US and regions in Canada and the distance between those two locations. For expositional purposes, the bilateral city data is shown in three different plots: comparisons within city pairs in

the US, within region pairs in Canada and between cities in the US and regions in Canada. The pattern observed in each of these plots is unmistakable: as distance between cities rise, the share of common identical goods between cities falls. Within the US, the share of common goods across cities is over 28 percent between New York and Philadelphia—the closest city pair in our data—and is less than 18 percent for goods between New York and Los Angeles—the two cities further apart. Within Canada, Ontario and Quebec share almost 60 percent of goods while British Columbia and Maritimes share less than 45 percent of the goods. In Table A2 in the appendix we show regressions of the share of common goods in terms of the simple count of the number of goods and in value terms against bilateral distance between cities.

The levels of the share of common goods within countries are not directly comparable between Figures 2A and 2B. This is because our data is based on different household sizes per city/region in the US and Canada, and because regions in Canada include several large cities. However, it is still surprising that despite the large sample of goods that are included in each city only around 25 percent of the UPCs are common between any two cities in the US. While this probably understates the true degree of overlap in the US because some UPCs might not have been purchased by the sample households included in our data but did exist in the city, it underscores the importance of compositional effects when comparing prices of similar “product categories” across cities within a country. Our sample of over 50,000 UPCs per city is around *40 times larger* than those used by the Bureau of Labor Statistics when computing regional price indexes.¹ This suggests that the amount of overlap in city or regional price indexes in national statistics data is quite small.

Figure 2C shows the importance of compositional effects across the border. A large number of the products sold in the US are not sold in Canada *in identical form*. In the typical bilateral city/region comparison between the US and Canada only 7.5 percent of the goods are common, this is less than one third the common set of goods available between city pairs of equal distance within the US (Figure 2A and 2C are directly comparable). This means that the composition of a random sample of goods sold in the US is likely to differ substantially with the composition of a sample of goods sold in Canada. By the same token, more proximate locations have more similar consumption bundles than distant locations.

¹ The BLS collects around 34,000 price quotes (for the same categories included in our database) over 23 different cities. This implies that they collect around 1,260 price quotes per city.

The fact that price indexes across regions or countries are largely composed of different goods would not be a problem for understanding the LOP or PPP if goods within categories are fairly homogenous. In this case, one could have a reasonable degree of confidence that similar goods would have similar prices or at least these prices would move together. The time series properties of disaggregated data have been examined extensively in Broda and Weinstein (2007) and Klenow and Kryvtsov (2007), so here we will just review a few key stylized facts uncovered in those papers. In Figure 3, we plot the kernel density of quarterly UPC relative price changes and quarterly UPC relative price changes after controlling for product group-time fixed effects. It is useful at this point to introduce some notation. Let $p_{ugc,t}$ be the log price of UPC u that belongs to product group g in city c and period t . We denote the relative log price of a UPC with respect to the largest Canadian province, Ontario, as $q_{ugc,t} = p_{ugc,t} - p_{ugOnt,t}$. The red line in Figure 3 shows the distribution of $\Delta q_{ugc,t}$ for all UPCs in all time periods in Canada. The typical quarterly UPC *absolute* price change is around 9 percent and the standard deviation is around 18 percent. This number is similar in magnitude to the typical price change of the quotes underlying the CPI calculation in the US.² It suggests that there is a large amount of price volatility over time within Canada and the US.

These numbers imply that there is vastly more volatility in the raw price data than in exchange rates. The typical quarterly exchange rate change among developed economies with flexible exchange rates is less than 2 percent (see Calvo and Reinhart (2004)). The large volatility of the raw price data relative to exchange rate data has an important implication for examining convergence. It implies that a large share of the fluctuations in the prices of individual goods across countries is likely to come from UPC specific shocks that are ignored at the aggregate level. As we will see in later sections, the distinction between idiosyncratic versus common price shocks will help us explain differences in the observed rates of convergence back to the LOP when we use disaggregate as opposed to aggregate data.

Figure 3 also includes direct information on how important are UPC idiosyncratic price shocks are to explain the volatility of prices over time. The blue line shows the kernel distribution of $\Delta \tilde{q}_{ugc,t}$ where $\tilde{q}_{ugc,t} = q_{ugc,t} - q_{gc,t}$ i.e. the log relative price of a particular UPC in a

² Klenow and Kryvtsov found that the median absolute price change for a price quotation used in the CPI was 13.3 percent *per month*. Similarly Broda and Weinstein found that the standard deviation of price changes of a UPC was 20 percent per quarter.

particular city and time once it has been purged of common city-time effects. As one can see from the plot, there is enormous dispersion of prices within product groups as both distributions lie almost on top of each other. In other words, common product-group and time factors play a tiny role in explaining the observed time-series volatility of UPC prices. The standard deviation of the UPC specific component of prices is close to 15 percent, which is almost identical to the standard deviation of the raw price changes. If we focus our attention on a relatively homogeneous good like fresh eggs, the standard deviation falls to 10 percent, but it is pretty clear that one cannot even treat a relatively homogeneous good like fresh eggs as a single item.

The preceding analysis suggests that even though goods may have identical prices, goods categories might exhibit very different average prices and price changes. Fortunately, the use of UPC data means that we can be incredibly precise about the prices that we are comparing. In Table 2, we compare the prices of individual UPCs across cities and regions in the fourth quarter of 2003. In the first panel, we focus on the US. Since we have data for 10 cities, we can make 45 bilateral comparisons of prices across city pairs. The middle and lower panels examine all bilateral comparisons between regions in Canada and between cities in the US and regions in Canada. As the first column indicates, we typically have 10,616 prices of common UPCs for every city pair in the US, 25,094 goods in the typical bilateral region comparison within Canada and 1,531 identical goods across countries. Columns 2-4 of Table 2 present medians, averages, and standard deviations of bilateral city comparisons for several sample statistics (in Table A3 in the appendix we present all city pair comparisons). In column 2, we first computed the median price differential for each city pair. In column 3, we compute the standard deviation of log relative prices of the same UPCs consumed in each city pair. Finally, in column 4, we computed the median absolute difference in the log prices for each UPC consumed in city pair.

The first interesting number presented in the table is the standard deviation of the median price differential in the city pairs. The standard deviation of 0.016 means that the typical price differential between cities in the US is very close to zero (upper panel). We repeat the same exercise for Canadian regions and obtain very similar results (middle panel). These results suggest that whatever price differentials exist within countries, they are distributed around zero. This is strong indication that absolute PPP holds in different locations within countries. This finding is present in our data in all quarters for which we have regional Canadian data. The average difference in prices of identical goods does seem to rise as we cross borders, but the rise

is quite modest (lower panel). The median price difference in the 4th quarter of 2003 for a given UPC in a US city relative to a Canadian region is only 1.9 percent higher on average.³ This result however, is not robust to the time period being studied, as large cumulative exchange rate movements over this period have made absolute PPP fluctuations vary from around 15 percent to 2 percent.

In columns 3 and 4 we present the standard deviation of the log relative prices and the median absolute price deviation. The table reveals that the typical standard deviation of log price differences between any two cities is 22.3 percent in the US and 18.7 percent in Canada. These numbers reveal something very important about the LOP: even within a country the standard deviation of prices of identical goods is typically 20 percent. To put this number in perspective, consider the results of Froot, Kim, and Rogoff's [1995] study of international violations of the law of one price. In that study, they concluded, "the volatility of law of one price deviations is both remarkably high (typically on the order of 20% or more per year for most commodities in most centuries) and remarkably stable over time." The important fact to bear in mind is that the LOP deviations that these authors found internationally are approximately the same magnitude as those we observe within countries. In other words, the prices of individual goods vary substantially across space regardless of whether two regions are in the same country or not.

This point notwithstanding, we can see that the dispersion of prices of individual goods vary slightly more when crossing the border. The lower panel of Table 2 shows that the standard deviation of prices of identical goods across the border is typically 26.7 percent, roughly 4 percentage points larger than within the US and 8 percentage points larger than within Canada. Results are similar using the typical absolute price difference between cities. However, we need to be cautious about interpreting this raw data because we need to adjust these numbers for the fact that cities within a country are likely to be closer together on average than cities in different countries.

One can also inspect the importance of the border visually in Figure 4. Here we plot the kernel densities of all relative prices across cities within the US, within Canada, and between the US and Canada. As the plot makes clear, prices in the US are a bit higher than prices in Canada, and there is evidence of greater dispersion in international prices than in domestic prices, but the

³ We adjust Canadian prices downwards by 7 percent because Canadian prices are inclusive of the VAT.

distributions are not radically different. Rather the border seems to add a small amount to the very large within-country dispersion in prices across cities. This creates some tension with the results of Engel and Rogers (1996), and is a point that we will need to explore more systematically.

In sum, the sample statistics reveal a number of important lessons for understanding international pricing. First, there are a vast number of goods in the market and the composition of consumption varies systematically with distance and across borders. This implies that one must take great care about how samples are constructed when comparing relative price movements across space and borders. Second, the prices of these goods vary substantially even for narrowly defined commodities. This implies that absolute deviations in the LOP will be quite sensitive to whether precisely the same goods are compared. Third, one should not equate the international violation of the law of one price with a barrier at the border. The data strongly suggests that there are substantial violations of the law of one price within countries and that these violations are of similar magnitudes as international violations. A corollary of this lesson is that one should not be surprised at the existence of LOP deviations – these happen all the time within countries – the more interesting question is how much larger international deviations are than the regional ones. Fourth, there is vastly more volatility in individual price quotes than in price indexes. This means that much of the price variation is eliminated when one focuses on price indexes. As we will see in the next few sections, each of these stylized facts will play a key role in understanding why absolute price convergence holds and why it has been so hard to find evidence in favor of it.

IV. The Width of the Border Redux

In order to understand the magnitude of international deviations of the LOP, we need to think of a benchmark. One of the simplest and most compelling reasons why prices may differ spatially is that it is difficult to transport goods. Thus, one might expect smaller LOP deviations in close cities than in distant cities. In their seminal work, Engel and Rogers [1996] developed this concept further by expressing border effects in terms of distance – a convention we will adopt here.

A simple way of computing the “width” of the border is to regress a measure of the price dispersion on the log of distance and a dummy variable that is one if the price difference is computed for a good purchased in cities that are located in different country. In this case one can compute the width of the border by dividing the border coefficient by the distance coefficient and then exponentiating. In Table 3, we present the results for a similar regression as that in Engel and Rogers. The only difference is that we use two different measures of price dispersion. First, we look at a price variance measure: the square of the log price difference of a UPC purchased in two different cities, i.e. $(r_{ugcc',t})^2 = (p_{ugc,t} - p_{ugc',t})^2$; second, we look at the absolute log price difference paid for the same UPC in two cities, i.e. $|r_{ugcc',t}| = |p_{ugc,t} - p_{ugc',t}|$.

Specifically, we run the following regression:

$$(9) \quad (r_{ugcc',t})^2 = \alpha_c + \beta \ln dist_{cc'} + \gamma Border_{cc'} + \varepsilon_{ugcc',t}$$

$$(10) \quad |r_{ugcc',t}| = \alpha_c + \beta \ln dist_{cc'} + \gamma Border_{cc'} + \varepsilon_{ugcc',t}$$

where α_c are city dummies, and standard errors are clustered by city pair. The “width” of the border adopted by Engel and Rogers is given by $\exp(\hat{\gamma} / \hat{\beta})$, where circumflexes indicate estimated parameters.

One of the problems of this approach is that this measure of the border is unitless, and hence one cannot strictly interpret it in terms of a mileage equivalent.⁴ However, the coefficient can be interpreted in terms of the *relative* distance between any two cities. This can most easily be understood in terms of an example. Suppose there are three cities: A , B , and C . Let $|r_{ugAC,t}|$ and $|r_{ugBC,t}|$ be the absolute relative price differences between city pairs AC and BC , respectively. If the distance between cities A and C is 2 log units more than the distance between cities B and C and the border coefficient is 2, then the impact of crossing the border would have the same impact on absolute relative price differences as comparing relative prices in city B with those in city A . Note that the critical factor driving the width of the border is not the distance from B to C but the relative differences from C of A and B . Thus, if there were two additional cities A' and B' that were twice as far from C as A and B respectively, the border effect would have an identical impact on relative prices as travelling from A' to B' . As this example makes clear, we should be

⁴ For instance, whether we measure distance in inches or in miles would imply the same γ and β coefficients.

speaking of the border in terms of relative distances from a given location. This point notwithstanding we will stick with convention for the purposes of comparability with prior research and sometimes express the “width” of the border in terms of miles. At times we will focus the magnitude of the border coefficient, which is a less colorful, but more meaningful measure of the border.

The results of this exercise are presented in Table 3. The first panel presents the raw regression results and the second panel presents results in which we weight the observations by the sales of the UPC.⁵ The weighted regression results are probably more reasonable because the forces of goods arbitrage are probably much greater for a good with a large amount of sales than for a good that is only purchased by a few people. In all regressions, distance contributes significantly to price dispersion and there is a positive and significant border effect. This is comforting because our priors strongly suggest that borders and distance interfere with the law of one price.

What is most interesting in the table, however, is our estimate of the impact of the border. If we look at the regressions with the absolute log price difference as the dependent variable, we see that the border introduces a price wedge of seven percent between the US and Canada. We can obtain some sense of how small this is compared to prior work by computing the “width” of the border. In the unweighted regressions, the width of the border ranges from 720 miles to 328 miles depending on the specification. By contrast the point estimate in Engel and Rogers was 75,000 miles for all goods and 3.8 million miles for food at home – the category closest to our sample of goods. Similarly Parsley and Wei (2001) estimate that the width of the border *vis à vis* Japan is 43 quadrillion miles. Of course, the unweighted estimates are likely to overstate the border for the reasons we have highlighted above. If we turn to the weighted regression results, we find that that width of the border ranges between 36 and 106 miles. In other words, Canada is not located midway between the Earth and the Moon – it’s really just a few miles north of Buffalo. We show in Figure A1 in the appendix that this result is robust to the quarter we use.

The fact that we find the border effect to be so small strikes us as both deeply comforting and confounding. On the one hand given Canada’s proximity to the US, the existence of a Free Trade Agreement, and the similarity of the economies suggests that we should expect small

⁵ We use the average value of consumption of each UPC between city pairs as a weight.

border effects. However, it is puzzling why we should find such a small border effect when so many other studies have not.

One possible explanation harks back to our earlier discussion of the heterogeneity of products within product categories. If categories like “fresh eggs” are very heterogeneous, then a basket of fresh eggs in one country is likely to contain very different eggs than a basket of eggs in another country. We have already seen that this compositional effect becomes more important with distance and when one crosses a border. We can now examine the importance of this effect in three stages. Our first task is to demonstrate that carefully aggregating the data does not affect the estimates of the border effect. In order to do this, we need to be precise about what goods and weights are used to compute city price indexes. We first define $I_{cc'}$ as the set of commonly consumed UPCs in city pair cc' . We first construct a common weighted index as a Geometric index of the relative prices of common goods within every product group in every bilateral city pair:

$$(11) \quad \text{Common Weight Index}_{gcc't} = \prod_{u \in I_{gcc'}} \left(\frac{P_{ugct}}{P_{ugc't}} \right)^{\frac{1}{2}(s_{ugc} + s_{ugc'})}$$

where s_{ugct} is the share of expenditure in product group g on UPC u in city c in time t . Note that the log of equation (11) can be expressed in terms of the actual log price difference of a UPC purchased in two different cities $\ln(\text{Common Weight Index}_{gcc't}) = \sum_{u \in I_{gcc'}} \frac{1}{2}(s_{ugc} + s_{ugc'}) r_{ugc't}$. The two key characteristics of this index is that it only uses prices for common UPC across cities and it only depends on the average share of consumption in the two cities and not on the city specific expenditure shares.

The second index we consider is the city-specific index:

$$(12) \quad \text{City-Specific Weight Index}_{gcc't} = \frac{\prod_{u \in I_{gcc'}} (P_{ugct})^{s_{ugc}}}{\prod_{u \in I_{gcc'}} (P_{ugc't})^{s_{ugc'}}$$

In contrast to the common-weight index, the city-specific index can vary with the market shares of expenditures in two locations even if the average expenditure level is the same. The distinction is important because it lets us examine whether simply allowing the weights of common goods to vary has an effect on the results. We would expect distance to have a different effect on this index if compositional effects are important.

Finally we form an all-goods price index defined below:

$$(13) \quad \text{All-Goods Index}_{gcc't} = \frac{\prod_{u \in I_{gc}} (P_{ugct})^{s_{ugc}}}{\prod_{u \in I_{gc'}} (P_{ugc't})^{s_{ugc'}}$$

where I_c is the set of UPCs available in city c . The major difference between this equation and equation (12) is that we now allow all goods in each city to enter the index, not just the common ones.

Our basic tests involve re-estimating the regressions in (9) - (10) using the log of the price indexes at the product group level instead of the log price differences of individual UPCs to see whether the simple act of aggregation creates a problem. As one can see in the first panel of Table 4, simply using common goods price indexes has almost no impact on our measure of the border. The estimated border effects do not move by much and the “width” of the border stays within 100 miles of the estimates that we obtained with the UPC-level data.

However, it is important to remember that the data used by researchers to examine border effects is not based on a common set of goods, but rather on non-overlapping samples of the goods available in each country. Panels 2 and 3 in Table 4 can help us understand the impact of using price indexes to assess the border effect. The second panel shows the impact that compositional effects through city-specific weights can have on the distance and border coefficients. The impact of distance on the square log price differences is over 5 times larger than in the first panel and the border dummy jumps by a factor of 3. The difference between the two panels can be traced directly to compositional effects. The prices of disaggregated goods categories may vary a lot even if the underlying prices hardly vary at all. The border dummy also rises to almost 3 times its value when common-weights are used. Since compositional effects tend to raise both the log distance and border coefficients the impact on the “width” of the border is not strongly affected by using city-specific weights.

The importance of distance and the border rises dramatically when we use an index composed of all goods. Now the distance and border coefficients rise by at least an order of magnitude. Interestingly, when goods that are specific to each city are included, the width of the border dummy jumps to literally astronomical values. Most of this is driven by an enormous jump in the border coefficient. The width of the border ranges from 23 million miles to 7.9 billion miles depending on the specification. The difference between this set of results and the

previous one arises solely from the fact that the composition of goods within a product group differs across the border sufficiently to affect the average price. This large border effect results in apparent rejections of the law of one price or PPP because the Canadians drink RC Cola and Americans drink Coca-Cola. While one may hope that RC Cola and Coca-Cola move together in the time series, there are many reasons to worry that this may not be the case. At the very least, one can see ample reasons why LOP might hold precisely, but the way in which aggregate indexes are formed produces failures of PPP. This establishes that it can be very misleading to estimate deviations from the law of price or PPP using even highly disaggregated product categories.

This explanation, however, is unsatisfactory to explain the results of Engel and Rogers (1996) because those results are based on the time-series volatility of price indexes as opposed to the dispersion of price *levels*. For instance, if prices in a product group all move together, it is possible for the levels to deviate across regions but the time series not to show a large border effect. In order to examine what role is played by aggregation in the results by Engel and Rogers we exploit the fact that we have time series data at the UPC level in the US National Sample and for each of six Canadian regions in the Canadian Regional Sample. Following Engel and Rogers, for each region pair, we compute the standard deviation of the relative log price changes of the goods common to that pair. In particular, we calculate $sd(\Delta r_{ugcc',t})$ where $\Delta r_{ugcc',t} = r_{ugcc',t} - r_{ugcc',t-1}$. This is the same statistic that Engel and Rogers use in their study but computed at the UPC level rather than at the product group level. We then regress these standard deviations on the log of distance between the regional pairs (counting the US as another region) and a border dummy, using the average distance between the Canadian region and our sample of cities as a proxy for the Canadian region's distance to the US. That is, we just use this time-series proxy for market segmentation as the dependant variable in regression (9).

The results from this exercise are presented in Table 5. At first glance, the results are quite similar to those of Engel and Rogers (1996) – we find that the standard deviation of relative inflation rates rises with distance and jumps discretely at the border. This result is suggestive of trade costs and border effects mattering for price arbitrage. However, what is most striking is the magnitude of the border. While Engel and Rogers found a border effect of 3.8 million miles for the “food at home” sector, we find a more modest border that is 3 miles wide. Thus the UPC

level data also suggests much smaller border effects even when we use the same proxy for market segmentation as Engel and Rogers.

But why do these results differ so much? Before we begin our investigation of the cause for the much smaller border effect, it is useful to first focus on why it is likely that disaggregated data would produce different results. The major difference between Engel and Roger's use of price indexes and our use of UPC level data is that price indexes are averages of individual price quotes. We have already seen in our analysis of the sample statistics that individual price movements exhibit enormous volatility in the time series but there is not much difference in the average price level across cities. Thus averaging the prices of UPCs together tends to eliminate much of the idiosyncratic variance of UPCs and leaves us with only the relatively small levels of variance of average prices across cities. Internationally, however, the impact of exchange rate fluctuations is not compressed by averaging because the impact is common to all UPCs in a country. This causes the border coefficient to fall less slowly than the distance coefficient. Since we divide by the distance coefficient when computing the border effect, *ceteris paribus*, this will tend to make the border appear wider.

We can see this formally by conducting the following exercise. Suppose that we can write the log relative price between foreign city c , and home city c' , for UPC, u , in product group, g , in time t as $r_{ugcc't}$. We then can decompose the change in the relative price as follows:

$$(14) \quad \Delta r_{ugcc't} = \delta_{cc't} + \delta_{et} + \varepsilon_{ugcc't},$$

where the δ 's correspond to city pair and exchange rate shocks, and $\varepsilon_{ugcc't}$ is the idiosyncratic shock to a UPC. Similarly if two cities are in the same country, we decompose the price movement using the same terms with the exception that $\delta_{et} = 0$. If we assume that all these terms are independent, then we can write

$$(15) \quad \text{Var}(\Delta r_{ugcc't}) = \sigma_{cc'}^2 + \sigma_e^2 + \sigma_\varepsilon^2$$

in the case when the cities are in different countries and

$$(16) \quad \text{Var}(\Delta r_{ugcc't}) = \sigma_{cc'}^2 + \sigma_\varepsilon^2$$

when the cities are in the same country. In this case the border effect (expressed in terms of ratios of variances instead of standard deviations) would be

$$(17) \quad \frac{\text{Var}\left(\Delta r_{ugcc't} \Big|_{\text{International}}\right)}{\text{Var}\left(\Delta r_{ugcc't} \Big|_{\text{Domestic}}\right)} = \frac{\sigma_{cc'}^2 + \sigma_e^2 + \sigma_\varepsilon^2}{\sigma_{cc'}^2 + \sigma_\varepsilon^2}.$$

However, if there are n UPCs in a product group and we first average the data before computing the variances, the ratio of the variances will be

$$(18) \quad \frac{\text{Var}\left(\Delta r_{ugcc't} \Big|_{\text{International}}\right)}{\text{Var}\left(\Delta r_{ugcc't} \Big|_{\text{Domestic}}\right)} = \frac{\sigma_{cc'}^2 + \sigma_e^2 + \frac{\sigma_\varepsilon^2}{n}}{\sigma_{cc'}^2 + \frac{\sigma_\varepsilon^2}{n}},$$

which is strictly larger than the expression given in equation (17) for $n > 1$. This suggests that if one computes border effects by comparing the variances of relative prices using price indexes, one will tend to find larger effects than if one uses the underlying micro-data. Moreover in datasets like ours, where the variance of the idiosyncratic shocks is likely to be large and the variance of bilateral city-pair shocks small, one would expect this effect to be substantial for large n .

In Table 6, we examine this aggregation bias by running the same regressions that we ran in Table 5, but first pooling the UPC level data to form product group averages and then computing the standard deviations in the movements of the product group level prices. We present two sets of results based on the two different ways of pooling the data given by equations (11) and (13). As one can see from the upper panel of this table, the width of the border estimated from regressions that use indexes based only on a common set of goods rises substantially (relative to Table 5). Averaging the data causes the width of the border to rise to 1000 to 100,000 miles depending on the specification.

Although these numbers are much larger, they are still smaller than the typical border effects of millions, if not quadrillions of miles, that often appear in studies. The lower panel of Table 6 shows the width of the border based on aggregate city-specific price indexes. A key distinction between these aggregate prices and those used in the upper panel is that each product-group index is an average of a much larger number of UPCs than in the upper panel. This is because the share of common goods across the border is less than 5 percent the size of the sample of goods in each region in Canada. As we noted earlier, this suggests that we might expect to see even larger border effects if we just formed indexes based on the set of UPCs consumed within a city in a particular product group. We verify that using indexes based on all

the UPCs within a city widens the border substantially. The border with Canada now rises to between 16 billion to 120 billion miles – still not in the quadrillion mile range but much further away than the 3 miles suggested by the UPC-level data.

Finally, the last two columns of the lower panel show the impact that the exchange rate fluctuations have on the measure of market segmentation based on aggregate price indexes. We replicate the results in columns 2 and 4 but drop the exchange rate from the relative price terms, i.e. we simply compute the US price index in US dollar terms and the Canadian price index in Canadian dollar terms. Not surprisingly we find no border effect in this case. This result underscores the importance of exchange rates shocks that are common across all UPCs when using aggregate data that collapses the UPC specific shocks.

V. Absolute Convergence Within and Across Countries

Having established that border effects are small, we now turn our attention to convergence. We have two objectives in this section. First, we want to estimate convergence rates using barcode data and second we want to explain why our results differ from those in other studies that use more aggregate data.

Our measure of relative prices is the log difference between the price of the UPC in a particular region and the price of the same UPC in Ontario. In modeling deviations of relative prices from their long-run levels we start by estimating the following regression:

$$(19) \quad q_{ugc,t} = \alpha_c + \beta q_{ugc,t-1} + \varepsilon_{ugc,t}$$

where α_c is a city-specific dummy and β denotes the speed of convergence. Under the null of no convergence, β is equal to one. In this case, a shock to $q_{ugc,t}$, i.e. $\varepsilon_{ugc,t}$, is permanent.

Convergence implies that β is less than one, with the approximate half-life of a shock to log prices given by $-\ln 2 / \ln \beta$.⁶ If β is less than one, the long-run level of relative prices is given by $\alpha_c / (1 - \beta)$. If $\alpha_c = 0$ and $\beta < 1$ then we can say that we observe absolute convergence in the data.

⁶ As Goldberg and Verboven (2005) note, this formula is only correct for AR1 processes. In general, the correct half life can be computed from the impulse response functions. However, we will follow the literature and drop the word “approximate” in future discussions of the half life.

This means that not only are shocks to relative prices transitory, but that eventually relative price differences between cities disappear. In the case where $\alpha_c \neq 0$ and $\beta < 1$ then we observe relative convergence in the data, i.e. shocks to $q_{ugc,t}$ are transitory but relative price differentials will persist.

The dummies α_c capture city fixed effects that account for non-time dependent price differences across cities (and countries). In addition to the speed of convergence, β , we are also interested in examining the absolute values of α_c . If these are zero or small (and β is less than one), then this would indicate that markets are not very segmented and that absolute price convergence is a good description of the data.

For estimation purposes we want to control for potential serial correlation in equation (19). For this reason, we augment equation (19) to include higher order auto-regressive terms as in Dickey-Fuller (1979):

$$(20) \quad q_{ugc,t} = \alpha_c + \beta q_{ugc,t-1} + \sum_{s=1}^S \gamma_s \Delta q_{ugc,s-1} + \varepsilon_{ugc,t}$$

where $\Delta q_{ugc,t-1} = q_{ugc,t-1} - q_{ugc,t-2}$ and S is the number of lags included in the regression. Since we are interested in studying the different convergence speeds of prices within and across countries, we allow for the convergence and autocorrelation terms to vary by country. Specifically, we estimate the following equation:

$$(21) \quad q_{ugc,t} = \alpha_c + \beta_w q_{ugc,t-1} + \beta_a q_{ugc,t-1} \times Border + \sum_{s=1}^4 \gamma_{ws} \Delta q_{ugc,s-1} \\ + \sum_{s=1}^4 \gamma_{as} \Delta q_{ugc,s-1} \times Border + \varepsilon_{ugc,t}$$

where *Border* is a dummy that takes the value of 1 when city c is not in Canada, β_w is the convergence parameter “within” countries, and $\beta_w + \beta_a$ is the convergence parameter “across” countries. In each of the specifications we run several tests: 1) whether $\alpha_c = 0$ within countries and $\alpha_c = 0$ for cities across countries; 2) whether $\beta_w = 1$, that is if there is a unit root within countries; 3) whether $\beta_w + \beta_a = 1$, that is if the data supports a unit root process across countries; and 4) whether $\beta_a = 0$, that is if the convergence rates within country are the same as across country.

Table 7, column 1, reports the results for equation (21) estimated on all the set of common UPCs between cities assuming a homogenous panel (i.e., $\alpha_c = 0 \forall c$). The coefficient

estimate for β is 0.789 with a standard error of 0.021. Since we have a limited time series dimension (12 – 16 quarter), it is inappropriate to employ conventional panel unit root tests that rely on large T asymptotics. Instead, we employ a unit root test for short panels developed by Harris and Tzavalis (1999). In the homogeneous panel case we can reject the unit-root test within and across borders at the 1 percent level. This suggests that prices revert back to their long-run level. In particular, the implied half-life for convergence is 2.9 quarters.

A notable feature of our data is that we can compare the rates of convergence back to PPP *across* as well as within countries. The second column allows for the β coefficient to vary within and across countries. In particular, we find that β_a , the difference in the autoregressive coefficient within and across the border to be around 0.06 and statistically significant. This suggests that while prices take longer to converge back to PPP when cities are across the border as opposed to within countries, the increase in the half-life of the shock is less than 2 quarters! Overall this implies a half-life for convergence of shocks across the border in this specification is 4.1 quarters.

Column 3 repeats the regression in column 2 but weights each UPC by how important they are in consumption in each pair of city.⁷ Half-lives for shocks within Canada double to 5.6 quarters as UPCs with large weights in consumption seem to have slightly slower convergence rates. The rate of convergence for UPCs across country also rises, but the difference between convergence rates between and within countries is less than 1 quarter. While we discuss the magnitude of city-specific effects below, when these are included in the regression convergence rates across countries rise to around 8 to 9 quarters, while within country convergence rates remain around 4 quarters. Overall, we find estimates for the rate at which PPP deviations diminish of between 3 to 4 quarters within borders and 4 to 9 quarters across borders. These numbers are broadly consistent with the micro price evidence on sticky prices. For example, Bills and Klenow (2007) find that half of domestic goods' prices last less than 4.3 months while the median duration in prices (including sales) in Nakamura and Steinsson (2007) is around 4.6 months. Gopinath and Rigobon (2007) find that price stickiness in US import prices can last up to 11 months.

⁷ The actual weight used is $w_{ugc} = 0.5 \times value_{ugc,t_0} + 0.5 \times value_{ugOm,t_0}$ where t_0 is 1999.

Given that we have price data on identical goods across cities within and across countries we can assess the economic magnitude of the deviations from absolute PPP within and across countries. As mentioned above, α_c / β defines the long-run level of $\ln q_{ugc,t}$. In columns (3) – (5) we compare how large are the absolute deviations from the PPP within and across countries. We find that the within Canada the average deviation from absolute PPP is statistically significant but small, between 0.9 percent and 1.5 percent. That is, any deviation in prices between regions in Canada and Ontario converge back to levels that imply that Ontario is around 2 percent cheaper than the average of the 5 other regions in Canada (column 6).

Furthermore, we find evidence that the absolute deviation between Canada and the US is approximately the same as that within regions in Canada. The absolute long-run values converge to levels that are between 0.6 percent to 3.3 percent more expensive in the US. Interestingly, deviations between Ontario and British Columbia and Alberta are equally as large as those between Ontario and the US. The hypothesis of absolute price convergence fails in a statistical sense within and across borders, but the magnitude of the failure is negligible from an economic standpoint. This is strong evidence in favor of a small role played by the border in terms of market segmentation.

VI. Aggregation and Non-Linear Convergence Rates

Once again the results should be seen as deeply comforting and confounding. On the one hand one should expect that in the absence of trade barriers and absolute price convergence should be a good description of the data when examining similar countries. However, no study has ever found this result before. Moreover, with the notable exception of Imbs et al (2005), studies have not investigated why the results are so dependent on the data. For example, Crucini and Shintani (2006) find faster half lives when using more disaggregated data than is typically found using aggregate data, but after rejecting the aggregation bias explanation of Imbs et al (2005), they do not offer an explanation reconciling these two findings. We now turn to trying to understand this puzzle

The evidence presented in Table 7 suggests that when individual product data is used, we find estimates for the rate at which PPP deviations diminish is between 3 to 6 quarters within

borders and 4 to 9 quarters across borders. However, we still have not addressed whether if we aggregate our data we would obtain slower rates of convergence that have plagued the existing literature. In particular, in the next two tables we not only assess whether half-lives estimated using aggregate data are large but also examine what are the reasons behind any difference between results at different level of aggregations.

Table 8 presents re-estimates equation (21) using product-group price indexes across cities instead of UPC price ratios. In particular, we run the following specification:

$$(22) \quad q_{gc,t} = \alpha_c + \beta_w q_{gc,t-1} + \beta_a q_{gc,t-1} \times Border + \sum_{s=1}^S \gamma_{ws} \Delta q_{gc,s-1} + \sum_{s=1}^S \gamma_{as} \Delta q_{gc,s-1} \times Border + \varepsilon_{gc,t}$$

where we define the price ratio at the product group level as

$$q_{gc,t} = \ln \left(\frac{P_{gc,t}}{P_{gOnt,t}} \right),$$

where $P_{gc,t}$ is a product-group price index. We will vary the method we use to compute this index to obtain a better understanding of how aggregation affects the data.

We consider two ways of computing these price indexes. First, we consider an index in which we allow all goods in each city to be averaged together.

$$(23) \quad \frac{P_{gc,t}}{P_{gOnt,t}} = \frac{\sum_{u \in I_{gc}} w_{ugc,0} P_{ugc,t}}{\sum_{u \in I_{gOnt}} w_{ugOnt,0} P_{ugOnt,t}}$$

where $w_{ugc,0}$ is the weight of UPC u in product group g in city c in 1999 and I_{gc} includes all the set of available UPCs in product group g in city c . Second, we build an index that aggregates only those goods that are common in Ontario and the region:

$$(24) \quad \frac{P_{gc,t}}{P_{gOnt,t}} = \frac{\sum_{u \in I_{gc-Ont}^{Com}} \tilde{w}_{ugc} P_{ugc,t}}{\sum_{u \in I_{gc-Ont}^{Com}} \tilde{w}_{ugc} P_{ugOnt,t}}$$

where $w_{ugc} \equiv \frac{1}{2} w_{ugc,0} + \frac{1}{2} w_{ugOnt,0}$.

Table 8 shows the convergence results under the these different aggregation schemes. For simplicity we will focus our discussion on columns 8 and 12 but results are similar using the comparison between other columns. Aggregation of the micro data produces significantly higher half lives. If we aggregate the data using only common goods, the convergence coefficient rises

from 0.85 (Table 7 column 6) to 0.95 (Table 8 Column 8). Despite the increase in half-lives, the rich panel nature of our data allows us to reject the presence of a unit-root in all cases (we also use the Harris and Tzavalis (1999) distributions). The implied half life of price shocks rises from 4 to 13 quarters within Canada and from 9 to 13 quarters across the border. If we form the index using all goods within the product group instead of just the common ones, the half lives jump to 138 quarters within Canada and 346 across the border! These half lives are essentially infinite considering that we no longer can reject a unit root.

These results suggest that whatever causes the discrepancy between aggregate results and those of micro data studies is present in our data. However, we can immediately rule out one source of this bias. Since we were consistent in the construction of the price indexes, we know that the difference between aggregate results and those obtained with the UPC-level variation is not due to differences in how aggregate indexes are constructed internationally. Compositional effects may explain why indexes comprised of disjoint samples of goods exhibit unit roots (e.g. the difference between panel 3 and panel 2), but they do not explain why the results in panel 2 differ from those in Table 7.

A second hypothesis for what might be driving aggregation bias has been suggested by Imbs et al (2005). However, it is difficult to see how it could apply here. Their explanation relies on the convergence coefficients varying systematically with the goods. Hence the aggregation bias can be solved by estimating different β 's for different goods. The bias we have identified arises solely from aggregating the data and is present even though we estimate one β in the UPC-level regressions and one β in the aggregate regressions.

Nonetheless, we can examine the importance of this form of aggregation bias in our data. In particular, the type of aggregation bias they study can be briefly explained using a general version of equation (19):

$$(25) \quad q_{ugc,t} = \alpha_c + \beta_u q_{ugc,t-1} + \varepsilon_{ugc,t}$$

where the main difference with (19) is related to the fact that the persistence coefficient is allowed to vary with each UPC u , i.e. β_u .⁸ For simplicity and without loss of generality we can define $\beta_u = \beta + \delta_u$, where $E(\delta_u) = 0$. In the case where the true model is that given by (25) but

⁸ Equation (25) is identical to the benchmark regression used in Imbs et al (2006).

instead equation (19) is estimated, then Imbs et al. (2006) argue that the estimated $\hat{\beta}$ from (19) has a bias. In particular, under certain conditions $E(\hat{\beta}) = \beta + \chi$ where $\chi > 0$.

Our data is particularly well suited to test for the existence of this type of aggregation bias. Since we are constrained by the time-series dimension of our panel, we start by assuming that the heterogeneity in persistence coefficients exists not at the UPC level but at the product group level. Specifically, this means we assume that $\beta_u = \beta_g \forall u \in g$ and $\beta_g = \beta + \delta_g$. That is, two UPCs in the same product group g share the same persistence coefficient, but UPCs in different product groups have different β .

Table 9 presents the results for the mean $\hat{\beta}_g$ and half-lives across all product group levels from (25) separately estimated for each product group, and the single $\hat{\beta}$ estimated from equation (19) using exactly the same set of UPCs. The first two columns of this table report the results for both weighted and un-weighted regressions. The average persistence coefficient across all product groups is 0.79, while the single persistent coefficient estimated from (19) is 0.81. This implies that allowing for heterogeneous coefficients at the product group level implies an average half-life for UPCs of around 2.9 quarters, while assuming homogenous coefficients across all product groups, implies the single estimated half-life to be 3.3 quarters. When we extend equations (25) and (19) to allow for both within and across coefficients, we find that the mean group estimate of the half-life across borders is 4.6 quarters while at the aggregate level the half-life is 6.6 quarters. Despite the small positive “aggregation bias”, standard tests reject the equality of the mean group coefficient with the homogenous coefficient. Hence, like Crucini and Shintani (2006), who also looked at the importance of this form of aggregation bias, we conclude, “our estimate of aggregation bias is modest, as anticipated by the simulation procedures used by Chen and Engel (2005).”

If the aggregation bias does not resolve the PPP puzzle in our data, how can we explain the large differences in estimated persistence at different levels of aggregation? An attractive explanation is the presence of strong nonlinearities in response to shocks. Nonlinearities have been found in other studies, e.g. Parsley and Wei (1996), but no one has conjectured that these nonlinearities can explain the difference between national and international convergence results.

We do so by first verifying the importance of nonlinearities in our data. We can estimate their importance by running the following regression

$$(25) \quad q_{ugc,t} = \alpha_c + \beta_1 q_{ugc,t-1} + \beta_2 q_{ugc,t-1} |q_{ugc,t-1}| + \sum_{s=1}^S \gamma_{ws} \Delta q_{ugc,s-1} + \varepsilon_{gc,t} ,$$

which is identical to the regression run in Parsley and Wei (1996). Here, if β_2 is less than zero, then this implies that larger shocks dissipate faster than small ones. In this case, the decay rate can be written as $\beta_1 - 1 + 2\beta_2|q_{ugc,t-1}|$ and the approximate half life is $-\ln(2)/\ln(\beta_1 + 2\beta_2|q_{ugc,t-1}|)$. We run this regression separately for within Canada relative prices and relative prices across the border. Table 10 reports the results of these regressions. β_2 is less than zero regardless of the data set used, indicating that there is strong evidence in favor of non-linear price adjustment both within countries and across the border. The importance of these nonlinearities can be assessed by considering a number of different-sized shocks. The standard deviation of $q_{ugc,t-1}$ is 26 percent which implies an approximate half life of two quarters within Canada and about 6 quarters across the border. This implies that the typical price shock dissipates quite rapidly. Smaller price shocks are obviously more persistent.

In particular, quarterly exchange rate shocks over this period had a standard deviation of 4 percent. Small shocks like this are estimated to have a half life of 4 quarters in Canada but 23 quarters (5.75 years) across the border. In other words, the UPC-level data is completely consistent with the observation that prices respond very slowly to exchange rate shocks. The nonlinearity of price adjustment implies that these shocks should be very persistent.

As attractive as this explanation seems, it still does not resolve why the convergence coefficient rises as we move to aggregate indexes. In order to see this result more clearly, it is useful to draw a picture that summarizes our findings. In particular the non-linearity implies that the persistence of relative price deviations will drop off as the absolute magnitude of the deviation rises. We portray this in Figure 5. The implication of this non-linearity is that the slope of a regression of current relative prices on past relative prices will depend on the amount of dispersion observed in the relative prices of the past period. The convergence coefficient, β , will be strongly influenced by this non-linearity because the OLS estimates will place a heavy emphasis on the observations where $|q_{ugc,t-1}|$ was large. However, if we aggregate the data, these large positive and price deviations are likely to cancel and hence the relative weight given to goods with small price deviations will rise. To the extent that these goods converge at a slower

rates, this means that the use of aggregated data such as a price index will produce estimates of the convergence coefficient that are larger than those produced using disaggregated data.

In Table 11 we present an illustrative example of this effect. In each time period there are two goods whose log relative prices converge from 30 to 21 to 15 and -30 to -21 to 0.15 (i.e. have a convergence coefficient of 0.7) and one good whose relative log price deviation ranges between 3 and -3 and has a convergence coefficient of 0.9. In this case if we estimated the convergence coefficient using the disaggregated data we would estimate a coefficient close to 0.7, but one would obtain an estimate of 0.9 if one first aggregated the data. In other words, the non-linearity of convergence rates would mean that the estimated half life in the rise from 2 quarters in the disaggregated data to 6.5 quarters in the aggregated data.

One way to see that this is what is driving the aggregation bias is to form our aggregates in such a way that we preserve much of the underlying volatility of the UPC level data to see if this causes us to recover our disaggregated estimates. In order to this, we form our product group price indexes according to equation (24) but only use 5 UPCs chosen at random to form the product group level price indexes. By using a small number of UPCs we allow the product group prices to be affected by large outliers. As one can see from the estimates, the rate of convergence in this table hardly differ from those of Table 7. The contrast with panel 2 of Table 8 is striking, however. Increasing the number of UPCs in the aggregate price index drives up the convergence coefficient significantly because the individual relative price deviations cancel out in the larger sample. As a result, the estimated convergence coefficient rises in all specifications and the corresponding half-lives rise as well.

VII. Conclusion

The use of barcode data reveals a very different picture of international price differences than what one sees with coarser data. In particular we find that the LOP and PPP hold in their absolute forms as well across the border as they do within countries. Moreover, the importance of distance for price differences is five to ten times larger in aggregate data than in barcode data. Much of this is driven by the fact that the set of common goods falls systematically with distance

leading price indexes to diverge because their composition diverges. Finally, we find that rates of price convergence within and across borders are fast and completely in line with micro studies.

Our study also explains why prior work has failed to identify these facts. In particular, our examination of the barcode data reveals that there is enormous heterogeneity in individual goods even when one examines goods that strike most researchers as homogeneous. This heterogeneity gives rise to very large relative price deviations that are lost when one examines price aggregates. The combination of using price indexes that are comprised of non-comparable goods and averaging away the large idiosyncratic price changes gives rise to a biased picture of how fast price differentials dissipate.

Finally, we are able to demonstrate that there are important non-linearities in the rate of price convergence. Large price differences disappear rapidly, but small ones persist. The non-linearity of price adjustment is present both when we look within and across borders. These results imply that the relatively small price differences generated by the typical exchange rate movement will tend to be quite persistent, but larger ones will be short lived. This may have important implications for understanding why prices sometimes seem to respond to exchange rate changes but other times do not. Obviously more work is needed to understand the implications of this for local versus producer currency pricing.

References

- Bils, Mark and Peter Klenow, "Some Evidence on the Importance of Sticky Prices," *Journal of Political Economy*, CXII (2004), 947–985.
- Broda, Christian and David Weinstein, "Product Creation and Destruction: Evidence and Price Implications," NBER Working Papers No.13041, 2007.
- Calvo, Guillermo and Carmen Reinhart, "Fear Of Floating," *The Quarterly Journal of Economics*, CXVII (II) (2002), 379-408.
- Chen, Shiu-Sheng and Charles Engel, "Does 'Aggregation Bias' Explain the PPP Puzzle?" *Pacific Economic Review*, X (2005), 49-72.
- Crucini, Mario, Telmer, Chris and Marios Zachariadis, "Understanding European Real Exchange Rates," *American Economic Review*, VC (III) (2005), 724-738.
- Crucini, Mario and Mototsugu Shintani, "Persistence in Law-Of-One-Price Deviations: Evidence from Micro-Data," Working Papers 0616, Department of Economics, Vanderbilt University, 2006.
- Cumby, Robert E., "Forecasting Exchange Rates and Relative Prices with the Hamburger Standard: Is What You Want What You Get with McParity?" NBER Working Papers: No. 5675, 1996.
- Dickey, D.A. and W.A. Fuller, "Distribution of the Estimators for Autoregressive Time Series with a Unit Root," *Journal of the American Statistical Association*, LXXIV (1979), 427–431.
- Engel, Charles, "Real Exchange Rates and Relative Prices: An Empirical Investigation," *Journal of Monetary Economics*, XXXII (I) (1993), 35–50.
- Engel, Charles and John H. Rogers, "How Wide Is the Border?" *American Economic Review*, LXXXVI (V) (1996), 1112–25.
- Froot, Kenneth A., Michael Kim, and Kenneth Rogoff, "The Law-of-One-Price over 700 Years," NBER Working Papers: No.1532, 1995.
- Ghosh, Atish and Holger Wolf, "Pricing in International Markets: Lessons from the Economist," NBER Working Papers: No. 4806, 1994.
- Giovannini, Alberto, "Exchange Rates and Traded Goods Prices." *Journal of International Economics*, XXIV (I/II) (1988), 45–68.

- Goldberg, Pinelopi and Frank Verboven, "Market integration and convergence to the Law of One Price: evidence from the European Car Market," *Journal of International Economics*, LXV (2005), 49-73.
- Gopinath, Gita and Roberto Rigobon, "Sticky Borders," NBER Working Paper No. 12095, 2007.
- Harris, Richard and Elias Tzavalis, "Inference for unit roots in dynamic panels where the time dimension is fixed," *Journal of Econometrics*, XCI (1999), 201-226.
- Haskel, Jonathan and Holger Wolf, "The Law of One Price—A Case Study," *Scandinavian Journal of Economics*, 2001, CIII (IV), 545–58.
- Imbs, Jean, Haroon Mumtaz, Morten O.Ravn and H el ene Rey, "Aggregation Bias Does Explain the PPP Puzzle," NBER Working Papers 11607, 2005.
- Imbs, Jean, Haroon Mumtaz, Morten O.Ravn and H el ene Rey, "PPP Strikes Back: Aggregation and the Real Exchange Rate," *The Quarterly Journal of Economics*, CXX (I) (2005), 1-43.
- Isard, Peter, "How Far Can We Push the Law of One Price?" *American Economic Review*, LXVI (V) (1977), 942-48.
- Klenow, Peter J. and Oleksiy Kryvtsov, "State-dependent or Time-dependent Pricing: Does it matter for recent U.S. inflation?", working paper, 2007.
- Knetter, Michael M., "Price Discrimination by U.S. and German Exporters." *American Economic Review*, LXIX (I) (1989), 198–210.
- Knetter, Michael M., "International Comparisons of Price-to-Market Behavior." *American Economic Review*, LXXXIII (III) (1993), 473–86.
- Lutz, Matthias, "Pricing in Segmented Markets, Arbitrage Barriers, and the Law of One Price: Evidence from the European Car Market," *Review of International Economics*, XII (III) (2004), 456–75.
- Nakamura, Emi and Jon Steinsson, "'Five Facts about Prices: A Reevaluation of Menu Cost Models", manuscript, Harvard University, 2007.
- Parsley, David C. and Shang-Jin Wei, "Convergence to the law of one price without trade barriers or currency fluctuations," *Quarterly Journal of Economics*, CXI (IV) (1996), 1211– 1236.
- Parsley, David C. and Shang-Jin Wei, "Explaining the Border Effect: The Role of Exchange Rate Variability, Shipping Costs, and Geography." *Journal of International Economics*, LV (I) (2001), 87-105.

Rogers, John H. and Michael Jenkins, "Haircuts or Hysteresis? Sources of Movements in Real Exchange Rates" *Journal of International Economics*, XXXVIII (III-IV) (1995), 339–60.

Rogoff, Kenneth. "The Purchasing Power Parity Puzzle," *Journal of Economic Literature*, XXXIV (II) (1996), 647–68.

Figure 1




Figure 2

Share of Common UPCs, Distance and Border

Figure 2A: Within Cities in the US




Figure 2C: Between Cities in US and Regions in Canada




Figure 2B: Within Regions in Canada




Figure 3




Figure 4




Figure 5




Table 1: Descriptive Statistics

	US National	US Cross-Section	Canada-Regional
Number of Cities/Regions	1	10	6
Number of Households per City/Region	55,000	300	2500
Time Period	1999Q1 - 2003Q4	2003Q4	2001Q1-2004Q4
Number of UPCs per City/Region	697,312	50,628	57784
Number of Product Groups per City/Region	123	118	156
Number of UPCs per Product Group per City/Region	5,669	429	370
Number of CPI Individual Quotes per ELI per City/Region	-	10	30

Cities included in the US: Boston, Chicago, Houston, Los Angeles, New York, Atlanta, Detroit, Philadelphia, Buffalo-Rochester, and Phoenix.

Regions included in Canada: Alberta, British Columbia, Manitoba, Maritimes, Ontario and Quebec.

Source: ACNielsen Homescan US and ACNielsen Homescan Canada.

Table 2: Law of One Price Deviations within City/Region Pairs in the U.S. and Canada

	Number of Common UPCs	Price Differences across Cities Common UPCs ONLY		
		Median	Standard Deviation	Median Absolute
	(1)	(2)	(3)	(4)
Upper Panel: U.S. - U.S.				
All 45 US city comparisons:				
Median	10,616	0.000	0.223	0.113
Average	10,730	-0.001	0.224	0.114
St. Deviation	1,303	0.016	0.012	0.013
Middle Panel: Canada - Canada				
All 15 Canadian region comparisons:				
Median	25,094	0.003	0.187	0.085
Average	25,980	0.007	0.181	0.083
St. Deviation	4,682	0.010	0.015	0.016
Lower Panel: U.S. - Canada				
All 60 U.S. City-Canada region comparisons:				
Median	1,531	0.021	0.267	0.161
Average	1,634	0.019	0.266	0.160
St. Deviation	328	0.020	0.008	0.008

Table 3: Deviations in the Prices of UPCs

Upper Panel: Unweighted Regression						
Dependent Variable	Square of log Price Difference			Absolute Log Price Difference		
	All	US-US	Can-Can	All	US-US	Can-Can
Data						
Log Distance	0.0047 [0.0006]**	0.0028 [0.0007]**	0.0083 [0.0006]**	0.0120 [0.0015]**	0.0068 [0.0016]**	0.0213 [0.00165]**
Border Dummy	0.0312 [0.0009]**			0.0694 [0.0020]**		
Observations	970338	482869	389701	970338	482869	389701
R-squared	0.02	0.00	0.01	0.03	0.00	0.01
"Width" of the Border	720			328		

Lower Panel: Weighted Regression						
Dependent Variable	Square of log Price Difference			Absolute Log Price Difference		
	All	US-US	Can-Can	All	US-US	Can-Can
Data						
Log Distance	0.0062 [0.0006]**	0.0024 [0.0008]**	0.0084 [0.0007]**	0.0182 [0.0019]**	0.0058 [0.0019]**	0.0245 [0.0023]**
Border Dummy	0.0290 [0.0013]**			0.0654 [0.0027]**		
Observations	970338	482869	389701	970338	482869	389701
R-squared	0.04	0.00	0.01	0.05	0.00	0.03
"Width" of the Border	106			36		

All Regressions include city dummies. Robust standard errors in brackets. All standard errors are clustered by city pair. ; * significant at 5% level; ** significant at 1% level.

Table 4: Border Effects for Product Group Level Price Indexes

Upper Panel: Common Weighted Index of Common Goods						
Dependent Variable	Square of log Price Difference			Absolute Log Price Difference		
	All	US-US	Can-Can	All	US-US	Can-Can
Data						
Log Distance	0.003 [0.001]**	0 [0.000]	0.004 [0.000]**	0.02 [0.003]**	0.004 [0.002]*	0.03 [0.002]**
Border Dummy	0.018 [0.002]**			0.064 [0.005]**		
Constant	-0.019 [0.007]*	0.002 [0.002]	-0.023 [0.003]**	-0.119 [0.033]**	0.028 [0.013]*	-0.148 [0.017]**
Observations	12471	5211	2333	12471	5211	2333
R-squared	0.12	0.02	0.08	0.15	0.04	0.14
"Width" of the Border	403			25		

Middle Panel: City-Specific Weighted Index of Common Goods						
Dependent Variable	Square of log Price Difference			Absolute Log Price Difference		
	All	US-US	Can-Can	All	US-US	Can-Can
Data						
Log Distance	0.016 [0.002]**	0.004 [0.001]**	0.024 [0.003]**	0.048 [0.006]**	0.011 [0.004]**	0.07 [0.006]**
Border Dummy	0.063 [0.005]**			0.109 [0.010]**		
Constant	-0.086 [0.017]**	0.017 [0.013]	-0.141 [0.019]**	-0.232 [0.051]**	0.094 [0.039]*	-0.365 [0.046]**
Observations	12471	5211	2333	12471	5211	2333
R-squared	0.06	0.02	0.1	0.12	0.06	0.17
"Width" of the Border	51			10		

Lower Panel: City-Specific Weighted Index Composed of All Goods						
Dependent Variable	Square of log Price Difference			Absolute Log Price Difference		
	All	US-US	Can-Can	All	US-US	Can-Can
Data						
Log Distance	0.162 [0.071]*	0.013 [0.004]**	0.15 [0.032]**	0.103 [0.024]**	0.023 [0.011]*	0.173 [0.034]**
Border Dummy	3.693 [0.100]**			1.746 [0.029]**		
Constant	-0.973 [0.453]*	0.016 [0.032]	-1.188 [0.289]**	-0.387 [0.156]*	0.16 [0.089]	-1.207 [0.308]**
Observations	12471	5211	2333	12471	5211	2333
R-squared	0.34	0.02	0.19	0.54	0.09	0.25
"Width" of the Border	7.95E+09			2.30E+07		

Robust standard errors in brackets; * significant at 5% level; ** significant at 1% level.

Table 5: Engel and Rogers at the UPC level

Dependent Variable	All UPCs			
	St. Deviation over time of the Log of Price Ratio between Cities			
	Within Canada	All	Within Canada	All
Weighted	No	No	Yes	Yes
Log Distance	0.01 [0.0010]**	0.009 [0.0008]**	0.014 [0.0018]**	0.012 [0.0013]**
Border Dummy		0.012 [0.0048]**		0.012 [0.0039]**
Observations	99444	116744	99444	116744
R-squared	0.01	0.01	0.01	0.01
"Width" of the Border		3.8		2.7

Robust standard errors in brackets; * significant at 5% level; ** significant at 1% level.

Table 6: Engel and Rogers at the Product Group level using only Common UPCs across Cities

Dependent Variable	All Product Groups - Common UPCs - Common Weights			
	St. Deviation over time of the Log of Price Ratio between Cities			
	Within Canada	All	Within Canada	All
Data	No	No	Yes	Yes
Weighted	No	No	Yes	Yes
Log Distance	0.01 [0.0017]**	0.004 [0.0019]*	0.003 [0.0007]**	0.002 [0.0009]*
Border Dummy		0.046 [0.0036]**		0.014 [0.0027]**
Observations	1213	4336	1213	4336
R-squared	0.01	0.07	0.01	0.02
"Width" of the Border		98716		1097

Dependent Variable	All Product Groups - All UPCs - City-Specific Weights					
	St. Deviation over time of the Log of Price Ratio between Cities					
	US \$	US \$	US \$	US \$	Canadian \$	Canadian \$
Data	Within Canada	All	Within Canada	All	All	All
Weighted	No	No	Yes	Yes	No	Yes
Log Distance	0.007 [0.0007]**	0.002 [0.0006]**	0.007 [0.0007]**	0.002 [0.0006]**	0.002 [0.0006]**	0.002 [0.0006]**
Border Dummy		0.051 [0.0000]**		0.047 [0.0000]**	-0.005 [0.005]	-0.003 [0.003]
Observations	1268	11941	1268	11941	11941	11941
R-squared	0.01	0.27	0.01	0.27	0.18	0.18
"Width" of the Border		1.19E+11		1.61E+10	.	.

Robust standard errors in brackets; * significant at 5% level; ** significant at 1% level.

Table 7: Convergence Rates at the UPC Level

Dependent Variable	Log of UPC Price Ratio Relative to Ontario					
	No No (1)	No No (2)	No Yes (3)	Yes No (4)	Yes No (5)	Yes Yes (6)
City Dummies						
Value Weights						
$q_{ugc,t-1}$	0.787 [0.002]	0.779 [0.002]	0.866 [0.005]	0.772 [0.020]	0.762 [0.002]	0.853 [0.006]
$q_{ugc,t-1}$ * Border		0.079 [0.006]	0.05 [0.020]		0.156 [0.008]	0.071 [0.024]
Dummy ALB				-0.013 [0.001]	-0.014 [0.000]	-0.020 [0.001]
Dummy BRC				-0.017 [0.000]	-0.018 [0.001]	-0.020 [0.002]
Dummy MAN				-0.009 [0.000]	-0.01 [0.001]	-0.014 [0.002]
Dummy MAR				0.004 [0.001]	0.003 [0.001]	-0.006 [0.002]
Dummy QUE				0.003 [0.001]	0.003 [0.001]	-0.005 [0.001]
Dummy US				-0.005 [0.005]	-0.030 [0.002]	-0.019 [0.007]
Constant	-0.005 [0.000]	-0.006 [0.000]	-0.012 [0.000]			
Observations	399879	399879	399879	399879	399879	399879
R-squared	0.39	0.39	0.61	0.4	0.4	0.62
Half-life Within Canada p-value (*)	2.9 0.000	2.8 0.000	4.8 0.000	2.7 0.000	2.6 0.000	4.4 0.000
Half-life Across Border p-value (*)	. .	4.5 0.000	7.9 0.000	. .	8.1 0.000	8.8 0.001
Long-Run Convergence Coefficient within Canada	-0.026	-0.026	-0.090	-0.029	-0.029	-0.088
Absolute Convergence Test within Canada (p-value)	0.000	0.000	0.000	0.000	0.000	0.000
Equality Test between Within and Across Absolute Convergence (p-value)				0.677	0.000	0.291

(*) P-value for a standardized normal coefficient test based on the asymptotic distribution estimated by Harris and Tzavalis (1999).

(†) This is the average of the dummies for Canada divided by the coefficient on $L1dlnp$. Standard errors are computed using the delta method.

Robust standard errors in brackets.

Table 8: Results from Aggregating UPC prices at the product group level

Dependent Variable	Common UPCs - Common Weights - Small Sample				Common UPCs - Common Weights - Large Sample				All UPCs - City-Specific Weights - Large Sample			
	Log of Product Group Price Ratio Relative to Ontario				Log of Product Group Price Ratio Relative to Ontario				Log of Product Group Price Ratio Relative to Ontario			
	No	No	Yes	Yes	No	No	Yes	Yes	No	No	Yes	Yes
City Dummies	No	Yes	No	Yes	No	Yes	No	Yes	No	Yes	No	Yes
Product Group Weights	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)	(10)	(11)	(12)
$Q_{ugc,t-1}$	0.821	0.843	0.82	0.83	0.873	0.967	0.848	0.947	0.989	0.999	0.986	0.995
	[0.028]	[0.053]	[0.030]	[0.057]	[0.014]	[0.016]	[0.016]	[0.019]	[0.006]	[0.003]	[0.006]	[0.003]
$Q_{ugc,t-1} * \text{Border}$	0.075	0.038	0.138	0.111	0.015	-0.034	0.135	0.039	-0.01	0.001	0.017	0.033
	[0.040]	[0.055]	[0.051]	[0.066]	[0.034]	[0.030]	[0.048]	[0.033]	[0.011]	[0.016]	[0.012]	[0.013]
Dummy ALB			0.013	0.015			0.01	0.009			0.002	0.007
			[0.003]	[0.006]			[0.002]	[0.003]			[0.003]	[0.003]
Dummy BRC			0.014	0.018			0.013	0.009			0.006	0.005
			[0.004]	[0.005]			[0.003]	[0.003]			[0.003]	[0.003]
Dummy MAN			0.003	0.003			0.006	0.006			0.001	0.003
			[0.004]	[0.006]			[0.003]	[0.003]			[0.003]	[0.003]
Dummy MAR			-0.003	0.003			0	0.004			-0.002	0
			[0.004]	[0.007]			[0.003]	[0.003]			[0.004]	[0.003]
Dummy QUE			-0.006	-0.001			-0.002	0.003			-0.001	0.003
			[0.004]	[0.007]			[0.003]	[0.003]			[0.003]	[0.003]
Dummy US			-0.04	-0.022			-0.04	-0.03			-0.041	-0.033
			[0.009]	[0.011]			[0.009]	[0.010]			[0.005]	[0.012]
Constant	0.002	0.006			0.003	0.006			-0.001	0.004		
	[0.001]	[0.002]			[0.001]	[0.001]			[0.001]	[0.001]		
Observations	6144	6144	6144	6144	6144	6144	6144	6144	6432	6432	6432	6432
Average UPCs per product group	5	5	5	5	80	80	80	80				
R-squared	0.63	0.71	0.64	0.71	0.73	0.86	0.72	0.87	0.97	0.98	0.97	0.98
Half-life Within Canada	4	4	3	4	5	21	4	13	63	.	49	138
p-value (*)	0.001	0.012	0.000	0.010	0.001	0.012	0.000	0.010	0.214	1.000	0.063	0.161
Half-life Across countries	6	5	16	11	6	10	40	49	33	.	.	.
p-value (*)	0.504	0.000	0.000	0.015	0.504	0.000	0.000	0.015	0.125	1.000	1.000	1.000
Long-Run Convergence Coefficient within Canada	0.011	0.039	0.022	0.044	0.024	0.182	0.0342	0.149	-0.091	2.628	0.114	0.879
Absolute Convergence Test within Canada (p-value)	0.093	0.000	0.000	0.000	0.002	0.035	0.000	0.011	0.410	0.628	0.083	0.152
Equality Test between Within and Across Absolute Convergence (p-value)			0.000	0.015	.	.	0.000	0.043

(*) P-value for a standardized normal coefficient test based on the asymptotic distribution estimated by Harris and Tzavalis (1999).

(†) This is the average of the dummies for Canada divided by the coefficient on L1dnp. Standard errors are computed using the delta method.

Robust standard errors in brackets; * significant at 5% level; ** significant at 1% level

Table 9: Aggregation Bias

	Within and Across		Within Country		Across the Border	
	Unweighted (1)	Weighted (2)	Unweighted (3)	Weighted (4)	Unweighted (5)	Weighted (6)
Persistence						
Mean Product Group Estimates	0.79	0.81	0.57	0.64	0.85	0.86
Aggregate	0.81	0.88	0.64	0.8	0.87	0.9
p-value*	0.000	0.000	0.000	0.000	0.000	0.000
Half-Lives						
Mean Product Group Estimates	2.9	3.3	1.2	1.6	4.3	4.6
Aggregate	3.3	5.4	1.6	3.1	5.0	6.6

Table 10: Non-Linearity in Price Adjustment within and Across the Border

Dependent Variable	Log of UPC Price Ratio Relative to Ontario			
	Yes	Yes	Yes	Yes
City Dummies				
Product Group Weights	Yes	Yes	Yes	Yes
	Within Canada		Across the Border	
$q_{ugc,t-1}$	0.779	0.845	0.866	0.984
	[0.002]	[0.004]	[0.008]	[0.019]
$q_{ugc,t-1} * q_{ugc,t-1} $		-0.221		-0.17
		[0.012]		[0.037]
Observations				
Average UPCs per product group	389343	389343	10536	10536
Approximate half-life for a one standard deviation price shock	2.8	2.2	4.8	6.3
R-squared	0.37	0.37	0.60	0.60

Table 11: A Simple Example of the Role of Non-Linearities

	Individual Good's Convergence Coefficient	Period 1	Period 2	Period 3
		Log Price Difference Relative to Ontario		
Good 1	0.7	-0.3	-0.21	-0.147
Good 2	0.7	0.3	0.21	0.147
Good 3	0.9	-0.03	-0.027	-0.024
Average		-0.01	-0.009	-0.0081
Aggregate Convergence	0.9			
Microdata Convergence	0.7008			

Table A1: Examples of Common UPCs across the Border

UPC	UPC Descriptor	Product Group Descriptor
6897829901	PLAYSTATION 2 RF ADAPTER 1S (#	AUDIO/VIDEO/COMPUTER UNITS
6897879500	XBOX UNI RF ADAPTER 1S (#79500	AUDIO/VIDEO/COMPUTER UNITS
1380300201	CANON POWERSHOT A10 DIG CAMERA	CAMERAS/FILM/ACCESSORIES
1821070001	NIKON COOLPIX 2000 DIGITAL CAM	CAMERAS/FILM/ACCESSORIES
5820038576	LUCERNE BUTTER UNSALTED 454GM	BUTTER & MARGARINE
5574227472	SMART CHOICE SOFT TUB 454 GM(#	BUTTER & MARGARINE
5980061302	NESTLE AFTER EIGHT BISCUIT CAR	COOKIES & SWEET BISCUITS
7241709129	CADBURY CARAMEL FINGERS 125 GM	COOKIES & SWEET BISCUITS
5610015728	PRINGLES REGULAR PLAIN 50 GM	SNACK FOODS
6041002521	LAYS CLASSIC PLAIN BIG GRAB 70	SNACK FOODS
5218132276	SAFETY 1ST FISH N ROD TUB TOY	TOYS
6487432633	MATTEL HOT WHEELS RIPPIN WHEEL	TOYS

Table A2: Examples of Common UPCs across the Border

	Share of Common UPCs (in terms of the count of UPCs)			Share of Common UPCs (in terms of the value of UPCs)		
	US-US	Can-Can	US-Can	US-US	Can-Can	US-Can
Ln(Distance)	-0.0301*** [0.0029]	-0.0460** [0.018]	-0.0184*** [0.0032]	-0.0301*** [0.0028]	-0.0238 [0.027]	-0.0186*** [0.0029]
Border Dummy			-0.134*** [0.0061]			-0.157*** [0.0056]
Constant	0.434*** [0.026]	0.972*** [0.16]	0.347*** [0.027]	0.469*** [0.025]	0.888*** [0.24]	0.387*** [0.023]
Observations	45	15	105	45	15	105
R-squared	0.82	0.88	0.97	0.8	0.83	0.97

Robust standard errors in brackets; *** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$

Appendix Table A3: Law of One Price Deviations within City/Region Pairs in the U.S. and Canada

Price Differences across Cities Common UPCs ONLY					Price Differences across Cities Common UPCs ONLY						
City/Region 1	City/Region 2	Number of Common UPCs (1)	Median (2)	Standard Deviation (3)	Median Absolute (4)	City/Region 1	City/Region 2	Number of Common UPCs (1)	Median (2)	Standard Deviation (3)	Median Absolute (4)
U.S. - U.S.						U.S. - Canada					
Boston	Chicago	10,362	0.000	0.226	0.117	Boston	Alberta	1737	0.018	0.267	0.152
Boston	Houston	10,235	0.017	0.212	0.107	Boston	British Columbia	1686	0.021	0.265	0.162
Boston	Los Angeles	9,119	0.000	0.242	0.134	Boston	Manitoba	1517	0.038	0.256	0.147
Boston	New York	11,503	0.000	0.223	0.115	Boston	Maritimes	1529	0.031	0.272	0.162
Boston	Atlanta	10,257	0.001	0.208	0.098	Boston	Ontario	2513	0.030	0.275	0.164
Boston	Detroit	10,863	0.000	0.220	0.110	Boston	Quebec	1616	0.036	0.275	0.175
Boston	Philadelphia	12,346	0.000	0.220	0.106	Chicago	Alberta	1569	0.025	0.255	0.145
Boston	Rochester	10,996	0.000	0.214	0.102	Chicago	British Columbia	1595	0.013	0.256	0.145
Boston	Phoenix	10,111	0.000	0.227	0.117	Chicago	Manitoba	1450	0.038	0.259	0.145
Chicago	Houston	11,102	0.038	0.221	0.123	Chicago	Maritimes	1407	0.034	0.274	0.155
Chicago	Los Angeles	9,773	0.000	0.234	0.120	Chicago	Ontario	2275	0.028	0.267	0.160
Chicago	New York	9,231	0.000	0.232	0.122	Chicago	Quebec	1442	0.029	0.274	0.171
Chicago	Atlanta	10,677	0.021	0.219	0.114	Houston	Alberta	1548	-0.013	0.249	0.153
Chicago	Detroit	12,798	0.000	0.222	0.106	Houston	British Columbia	1552	-0.028	0.250	0.149
Chicago	Philadelphia	11,213	0.000	0.226	0.112	Houston	Manitoba	1408	-0.003	0.254	0.144
Chicago	Rochester	10,466	0.000	0.214	0.102	Houston	Maritimes	1375	-0.003	0.267	0.152
Chicago	Phoenix	10,996	0.000	0.227	0.112	Houston	Ontario	2191	-0.010	0.269	0.163
Houston	Los Angeles	10,425	-0.039	0.241	0.141	Houston	Quebec	1450	-0.007	0.264	0.161
Houston	New York	8,910	-0.062	0.235	0.143	Los Angeles	Alberta	1558	0.007	0.257	0.147
Houston	Atlanta	13,209	0.000	0.193	0.083	Los Angeles	British Columbia	1558	-0.001	0.262	0.162
Houston	Detroit	12,322	-0.023	0.213	0.113	Los Angeles	Manitoba	1356	0.025	0.256	0.156
Houston	Philadelphia	10,823	-0.013	0.214	0.109	Los Angeles	Maritimes	1337	0.027	0.267	0.154
Houston	Rochester	10,074	-0.018	0.215	0.109	Los Angeles	Ontario	2210	0.021	0.279	0.169
Houston	Phoenix	12,853	-0.019	0.218	0.115	Los Angeles	Quebec	1437	0.018	0.272	0.158
Los Angeles	New York	8,346	0.000	0.252	0.136	New York	Alberta	1514	0.035	0.267	0.159
Los Angeles	Atlanta	9,494	0.029	0.239	0.133	New York	British Columbia	1518	0.038	0.271	0.169
Los Angeles	Detroit	10,116	0.002	0.237	0.124	New York	Manitoba	1358	0.057	0.257	0.166
Los Angeles	Philadelphia	9,361	0.002	0.245	0.135	New York	Maritimes	1401	0.067	0.267	0.167
Los Angeles	Rochester	8,449	0.000	0.236	0.124	New York	Ontario	2313	0.056	0.269	0.168
Los Angeles	Phoenix	12,752	0.000	0.222	0.102	New York	Quebec	1522	0.061	0.273	0.173
New York	Atlanta	8,963	0.043	0.229	0.131	Atlanta	Alberta	1383	-0.017	0.251	0.154
New York	Detroit	9,964	0.001	0.232	0.118	Atlanta	British Columbia	1363	-0.016	0.257	0.159
New York	Philadelphia	12,893	0.003	0.226	0.116	Atlanta	Manitoba	1234	0.004	0.257	0.151
New York	Rochester	9,723	0.000	0.233	0.119	Atlanta	Maritimes	1241	0.014	0.266	0.163
New York	Phoenix	8,684	0.000	0.240	0.128	Atlanta	Ontario	1982	0.001	0.273	0.168
Atlanta	Detroit	12,539	-0.005	0.208	0.105	Atlanta	Quebec	1345	0.000	0.265	0.163
Atlanta	Philadelphia	11,280	0.000	0.212	0.098	Detroit	Alberta	1756	0.007	0.262	0.152
Atlanta	Rochester	10,616	0.000	0.209	0.094	Detroit	British Columbia	1755	0.010	0.270	0.161
Atlanta	Phoenix	11,464	-0.007	0.218	0.111	Detroit	Manitoba	1608	0.022	0.256	0.151
Detroit	Philadelphia	11,984	0.000	0.221	0.107	Detroit	Maritimes	1617	0.034	0.270	0.159
Detroit	Rochester	11,593	0.000	0.214	0.096	Detroit	Ontario	2587	0.023	0.276	0.163
Detroit	Phoenix	11,603	0.000	0.224	0.111	Detroit	Quebec	1662	0.024	0.267	0.164
Philadelphia	Rochester	12,196	0.000	0.214	0.100	Philadelphia	Alberta	1624	0.024	0.254	0.151
Philadelphia	Phoenix	10,510	0.000	0.231	0.119	Philadelphia	British Columbia	1616	0.021	0.268	0.163
Rochester	Phoenix	9,675	0.000	0.226	0.113	Philadelphia	Manitoba	1464	0.034	0.260	0.165
All 45 US city comparisons:						All 60 Uscity-Canadian region comparisons:					
Median		10,616	0.000	0.223	0.113	Median		1,531	0.021	0.267	0.161
Average		10,730	-0.001	0.224	0.114	Average		1,634	0.019	0.266	0.160
St. Deviation		1,303	0.016	0.012	0.013	St. Deviation		328	0.020	0.008	0.008
Canada - Canada											
Alberta	British Columbia	29014	0.000	0.160	0.063	Rochester	Ontario	2215	0.026	0.282	0.177
Alberta	Manitoba	27824	0.000	0.154	0.056	Rochester	Quebec	1455	0.037	0.274	0.170
Alberta	Maritimes	22004	0.022	0.188	0.096	Phoenix	Alberta	1629	-0.015	0.258	0.154
Alberta	Ontario	30995	0.003	0.187	0.085	Phoenix	British Columbia	1667	-0.010	0.268	0.156
Alberta	Quebec	22359	0.005	0.193	0.094	Phoenix	Manitoba	1483	0.015	0.261	0.156
British Columbia	Manitoba	25094	0.007	0.168	0.071	Phoenix	Maritimes	1410	0.016	0.268	0.156
British Columbia	Maritimes	20286	0.031	0.196	0.106	Phoenix	Ontario	2303	0.006	0.278	0.163
British Columbia	Ontario	29281	0.016	0.194	0.096	Phoenix	Quebec	1532	0.002	0.282	0.169
British Columbia	Quebec	21126	0.017	0.200	0.103						
Manitoba	Maritimes	20879	0.008	0.189	0.092						
Manitoba	Ontario	28757	0.000	0.185	0.083						
Manitoba	Quebec	20994	0.000	0.192	0.089						
Maritimes	Ontario	30914	0.000	0.168	0.066						
Maritimes	Quebec	24800	0.000	0.171	0.073						
Ontario	Quebec	35374	0.000	0.165	0.068						
All 15 Canadian Region comparisons:											
Median		25,094	0.003	0.187	0.085						
Average		25,980	0.007	0.181	0.083						
St. Deviation		4,682	0.010	0.015	0.016						

Figure A1

