

Measuring the Cost of Living in Mexico and the US

David Argente

Chang-Tai Hsieh

Munseob Lee*

October 2021

Abstract

We use a dataset with prices and spending on consumer packaged goods matched at the barcode-level across the US and Mexico to measure the price index in Mexico relative to the US. Mexican prices relative to the US are 23% *lower* compared to the International Comparisons Project’s (ICP) price index. We decompose the 23% gap into the biases from imputation, sampling, quality, and variety. Quality bias *increases* Mexican prices by 48%. Imputation, sampling, and variety bias *lowers* Mexican prices by 11%, 13%, and 33%, respectively.

JEL Classification Numbers: E01, E31, O47

Keywords: Price Index, Purchasing Power Parity, International Price Comparison

1 Introduction

Indexes of prices across countries are a vital ingredient in estimates of standards of living and real output across countries. The most widely used price indexes are those by the International Comparison Program (ICP). The ICP collects prices of more than a thousand specific products (“items”) in multiple countries which it then aggregates into price indexes of 155 broad product categories (“basic headings”).

*David Argente, Pennsylvania State University, dargente@psu.edu. Chang-Tai Hsieh, University of Chicago and NBER, chsieh@chicagobooth.edu. Munseob Lee, University of California San Diego, munseoble@ucsd.edu. We thank Robert Inklaar, Esteban Rossi-Hansberg and Ina Simonovska for helpful comments. Researchers own analyses calculated (or derived) based in part on data from Nielsen Consumer LLC and marketing databases provided through the NielsenIQ Datasets at the Kilts Center for Marketing Data Center at The University of Chicago Booth School of Business. The conclusions drawn from the NielsenIQ data are those of the researchers and do not reflect the views of Nielsen. Nielsen is not responsible for, had no role in, and was not involved in analyzing and preparing the results reported herein. Churn Ken Lee, Jake Orchard, and Jihoon Sung provided excellent research assistance.

Despite their widespread use, it is well known that there are four potential biases in the prices provided by the ICP. The first is imputation bias. The ICP’s surveyors are unable to collect the prices of many items. The ICP imputes prices for the items with missing prices using data from other countries. The second is sampling bias. The ICP calculates the price of an item as an average of store-level prices weighted by each store’s *total* sales, but this procedure may not yield prices that reflect what consumers actually pay. Third, products differ in quality across countries, and it is possible that the ICP matches lower quality items in poor countries with higher quality ones in richer countries. Fourth, products available in one country are not available in other countries and no adjustment is made for potential differences in the availability of products across countries.

These potential biases have been brought up by many authors. For example, [Deaton \(2010\)](#) discusses how the ICP imputes missing prices, [Deaton and Heston \(2010\)](#) bring up the potential bias due to sampling and quality, and [Feenstra, Xu and Antoniadis \(2020\)](#) address the issue of variety bias in the ICP. However, although these biases are well understood conceptually, we have very little evidence on their empirical magnitude. Our goal in this paper is to use a new dataset to measure these biases empirically for the specific case of the price index of nondurable goods in Mexico relative to the US.

Our data is the Nielsen Consumer Panel data for Mexico and the US (the Nielsen data). The Nielsen data collects data on spending on nondurables for 40-60 thousand households in the US and 6 thousand households in Mexico. Households in the two countries use in-home scanners or diaries to record their purchases of packaged goods. The Nielsen data includes information on prices and quantities of nondurable goods (identified by a 12-digit barcode) purchased by each household on each shopping trip and from each retail store visited. The nondurables in the Nielsen data account for 60-65% of total nondurable spending in the two countries.

The richness of the Nielsen data allows us to measure potential biases in the ICP for nondurable goods. First, the Nielsen data has better coverage of items across both countries within the basic headings it covers. As a result, we can construct a price index similar to the one produced by the ICP but that does not rely on price imputations. We find that the ICP’s imputation of missing prices overstates prices by 11% in Mexico compared to the US.

Second, the Nielsen data has information on prices and expenditures on all products for a representative sample of households in each country. The Nielsen data also identifies the retail store where each purchase was made. We can therefore construct two price indices, one using weights from household expenditures and another based on a store’s total sales. Mexican prices (compared to the US) aggregated from consumer expenditure weights are 13% lower compared to prices aggregated from weights that reflect a store’s total sales. This

gap comes from the fact that Mexican households shop more frequently and are more likely to purchase only lower priced items in a given store compared to American households. Thus, using total sales in a store as weights overstates average prices in both countries, but more in Mexico compared to the US.

Third, the Nielsen data includes barcode information, which we use to identify more than 5,000 identical barcodes in the two countries. Mexican prices (compared to the US) among products with the same barcode are 48% *higher* compared to the price index of “comparable” products calculated following the ICP’s methodology. This suggests that, as conjectured by [Deaton and Heston \(2010\)](#), the ICP may match low-quality products in low-income countries with high-quality products in high income countries.

Fourth, since we observe all the purchases made by households, we can measure the importance of products available to Mexican consumers but not to American consumers, and vice versa. We find that Mexican varieties missing in the US market matter more than US products missing in Mexico. When we take into account the differences in the availability of varieties in the two countries, effective prices in Mexico are 33% lower compared to the US. The net effect of all four adjustments, for imputation, sampling, quality, and variety, lowers Mexican prices relative to the US by 23% compared to the ICP.

This paper builds on recent work that use alternative micro-data to estimate price indexes across countries. Specifically, [Cavallo, Diewert, Feenstra, Inklaar and Timmer \(2018\)](#), [Cavallo, Feenstra and Inklaar \(2021\)](#), and [Simonovska \(2015\)](#) use online data and [Feenstra, Xu and Antoniadis \(2020\)](#) use scanner data of toothpaste, laundry detergent, personal wash items and shampoo to estimate prices across countries. Our contribution is to bring new data to this long-standing question and to empirically measure multiple potential biases in the ICP using this new data. Specifically, our data contains detailed data on both prices and quantities on the *majority* of nondurables for a representative sample of *consumers*.

The four biases we document for the price index in Mexico relative to the US are also potentially present in measures of inflation. For example, [Nordhaus \(1997\)](#), [Bils and Klenow \(2001\)](#), and [Bils \(2009\)](#) document the bias in price indexes over time because quality improvements are not fully taken into account. [Broda and Weinstein \(2010\)](#) quantify the size of the variety bias in the US CPI over time. [Coibion, Gorodnichenko and Hong \(2015\)](#) studies the role of store-switching and [Chevalier and Kashyap \(2019\)](#) document the importance of price discrimination in the bias in the CPI’s sampling weights. We are not aware of a study that measures the importance of all these biases in the price index over time taken together in one consistent dataset.¹

¹There is also a literature that uses Engel’s law to measure price indexes over space and time without the need of price data. See [Hamilton \(2001\)](#), [Costa \(2001\)](#), and [Nakamura, Steinsson and Liu \(2016\)](#), who quantify biases in the CPI of the US and China. [Almås \(2012\)](#) and [Atkin, Faber, Fally and Gonzalez-Navarro](#)

The rest of the paper is organized as follows. Section 2 describes the data. In Section 3 we develop our price index. Section 4 presents our results and decomposes the gap between our price index and the ICP into the contribution of imputation, sampling, quality, and variety. The last section concludes.

2 Data Description

In this section, we describe the data collected by the ICP, the Nielsen data, and the matched data we constructed from the Nielsen data matched to the items covered in the ICP.

2.1 International Comparison Program (ICP)

The ICP is a statistical initiative that collects prices on more than thousand detailed products (“items”) around the world. The prices of different items are aggregated into 155 broad product categories (“basic headings”) that cover all the components of GDP. Approximately 53 basic headings refer to goods, of which 33 are non-durables. The publicly available data has price indices for the basic headings. We use the restricted 2011 ICP micro-data for Mexico and the United States, which has the prices at the item level.

The ICP specifies the characteristics of a “representative product” for each item. The specification of a representative product includes quantity and packaging (e.g. 250 milliliters of milk), source (e.g. produced domestically or imported), seasonal availability (e.g. year-round or only seasonal), product characteristics, and brand. For example, the representative product for the item “Baby Diapers” is a well-known brand, containing between 18 to 24 pieces, either classic or basic type, with a size between 4 and 9.5 kg, and with a multi-pack package.

Prices of the representative product of each item are collected from a sample of retail establishments chosen based on their total sales. Prices are collected from products in a retail establishment that meet the specifications of the representative products. The ICP then calculates the average price of an item as a weighted average of the unit price of the representative product across all the sampled stores, where the weights are the total sales of each store.

There are two points to note about the ICP’s sampling. First, we do not know the exact product chosen as the representative product in each store. Although the goal is to price the same product in multiple stores, the unavoidable problem is that stores differ in the products they sell. Therefore, it is likely that surveyors choose different products as the representative

(2020) apply this methodology to estimate welfare differences across income groups.

product of an item in different stores. Second, the weights used to aggregate prices in each store are calculated using the total sales of each store instead of the sales of an item in each store. The problem, as we will see later, is that total sales of a store as a share of total sales in all stores can be very different from weights calculated from the sales of an item in a store as a share of total sales of the item in all stores. For example, whole milk can account for a large share of sales in grocery stores but is relatively unimportant in sales for gas station convenience stores.

The ICP also collects national accounts expenditures for each basic heading but not for the items within each basic heading. These expenditures are used to aggregate basic headings into an aggregate price index. Within each basic heading, the ICP does not have expenditure weights at the item level. Instead, it classifies each item as “important” or “less-important.” For the nondurables in Mexico and the US that we consider, the ICP classifies the majority of these items as “important.”

2.2 Mexico and US Nielsen

The Nielsen data for the US tracks the shopping behavior of 40,000 to 60,000 households in 48 contiguous states plus Washington D.C. Each household uses in-home scanners to record their purchases. The US data contain slightly under one million distinct 12-digit barcodes. For each barcode, the data contains information on the brand, size, packaging, and other rich sets of product features. We combine the information on the price paid by the consumer with Nielsen’s data on the products’ characteristics to calculate the *unit price* of each product (e.g., price per ounce). In what follows, we use “price” as a shorthand for a product’s unit price.

The data also contains information on each purchasing trip the panelist makes, including information on the retailer, the retailer’s location, the date of the transaction, and the expenditures and prices of each barcode purchased in each store. Furthermore, the data have demographic variables such as age, education, annual income, marital status, and employment that are updated annually based on surveys sent to the households. Nielsen constructs projection weights that make the sample representative of the US urban population that we use in our calculations.

The Nielsen data for Mexico tracks the shopping behavior of 6,000 households for the years 2012-2013.² The sample is representative of all cities over 50,000 people and covers 55 cities in Mexico. Instead of using in-home scanners, households record their spending in

²Since the ICP sampling is conducted every 6 years, and the Mexican scanner data is only available from 2012-2013, we focus on the cross-sectional patterns of the 2011 ICP estimates.

diaries that are collected biweekly by Nielsen. Nielsen’s data for Mexico contains approximately 55,000 distinct barcodes.

The Mexican data contains detailed information on each shopping trip (date, store, amount spent), transaction level information for each product purchased (quantity, price, deals, coupons), as well as detailed product level characteristics (brand, size, packaging, flavor). As in the US data, we calculate the unit price from the products’ characteristics and the price paid by the consumer. The data also include demographic variables at the household level such as the occupation of the household members, education, age, and family size. As in the case of the US, Nielsen constructs projection weights that make the sample representative of the Mexican urban population that we use in our calculations.

2.3 Nielsen Matched to ICP Items

We use the ICP’s definition of an item along with Nielsen’s description of the characteristics of each barcode to assign barcodes to items in the ICP. We match the barcodes in Nielsen in Mexico and the US to 71 ICP items in 18 basic headings.³ These 71 items account for 60% of aggregate expenditures on non-durables in Mexico and 65% in the US. The basic headings of nondurables in the ICP not covered by the Nielsen data are products without barcodes such as fresh meats and fruits.

In our matched sample of 71 ICP items, there are 45 items whose prices are collected by the ICP in Mexico and the US. Table 1 lists these 45 items along with the number of barcodes in Nielsen we classify under each item. The 2011 ICP imputes the prices for Mexico and/or the US for the remaining 26 items. These items are listed in Table 2, again with the number of Nielsen barcodes we match to each item.

We also match barcodes in the US and Mexico, using the fact that the two countries use the same barcode system.⁴ The third row in Table 1 and 2 shows the number of matched

³There are a total 117 items in the ICP in the 18 basic headings we match with the Nielsen data. We aggregated the ICP items that we cannot distinguish from the barcode description in the Nielsen data. For example, the description of the characteristics of barcodes of breads do not distinguish between “White Bread (Not Sliced)” and “Sliced White Bread” (two distinct items in the ICP) so we aggregate these two ICP items into one. This aggregation reduces the number of ICP items we match to Nielsen to 104. The “new” items after we aggregate are: White Bread (not sliced and sliced), Dried/Instant Noodles (dried and instant noodles), Milk, un-skimmed (pasteurized and ultra-pasteurized), Cheese (cheddar, processed, camembert, and gouda), Tomato Paste (small and large), Ice Cream (cornetto-type and packed), Black Tea (bags and loose leaves), Canned Beer (domestic and imported), Coffee Roasted (Arabica and Robusta), and Detergent Powder (washing machine and hand wash). Among these 104 items, 6 of them are not covered by Nielsen and 27 are priced only in the US Nielsen data but not in the Mexican data. We end up with a total of 71 ICP items for which we have Nielsen prices in Mexico and the US. We find no significant difference in ICP’s Mexican prices (compared to US prices) between matched and unmatched items.

⁴Both countries use the Universal Product Code (UPC), which consists of 12 numeric digits. The organization that assigns these barcodes (i.e. GS1), is present in over 115 countries.

barcodes for each ICP item. It is clear that the number of common barcodes is small share of the barcodes sold in each country. The majority of barcodes sold in Mexico are not sold in the US, and vice versa.

Table 1: Items with Prices in the ICP

	# barcodes MX	# barcodes US	# common
Cornflakes	1312	5813	263
White Bread	159	1683	42
Roll	77	4662	16
Sandwich biscuits/cookies	568	709	3
Butter biscuits	18	26	0
Flavored biscuits/cookies sweet	1803	13124	134
Spaghetti	147	1971	21
Dried/Instant Noodles	401	1169	60
Canned tuna without skin	328	864	5
Milk, un-skimmed	763	1521	15
Milk, low-fat, Pasteurized	46	3486	9
Yoghurt, plain	1210	6207	13
Sour cream	206	909	5
Cheese	1445	15713	99
Cream cheese	72	1113	19
Potato chips	725	5043	118
Tomato Paste	104	1022	20
Tinned sweet corn/Maize	186	877	37
White sugar	360	375	3
Strawberry/Apricot Jam	164	901	12
Orange marmalade	10	204	2
Chocolate bar	882	3406	75
Ice Cream	547	8552	93
Cooking salt	235	1468	13
Tomato ketchup	167	627	31
Chicken Extract (bouillon/stock cube)	258	208	11
Baby food	45	1148	4
Cocoa Powder, Tin	201	558	8
Instant coffee	362	463	31
Black Tea	15	2307	0
Mineral water	920	4095	43
Carbonated Soft Drink (Small)	267	4376	47
Carbonated Soft Drink (Large)	1495	7669	41
Apple juice	204	1742	27
Canned Beer	231	1025	22
Bottled Beer	289	2196	23
Detergent Powder	821	710	41
Liquid Window Cleaner	70	355	4
Kitchen paper roll	481	1567	57
Dishwashing detergent	189	1405	57
Tooth paste, tube	420	1275	75
Shower gel	719	7457	133
Regular sanitary pad/napkin	554	1340	20
Shampoo	2164	4321	199
Toilet paper - Multipack	1029	2042	49

Note: The table reports the number of barcodes in Mexico and the US in the 45 items whose prices are collected by the ICP.

Table 2: Items with Imputed Prices in the ICP

	# barcodes MX	# barcodes US	# common
Wheat Semolina (Suji)	8	2	0
Oats, rolled	222	1898	19
Whole wheat bread	192	1988	65
Pita bread	2	390	0
Salted crackers	225	2961	32
Short pasta	210	50	0
Vermicelli (Angel Hair)	8	30	0
Macaroni	54	3307	23
Milk, condensed	56	132	2
Milk, powdered	140	567	2
Green Olives (with stones)	155	2354	22
Tinned green peas	71	709	4
Tinned Button Mushrooms	101	618	4
Brown sugar	17	652	1
Pineapple Jam	20	72	0
Natural honey, Mixed blossoms	198	2721	23
Thin Soya Sauce	73	337	13
Chili sauce	443	3095	33
Coffee Roasted	319	5574	43
Tea, green	12	1281	0
Orange juice	285	2634	33
Lemonade	52	1029	3
All purposes household cleaner	1352	4211	57
Deodorant roll-on for men	638	119	13
Toilet Soap	910	1787	69
Baby diapers	957	2399	45

Note: The table reports the number of barcodes in Mexico and the US in the 26 items where the ICP imputes the price of the item in Mexico or the US (or in both countries) .

3 An Exact Price Index

This section derives the ideal price index. We also describe how we use the Nielsen data to construct a price index that mimics the procedure followed by the ICP.

3.1 Exact Price Index

Following the ICP, we assume that utility is a function of consumption in the basic headings, and that basic headings are in turn an aggregate of expenditures of specific items.

Specifically, the utility of a representative household in a country is given by:

$$\mathbb{U} = \left[\sum_b \left(\sum_i C_{ib}^{\frac{\eta_b-1}{\eta_b}} \right)^{\frac{\eta_b}{\eta_b-1} \frac{\gamma-1}{\gamma}} \right]^{\frac{\gamma}{\gamma-1}} \quad (1)$$

where C_{ib} is consumption of item i in basic heading b . When possible, we omit the country index in the notation. The parameters γ and η_b denote the elasticity of substitution across basic headings and items within basic headings, respectively. The set of basic headings and items are the same for all countries in the ICP. Therefore, the indexes for basic headings and items do not carry an index for the country.

The ICP assumes that an item consists of a single representative product, but there are multiple barcodes that satisfy the ICP definition of the representative product of an item. For example, there are 463 barcodes in the US and 362 in Mexico that meet the ICP's definition of the representative product for "instant coffee" (see Table 1). To account for the varieties in each item, we assume consumption of an item in a country (say Mexico, indexed by m) is itself a CES aggregate of individual *barcodes* indexed by k :

$$C_{ib}^m = \left(\sum_{k \in m} (\varphi_{kib} C_{kib}^m)^{\frac{\sigma_{ib}-1}{\sigma_{ib}}} \right)^{\frac{\sigma_{ib}}{\sigma_{ib}-1}}. \quad (2)$$

Here C_{kib}^m denotes total physical units and φ_{kib} the quality of barcode k , where the parameter σ_{ib} denotes the elasticity of substitution between the barcodes. Note that the aggregation over the barcodes is only over the set available in each country.⁵

The exact price index of item i in basic heading b in Mexico, taking the US (indexed by u) as the numeraire, is then given by:

$$\text{EPI}_{ib} = \left(\frac{\lambda_{ib}^m}{\lambda_{ib}^u} \right)^{\frac{1}{\sigma_{ib}-1}} \times \prod_{k \in \text{common}} \left(\frac{p_{kib}^m}{p_{kib}^u} \right)^{\omega_{kib}} \quad (3)$$

The first term in equation 3 is the ratio of the share of total spending on the common barcodes in Mexico relative to the US, where the spending shares are defined as:

$$\lambda_{ib}^m \equiv \frac{\sum_{k \in \text{common}} p_{kib}^m C_{kib}^m}{\sum_{k \in m} p_{kib}^m C_{kib}^m} \quad \text{and} \quad \lambda_{ib}^u \equiv \frac{\sum_{k \in \text{common}} p_{kib}^u C_{kib}^u}{\sum_{k \in u} p_{kib}^u C_{kib}^u}$$

⁵In Appendix ?? we extend the utility function at the item level to allow for non-homothetic preferences. There we show that the bias in the ICP price index between Mexico and the US is almost the same when we account for non-homothetic preferences as in our baseline case with homothetic preferences.

The second term is the geometric mean of the ratio of the price in Mexico relative to the US of barcode k for the barcodes common to the two countries, weighted by the logarithmic mean of the expenditure shares of the barcode.⁶ The price of a barcode in equation 3 is a weighted average of the price of the same barcode in all the stores in a country

$$p_{kib}^m = \prod_{s \in m} (p_{skib}^m)^{\phi_{skib}^m} \quad \text{and} \quad p_{kib}^u = \prod_{s \in u} (p_{skib}^u)^{\phi_{skib}^u}$$

where p_{skib}^m is the price of barcode k in store s (in Mexico) and the weights are the share of spending on each barcode in the store.⁷

The exact price index of item i in basic heading b requires at least one common barcode between Mexico and the US. For 63 of the 71 items in the matched data, we observe at least one common barcode. For the remaining 8 items that do not have a common barcode, we account for these items with a variety correction term at the item level. Specifically, define I_C as the set of items with common barcodes. The exact price index for the basic heading can be defined as:

$$\text{EPI}_b = \left(\frac{\lambda_b^m}{\lambda_b^u} \right)^{\frac{1}{\eta_b - 1}} \times \prod_{i \in I_C} \text{EPI}_{ib}^{\omega_{ib}^*} \quad (4)$$

where ω_{ib}^* is logarithmic mean of the expenditure share of each item and the spending shares are defined as:

$$\lambda_b^m \equiv \frac{\sum_{i \in I_C} p_{ib}^m C_{ib}^m}{\sum_i p_{ib}^m C_{ib}^m} \quad \text{and} \quad \lambda_b^u \equiv \frac{\sum_{i \in I_C} p_{ib}^u C_{ib}^u}{\sum_i p_{ib}^u C_{ib}^u}$$

The second term in equation 4 is the weighted geometric average of the price index of the items that can be priced from equation 3. The first term in equation 4 is ratio of the spending shares on items with common barcodes in the two countries. This term captures the weighted geometric average of the price index of the items that can *not* be priced from equation 3.

Finally, the aggregate exact price index is a weighted geometric average of the exact price index of each basic item:

$$\text{EPI} = \prod_b \text{EPI}_b^{\omega_b} \quad (5)$$

⁶The logarithmic mean is $\omega_{kib} \equiv \frac{s_{kib}^m - s_{kib}^u}{\ln s_{kib}^m - \ln s_{kib}^u} / \sum_{k \in \text{common}} \frac{s_{kib}^m - s_{kib}^u}{\ln s_{kib}^m - \ln s_{kib}^u}$ where $s_{kib}^m \equiv \frac{p_{kib}^m C_{kib}^m}{\sum_{k \in m} p_{kib}^m C_{kib}^m}$
and $s_{kib}^u \equiv \frac{p_{kib}^u C_{kib}^u}{\sum_{k \in u} p_{kib}^u C_{kib}^u}$

⁷Note that in the data we observe multiple prices for the same barcode purchased from multiple stores. Appendix ?? shows that the ideal price of a barcode can be derived as a CES price index from a consumer's discrete choice problem and that a first order approximation to this index is the Cobb-Douglas price index.

where ω_b is the logarithmic mean of the expenditure shares of each basic heading and \mathbb{EPI}_b is given by equation 4.

3.2 A “Pseudo-ICP” Price Index

There are two differences between the ideal price index we will calculate and the ICP’s price index. First, the underlying data (the Nielsen data) is different from the data used by the ICP. Second, there is the difference in methodology. To isolate the effect of methodology, we use the Nielsen data to construct a price index that mimics the ICP. We call this a “Pseudo-ICP” price index.

We proceed in four steps. First, for each ICP item, we identify the set of barcodes in each store that meet the ICP’s specifications for the representative product.⁸ From these set of barcodes, we pick the barcode with the largest volume of sales in each store as the representative product of the item in the store. The representative product of an item is therefore not the same in each store. We obviously do not know the exact product in each store surveyed by the ICP’s price surveyors, but we believe this procedure is a good approximation of how the ICP chooses the representative product of an item in each store. The price of an item is then the geometric mean of the price of the store-specific representative product weighted by the store’s *total* sales.

Second, the ICP does not collect prices for 26 items in Table 2. Instead it imputes the prices for these items using data from other countries. We estimate the imputed prices by comparing the basic heading price indexes reported by the ICP with the geometric mean of the item-level prices they do collect. If, for example, a basic heading price reported by the ICP is higher than the geometric mean of item-level observed prices, we infer that the imputed prices for the average item with a missing price are higher than those of other items in the same basic heading.

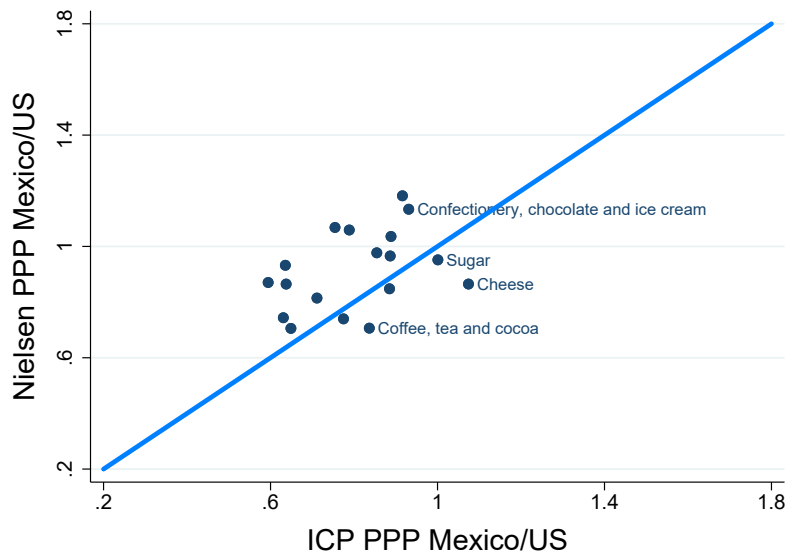
Third, the price index of a basic heading is an equally weighted geometric average of the price of the items in each basic heading.⁹ Figure 1 shows the scatterplot of the Pseudo-ICP price index of a basic heading (Mexico relative to the US) vs. the price index published by the ICP. The correlation is not one but the pseudo ICP price index generally aligns with the ICP’s price index. The median of the pseudo ICP price index at the basic heading level is 0.90 whereas the median of the actual ICP price index is 0.81. The regression coefficient is 1.12 and statistically significant at the 1 percent level. At the item level, the median of the

⁸We also mimic the fact that the ICP survey only covers four cities in Mexico (Mexico City, Guadalajara, Monterrey, and Puebla), while it covers all urban areas in the US.

⁹In principle the ICP classifies items as “important” and “less important” and gives “important” items more weight when aggregating to the basic heading level. However, most of the nondurable ICP items in Mexico and the US that we match to the Nielsen data are classified as “important.”

pseudo ICP price index is 0.85 whereas the median of the numbers published by the ICP is 0.73. A linear regression across the prices at the item level yields a coefficient is 1.10 and is also statistically significant at the 1 percent level.

Figure 1: Pseudo vs. Actual ICP Price Index



Notes: The figure plots the pseudo-ICP index of each 18 basic heading against the index published by ICP.

The final step is to aggregate the price of a basic heading into an aggregate price index. The ICP calculates the geometric average of the price index of a basic heading across all basic headings. In the case of the price index of Mexico vs. the US, the price indexes of the basic headings are first averaged using Mexico's weights, then averaged using the US weights, and then the geometric mean of the two is taken. The result is a Fisher index, which facilitates multilateral comparisons since such indexes are transitive (i.e. price comparisons between two countries are the same whether it is computed directly or indirectly through a third country) and base country-invariant (i.e. price comparisons between two countries are the same regardless of the choice of base country). Since we are only comparing two countries, we aggregate the price index at the basic heading level using the logarithmic weight of each basic heading. This is the theoretically consistent way to aggregate prices with a CES utility function.

4 The Price of Nondurables in Mexico vs. the US

In this section, we calculate the exact price index from the Nielsen data. We decompose the gap between the exact and the pseudo-ICP price index into the bias due to imputation, sampling, quality, and variety.

We need the elasticity of substitution between barcodes within each item. We follow Feenstra (1994)’s procedure as extended by Broda and Weinstein (2006) and Broda and Weinstein (2010). The procedure consists of estimating a demand and supply equation for each barcode using the information on prices and quantities of each barcode from the Nielsen data.¹⁰ The average of the elasticity of substitution across barcodes we obtain is 6.56 with standard deviation of 3.13.¹¹

We use these estimates of the elasticity of substitution between barcodes, along with the Nielsen data, to estimate the exact price index of each item from equation 3. We then use equation 4 to aggregate the price of an item using logarithmic weights of each item in a basic heading, which we then aggregate into an aggregate price index using logarithmic weights for the basic heading as equation 5. Column 1 in Table 3 shows that the exact price index is 0.65, which implies that Mexican prices are 35% lower than in the US.

The second column shows the pseudo-ICP price index. Here we take the weighted average of the pseudo ICP price of the basic heading (shown in figure 1) using the same logarithmic weights for each basic heading that we used to aggregate the exact price index of each basic heading. The aggregate pseudo-ICP price index, shown in column 2 in Table 3, is 0.86 which suggests that Mexican prices are only 14% lower than in the US. The exact price index calculated from the same data suggests that Mexican prices are 35% lower, which indicates an aggregate bias of 23%.

Table 3: Cost of Living in Mexico vs. the US

Aggregate Price Index		Bias due to:				Aggregate
Exact	Pseudo-ICP	Imputation	Sampling	Quality	Variety	Bias
0.65	0.86	0.89	0.87	1.48	0.67	0.77

Notes: The table reports the aggregate exact price index, pseudo ICP price index, and the gap between the exact and ICP index due to imputation, sampling, quality, and variety. Aggregate bias is the product of the bias due to imputation, sampling, quality, and variety.

¹⁰See Appendix ?? for more details.

¹¹We check the robustness of these results by calculating price indexes and biases with the common elasticity of substitution, 6. The estimated exact price index (0.64) is very close to the item-specific elasticity of substitution (0.65). See Appendix ?? for more details.

Recall that we use the same logarithmic weights at the basic heading level to aggregate the exact and pseudo ICP price indexes for each basic heading. Therefore, the difference between the aggregate exact and the ICP price index in Table 3 comes from the price index at the basic heading level. In turn, the gap between the two indexes at the basic heading level is the product of the biases due to sampling, quality, and variety for each basic heading:

$$\frac{\text{EPI}_b}{\text{ICP}_b} = \text{Imputation Bias}_b \times \text{Sampling Bias}_b \times \text{Quality Bias}_b \times \text{Variety Bias}_b$$

We now quantify each of these biases.

Imputation Bias: Out of the 71 items in our sample, the ICP does not collect prices for 26 of them from either Mexico or the US. The ICP runs a set of country-product dummy (CPD) regressions to impute the missing values for items whose prices cannot be found (Deaton, 2010). We revisit this imputation and quantify the size of the bias. As discussed in the previous section, we back out imputed prices by comparing a basic heading price reported by the ICP and geometric mean of item-level prices within a basic heading.

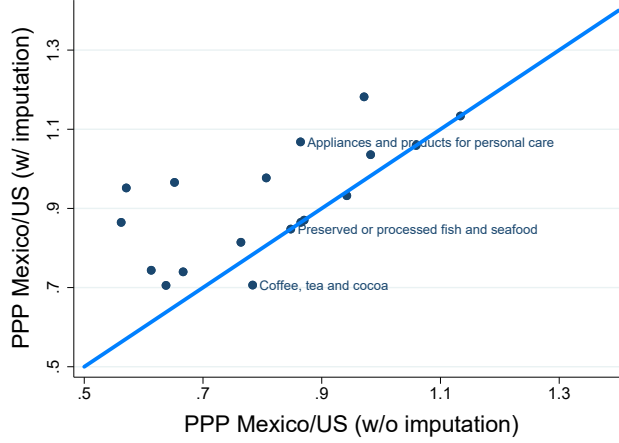
The imputation bias is defined as follows:

$$\text{Imputation Bias}_b \equiv \left[\prod_i (\tilde{p}_{ib}^m / \tilde{p}_{ib}^u)^{\frac{1}{N_b}} \right] / \text{ICP}_b \quad (6)$$

where ICP_b is the pseudo ICP that uses imputed prices for items not collected from either Mexico or the US and \tilde{p}_{ib} is the price of the item computed from the Nielsen data. There is no imputation bias if the ICP collects prices for all the items.

Figure 2 plots the ratio of US to Mexican prices of a basic heading with imputed item level prices on the y-axis against the price ratio calculated from the Nielsen data on the x-axis. The observations on the 45-degree line are those for which the ICP has prices for all items within the basic heading. For these basic headings the imputation bias is one. However, on average the majority of the observations are above the 45-degree line: the ratio of Mexican to US prices with imputed item level prices is typically higher than the price ratio calculated from actual price data.

Figure 2: Imputation Bias for Each Basic Heading



Notes: The figure plots the PPP with and without item-level imputation for each of the 18 basic headings.

Table 3 shows the net effect of imputation on the aggregate price index. Specifically, the third column shows the imputation bias calculated from equation 6. Imputation lowers Mexican prices by 11% compared to the ICP's index.

Sampling Bias: There are three sources of sampling bias. First, the ICP samples items only from the four cities in Mexico (i.e. Mexico City, Guadalajara, Monterrey, and Puebla). Second, at the item level, the ICP uses an equal-weights geometric average because they lack data on expenditures at the item level. Third, the ICP aggregates prices from each store using total sales of the store instead of the sales of the item in the store as weights. The sampling bias is the product of these three biases:

$$\text{Sampling Bias}_b \equiv \left[\prod_i \left(\frac{\bar{p}_{ib}^m / \bar{p}_{ib}^u}{\tilde{p}_{ib}^m / \tilde{p}_{ib}^u} \right)^{\frac{1}{N_b}} \right] \times \left[\frac{\prod_i \left(\frac{\bar{p}_{ib}^m}{\bar{p}_{ib}^u} \right)^{\omega_{ib}}}{\prod_i \left(\frac{\bar{p}_{ib}^m}{\bar{p}_{ib}^u} \right)^{\frac{1}{N_b}}} \right] \times \left[\prod_i \left(\frac{\hat{p}_{ib}^m / \bar{p}_{ib}^m}{\hat{p}_{ib}^u / \bar{p}_{ib}^u} \right)^{\omega_{ib}} \right] \quad (7)$$

where \tilde{p}_{ib}^m is the average price of an item computed as the weighted average of store level prices, where the weights are the total sales of the store considering only the main cities in Mexico, \bar{p}_{ib}^m is the same weighted average considering all cities and \hat{p}_{ib}^m is the weighted average where the weights are the sales of the particular item in the store.

The first term denotes the bias from the geographical coverage of the ICP survey in Mexico. By sampling only from the four major cities in Mexico, the ICP underestimates the average price of items by approximately 1%.¹² The four cities the ICP samples in Mexico

¹²Consistent with the coverage of our data, the ICP samples items from urban areas in all US states (i.e.

are among the largest in terms of population and their average household income is close to that of the median city in the country (i.e. choosing these four cities is close to choosing a representative city for Mexico.) This, combined with the fact that the ICP only samples urban cities in the US as well, yields a small geographic bias.

The second term in equation 7 is the sampling bias discussed by [Deaton and Heston \(2010\)](#). It denotes the bias from using the simple average instead of the weighted average of the item level price, and its magnitude depends on the covariance of the expenditure weights of each item and its price.¹³ The covariance is negative if people spend more on items with lower prices. In the Nielsen data, the difference in these covariances between Mexico and the US is essentially zero. Thus, at least in the case of Mexico relative to the US, there is little bias from not using expenditure weights at the item level.


The third term in equation 7 is the bias from aggregating store-level prices using the total sales of the store as weights instead of the sales of the item in the store. For example, a fruit vendor may also sell milk, but the vendor's total sales depend mostly on fruit sales instead of milk. The resulting bias in the price index depends on whether consumers in Mexico purchase more products at stores where these products are cheaper compared to consumers in the US.

Figure 3 shows that this is indeed the case. It plots the distribution of prices for each item where the items' average price is averaged over the store specific price using the stores' total sales and the sales of each item in the store. For both countries, the distribution weighted using store-item weights is shifted to the left compared to the distribution using store level weights. Furthermore, the gap between the two distributions is larger for Mexico compared to the US. Weighting store specific prices using a store's total sales overstates the average price paid by consumers in the two countries but more so in Mexico.

$\tilde{p}_{ib}^u = \bar{p}_{ib}^u$).

¹³Proposition ?? in the Appendix ?? formally shows this relationship and empirically estimates the covariance using a reduced form approach.

Figure 3: Distribution of Prices using Store vs. Store-Item Weights




Notes: The figure shows the distribution of the price of an item calculated as a geometric average of the price of the item at the store level using either the store's total sales or the sales of the item in the store as weights.

Figure 4 shows why the gap between the two price distributions is larger for Mexico. The left panel in Figure 4 plots the demeaned price of a barcode paid by a household in a shopping visit against the number of shopping visits of the household in the Nielsen data.¹⁴ Households in Mexico and the US that shop more frequently pay lower prices for the same product compared to households that shop less frequently. The right panel in Figure 4 shows the distribution of the number of shopping trips per household. A typical US household makes one shopping trip per week whereas in Mexico the number is five. The average Mexican household shops more intensely than the typical American household, and thus buys more of the cheaper products in a given store.

Table 3 shows the effect of the difference in shopping behavior between Mexico and the US on the price index. Specifically, the third column shows the weighted mean of the sampling bias at the basic heading level calculated from equation 7, where the weights are the logarithmic share of each basic heading. The difference in shopping intensity between Mexico and the US lowers Mexican prices by 13% compared to the ICP because the ICP's sampling weights ignores the effect of shopping intensity.

¹⁴The number of shopping visits is the number of retail outlets a household makes purchases from in a week.

Figure 4: Shopping Intensity and Prices in Mexico and the US



Notes: Left panel plots the demeaned price paid for a product vs. the demeaned number of shopping trips. The demeaned product price is the residual from a regression of the log price of a barcode on category, store, and quarter fixed effects. The demeaned number of shopping trips is the residual of the log number of shopping trips on category, store, and quarter fixed effects. Right panel shows the distribution of the number of shopping trips per household. Number of shopping trips is the number of stores visited by a household in a week.

Quality Bias: In 63 out of the 71 items in our matched sample, there is at least one common barcode sold in the two countries. With at least one common barcode within an item, we can calculate quality and variety biases separately.¹⁵ The quality bias in the ICP reflects the effect of matching high-quality goods in one country with a low-quality good in another country. We measure quality bias of an item as:

$$\text{Quality Bias}_{ib} \equiv \prod_{k \in \text{common}} \left(\frac{p_{kib}^m}{p_{kib}^u} \right)^{\omega_{kib}} \bigg/ \left(\frac{\hat{p}_{ib}^m}{\hat{p}_{ib}^u} \right) \quad (8)$$

The denominator in equation 8 is the price gap between Mexico and the US at the item level computed as the average price of the “representative product” in each store. The numerator is the price gap where the “representative product” in a store is chosen from the set of barcodes sold in both Mexico and the US.¹⁶ Quality bias is greater than one if

¹⁵The relative magnitude of the variety bias and the quality bias could be affected if we identify too few barcodes in common, particularly if there are a significant number of products among the unmatched barcodes that are in fact the same products. Nonetheless, the aggregate bias is the *product* of the variety and the quality biases. Therefore, as long as products with the same barcode are in fact the same product in the two countries, the aggregate bias is not affected by the possibility that some of the unmatched products may in fact be the same product.

¹⁶Both the numerator and the denominator in equation 8 weights the price of the item in each store by the sales of the item in the store.

the representative product of an item in a store in Mexico are of lower quality than the representative products of the same item in the US.

Table 4 shows that this is indeed the case. The first row follows the ICP’s methodology and chooses one representative product per store to calculate the average price of an item. The first column shows the average percent difference between Mexico and the US in the price of an item, which is 28.0%. The second column picks the representative product of an item in the store chosen from the barcodes common to the two countries. The average price gap using products with same barcode is close to zero. The second row provides a similar number but where the average price of an item is calculated as a weighted average of the price of *all* the barcodes in the item. The first column considers all the barcodes in an item; the second column only considers barcodes sold in the two countries. Average prices calculated from all the barcodes are 42.1% lower in Mexico. The second column restricts to barcodes sold in the two countries. Here, the average price in Mexico is very close to the average price in the US.

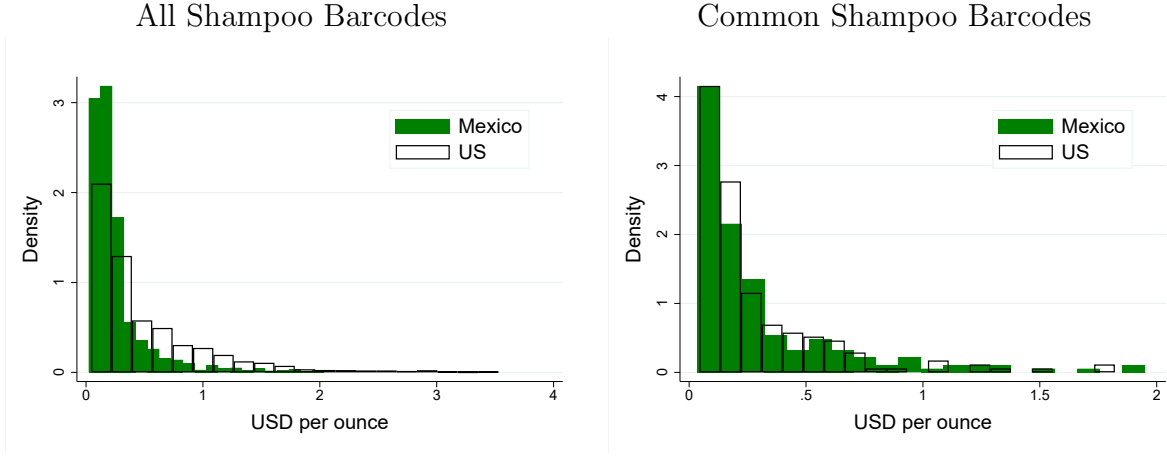
Table 4: Price Gap between Mexico and US for All vs. Common Barcodes

Sampling	All barcodes	Common barcodes
Representative product	-28.0 %	-0.5 %
All products in item	-42.1 %	0.5 %

Notes: The table reports average percent difference in item-level prices between Mexico and the US. We take the weighted average with item-level expenditure weights. Column 1 considers all barcodes available in the two countries within each item. Column 2 restricts the sample to barcodes sold in Mexico and the US. Row 1 chooses one representative product per store. Row 2 considers all the barcodes and calculates the weighted average of the prices of chosen products.

Figure 5 shows the distribution of the price of shampoo. The left panel shows distribution of the unit price of shampoo in the US and Mexico for all shampoo barcodes. The right panel shows the distribution of the unit price of shampoo only for shampoo barcodes sold in both countries. As can be seen, the average shampoo in the US is more expensive than in Mexico, but the difference is much smaller for shampoos sold in both countries.

Figure 5: Price Distribution for All and Common Shampoo Barcodes



Notes: Left panel shows the price distribution of barcodes of all shampoo products in the US and Mexico. Right panel plots the price distribution of shampoo products sold in both countries.

Table 3 shows the aggregate effect of the patterns shown in Figure 5 and Table 4 on aggregate quality bias (for the 63 items for which this calculation is possible). The effect is large: quality bias *increases* Mexican prices by 48% relative to the US.

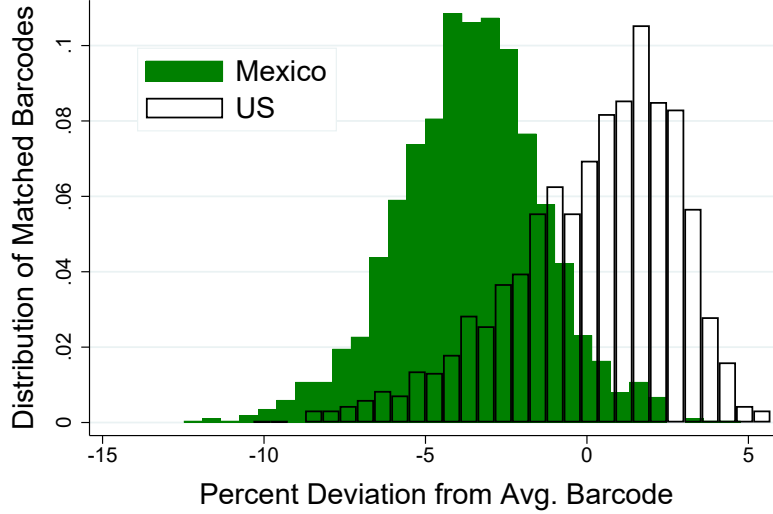
Variety Bias: We have already seen that the majority of Mexican barcodes are not available in the US market, and vice versa. The ICP price index is biased due to missing varieties if the importance of Mexican barcodes not available in the US is different from the importance of American barcodes not available in Mexico. We measure variety bias (of an item) by:

$$\text{Variety Bias}_{ib} \equiv \left(\frac{\lambda_{ib}^m}{\lambda_{ib}^u} \right)^{\frac{1}{\sigma_{ib}-1}}. \quad (9)$$

This is the ratio of the expenditure share of common barcodes in Mexico relative to the expenditure share of common products in US as in [Feenstra \(1994\)](#).

Figure 6 plots the distribution of share of spending on common barcodes relative to the spending share on the average barcode in the country. US households spend more on typical common barcodes while the opposite is true in Mexico. This means that missing Mexican varieties are more important compared to missing American varieties.

Figure 6: Distribution of Revenue of Barcodes Sold in Both Countries



Notes: The figure shows the distribution of revenue generated by the common barcodes relative to the average product in each country.

Barcodes that are present in Mexico and not present in the US are both very popular in Mexico (in terms of expenditures) and also very cheap relative to those sold in both countries. In the case of cereals, for instance, unmatched barcodes in Mexico are mainly produced by Mexican firms at a significantly lower price per unit.

Table 3 summarizes the effect of variety bias on the price index (for items with common barcodes in the two countries). We estimate that varieties only available to Mexican consumers lowers the effective prices paid by Mexican consumers by 33% compared to the US.¹⁷

Bias from Items Without Common Barcodes: Finally, 8 of the 71 items in our sample do not have common barcodes across the two countries. For these items, we cannot separately calculate quality and variety biases, but we can measure the *product* of these two biases on

¹⁷Our results from comparing Mexico and the US differs from those found in studies using within-country variation (e.g. Handbury and Weinstein, 2015). Intuitively, the variety bias captures the relative share of expenditures in common goods across locations. Within countries, households residing in a low income state consume many of the same goods that are available in higher income states. Across countries, variety bias is driven by country-specific factors and, in the case of Mexico and the US, the importance of local brands in Mexican households' consumption is remarkable. As a result, the share of expenditures in common goods for Mexico is substantially lower than that of the US, even if the US has a higher average household income.

the price of a basic heading from the following equation:

$$\begin{aligned} & \text{Quality and Variety Bias From Items Without Common Barcodes}_b \\ &= \left(\frac{\lambda_b^m}{\lambda_b^u} \right)^{\frac{1}{\eta_b-1}} \times \frac{\prod_{i \in I_C} \left(\frac{\hat{p}_{ib}^m}{\hat{p}_{ib}^u} \right)^{\omega_{ib}^*}}{\prod_i \left(\frac{\hat{p}_{ib}^m}{\hat{p}_{ib}^u} \right)^{\omega_{ib}}} \end{aligned} \quad (10)$$

where as before I_C is the set of items with common barcodes. In our data, equation 10 is very close to one for all the basic headings with at least one item that does not have a common barcode in the two countries. So at least for these 8 items for which we need to calculate equation 10 the net effect of quality and variety biases does not change the price index of Mexican goods relative to the US. Note that this is not the case for the items for which we can separately estimate the effect of quality and variety biases. For these 63 items, the net effect of these two biases decreases Mexican prices by 1% ($1.48 \times 0.67 = 0.99$) compared to the US.

Price Index Bias for Basic Headings: The implication from Table 3 is that the ICP overstates Mexican prices because of imputation, sampling and variety bias, but understates Mexican prices because of quality bias. The net effect – aggregate bias in Table 3 – is that the exact price index is about 23% lower compared to the pseudo-ICP index.

Table 5 shows the bias in the ICP for the basic headings.¹⁸ The first message is that there is a large amount of heterogeneity in the aggregate bias. The exact price index is similar to the ICP for “sugar” and “coffee, tea and cocoa.” At the other extreme, the ICP price index is significantly lower than the ideal price index for “food products nec” and “non-durable household goods”. This heterogeneity suggests that one should be careful about drawing strong inferences from prices obtained from a narrow set of products.

The relative importance of the four biases also differs quite bit across products. Imputation bias is large for basic headings with a large number of missing prices in the ICP such as “Pasta products” where 4 out 6 items are missing either from Mexico or the US. Sampling bias is very large for “appliances and articles for personal care” and virtually one for “fresh milk.” Mexican households appear to be more price sensitive in where they buy appliances and articles for personal care compared to fresh milk. Quality bias is basically one for fresh milk, sugar, and beer: the price of the average fresh milk, sugar, and beer product is about the same as fresh milk, sugar, and beer sold both in Mexico and the US. On the other hand, quality bias is very large for cereals, drinks and juices, non-durable household goods, and

¹⁸The product of the quality and variety terms from items without common barcodes are very close to one for all basic headings.

personal care products. Finally, Mexican varieties are more important than American varieties for Mexican households in almost every basic item (the variety bias term is less than 1), but the exception is sugar. For this basic heading, American varieties not sold in Mexico are slightly more important than Mexican barcodes not sold in the US.

Table 5: Biases in ICP for Each Basic Heading

Basic Heading	Biases				Aggregate Bias
	Imputation	Sampling	Quality	Variety	
Other cereals, flour and other products	0.83	0.53	1.56	0.55	0.38
Bread	1.01	0.95	0.89	0.64	0.55
Other bakery products	0.65	1.27	1.77	0.46	0.67
Pasta products	0.68	1.02	1.41	0.83	0.80
Preserved or processed fish and seafood	1.00	0.94	0.94	0.21	0.19
Fresh milk	1.00	1.04	0.98	0.38	0.38
Preserved milk and other milk products	0.82	1.02	1.41	0.67	0.79
Cheese	1.00	0.88	1.46	0.58	0.75
Frozen, preserved or processed vegetables	0.82	0.95	0.84	0.60	0.39
Sugar	0.60	1.44	1.10	1.06	1.01
Jams, marmalades and honey	0.95	0.69	1.05	0.69	0.47
Confectionery, chocolate and ice cream	1.00	0.75	1.24	0.92	0.86
Food products nec	0.90	0.94	1.61	0.98	1.34
Coffee, tea and cocoa	1.11	0.97	1.74	0.56	1.05
Mineral waters, soft drinks, and juices	0.94	0.73	1.71	0.56	0.65
Beer	1.00	0.96	1.06	0.87	0.88
Non-durable household goods	0.90	0.99	2.81	0.60	1.50
Appliances and products for personal care	0.81	0.70	1.54	0.66	0.57

Notes: The table reports imputation bias, sampling bias, quality bias, and variety bias for each basic heading. Aggregate bias is defined as the product of the imputation, sampling, quality, and variety biases.

5 Conclusion

The construction of cross-country price indexes is of crucial importance to compare living standards between countries and to measure global inequality. The ICP has taken on the important and heroic exercise of measuring these prices, but faces severe data limitations. In this paper, we construct a data set for two countries that allows us to address some of these data limitations; namely, the fact that the ICP has incomplete information on item across countries, that they do not have expenditure information to weight items appropriately, that they cannot compare exactly the same item across countries, and that they do not have information on differences on the set of products available in each country. Using our alternative data, we estimate that Mexican real consumption is larger relative to the

United States than was previously estimated. We identify the imputation of prices and the heterogeneity in shopping behavior, quality of products, and variety availability as important sources of bias in international price comparisons. Overall, our results show that the real non-durable consumption inequality across the US and Mexico is 23% lower than that predicted by the ICP estimates.

There are several generalizable lessons from our study for international price comparisons. First, aggregate bias must be estimated considering all biases jointly. Biases often imply adjustments in opposite directions, so addressing them in isolation could lead to drastically different conclusions about the comparison of the standards of living across countries. Second, because average prices are correlated with income per capita, the magnitude of the quality bias increases when comparing unequal countries. Third, the relevance of the variety bias is not correlated with income, but instead with the importance of *domestic* varieties in overall consumption. For this reason, using within-country variation yields the opposite results to using cross-country variation. Fourth, the bias from the limited geographical coverage of the ICP survey in Mexico is not large, indicating that coverage is less relevant if the sampling for each country only focuses on urban areas, as it does in the case of Mexico and the US. Fifth, the bias from using the simple average instead of the weighted average of the item level price is not large; the covariance of the expenditure weights of each item and its price is not large for neither the US or Mexico. Sixth, a country's shopping intensity is correlated with the relevance of the sampling bias. This is because weighting store specific prices using a store's total sales overstates the average price paid by consumers in all countries. Finally, we find a great deal of heterogeneity in the size of aggregate bias across different product categories. As more product-level data becomes available, we will be able to get better estimates for aggregate price index beyond consumer packaged goods.

References

- Almås, Ingvild, “International income inequality: Measuring PPP bias by estimating Engel curves for food,” *American Economic Review*, 2012, *102* (2), 1093–1117.
- Atkin, David, Benjamin Faber, Thibault Fally, and Marco Gonzalez-Navarro, “Measuring Welfare and Inequality with Incomplete Price Information,” *NBER Working Paper No. 26890*, 2020.
- Bils, Mark, “Do higher prices for new goods reflect quality growth or inflation?,” *The Quarterly Journal of Economics*, 2009, *124* (2), 637–675.
- and Peter J Klenow, “Quantifying Quality Growth,” *American Economic Review*, 2001, pp. 1006–1030.
- Broda, Christian and David E Weinstein, “Globalization and the Gains from Variety,” *The Quarterly Journal of Economics*, 2006, *121* (2), 541–585.
- and —, “Product creation and destruction: Evidence and price implications,” *American Economic Review*, 2010, *100* (3), 691–723.
- Cavallo, Alberto, Robert C Feenstra, and Robert Inklaar, “Product Variety, the Cost of Living and Welfare Across Countries,” Technical Report, National Bureau of Economic Research 2021.
- , W Erwin Diewert, Robert C Feenstra, Robert Inklaar, and Marcel P Timmer, “Using Online Prices for Measuring Real Consumption across Countries,” *AEA Papers and Proceedings*, 2018, *108*, 483–87.
- Chevalier, Judith A and Anil K Kashyap, “Best prices: Price discrimination and consumer substitution,” *American Economic Journal: Economic Policy*, 2019, *11* (1), 126–59.
- Coibion, Olivier, Yuriy Gorodnichenko, and Gee Hee Hong, “The cyclicalities of sales, regular and effective prices: Business cycle and policy implications,” *American Economic Review*, 2015, *105* (3), 993–1029.
- Costa, Dora L, “Estimating real income in the United States from 1888 to 1994: Correcting CPI bias using Engel curves,” *Journal of Political Economy*, 2001, *109* (6), 1288–1310.
- Deaton, Angus, “Price indexes, inequality, and the measurement of world poverty,” *American Economic Review*, 2010, *100* (1), 5–34.

- and Alan Heston, “Understanding PPPs and PPP-based national accounts,” *American Economic Journal: Macroeconomics*, 2010, 2 (4), 1–35.
- Feenstra, Robert C, “New product varieties and the measurement of international prices,” *American Economic Review*, 1994, pp. 157–177.
- , Mingzhi Xu, and Alexis Antoniadis, “What is the Price of Tea in China? Goods Prices and Availability in Chinese Cities,” *The Economic Journal*, 2020, 130 (632), 2438–2467.
- Hamilton, Bruce W, “Using Engel’s Law to estimate CPI bias,” *American Economic Review*, 2001, 91 (3), 619–630.
- Handbury, Jessie and David E Weinstein, “Goods Prices and Availability in Cities,” *Review of Economic Studies*, 2015, 82, 258–296.
- Nakamura, Emi, Jón Steinsson, and Miao Liu, “Are Chinese growth and inflation too smooth? Evidence from Engel curves,” *American Economic Journal: Macroeconomics*, 2016, 8 (3), 113–44.
- Nordhaus, William D, “Do Real-Output and Real-Wage Measures Capture Reality? The History of Lighting Suggests Not.,” in Timothy F. Bresnahan and Robert J. Gordon, eds., *The Economics of New Goods*, University of Chicago Press, 1997, pp. 29–66.
- Simonovska, Ina, “Income differences and prices of tradables: Insights from an online retailer,” *Review of Economic Studies*, 2015, 82 (4), 1612–1656.