Using the Retail Distribution to Impute Expenditure Shares*

Alexis Antoniades
Georgetown University
Robert C. Feenstra
University of California, Davis
and National Bureau of Economic Research
Mingzhi (Jimmy) Xu
INSE at Peking University
and National Bureau of Economic Research

April 20, 2019

Abstract

Online price data provide a new and rich source of information. But in the absence of information on quantities or expenditure, researchers will treat equally the prices and price behavior across all products within a product group. In doing so, they introduce significant measurement error and potential bias. In this paper, we address this limitation by presenting a simple methodological innovation that allows researchers to impute expenditure when expenditure is not known. With a retail model based on standard assumptions, we show that measures of the retail distribution — which can be computed solely from price observations — provide a good and theoretically consistent proxy for expenditure. Through a series of simulations that use scanner-level price and quantity information on about 85% of the Fast Moving Consumer Goods sold in six Gulf countries, we show that treating all products equally introduces substantial measurement error and bias in the calculation of price stickiness, inflation, and international price differences. But we also show that adding information on the retail distribution reduces measurement error substantially in each of these exercises. Our findings also have important implications for the work of the International Comparison Program.

Key Words: Scanner data, online price data, ICP, retail distribution, measurement error
JEL Classification: C81, E01, E31, E37

*We thank ICP/PPP and ESCOE conference participants. This research was made possible through the support of an NPRP grant from the Qatar National Research Fund. Calculations on US indexes are based on data from the Nielsen Company (US), LLC, and marketing databases provided by the Kilts Center for Marketing Data Center at the University of Chicago Booth School of Business ©The Nielsen Company. Financial support from the China Government Scholarship (CSC) for a doctoral study and NBER post-doctoral fellowship is gratefully acknowledged. Please direct correspondence to: Alexis Antoniades, Georgetown University, aa658@georgetown.edu.

1 Introduction

Online price data offer a promising and rich new source of information for informing economic studies. Exemplified by the work of Cavallo and Rigobon (2016) in the Billion Prices Project (BPP), these data are being increasingly used to answer questions on price behavior and market structure, and to assess differences in the cost of living across time and space.[1] Because these data are up to date, easy to obtain, and cover a very large number of retailers, locations, and products, their use in empirical work can be found across several economics disciplines.

For all their promise, these data scraped from retailers' websites face a major limitation: they offer no information on expenditure or quantities. Researchers can get immediate access to the prices of all products sold (and of some not sold) from a retailer's website, but they do not know what consumers actually buy and in what quantities. Because expenditure data are not available, at the product group levels these studies treat each price observation equally by (implicitly) assuming that spending per product is evenly distributed within a product group. But not all goods are made equal. While consumers purchase a variety of products, they exhibit strong preferences for only a small subset of the available products and brands within any narrowly defined product group. The sales distribution is so skewed that, based on our own calculations, total sales of the top 2% of grocery products per product category can account for as much as the total sales of the bottom 96%.[2] For these products, we observe that (i) sales are hundreds of times higher than sales of other products in the same category, (ii) prices are less sticky (i.e., they change two to three times more often), and (iii) price differentials across retailers are much smaller due, perhaps, to the fact that consumers are more likely to assert store expensiveness by comparing prices of the most well-known products rather than random ones. Since a minority of products accounts for the majority of sales, and since prices for these products may behave differently, the tendency to average prices and price behavior across all products introduces significant measurement error and potential bias.

The unavailability of expenditure data at the basic heading level is also an issue that has come up in the work of the International Comparison Program (ICP). Labeled as "the largest and most complex international statistical activity in the world," the ICP aims to measure the cost of living across the world by computing purchasing power parities (PPPs).[3] In the 2011 round, its latest, the ICP collected price data from across 199 countries and regions. These data are used to provide direct

[1]For information on the BPP, see Cavallo and Rigobon (2016). Cavallo and Rigobon (2011) and Cavallo (2018) use the data to study the distribution of price changes. Coibion, Gorodnichenko, and Hong (2015) study sources of price rigidity using online price data, while Cavallo (2013) uses data collected online to compare estimated inflation measures with official statistics. Cavallo (2017) finds that online and offline price data are similar in most countries and have similar behavior patterns. References within those studies provide information on additional work that uses online price data to study price behavior.

[2]Computations are based on sales of Fast Moving Consumer Goods (FMCGs) between 2006 and 2011, inclusive, in Qatar, United Arab Emirates, Oman, Bahrain, Kuwait, and Saudi Arabia. Scanner price and quantity data for 30 product categories provided by Nielsen are used for computations. A detailed description of the data follows in section 3. The results are similar if scanner data from other regions (e.g., US, Canada, EU) are used.

[3]See Vogel and Hamadeh (2013), "World Bank to publish Purchasing Power Parities in December 2013": http://blogs.worldbank.org/opendata/world-bank-publish-purchasing-power-parities-december013 .

comparisons of well-being, compare growth rates by sector, report price levels, and assess poverty rates. Moreover, PPP-based GDP is used by the International Monetary Fund to determine voting rights, quota subscriptions, and financing amounts for its country members. The data are also used by the IMF to produce the World Economic Outlook database. Because in the data collection exercise of the ICP only prices are reported, which are then averaged to produce price aggregates at the basic heading level, measurement error is introduced. Similar to the studies that use online price data, in the ICP no weight information is collected to reflect the quantities of disaggregate products sold, so all products within the same basic heading are treated equally.[4]

In this paper, we propose a simple methodological innovation that enables researchers to impute expenditure shares when only prices are observed so that price observations can be weighted by importance. Our approach, which is easy to implement and does not require additional resources in terms of time and cost, can be applied to the ICP framework and (retroactively) to existing price data collected online, such as those collected by Cavallo and Rigobon (2016) in the BPP. By imputing market shares, we show in the paper that measurement error is reduced substantially in several broad applications and, therefore, we anticipate that the benefits from adopting the proposed methodology will significantly benefit the ICP, scholars, and statistical agencies worldwide looking to use online price data for CPI estimations.

To illustrate how our approach works, consider a typical dataset of prices collected online, either by scraping data from apps that record offline prices across retailers (e.g., Feenstra, Xu, and Antoniades (forthcoming)) or by scraping data from various retailers' websites (e.g., Cavallo and Rigobon (2011); Cavallo (2013); Cavallo and Rigobon (2016); Cavallo (2017, 2018)).[5] Such datasets contain information on prices but not on expenditure, at least not directly. Indirectly, however, these datasets yield important information on expenditure shares from the number of non-missing price quotations (selling points) per product. Because these datasets contain multiple prices per product — reflecting the many outlets where a product is available at — but are not balanced as many products are sold in only some outlets and not in others, we can easily obtain measures of the retail distribution that can then be used to impute expenditure shares. Specifically, using only price information, we can construct a metric for the retailer distribution by dividing the number of outlets carrying a particular product over the total number of outlets in our sample. We call this metric the *numeric distribution* (*ND*). We can also construct a measure of the distribution that takes into account the size of each retailer, where the number of price observations at a given point in time (i.e., products) per retailer is used to proxy for size. We call this metric the *weighted distribution* (*WD*).

We now have sufficient information to impute market shares for each product by assuming an exponential relation between market share and product distribution. That is, we assume that products available at more retailers have higher sales and that any additional distribution point gained

[4]In both the ICP and in studies that use online data, expenditure data from surveys and the CPI are used across categories to aggregate data up. Our focus here is the aggregation that takes place within a basic heading where no weight is applied.

[5]Similarly, one may also consider a dataset of prices collected by the ICP through surveying of local retailers.

raises market share by more than a point. We discuss the parameters used in the exponential relation later in the paper. These imputed market shares can then be used to weight price observations and reduce measurement errors in calculations, estimations, or both.

To corroborate the recommendation of using retail distribution to approximate product sales when sales are absent, we present a micro-founded theory of retailers that generates the observed convexity in the sales-distribution measure relation. Our model, which is based on standard assumptions (Hottman, Redding, and Weinstein (2016); Feenstra, Xu, and Antoniades (forthcoming)), provides a theoretical framework that accounts for the convex relation and shows it to be robust to alternative market structure settings, as long as different retailers charge manufacturers a different stocking fee.[6] This convexity — well documented in the marketing literature for over half a century — implies that simply knowing the share of outlets carrying a product may be sufficient to allow researchers to impute market share, even when expenditure or quantities are unknown.

We evaluate the performance of our approach in terms of reducing measurement error by considering three important applications: (i) measuring the frequency and magnitude of price changes, (ii) measuring inflation, and (iii) measuring international price differences. These applications are chosen because they cover a large, important, diverse, and an active body of work that is of great interest to academics, policy makers, and practitioners. For these exercises, we use scanner data on Fast Moving Consumer Goods (FMCGs) sold across the Gulf Cooperation Council (GCC) countries between 2006 and 2011.

First, we compute or estimate the measures of interest using actual prices *and* quantities. We set the outcome of this estimation as the benchmark because it is based on the prices consumers pay and the quantities they purchase. Next, we repeat the estimation by ignoring the expenditure information so that all products are treated (and weighted) equally. This approach mimics the approach taken by researchers working with online and ICP data. By comparing the new estimate with the benchmark, we are able to identify and quantify measurement error when expenditure or quantities are unknown and all products within a product group are treated equally. Finally, to test whether our proposed methodology of imputing market shares from price observations works in reducing measurement error, we repeat the exercise one last time, but this time each observation is weighted by an imputed market share derived from estimating retail distribution from the price data.

In the first application, we show that using only prices understates the true frequency of price changes by 25%. This happens because, as we document here, products with higher sales experience more frequent price changes. When only price data are used to compute the frequency of price changes, the prices of a handful of goods that account for the majority of sales and experience frequent price changes are marginalized by the vast number of all other goods with infrequent price changes that account for very few sales. To correct for this measurement error, we repeat the exercise by ignoring again expenditure information (which would not have been known to researchers

[6] The heterogeneity in the stocking fee is supported by various evidence in the marking literature, such as Rao and Mahi (2003); Kuksov and Pazgal (2007).

working with online price data), but instead we use information on retail distribution collected from price data to proxy expenditure shares. This approach reduces measurement error by 71%.

In the second application, we compute inflation rates for each of the GCC countries between January 2006 and December 2011 for FMCG products. By not using expenditure information, we understate inflation in each country by about 30% over the six-year period. Specifically, the democratic measure of inflation yields 4.4 percentage points lower inflation over the period than the plutocratic measure that takes into account actual expenditure by consumers. However, when we compute inflation using prices weighted by an imputed measure of expenditure share derived from retail distribution metrics, inflation is no longer understated: the downward bias/measurement error is reduced by 73%.

In the third and final application, we compute PPPs for the Gulf countries. Using information from the confidential World Bank ICP survey used to collect prices in 2011, we employ scanner data to simulate the ICP exercise without expenditure information, with expenditure information, and with information on prices and distribution but not on expenditure. An interesting aspect of this exercise is that by setting a lab-type experiment, we are also able to consider and evaluate several decision rules that are relevant to the ICP. For instance, we experiment with altering the number of outlets surveyed (10, 20, and 50) in each country. We also consider alternative practical rules when two or more items at a store fit the same product definition provided in the ICP product list (for example, take the minimum, maximum, average, median, and most-important-item prices or a random price among all products that fit the definition).[7] We elaborate more on all these important aspects of the exercise in section 5. In terms of measurement error, we find that using only prices at the basic heading level and excluding information on expenditure vastly overstates actual price differences across the GCC countries. Specifically, while we estimate prices to differ by 6% on average among the GCC countries when both prices and expenditure information are included in the estimation, we estimate differences to be as high as 18% when expenditure information is excluded from the calculations. In contrast, when information on the numeric distribution and weighted distribution is used, we find prices to differ by 9% and 7%, respectively, hence reducing measurement error by 75% over the case in which only prices are used.

To summarize, in the absence of any information on quantities or expenditure, using data on retail distribution provides a very good proxy for expenditure. Even with noisy data on retail distribution, the convexity between distribution and market share allows us to successfully separate the most important products from the rest. And as the applications above confirm, the returns of such strategy in terms of reducing measurement error and potential bias are substantial.

In the following section, we review the literature on retailer distribution and market share, and illustrate their relation through a simple exercise using prices collected online. In section 3, we present the data for analysis and discuss the theoretical framework in section 4. We devote section 5 to

[7] In practice, the price of the most-important-item is the price of the product with the largest expenditure share. The most-important-item can also be chosen based on the amount of shelf space devoted to it.

applications using our proposed methodology, and we conclude in section 6.

2 Retail Distribution and Market Share

To measure the retail distribution, the following metrics have been widely adopted in the marketing field (see the references just below):

$$\text{Numeric or Physical Distribution (\%)} = \frac{\text{Number of outlets carrying product}}{\text{Total number of outlets}} \tag{1}$$

$$\text{All Commodity Volume, ACV (\%)} = \frac{\text{Total sales of outlets carrying product}}{\text{Total sales of all outlets}} \tag{2}$$

$$\text{Product Category Volume, PCV (\%)} = \frac{\text{Total category sales of outlets carrying product}}{\text{Total category sales of all outlets}} \tag{3}$$

Numeric distribution (*ND*), also known as physical distribution, reports the share of outlets carrying a particular product. It is the least data-intense measure of the three metrics, but it does not distinguish between stores with high sales and low sales. *All commodity volume* (*ACV*) and *product category volume* (*PCV*) take into account variation in store size, but require more data, namely, expenditure.

Several studies in the marketing literature have found strong evidence of a convex relation between the retail distribution and market share, in both the cross section and the time series. Nuttall (1965) studied confectioneries; Mercer (1992) cigarettes in England and Scotland; Farris, Olver, and De Kluyver (1989) tortilla chips and instant coffee in the US; Borin, Vranken, and Farris (1991) shampoo in Japan. In 1995, Reibstein and Farris (1995) used scanner data from the IRI's 1988 Info Supermarket Review to test for convexity in 12 randomly chosen US grocery store categories. They first sketched a theoretical outline to provide some conceptual foundation for the hypothesized convex relation.[8] They then tested the following logistic function that resulted from their model:

$$MS = \beta_0 \times \frac{ACV^{\beta_1}}{(1 - ACV)^{\beta_2}} \tag{4}$$

For all categories but frozen pizza, they confirmed that a convex relation characterizes retail distribution and market share in the cross section. In the time series, the evidence was not as strong. A decade later, Kruger and Harper (2006) from IRI expanded the Reibstein and Farris (1995) analysis by testing for the presence of convexity in 263 US product categories and 817 product groups over a period of 22 quarters between 2000:Q1 and 2005:Q2. They found evidence of convexity in 95% of

[8] Reibstein and Farris (1995) attributed convexity to the presence of customer loyalty, to search costs, and to uncompromised choice from the unavailability of competing brands. According to their model, distribution gives access to consumers who are loyal to a particular product/brand, but also to consumers whose preferred product/brand is not available. With perfect brand loyalty or no search costs (or both), the relation between market share and distribution would be linear. But because search costs are non-negative and brand loyalty not perfect, the relation becomes convex.

the cases tested. Similarly to the studies above, we also tested for and confirmed the presence of convexity between market share and retail distribution on 30 product categories of FMCGs for each of the six GCC countries used in our sample.

The evidence suggests that in the absence of any information on expenditure, retail distribution can be used to impute market shares. However, to ensure that such an approach works well with the datasets we have in mind — namely, those that come from online sources and the ICP price surveys — a key prerequisite is to be able to produce reliable measures of retail distribution solely from price data. In the absence of expenditure data, *numeric distribution* (*ND*) can be computed from online price data, but *ACV* and *PCV* cannot. To account for variation in outlet size, we propose an alternative metric that uses outlet product variety as an indicator for outlet size. Specifically, we define the *weighted distribution* (*WD*) as:

$$\text{Weighted Distribution, WD (\%)} = \frac{\text{Total products of outlets carrying product}}{\text{Total products of all outlets}} \tag{5}$$

Counting the number of products offered by a store is a good indicator of its sales. As we show in the data section that follows, large stores carry more products, more brands, and more products per brand. Therefore, even if sales are not known, knowing the number of products sold is enough to help us distinguish between large and small stores.[9]

2.1 An Illustration

Do online price data exhibit the same convex relation between market share and retail distribution as documented in the studies using scanner data? In theory, they should because online price data are based on products that are available in stores. However, as far as we know, no study has formally put this hypothesis to the test.[10] Before we can apply our methodology, we must then first check whether the convexity exists in online data. To check, we conduct a simple test that matches online price data with actual expenditure data. The online price data allow us to construct the two measures of retail distribution discussed above, and the scanner data allow us to check whether the relation between the constructed retail distribution measures and actual market shares is convex. Our data source for this exercise is the online price data for toothpaste, personal wash items, shampoo, and laundry detergent products sold across 22 cities in China in 2014, which come from a mobile application that lists the (offline) prices of a large number of products available in several stores in each city.[11]

The data contain no information on expenditure. The number of retailers included in the dataset

[9] In the ICP price surveys it is not feasible to report the total number of products sold in each store. Nonetheless, price auditors can still report information on size based on store type (hypermarket, grocery, self-service) or simply by reporting the number of checkout counters.

[10] Cavallo (2017) does provide the first large-scale study comparing prices between online and offline stores across many countries and finds that they are identical about 72% of the time. But he does not (and cannot due to the lack of expenditure data) test whether in online price data a convex relation exists between retail distribution and market share , which is what we need to confirm here.

[11] More information on the dataset is provided in Feenstra, Xu, and Antoniades (forthcoming).

varies by city and ranges from 3 to 12. No information on retailers inclusion or exclusion exists. Furthermore, we do not know if we can proxy for retailer size with the number of listed products per retailer in the online application. If no price for a given barcode is reported at a particular store, this means either that the product is not sold there or that it is sold but the price is not uploaded. Therefore, counting the number of products per store may not be a good proxy for actual store size. These shortcomings mean that our measures of *ND* and *WD* will be very noisy at best. At worse, they will be very poor approximations of actual retail distribution and will make it hard for us to identify a convex relation from our sample, even if one exists in the population. In addition to the scraped price data from the phone app, we also use scanner data for these four categories provided by Nielsen China. The scanner data provide barcode-level price and quantity information for each city, but prices are averaged across time (weeks) and space (retailers) in each city. Therefore, we can compute market (expenditure) share for each barcode but not retail distribution.

We use the mobile app dataset to compute the two measures of retail distribution, and we use the Nielsen dataset to compute market share for each product. We then merge the two datasets, and for the majority of the barcodes found in the mobile application, we can now observe both distribution and market share information. Finally, we allocate the products into bins based on market share and take the median retail distribution (*ND* or *WD*) across all products in each market share bin.

Scatter plots of distribution and market share for each product category are reported, first using the numeric distribution (*ND*) (Figure 1, panel (a)) and then using the weighted distribution (*WD*) (Figure 1, panel (b)). The figure confirms that the relation between market share and computed retail distribution is convex, despite the data issues discussed above that may have compromised our ability to measure retail distribution accurately. The implications are significant. Scholars can exploit the convex relation to impute market share from distribution information obtained only from prices, and they can then use imputed market shares to weight products by importance. We elaborate more on this next as we consider three important applications. But first, we present the data that are used for the rest of the paper.

3 Data Description

So far we have (i) drawn attention to the convex relation that exists between retail distribution and market share, (ii) claimed that the relation can be exploited to impute market share from online (and ICP) data so that measurement error can be reduced, and (iii) shown that the convex relation exists in online price data. Next, we validate our claim that the proposed approach reduces measurement error by applying it to three important applications. We first discuss the data and then the applications in more detail.

The data come from AC Nielsen and cover sales of FMCGs in six GCC countries: Bahrain, Qatar, Oman, Kuwait, United Arab Emirates, and Saudi Arabia. The price and quantity information for thousands of products (barcodes) across 30 product categories between January 2006 and December

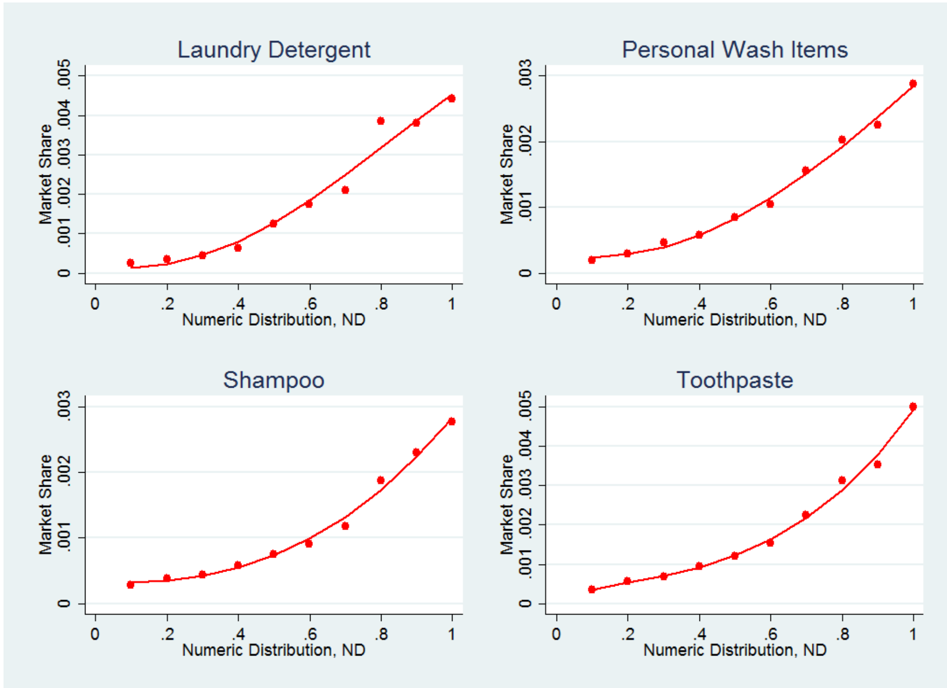

(a) Numeric Distribution

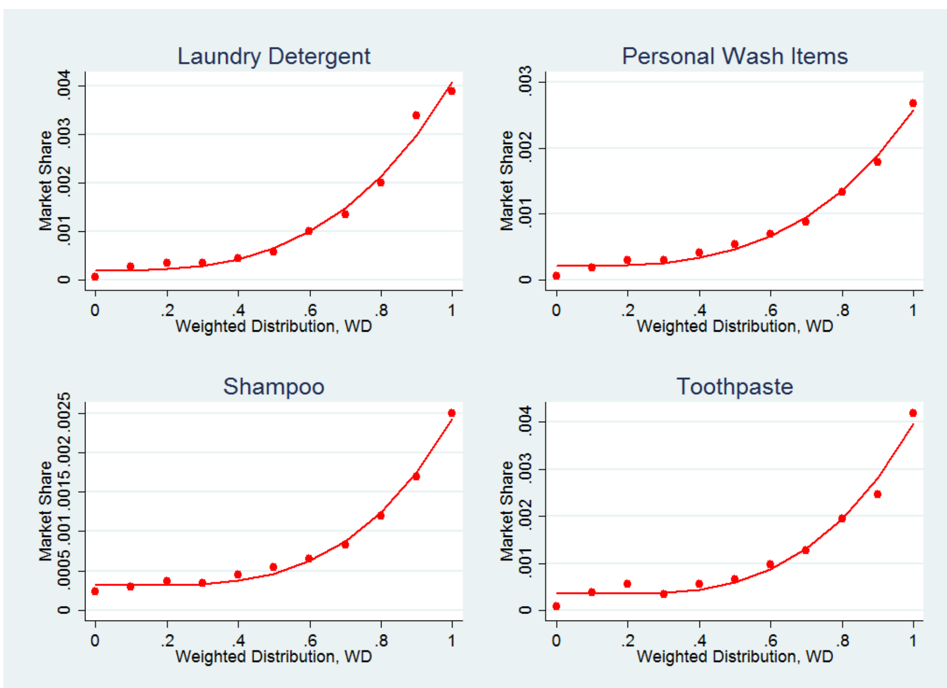

(b) Weighted Distribution

Note: We use price data from an online app in China to compute retail distribution (*ND* and *WD*) and Nielsen scanner data to compute market shares for products in the laundry detergent, shampoo, personal wash items, and toothpaste categories. Products are allocated into bins based on retail distribution (x-axis) and the average market share of all products within a bin is plotted on the graph, against measures of the numeric distribution (*ND*; panel (a)), and the weighted distribution (*WD*; panel (b)). For more information on the data, see Feenstra, Xu, and Antoniades (forthcoming).

Figure 1: Measures of Retail Distribution and Market Share, Chinese Data

2011 are provided.[12] The frequency is monthly or bi-monthly, and according to Nielsen, these data cover about 85% of all the FMCGs consumed in the GCC countries.

Table 1: Descriptive Statistics for the GCC Nielsen Data

Country	Categories	Products	Brands	Retailers	Start Date	End Date
Bahrain	30	24,259	2,168	311	6-Jan	11-Dec
Kuwait	30	37,660	3,052	285	6-Jan	11-Dec
Oman	30	40,165	3,442	614	6-Jan	11-Dec
Qatar	30	24,150	1,474	267	6-Jan	11-Dec
Saudi Arabia	30	34,447	3,030	3,398	6-Jan	11-Dec
United Arab Emirates	30	43,038	3,650	976	6-Jan	11-Dec

Note: Data are provided by Nielsen and cover sales of FMCGs between 2006 and 2011 in the GCC countries. The frequency of the data is monthly or bi-monthly, and price and quantity information is given at the level of the retailer in each period.

Three important characteristics of the dataset are worth highlighting. First, data are provided for each store, across thousands of stores. By analyzing the data at the store level, we are able to provide stylized facts on retailers. We are also able to accurately measure the retail distribution. Second, prices are reported during the day of the audit in each period. They are not averaged across all days within a period. This practice allows us to measure the frequency and magnitude of price changes across periods without measurement error.[13] Third, most of the products we study are imported, many of the consumers in these markets are expatriates (as many as 85% in Kuwait, UAE, and Qatar), and several international retailers operate in the markets. This suggests that the findings we present below can be generalized with some confidence to other economies outside the Gulf.[14]

Descriptive statistics for the dataset are provided in Table 1. In total, the dataset provides price and quantity information on 203,719 products sold in 5,851 outlets over a period of six years. Qatar and Bahrain are the smallest economies in terms of population, and Saudi Arabia the largest. The majority of products do not exist across all periods and all outlets. In each category, we observe that sales of FMCGs are highly concentrated in a handful of products per category.

[12] The categories are beans, blades, bouillon, cereals, cheese, chewing gum, chocolate, cigarettes, cooking oil, carbonated soft drinks, deodorants, detergents, dish wash, energy drinks, fabric conditioners, insecticides, juices, liquid cordials, male grooming, milk, milk powder, powder soft drink, shampoo, skincare, skin cleansing, sun care, tea, toothbrush, toothpaste, and water.

[13] In many cases, Nielsen provides price data that are averaged across the period of interest (e.g., week or month). This practice prevents researchers from accurately measuring the frequency and magnitude of price changes. For more on the time averaging measurement error in the Nielsen data, see Cavallo and Rigobon (2016).

[14] Antoniades and Zaniboni (2016) make a similar point. The authors use a subset of this dataset to study retailers' pass-through into consumer prices in the United Arab Emirates. They measure one-year pass-through to be 20%, which they find to be similar to estimates obtained using micro data in advanced economies.

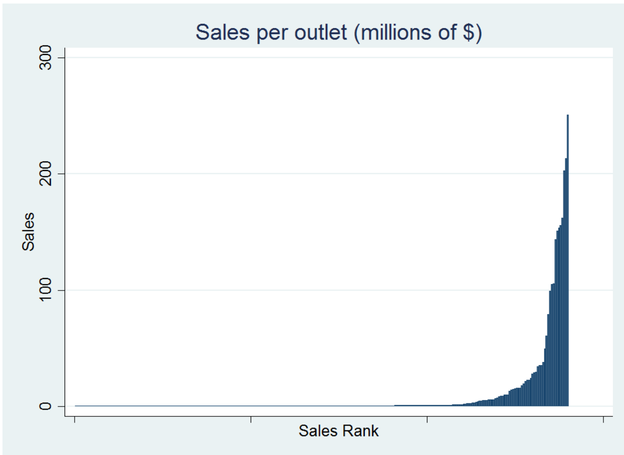

(a)

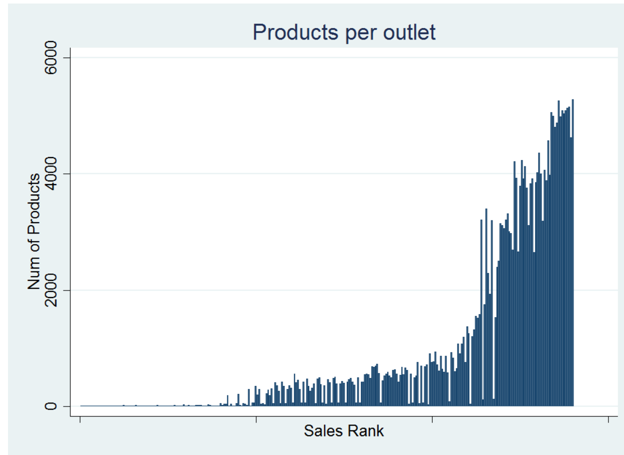

(b)

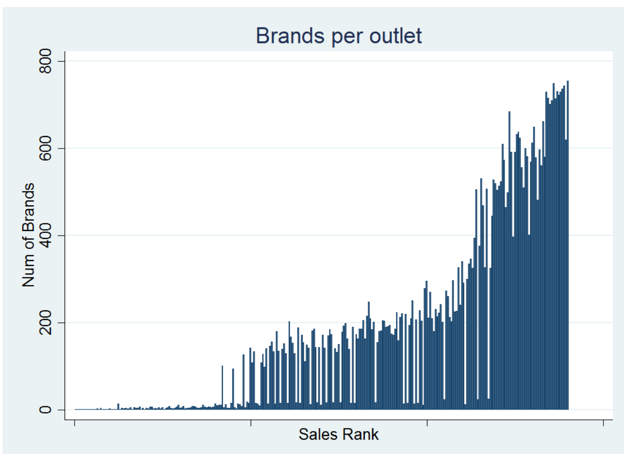

(c)

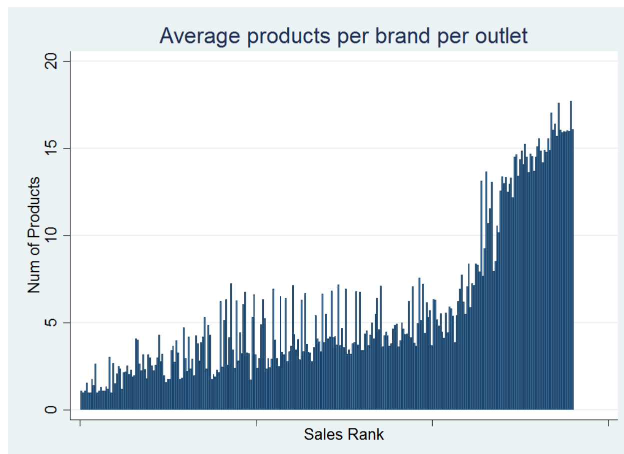

(d)

Note: Retail outlets in UAE are ranked based on average monthly sales (panel (a)). Monthly sales are computed from Nielsen scanner data across 30 product categories between January 2006 and December 2011. Products per outlet (panel (b)), brands per outlet (panel (c)), and average products per brand per outlet (panel (d)) are shown for each retail outlet while maintaining the ranking. The horizon axis denotes the store ranks based on store sales, in order of increasing sales from left to right.

Figure 2: Facts on Retailers' Heterogeneity, UAE

In the previous section, we made the assertion that counting the number of available products is a good proxy for store size. To provide support for this assertion, in Figure 2 we plot average monthly sales by store in UAE in US dollars on the vertical axis with store ranking on the horizontal axis. Rankings are based on sales, and stores are ranked from smallest to largest. Out of 976 outlets available in the sample, about a couple of dozen stores account for the majority of sales. The rest are small outlets with low sales. Next, we plot the average number of products, brands, and products per brand sold by each store each month, while maintaining the size ordering of outlets (Figure 2, remaining quadrants). We conclude that substantial variation in store size exists within a country and that large stores offer more products, brands, and products (varieties) per brand. The results for the other five countries are identical and omitted for brevity.[15]

4 Theoretical Foundation of Convexity

Having reviewed evidence from the literature and the data for the convex relation between market share and retail distribution, in this section we propose a theoretical model that provides some micro-foundations to account for such a pattern. The theory, based on a standard set of assumptions, characterizes both manufacturers' and retailers' decisions under alternative market structure settings. We show that assuming heterogeneity in the fixed costs paid by manufacturers to retailers is sufficient to generate the convex relation between sales and the distribution measure, which is robust to alternative market structures.

The Consumer

We study a closed economy, but our analyis could readily be extended to an open economy. The consumer's utility depends on the consumption of differentiated varieties, which are purchased from a set of retailers. Each manufacturer produces a single variety for simplicity, and they choose to which retailers they sell their product. We index manufacturers with j or ϕ, and retailers with r. The utility function follows Hottman, Redding, and Weinstein (2016); Feenstra, Xu, and Antoniades (forthcoming), and is assumed to be nested CES, as follows:

$$U = \left(\int_{r \in \Omega} X_r^{\frac{\eta-1}{\eta}} \right)^{\frac{\eta}{\eta-1}}, \; X_r = \left(\int_{j \in J_r} x_{rj}^{\frac{\sigma-1}{\sigma}} \right)^{\frac{\sigma}{\sigma-1}}, \; \sigma > \eta \tag{6}$$

where η and σ denote the elasticity of substitution across retailers and across varieties within retailers. The collection of varieties within retailer r is J_r, and the set of retailers is denoted as Ω. The demand for variety j served in r is,

$$x_{rj} = p_{rj}^{-\sigma} P_r^{\sigma-\eta} P^{\eta-1} Y. \tag{7}$$

The term $P_r^{\sigma-\eta} P^{\eta-1} Y$ reflects the total demand (in terms of market size) of retailer r, which will depend on the economy-wide total income (Y), as well as the price indexes given by:

$$P_r = \left(\int_{j \in J_r} p_{rj}^{1-\sigma} \right)^{\frac{1}{1-\sigma}}, \; P = \left(\int_{r \in \Omega} P_r^{1-\eta} \right)^{\frac{1}{1-\eta}} \tag{8}$$

The Suppliers

Two types of firms function as suppliers: manufacturers and retailers. Each manufacturer produces a single product and sells it to retailers, as already noted, while consumers purchase consumption

[15] In our data, we also find evidence that product prices fall as outlet size increases. That is, we find that larger stores have lower prices.

goods from retailers. In the subsequent analysis, both retailers and manufacturers are assumed to be profit maximizers that employ their optimal strategies simultaneously.

Manufacturers and retailers are heterogeneous in the model, and we denote them as ϕ and r, respectively. The manufacturers differ in productivity (ϕ). Retailers differ in terms of the fixed costs ($f_r > 0$) that they charge to manufacturers, which we call the *slotting fee*. For simplicity, we treat the slotting fee as exogenous (i.e., not chosen by retailers) and paid by manufacturers, so that it becomes a fixed cost for manufacturers. [16] Both manufacturers' productivity and retailers' slotting fees are exogenous in the model. The total measure of manufacturers is M, and their productivities are *i.i.d.* distributed with a $c.d.f.$ of $G(\phi)$.

There are many retailers, and the measure of retailers serving the economy is fixed and denoted by N. We line up retailers and rank them in order of their slotting fees from low to high. To simplify the following analysis, we treat retailers as if they are continuous, and we index them in relative terms (i.e., $r \in [0,1]$) where a retailer of $r = 0$ has the lowest fixed cost and a retailer of $r = 1$ has the highest fixed cost ($\partial f_r / \partial r > 0$). We study the equilibrium in which manufacturers will prefer to sell in retailers with lower fixed costs. That is, we assume that manufacturers go to retailers with the lowest fixed costs first and then to those with increasing higher fixed costs until it is no longer profitable to sell to other retailers. Let $r_\phi \in [0,1]$ denote the scope of the retailers to which manufacturer ϕ is possibly able to sell, and we formalize this assumption as follows.[17]

Assumption 1: *The manufacturer lines up retailers according to their slotting fees and sells to the lower-slotting-fee retailers $[0, r_\phi]$ until the manufacturer's additional profit goes to zero at r_ϕ.*

The *numeric distribution* of the product produced by manufacturer ϕ is exactly $r_\phi \equiv N_\phi / N$, where N_ϕ denotes the largest discrete index of retailers that manufacturer ϕ could serve. We assume retailers and manufacturers make their optimal decisions simultaneously to maximize profits; that is, retailers set retail prices taking wholesale prices as given, and manufacturers choose wholesale prices taking retailers' markups as given.

Manufacturers observe the pricing rule of the retailers and are aware that their pricing rule will affect the market outcome. Given the production efficiency ϕ, the marginal cost of this manufacturer is w/ϕ where w is labor wages. Manufacturer ϕ maximizes profit by choosing its prices $q_{r\phi}$ for the retailers $[0, r_\phi]$ to which it sells its product:

$$\pi_\phi \equiv \max_{q_{r\phi}} \int_0^{r_\phi} \pi_{r\phi} dr = \max_{q_{r\phi}} \int_0^{r_\phi} \left(q_{r\phi} x_{r\phi} - f_r \right) dr, \tag{9}$$

[16] The marketing literature refers to the fixed cost (f_r) as the *slotting fee* (or fixed trade spending), a fee charged to manufacturers by retailers in order to have manufacturers' products placed on retailers' shelves. It has also been well established that slotting fees differ across retailers (Rao and Mahi (2003); Kuksov and Pazgal (2007)). Retailers' slotting fees could reflect some other factors out of their control that affect manufacturers' willingness to sell goods in them (e.g., poor locations, traffic or logistics could increase such fixed costs), and we assume those obstacles are borne by the manufacturers.

[17] In the general scenario, multiple equilibria are possible, and we need this assumption for tractability in the analysis of the model.

where $\pi_{r\phi}$ is manufacturer ϕ's profit collected from retailer r, $x_{r\phi}$ is the demand for product ϕ by retailer r, f_r denotes the entry fee charged by retailer r to allow a manufacturer to sell on its shelves, and r_ϕ indicates the scope of the retailers that manufacturer ϕ is possibly able to serve. Manufacturers set wholesale prices taking retailers; markups as given. As shown in (9), $q_{r\phi}$ denotes the wholesale price, and the final price paid by consumers would be $p_{r\phi} = \mu_r q_{r\phi}$ where μ_r is the markup charged by retailer r. The pricing rule of retailers is specified later, and manufacturers take it as given and are aware that their wholesale prices will affect the market price $p_{r\phi}$. The first order condition with respect to $q_{r\phi}$ solves for the optimal prices:

$$q_{r\phi} = \frac{\sigma}{\sigma - 1}\frac{w}{\phi}, \ \forall r \in [0, r_\phi]. \tag{10}$$

We solve for the profit generated by selling to retailer r as:

$$\pi_{r\phi} = \frac{1}{\sigma - 1}\left(\frac{\sigma - 1}{\sigma}\right)^{\sigma} YP^{\eta-1}w^{1-\sigma} \times \mu_r^{-\sigma}P_r^{\sigma-\eta} \times \phi^{\sigma-1} - f_r.$$

The cutoff productivity ϕ_r of the manufacturer just able to make a profit by selling to retailer r while paying the slotting fee f_r is computed by setting $\pi_{r\phi}$ equal to zero:

$$\frac{1}{\sigma - 1}\left(\frac{\sigma - 1}{\sigma}\right)^{\sigma} YP^{\eta-1}w^{1-\sigma} \times \mu_r^{-\sigma}P_r^{\sigma-\eta} \times \phi_r^{\sigma-1} = f_r. \tag{11}$$

As multiple equilibria are possible in the general scenario, we employ Assumption 1 to focus on the equilibrium in which the retailers embedded with lower slotting fees always host more manufacturers (i.e., if $f_{r_1} < f_{r_2}$ then $\phi_{r_1} < \phi_{r2}$).[18] Then in equilibrium, only manufacturers with productivity ϕ greater than ϕ_r sell to retailer r. With the mass of manufacturers denoted as M, the measure of manufacturers serving retailer r is $M\left(1 - G\left(\phi_r\right)\right)$.

We are now more specific about the distribution of manufacturers' productivity ϕ in the economy. We assume that ϕ follows a Pareto distribution with a *c.d.f.* of $G\left(\phi\right) = 1 - \left(\bar{\phi}/\phi\right)^k, \ \phi \geq \bar{\phi}$, with $k > \sigma - 1$. We can use this distribution to solve for the price index P_r as defined in (8):

$$\begin{aligned} P_r &= \left[M\int_{\phi_r}^{+\infty} p_{r\phi}^{1-\sigma} g\left(\phi\right) d\phi\right]^{1/(1-\sigma)} \\ &= \frac{\sigma}{\sigma - 1}\left(\frac{k}{k - \sigma + 1}\right)^{\frac{1}{1-\sigma}} \bar{\phi}^{\frac{k}{1-\sigma}} M^{\frac{1}{1-\sigma}} w \times \mu_r \phi_r^{\frac{k-\sigma+1}{\sigma-1}}, \end{aligned} \tag{12}$$

Substituting (12) back to (11), we could solve the cutoff of productivity ϕ_r:

$$\phi_r^{\epsilon_1} = A_1 f_r \mu_r^{\eta}, \tag{13}$$

[18] In the equilibrium we studied, the market power (markup) of low-slotting-fee supermarkets cannot be large enough to overturn the advantage for manufacturers to sell products in them (due to low slotting fees). Otherwise, there may not exist a positively monotone pattern between (ϕ_r, f_r). That is, manufacturers may choose supermarkets with slightly higher slotting fees to avoid the profit reduction resulting from the high markup of a low-slotting-fee supermarket.

where ϵ_1 and A_1 are defined as:

$$\epsilon_1 \equiv \frac{(k-1)(\sigma-\eta)+\sigma\eta+1}{\sigma-1},$$
$$A_1 \equiv (\sigma-1)\left(\frac{\sigma}{\sigma-1}\right)^{\eta}\left(\frac{k}{k-\sigma+1}\right)^{\frac{\sigma-\eta}{\sigma-1}}\bar{\phi}^{\frac{k(\sigma-\eta)}{\sigma-1}}M^{\frac{\sigma-\eta}{\sigma-1}}w^{\eta-1}P^{1-\eta}Y^{-1}.$$

The observed sales ($p_{r\phi}x_{r\phi}$) of product ϕ through retailer r would be:

$$R_{r\phi} = \sigma\phi^{\sigma-1}A_1^{\frac{1-\sigma}{\epsilon_1}}f_r^{\epsilon_2}\mu_r^{1-\frac{\eta(\sigma-1)}{\epsilon_1}}, \tag{14}$$

where the equality uses (13). It can be easily shown that $\epsilon_2 \equiv 1-\frac{\sigma-1}{\epsilon_1} > 0$ given the imposed restriction that $k > \sigma - 1$. As the last step, we derive the total sales of product ϕ in the economy, where we also change notation from r_ϕ to n to denote the numeric distribution:

$$\begin{aligned} R_\phi &= \int_0^{r_\phi} R_{r\phi}dr \\ &= \int_0^{n} R_{r\phi}dr \\ &= \sigma\phi^{\sigma-1}A_1^{\frac{1-\sigma}{\epsilon_1}}\int_0^{n} f_r^{\epsilon_2}\mu_r^{1-\frac{\eta(\sigma-1)}{\epsilon_1}}dr. \end{aligned} \tag{15}$$

Proposition 1. *Under Assumption 1, and if retailers charge the same markups to consumers (i.e., $\mu_r = \mu, \forall r \in [0,1]$), product sales are convex in the numeric distribution, defined as $n \equiv N_\phi/N$.*

Proposition 1 is easily proved by taking the first and second derivatives of product sales R_ϕ with respect to *numeric distribution n* (see Appendix A1). It corresponds to a preliminary scenario in which retailers do not take their market shares into consideration when setting their retail prices, i.e. they do not see themselves as *multi-product* sellers. We next examine the case in which retailers optimally charge differing markups.

Product Sales with Variable Retailer Markups

In the more general case, the markups charged by retailers will differ. Retailers choose their prices for the range of products, taking into account that a change in any prices will affect their market shares for all their products. We first consider the case in which retailers fail to realize that the pricing rules could also affect the entry of manufacturers and hence profits. Let us call this case a "shortsighted" retailer. Manufacturers have to overcome the exogenous slotting fee to sell to a retailer, which implies that only manufacturers with productivity above the threshold can sell in that retailer. The profit maximization problem for retailer r is:

$$\max_{p_{rj},j\in J_r}\left[\sum_{j\in J_r}\left(p_{rj}-q_{rj}\right)x_{rj}\right]\Leftrightarrow\max_{p_{r\phi},\phi>\phi_r}\left[M\int_{\phi_r}^{+\infty}\left(p_{r\phi}-q_{r\phi}\right)x_{r\phi}g(\phi)d\phi\right] \tag{16}$$

where $p_{r\phi}$ is the retail price and $q_{r\phi}$ is the wholesale price of product ϕ. This problem is solved in Feenstra, Xu, and Antoniades (forthcoming), and the pricing rule of retailer r is:

$$p_{r\phi}=\mu_r q_{r\phi},\ \ with\ \ \mu_r\equiv 1+\frac{1}{(\eta-1)(1-s_r)},\ \forall\phi>\phi_r, \tag{17}$$

where s_r is the market share of retailer r over all its products sold and μ_r is retailer r's markup, which is equal across products sold by that retailer. Bigger retailers (larger s_r) would charge a higher markup.

Proposition 2. *When retailers are shortsighted, retailers' markups positively depend on their market shares as in* (17)*, and product sales are convex in the numeric distribution if:*

$$k\geq 1+\frac{\eta(\sigma-1)^2-\sigma\eta-1}{\sigma-\eta}.$$

The proof of Proposition 2 is in Appendix A2, and the above condition is sufficient for convexity. For cases outside the range as indicated in Proposition 2, we find that the convex relation between market sales and the *numeric distribution* still holds empirically, as we shall demonstrate below.

Next, we study the case of farsighted retailers, that is, retailers who are aware that their retail prices would affect both the intensive margin of sales (the sales conditional on the measure of manufacturers selling in those retailers) and the extensive margin of sales (the measure of the manufacturers selling in those retailers). Retailer r chooses a retail markup to maximize profit:

$$\max_{\mu_r}\left[M\int_{\phi_r}^{+\infty}\left(p_{r\phi}-q_{r\phi}\right)x_{r\phi}g(\phi)d\phi\right].$$

Given that $p_{r\phi}=\mu_r q_{r\phi}$ and$p_{r\phi}q_{r\phi}=\sigma f_r\phi_r^{1-\sigma}\phi^{\sigma-1}$, with $g(\phi)=k\bar{\phi}^k\phi^{-k-1}$, we can integrate retailer r's profit to obtain:

$$\max_{\mu_r}\left[\frac{\sigma kM\bar{\phi}^k}{k-\sigma+1}f_r\left(\mu_r-1\right)\phi_r^{-k}\right],$$

which could be further simplified given (13) as:

$$\max_{\mu_r}\left[\frac{\sigma kM\bar{\phi}^k A_1^{-\frac{k}{\epsilon_1}}}{k-\sigma+1}f_r^{1-\frac{k}{\epsilon_1}}\left(\mu_r-1\right)\mu_r^{-\frac{\eta k}{\epsilon_1}}\right]. \tag{18}$$

The first order condition of (18) with respect to μ_r implies that:[19]

$$\mu_r=1+\frac{1}{\eta k/\epsilon_1\left[\eta-(\eta-1)s_r\right]-1}, \tag{19}$$

[19]The derivation also takes into account that $\partial\ln P/\partial\ln\mu_r$=$\partial\ln P/\partial\ln P_r=s_r$, given that $\partial\ln P_r/\partial\ln\mu_r=1$.

where $\epsilon_1 \equiv \frac{(k-1)(\sigma-\eta)+\sigma\eta+1}{\sigma-1}$. To guarantee a meaningful markup $\mu_r > 1$, we require $\eta k/\epsilon_1 > 1$, which implies that:[20]

$$k > 1 + \frac{\eta + 1}{\sigma(\eta - 1)}. \tag{20}$$

Similar to the pricing rule for shortsighted retailers in (17), the markup of a farsighted retailer also positively depends on its market share. Therefore, we derive a proposition similar to Proposition 2.[21]

Proposition 3. *When retailers are farsighted, retailers' markups positively depend on their market shares as in* (19)*, and product sales are convex in the numeric distribution if:*

$$k \geq 1 + \frac{\eta(\sigma-1)^2 - \sigma\eta - 1}{\sigma - \eta}$$

The proof of Proposition 3 follows the similar steps in the proof of Proposition 2. Thus, we have completed the theoretical foundation to explain the observed sales pattern, which however provides with sufficient conditions. Nevertheless, we can go beyond model parameters and develop some inferences about the convex relationship between product sales and numeric distribution based on data.

Proposition 4. *Under Assumption 1, and when markups positively depend on market shares, if retailers' sales rank satisfies $s_{r_2} < s_{r_1}$ for $r_2 > r_1 \in [0,1]$ so that retailers hosting more products also have bigger total sales, product sales are convex in the numeric distribution.*

The proof of Proposition 4 is in Appendix A3. The condition in Proposition 4 that retailers hosting more products also have bigger total sales is not trivial, though it is the case on average in the data (see Figure 2). Conditional on entry, incumbent manufacturers will sell more to overcome higher fixed costs. In the case in which there is a substantial number of big manufacturers, the deterring effect of a high slotting fee on entry would be mitigated. In turn, the high slotting fee would bring more sales that are generated by incumbent manufacturers, and this would potentially break the positive relationship between the number of products a retailer hosts and its total sales.

The Model Simulation

To provide an overview of how well the model generates the convex relation between product market share and retail distribution, we perform a simulation exercise for the case in which retailers are shortsighted.[22] In the simulation, we simulate the sales and *numeric distribution* of a large number of products under three scenarios, and one of them ($k = 16$) corresponds to the case in which restriction of model parameters in Proposition 2 and 3 is satisfied. Our purpose is to demonstrate how our

[20] In the extreme case in which there is only one retailer, the markup is $\mu_r = \eta k/\epsilon_1 / (\eta k/\epsilon_1 - 1)$.

[21] Proposition 3 implicitly assumes that model parameters satisfies (20).

[22] The pattern for farsighted retailers remains similar, as is also discussed in Proposition 3. The detailed procedure for simulation is provided in Appendix A4.

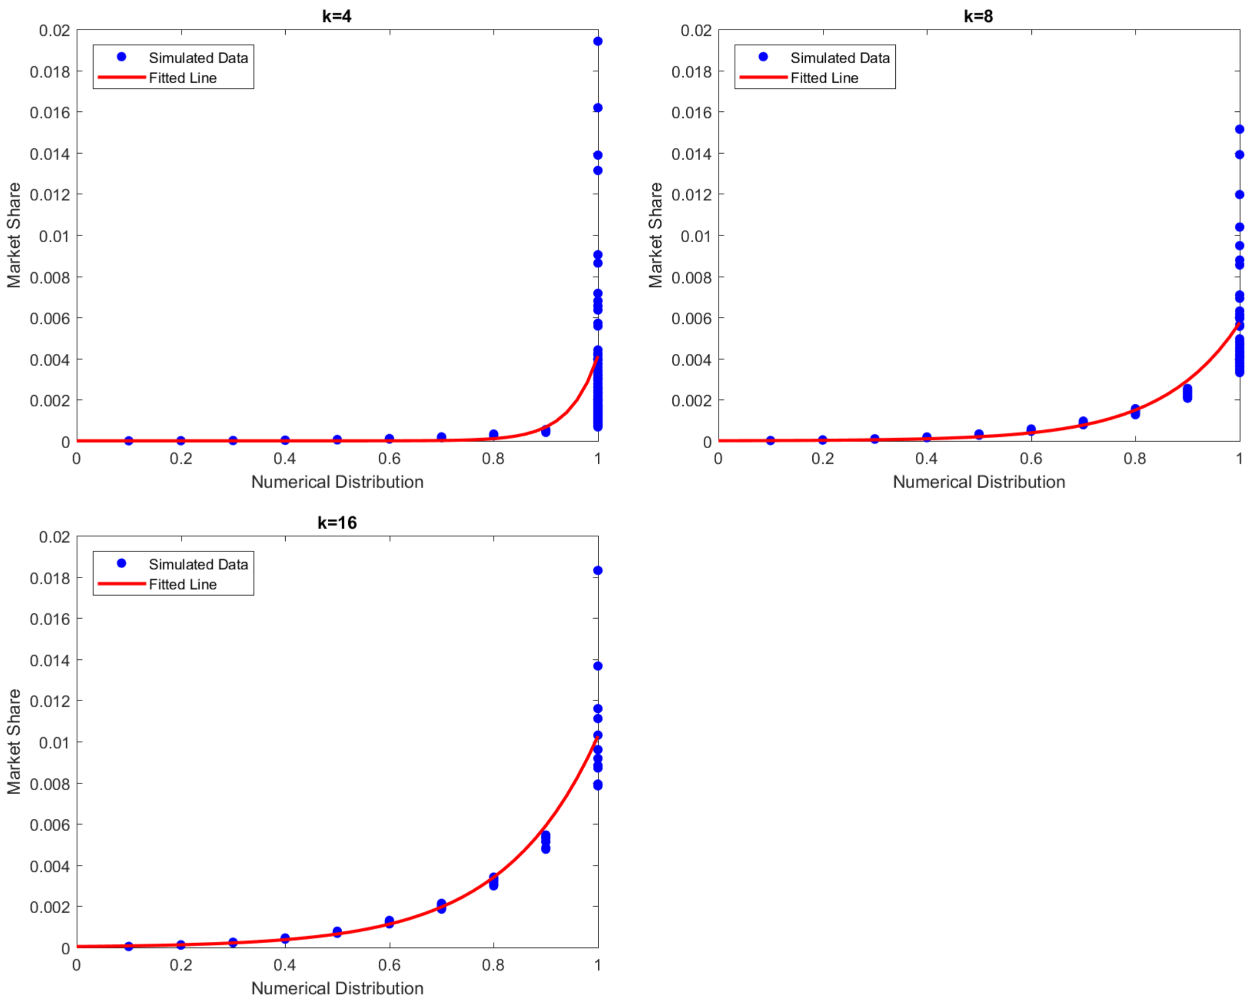

Note: $k = 16$ satisfies parameter restriction in Proposition 2 and 3.

Figure 3: Convexity between Sales and Numeric Distribution

model can replicate the convex relationship between sales and the *numeric distribution*, and investigate whether the convex relationship is robust to various candidate parameters of the distribution of productivity k with the minimum constraint $k > \sigma - 1$.

To give a brief idea of the procedure, setting parameters to satisfy the restriction, we simulate the economy in which consumers, manufacturers and retailers are specified by (6), (9), and (16). In practice, we specify the fixed costs as $f_r = \gamma e^{\theta r}$ ($\gamma > 0$ and $\theta > 1$) and simulate 10,000 draws u from a uniform distribution from 0 to 1. The corresponding Pareto productivity draws are $\phi = (1-u)^{-\frac{1}{k}}\bar{\phi}$. Given the functional forms, we solve the model by solving for the equilibrium retailer markups. Figure 3 presents the simulation results by values of k. In all three scenarios, we observe a convex relationship between product market share and the *numeric distribution*.[23]

[23] Analogously, the alternative measure of the *weighted distribution* could be shown to perform similarly to the *numeric distribution*. As the *numeric distribution* requires less information than the *weighted distribution* in practice, implementing it is more feasible.

To summarize, in this analysis, we present a micro-foundation for the observed convexity in the sales-distribution measure relation. Our model is based on the standard assumptions in the literature. We show that the implied convexity pattern is robust to various market structure settings, as long as the fixed costs incurred by manufacturers to sell in retailers vary across retailers. Our theoretical results further corroborate the robustness of using retail distribution to approximate product sales when they are absent.

5 Applications

5.1 Predicting Product Sales

We proceed to illustrate how we can exploit the relation between retail distribution and market size to obtain a proxy for expenditure when expenditure is not observed but prices are. We apply this to three different applications: (i) measuring the frequency and magnitude of price changes, (ii) measuring price levels and inflation, and (iii) measuring international price differences/purchasing power parities (PPPs).

While each application differs in nature, the core of the exercise is the same and consists of three steps. In the first step, we use actual price and quantity (expenditure) information from the Nielsen dataset to compute or estimate a measure of interest, such as inflation or the frequency of price changes. We set the outcome of this estimation to be the benchmark against which the results of the alternative estimations will be compared.

In the second step, we compute or estimate the same measure of interest, but this time we use only prices and treat all observations equally. This estimation mimics the approach of the ICP and studies employing online price data that lack information on expenditure and treat all products equally. We then compare these estimates to the benchmark case. Any difference between these two measures is due to measurement error (or bias) that arises from being unable to properly weight items by importance. Comparing the two estimates enables us to quantify how important such measurement error may be.

Finally, in the third step, we again exclude any information on quantities and expenditure, but instead we use measures of retail distribution extracted solely from price data to impute market shares. These imputed market shares are then used as weights in the estimations. The outcome of this estimation allows us to test whether our proposed approach, namely, using retail distribution as a proxy for expenditure, reduces measurement error and by how much. We first impute shares using *ND* and then using *WD*.

Table 2: Market Share and Distribution Regression Results

	Dependent variable: ln(market share)			
	GCC		U.S.	
	(1)	(2)	(3)	(4)
Numeric Distribution, ND	4.861***		5.246***	
	(0.001)		(0.004)	
Weighted Distribution, WD		4.984***		5.268******
		(0.000)		(0.003)
Constant	-7.259***	-8.055***	-10.390***	-10.601***
	(0.001)	(0.001)	(0.002)	(0.002)
Observations	105,383,354	105,383,354	2,118,960	2,118,960
R-squared	0.398	0.549	0.356	0.429
Period FE	YES	YES	YES	YES
Country FE	YES	YES	-	-

Note: Nielsen scanner data are used to compute measures of retail distribution and market share. ln(market share) is then regressed on either measure of distribution and on additional controls. The U.S. Nielsen data only have one year in regression. Estimated coefficients from these regressions are used to characterize the convex relation between the two measures in the simulations described in section 5. Standard errors are in parentheses, *** $p<0.01$, ** $p<0.05$, * $p<0.1$.

The first requirement behind our approach is to impute market (expenditure) shares from computed measures of retail distribution, namely, *ND* and *WD*. We do that by exploiting the convex relation between these two variables and imposing the following functional forms:

$$\text{Predicted Market Share} = \exp\left(a + bND\right) \tag{21}$$

or

$$\text{Predicted Market Share} = \exp\left(\tilde{a} + \tilde{b}WD\right) \tag{22}$$

where *ND* and *WD* are measures of the numeric and weighted distributions, respectively.

In general, the coefficients will not be known to the econometrician. Here, we employ the Nielsen data to estimate these coefficients. Specifically, we first measure the numeric distribution, weighted distribution, and market share for each product across all categories, countries, and time periods. Then we pool the data and regress product retail distribution on log market share in order to obtain the coefficients of interest. The results are reported in Table 2. Columns 1 and 2 present the estimates from regressing ln(market share) on *ND* and *WD*, respectively, using GCC data. Period and country fixed effects are also included in the estimation. For comparison purposes, estimates using US Nielsen data for the same product categories are provided. We observe that regardless of whether GCC or US data are used, the convexity coefficients are very similar and close to 5. In Appendix B3 we provide regression estimates for each country-category pair across the 30 product categories and seven countries (six GCC countries and the US).[24]

[24]To ensure that our results below are not sensitive to the specific coefficient of convexity employed, we replicate the

5.2 Frequency and Magnitude of Price Changes

We first compute the frequency and magnitude of price changes. We consider two methodologies for computing price changes: counting gaps in the price line, and carrying forward the last observed regular price through sale and stockout periods (for gaps of six months or less). We also consider estimates with or without sales included. For sales, we use a basic definition that identifies sales from a V-shape behavior in price. These measures are widely used in the literature. For more information, see Nakamura and Steinsson (2008).

We first plot the frequency and magnitude of price changes in UAE by allocating products into bins based on their market share. For each market share bin, we take the average frequency, the average price change (conditioning on a price change taking place), the average price change for price increases, and the average price change for price decreases across all categories. We report these scatter plots for the non-sales, no-carry-forward case in Figure 4.[25]

We observe that the magnitude of price change does not depend on market share (Figure 4, panels (b), (c), and (d)) but frequency does (panel (a)). Products with high sales experience more frequent price changes. This suggests that in the absence of any information on quantities or expenditure that would allow us to distinguish between important and non-important products, averaging the frequency of price changes across all products understates the degree of price stickiness in the economy.

Table 3 confirms that hypothesis. The frequency and magnitude of price changes in UAE are computed under four alternative specifications: (i) use expenditure information as weights (column (1), benchmark); (ii) exclude expenditure information and weight all products equally (column (2)); (iii) exclude expenditure information but use retail distribution (*ND*) to impute expenditure shares and weight all products (column (3)); and (iv) exclude expenditure information but use retail distribution (*WD*) to impute expenditure shares and weight all products (column (4)). Panel A reports computations for the frequency of price changes, and panel B for the magnitude.

When both prices and quantities are taken into consideration, we find that prices change 28% of the time. However, when information on expenditure is not known and all products are treated equally, prices appear to be stickier: they change only 22% of the time (a downward bias of one-fifth in the estimated price flexibility). Including retail distribution reduces measurement error (and the downward bias) substantially. When *ND* and *WD* are used to impute expenditure shares, we find

estimations in each of the three applications by allowing the convexity coefficient to vary from 3 to 7 when we impute expenditure shares from retail distribution metrics. We find that the results are robust to the alternative specifications. We conclude that including measures of retail distribution to impute expenditure shares reduces measurement error for a very generous range of coefficients, as long as the functional form used to link distribution and market share maintains convexity. In addition, we also considered the logistic function in equation (4) first proposed by Reibstein and Farris (1995) and repeated all the analysis using their functional form. The results, which are available upon request, are very similar and omitted for brevity. We chose the simplest version of the convexity function shown in equations (21) and (22) as it requires estimating only one convexity coefficient instead of two, and because the functional form in the paper produces a better fit than the logistic model in the applications below.

[25] Scatter plots for the other cases and countries are identical and omitted for brevity. Scatter plots are also identical when the median instead of the average frequency in each bin is computed .

that prices change 26% and 24% of the time, respectively. The results from alternative measures of frequency (rows 2, 3, and 4) provide the same conclusion: using retail distribution as a proxy for expenditure reduces measurement error and potential bias.

Unlike the case of the frequency, the magnitude of price changes does not depend on market share (see Figure 4, panels (b), (c) and (d)). Therefore, we do not expect to see substantial variation in the estimation under alternative estimation methods. Indeed, this is confirmed in panel B of Table 3, where the estimated magnitude of price changes does not change across columns.

Table 3: Frequency and Magnitude of Price Changes by Measures of Price Weight

	Benchmark (P&Q) (1)	P (2)	*ND* (3)	*WD* (4)
Panel A: Frequency of Price Changes				
(i) With Sales				
Continguous observations	0.28	0.22	0.26	0.24
Carrying regular price forward during sales and stockout	0.27	0.2	0.24	0.23
(ii) Without Sales				
Continguous observations	0.23	0.18	0.21	0.2
Carrying regular price forward during sales and stockout	0.22	0.17	0.2	0.19
Panel B: Magnitude of Price Changes				
All Changes*	0.02	0.03	0.03	0.03
Price Increases	0.07	0.07	0.07	0.07
Price Decreases	-0.06	-0.06	-0.06	-0.06

Note: Data are provided by Nielsen and cover sales of FMCGs between 2006 and 2011 in the GCC countries. The frequency of the data is monthly or bi-monthly, and price and quantity information are given at the level of the retailer in each period. * indicates that the computation is conditioning on a price change taking place.

5.3 Price Levels and Inflation

We measure inflation in each of the GCC countries using bi-monthly scanner data for 30 product categories of FMCGs between January 2006 and December 2011. To measure inflation, we estimate the model below with and without weights:

$$\ln p_{it} = \alpha_t + \beta_i + \epsilon_{it},$$

where i identifies products, α_t denotes the time fixed effect, and β_i denotes the product (barcode) fixed effect. In the versions of the estimation that includes weights, we use actual market shares, market shares imputed from the retail numeric distribution metric (*ND*), or market shares imputed

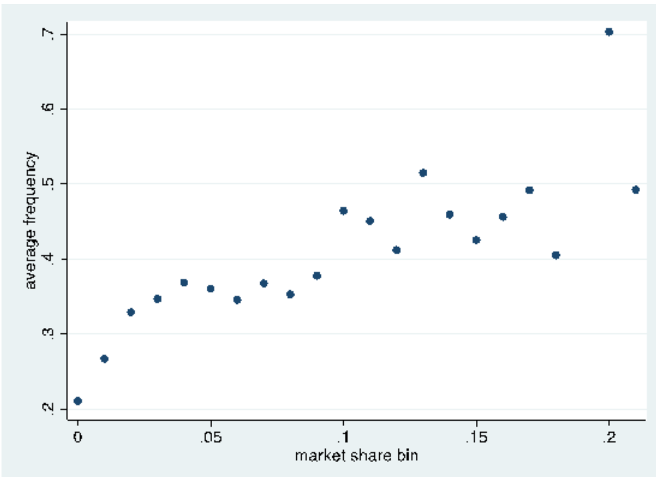

(a) Frequency of Price Changes

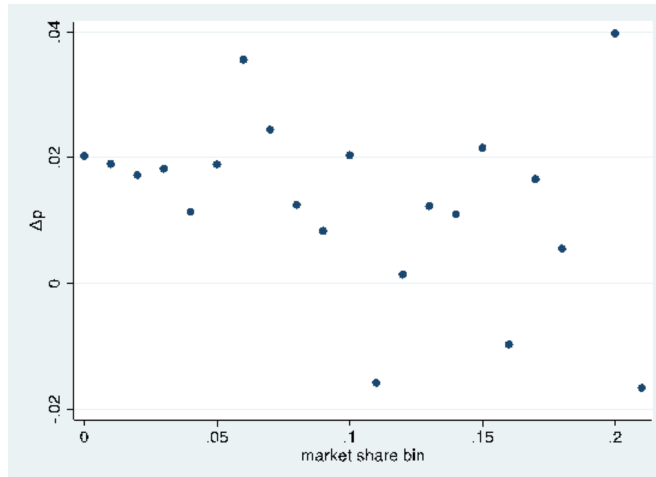

(b) Magnitude of Price Changes

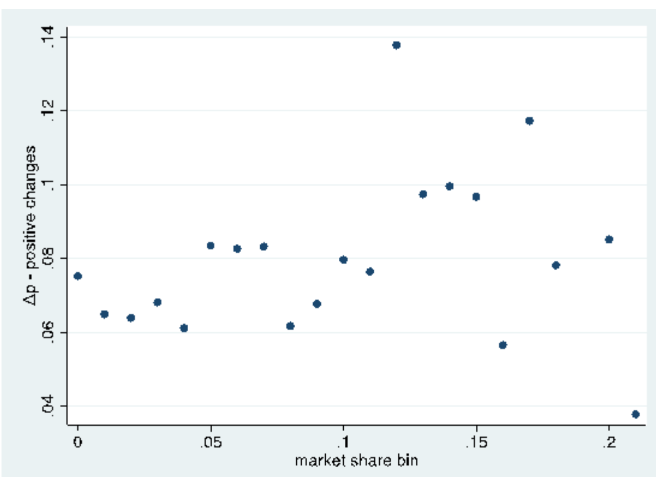

(c) Magnitude of Positive Price Changes

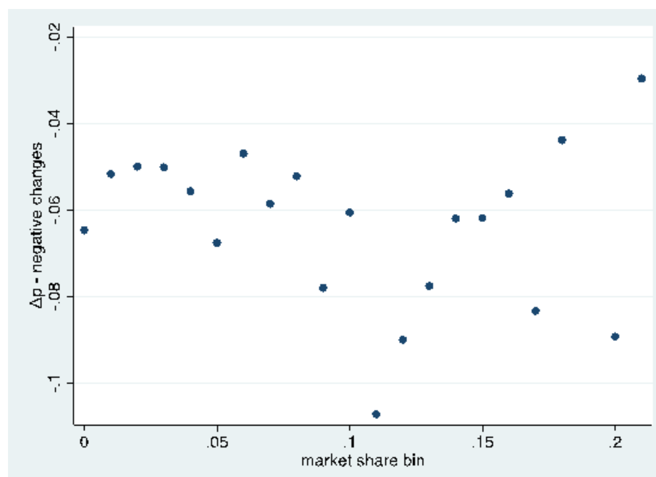

(d) Magnitude of Negative Price Changes

Note: We use Nielsen scanner data in UAE between January 2006 and December 2011 to compute the market share of each product in its product category and region. We also compute the frequency and magnitude of price changes for each product. Products are allocated into bins based on market share, and the average frequency and average magnitude of price changes in each bin are plotted. While the magnitude of price changes does not seem to depend on market share, the frequency does. Products with higher sales experience more frequent price changes. Panel (b) is conditioning on a price change taking place.

Figure 4: Frequency and Magnitude of Price Changes

from the retail weighted distribution metric (*WD*). For each product (barcode), prices are averaged across all outlets during that period.[26]

The estimated inflation series based on the coefficients of the time fixed effects are plotted in Figure 5 for all countries and all alternative estimations. For each country, measuring inflation using only prices understates the true increase in prices over the sample period; the democratic measure of inflation that does not take into account expenditure weights understates the true increase in prices by about 4.4 percentage points (Table 4, column (1)). This represents a deviation of *nearly one-third* from the actual inflation rate over the five-year period. However, imputing expenditure shares from retail distribution metrics, and using those shares to weight the data, substantially improves the performance of the inflation estimator: the deviation from the benchmark case is now reduced to 0.8% for the case of *ND* and 0.4% for the case of *WD*.

Columns (4) and (5) in Table 4 report the performance of estimation using *ND* and *WD* relative to the one without. Specifically, we first compute the root mean square error (RMSE) of deviations from the benchmark case across all periods. We then divide the RMSEs of the two measures using retail distribution information over that using only prices. A ratio below 1 suggests that using retail distribution measures improves the fit. Indeed, as shown in these columns, using *ND* decreases the gap from the benchmark case by *almost half* on average, and *WD* by *almost three-quarters*.

Table 4: Summary of Inflation Measures

Country	Total Gap			RMSE ratio	
	No wgt	ND	WD	ND	WD
	(1)	(2)	(3)	(4)	(5)
Bahrain (bah)	-4.9%	-1.4%	-0.6%	0.50	0.25
Saudi Arabia (ksa)	-4.6%	-1.2%	-1.1%	0.60	0.26
Kuwait (kuw)	-4.4%	1.1%	0.6%	0.57	0.52
Oman (omn)	-4.6%	-1.9%	-1.1%	0.56	0.33
Qatar (qtr)	-4.4%	-0.2%	-0.3%	0.28	0.13
United Arab Emirates (uae)	-3.6%	-0.9%	-0.1%	0.67	0.16
Average	-4.4%	-0.8%	-0.4%	0.53	0.27

Note: We use scanner data to compute the growth in prices between January 2006 to December 2011. Four alternative estimation methods are employed: (i) using prices and expenditure information to weight the data; (ii) using only prices and no weights; (iii) using prices and weights based on imputed market shares from *ND*; and (iv) using prices and weights based on imputed market shares from *WD*. The difference between the estimated price level at the end of the period for each method and the benchmark is reported in columns (1) to (3). We also report the share of the root mean square error (RMSE) of versions (iii) and (iv) to the RMSE of version (ii) in columns (4) and (5), respectively. The RMSE is based on deviations from the benchmark case. A coefficient less than 1 indicates a better fit relative to the case of no weights in estimation.

[26] We also experimented with taking the median, min, max, and a random price of a barcode across all outlets at which the product is sold. The results do not change and are available by the authors upon request.

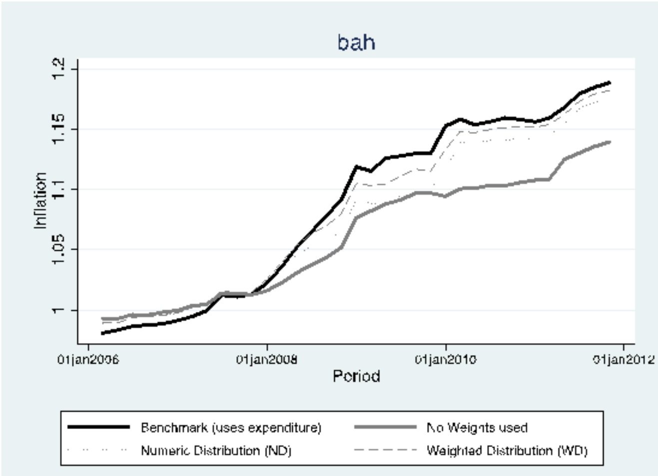

(a) Numeric Distribution

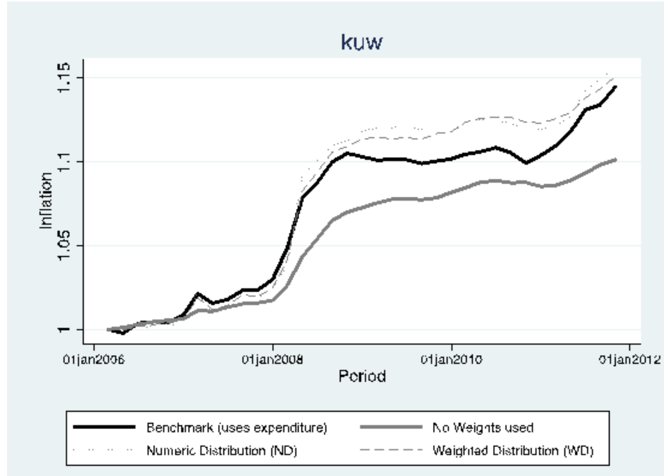

(b) Weighted Distribution

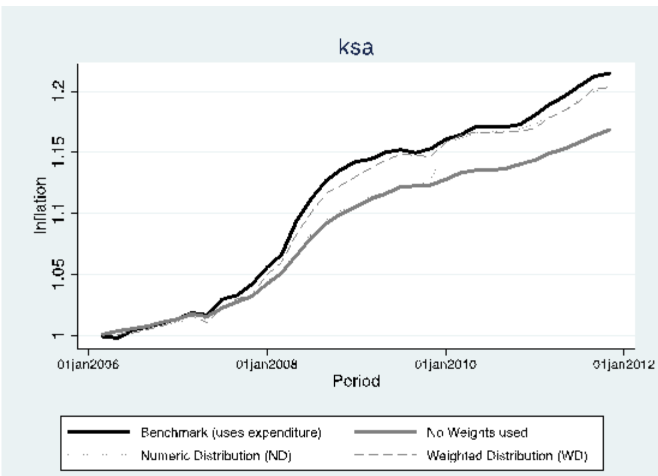

(c) Numeric Distribution

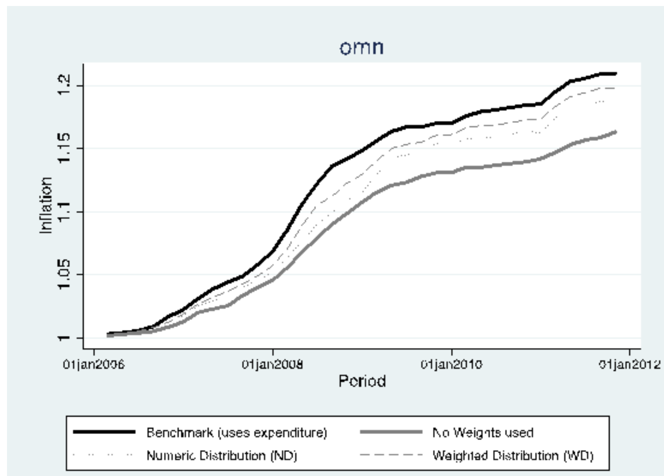

(d) Weighted Distribution

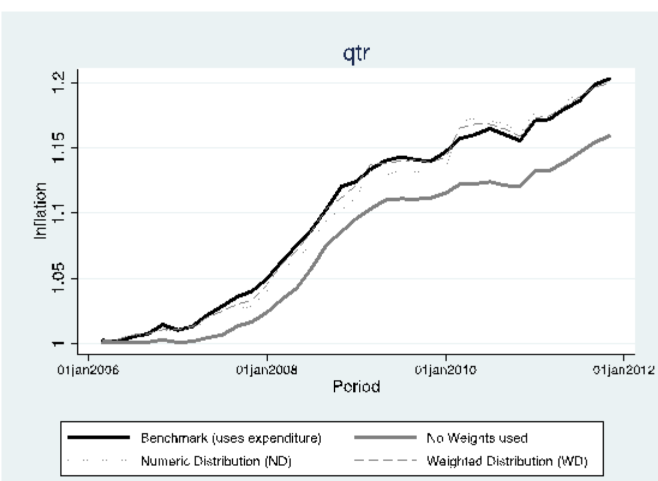

(e) Weighted Distribution

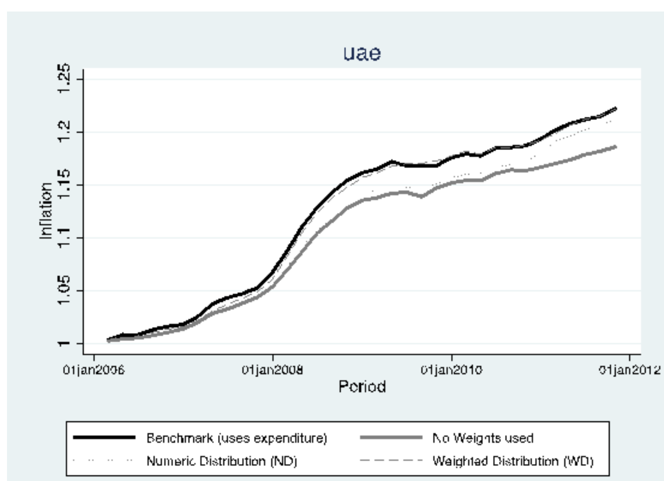

(f) Numeric Distribution

Note: We compute inflation measures between January 2006 and December 2011 for each of the GCC countries using four alternative specifications: (i) use prices and expenditure (solid black line), (ii) use only prices (solid grey line), (iii) use prices and information on numeric distribution (dotted grey line), and (iv) use prices and information on retail distribution (dashed grey line). In all cases, when information on expenditure is ignored and all products are treated equally, inflation is understated. However, using information on retail distribution obtained from prices reduces the measurement bias significantly.

Figure 5: Inflation Measures

5.4 Implications for the International Comparison Program

The International Comparison Program (ICP), a collaboration between the World Bank and national statistical agencies, is an initiative under the United Nations with the mandate to measure the relative cost of living across the world. Every few years, the World Bank puts together and distributes an extensive price survey to statistical agencies worldwide. The survey is broken down into product groups (e.g., "Bread and Cereals"; "Miscellaneous goods and services") and each product group into several basic headings (e.g., "Other cereals, flour, and other products"; "Appliances, articles and products for personal care"). Each basic heading contains a list of very detailed product definitions (e.g., " Cornflakes Kellogg's 500 gram, range 250-600 gram, milled corn (maize) pre-packed, ready to eat cereals, sugar and/or other ingredients"; "Tooth paste, tube, 80 mL, range 50-100 mL, Colgate, Classic Total, exclude whitening"). Its content varies by year and region and is highly confidential. Statistical agencies are asked to price each item in a number of stores and report the average price. Product prices are then used by the World Bank to compute price levels for each basic heading, for each product group, and for the overall basket. At the basic heading level, because expenditure information for each product is not known, prices are averaged across all products by assuming identical weights. Then prices for each basic heading are aggregated up using expenditure information from the components that make up the national CPI data. [27]

As straightforward as this exercise sounds, it presents an extremely daunting undertaking in terms of methodology and administration. Various issues arise in the pre-survey (e.g., how to construct the baskets), during the survey (e.g., how to price), and the post-survey (e.g., how to aggregate) stages. Rightly, the World Bank characterizes the ICP as the largest and most complex statistical exercise in the world. [28]

In this third and final exercise, we consider the averaging of prices at the basic heading level and ask whether this practice introduces measurement error and potential bias. Because prices of the most important items may converge faster across retailers and across countries for consumers paying more attention, taking an unweighted average price across all items within the basic heading can upwardly bias measures of differences in the cost of living across countries.

To check this, we use the scanner data to simulate the ICP under alternative scenarios that are described below. We begin by extracting from the World Bank 2011 confidential ICP survey the product definitions that overlap with the Nielsen scanner FMCG data. These are: (i) blades — 2 definitions, (ii) cereals — 4 definitions, (iii) detergents — 2 definitions, (iv) juices — 4 definitions, and (v) toothpaste — 1 definition. Examples of definitions selected are "Cornflakes Kellogg's 500 gram, range 250-600 gram, milled corn (maize) pre-packed, ready to eat cereals, sugar and/or other ingredients"

[27] We would like to thank the World Bank, especially Nada Hamadeh, for sharing with us a copy of the 2011 product survey. More information on the 2011 ICP survey is available at The World Bank, International Comparison Program (ICP), http://siteresources.worldbank.org/ICPEXT/Resources/ICP_2011.html.

[28] Two recent papers highlight key challenges in the ICP methodology. Deaton and Aten (2017) discuss the challenge of linking countries and regions together, while Inklaar and Rao (2017) compare and contrast alternative measurement methodologies. Antoniades (2016) succinctly captures key challenges in the collection of raw data with the 4Rs: the challenges of finding: (i) the right product, (ii) the right weight, (iii) the right price/retailer, and (iv) the right variety.

and "Tooth paste, tube, 80 mL, range 50-100 mL, Colgate, Classic Total, exclude whitening." For confidentiality purposes, we omit reporting the remaining 11 definitions.

While only 13 product definitions survive the matching, a total of 2,069 barcode products are selected. This happens because multiple varieties (barcodes) of the same product match the same ICP product description. Colgate Total 100ml, Colgate Total 100ml PD, Colgate Total 100ml Pump, Colgate Total 12 100ml, Colgate Total Fresh Stripe 100ml, and their 50ml variations all satisfy the ICP product definition "Tooth paste, tube, 80 mL, range 50-100 mL, Colgate, Classic Total, exclude whitening".[29] The availability of multiple varieties of the same product, which we call variety bias, poses an important challenge for price auditors as they have to pick one of potentially several different prices. In the simulation, we experiment with alternative pricing rules discussed below.

With the construction of the survey completed, the simulation is broken down into two parts: data collection and estimation. In the data collection part, we provide the simulation with a set of rules that mimic the actual process. Specifically, we first input the number of stores to be audited (n=10, 20, or 50). We then ask the simulation to pick those n stores out of the universe of outlets in our sample by selecting large stores first. If all supermarkets/hypermarkets are exhausted, the algorithm randomly picks the remaining from the population of groceries and mini-markets. This is an important stage as prices vary across stores, with the largest stores offering lower prices. Next, we ask the simulation to randomly pick a date for the audit out of the six bi-monthly periods in 2011. Finally, we give guidance as to which price must be quoted if multiple varieties of a product at a store satisfy the same definition. The six alternative rules are: (1) take an average price, (ii) take the median price, (iii) pick a price at random, (iv) pick the lowest price, (v) pick the highest price, and (vi) pick the price of the item you think is the most important based on sales (which can be asserted from shelf space). While the average and median rules are less practical, for the majority of cases in which a handful of varieties exist, they can easily be computed on the spot.

Once prices are collected in each country, the country-product-dummy (CPD) regression is estimated across the 13 product definitions and countries:

$$\ln p_{ic} = \alpha_c + \beta_i + \epsilon_{it},$$

where product is indexed by i and country is indexed by c. The variables α_c and β_i capture country and product fixed effects, respectively. PPPs, relative to the numeraire (in this case, Bahrain) are obtained from the exponent of country fixed effects. For example, if the exponent of the KSA coefficient is 1.2, then prices in Saudi Arabia are 20% higher than in Bahrain.

Three versions of the equation above are estimated. Version 1, the benchmark, considers both price and expenditure information so that each observation is weighted by importance. Version 2 mimics the ICP by omitting expenditure information and treating all products equally. Versions 3 and 4 omit expenditure information but use information on the numeric (*ND*) and weighted (*WD*)

[29] For a full list of products that fit the same definition in the case of toothpaste and cereals, see Table A.3 in the appendix.

distribution, respectively, to weight the data.

To ensure that the results are not sensitive to the random selection of time period and outlets, the exercise is repeated 50 times. Each time, average PPP differences across the GCC countries are collected, and the median differences across these 50 iterations are reported in Table 5 for alternative specifications. The first column indicates the rule specified for dealing with variety bias (explained above). The second column lists the number of outlets audited. The next four columns report the estimation results under the four alternative methods.

Table 5: Measuring International Price Differences

Estimation Type	Outlets Audited	Average PPP difference among the GCC countries by Weights Used			
		(1) Expenditure (benchmark)	(2) None	(3) *ND*	(4) *WD*
avg	10	0.06	0.18	0.09	0.07
avg	20	0.05	0.17	0.06	0.06
avg	50	0.08	0.2	0.06	0.09
max	10	0.06	0.18	0.11	0.08
max	20	0.05	0.21	0.07	0.09
max	50	0.08	0.27	0.12	0.16
med	10	0.06	0.09	0.06	0.01
med	20	0.05	0.12	0.04	0.02
med	50	0.08	0.16	0.05	0.08
mii	10	0.06	0.15	0.06	0.08
mii	20	0.05	0.18	0.08	0.11
mii	50	0.08	0.17	0.05	0.07
min	10	0.06	0.11	0.05	0.05
min	20	0.05	0.11	0.04	0.05
min	50	0.08	0.07	0.05	0.03
random	10	0.06	0.12	0.05	0.06
random	20	0.05	0.13	0.06	0.09
random	50	0.08	0.16	0.04	0.1

Note: Data are provided by Nielsen and cover sales of FMCGs between 2006 and 2011 in the GCC countries. The frequency of the data is monthly or bi-monthly, and price and quantity information are given at the level of the retailer in each period. In case multiple varieties of a product at a store satisfy the same definition of ICP product, we use alternative criteria to quote price: (i) take an average price (denoted as "avg"), (ii) take the median price (denoted as "med"), (iii) pick a price at random (denoted as "random"), (iv) pick the lowest price (denoted as "min"), (v) pick the highest price (denoted as "max"), and (vi) pick the price of the most-important-item based on sales (denoted as "mii").

For instance, the first row of the table lists average PPP differences for the scenario in which 10 outlets are audited and price auditors are asked to report the average price in case multiple products satisfy the same PPP product description. When both prices and expenditure information are used (version 1 — benchmark), prices in the GCC differ by 6%. However, when only prices are used and all products are treated equally, prices differ by 18%. Regardless of how many outlets are audited or which rule is used to deal with variety bias, excluding weights overstates PPP differences by a very large margin.[30] However, when information on numeric and weighted distribution is used to project expenditure shares (versions 3 and 4), estimated average PPP differences are in line with the benchmark case.

To summarize, the main lessons from this exercise are that (i) treating all products equally overstates the true cost of living, (ii) increasing the sampling size does not improve or worsen estimates, (iii) the results for alternative rules to deal with variety bias are similar, and (iv) projecting expenditure shares from retail distribution reduces measurement bias substantially.[31]

6 Conclusion

The availability of data on prices that can be collected online presents a new opportunity for researchers to study prices and price behavior. Yet, as we documented in this paper, the unavailability of information on quantities or expenditure introduces substantial measurement error, and in many cases, bias. By treating all prices equally, researchers may understate the cost of living, overstate price stickiness, and overstate price differences.

To overcome the challenge imposed by the lack of expenditure data, we propose that researchers use information on retail distribution to impute expenditure shares. By exploiting the convexity that characterizes the relation between retail distribution and market share, one can build measures of retail distribution solely from price data and use these measures to back out expenditure shares. Our approach, which we motivate through evidence in the literature but also a micro-founded framework, works because it helps researchers to identify the most important items within a product group, and thus allows them to weight the data accordingly in the estimation.

We illustrate that the proposed approach works well by reducing measurement error substantially when measuring the frequency of price changes, inflation, and international price differences. Adopting the methodology will benefit those working with price data scraped from retailers' websites or from online applications and will also benefit those working on the International Comparison Program.

[30] For robustness, we also report estimation results when the average (instead of the median) across the 50 iterations is computed. The results, which are available in the appendix, are identical to those reported in Table 5.

[31] Antoniades (2016) finds that there is more variation in prices across retailers than there is across varieties of the same product definition within retailers. That is, prices of a specific Colgate brand will vary substantially across retailers, but prices of Colgate varieties within a particular retailer will not be that different. The implication is that dealing with variety bias is not that important for the ICP as it is to decide what the best sample of retailers to audit is. This finding explains why alternative pricing rules in Table 5 do not yield different results.

References

ANTONIADES, A. (2016): "The Future of Price Statistics: Innovations in Data, Technology, and Methods," .

ANTONIADES, A., AND N. ZANIBONI (2016): "Exchange Rate Pass-Through Into Retail Prices," *International Economic Review*, 57(4), 1425–1447.

BORIN, N. A., C. V. VRANKEN, AND P. W. FARRIS (1991): "A pilot test of discrimination in the Japanese distribution system," *Journal of Retailing*, 67(1), 93.

CAVALLO, A. (2013): "Online and official price indexes: Measuring Argentina's inflation," *Journal of Monetary Economics*, 60(2), 152–165.

——— (2017): "Are online and offline prices similar? evidence from large multi-channel retailers," *American Economic Review*, 107(1), 283–303.

——— (2018): "Scraped data and sticky prices," *Review of Economics and Statistics*, 100(1), 105–119.

CAVALLO, A., AND R. RIGOBON (2011): "The distribution of the Size of Price Changes," Discussion paper, National Bureau of Economic Research.

——— (2016): "The billion prices project: Using online prices for measurement and research," *Journal of Economic Perspectives*, 30(2), 151–78.

COIBION, O., Y. GORODNICHENKO, AND G. H. HONG (2015): "The cyclicality of sales, regular and effective prices: Business cycle and policy implications," *American Economic Review*, 105(3), 993–1029.

DEATON, A., AND B. ATEN (2017): "Trying to Understand the PPPs in ICP 2011: Why are the Results so Different?," *American Economic Journal: Macroeconomics*, 9(1), 243–64.

FARRIS, P., J. OLVER, AND C. DE KLUYVER (1989): "The relationship between distribution and market share," *Marketing Science*, 8(2), 107–128.

FEENSTRA, R. C., M. XU, AND A. ANTONIADES (forthcoming): "What is the Price of Tea in China? Towards the Relative Cost of Living in Chinese and US Cities," *The Economic Journal*.

HOTTMAN, C. J., S. J. REDDING, AND D. E. WEINSTEIN (2016): "Quantifying the sources of firm heterogeneity," *The Quarterly Journal of Economics*, 131(3), 1291–1364.

INKLAAR, R., AND D. RAO (2017): "Cross-country income levels over time: did the developing world suddenly become much richer?," *American Economic Journal: Macroeconomics*, 9(1), 265–90.

KRUGER, M., AND B. HARPER (2006): "Market Share and Product Distribution: Re-Tested and Extended," .

KUKSOV, D., AND A. PAZGAL (2007): "Research note: the effects of costs and competition on slotting allowances," *Marketing Science*, 26(2), 259–267.

MERCER, A. (1992): *Implementable Marketing Research*. Prentice Hall.

NAKAMURA, E., AND J. STEINSSON (2008): "Five facts about prices: A reevaluation of menu cost models," *The Quarterly Journal of Economics*, 123(4), 1415–1464.

NUTTALL, C. (1965): "The relationship between sales and distribution of certain confectionery lines," *Commentary*, 7(4), 272–285.

RAO, A. R., AND H. MAHI (2003): "The price of launching a new product: Empirical evidence on factors affecting the relative magnitude of slotting allowances," *Marketing Science*, 22(2), 246–268.

REIBSTEIN, D. J., AND P. W. FARRIS (1995): "Market share and distribution: A generalization, a speculation, and some implications," *Marketing Science*, 14(3_supplement), G190–G202.

Online Appendix – Not for Publication

A. Theory Appendix

A1. Proof of Proposition 1

Since retailers charge the same markups, we denote it as $\mu_r = \mu \ \forall r \in [0,1]$. Product sales of ϕ can be written as:

$$R_\phi = \sigma\phi^{\sigma-1}A_1^{\frac{1-\sigma}{\epsilon_1}}\mu^{1-\frac{\eta(\sigma-1)}{\epsilon_1}} \times \int_0^n f_r^{\epsilon_2} dr.$$

The first and second derivative of R_ϕ with respect to n are ($\epsilon_2 > 0$):

$$\frac{\partial R_\phi}{\partial n} = \sigma\phi^{\sigma-1}A_1^{\frac{1-\sigma}{\epsilon_1}}\mu^{1-\frac{\eta(\sigma-1)}{\epsilon_1}} f_n^{\epsilon_2} > 0, \quad \frac{\partial^2 R_\phi}{\partial n^2} = \epsilon_2\sigma\phi^{\sigma-1}A_1^{\frac{1-\sigma}{\epsilon_1}}\mu^{1-\frac{\eta(\sigma-1)}{\epsilon_1}} f_n^{\epsilon_2-1}f_n' > 0,$$

where the first inequality holds given that there is no negative term, and the second inequality holds given that fixed cost f_r increase in r.

A2. Proof of Proposition 2

We rewrite (15) as:

$$\begin{aligned} R_\phi &= \sigma\phi^{\sigma-1}A_1^{\frac{1-\sigma}{\epsilon_1}-\frac{1}{\eta}\left[1-\frac{\eta(\sigma-1)}{\epsilon_1}\right]} \times \int_0^n f_r^{\epsilon_2-\frac{1}{\eta}\left[1-\frac{\eta(\sigma-1)}{\epsilon_1}\right]}\phi_r^{\frac{\epsilon_1}{\eta}\left[1-\frac{\eta(\sigma-1)}{\epsilon_1}\right]} dr \\ &= \sigma\phi^{\sigma-1}A_1^{\frac{1-\sigma}{\epsilon_1}-\frac{1}{\eta}\left[1-\frac{\eta(\sigma-1)}{\epsilon_1}\right]} \times \int_0^n f_r^{1-\frac{1}{\eta}}\phi_r^{\frac{\epsilon_1}{\eta}\left[1-\frac{\eta(\sigma-1)}{\epsilon_1}\right]} dr, \end{aligned}$$

where the first equality uses $\mu_r = A_1^{-\frac{1}{\eta}}f_r^{-\frac{1}{\eta}}\phi_r^{\frac{\epsilon_1}{\eta}}$ as implied by (13), and the second equality uses $\epsilon_2 \equiv 1 - \frac{\sigma-1}{\epsilon_1}$. The first and second derivative of R_ϕ with respect to n satisfy:

$$\frac{\partial R_\phi}{\partial n} = \phi^{\sigma-1}A_1^{\frac{1-\sigma}{\epsilon_1}-\frac{1}{\eta}\left[1-\frac{\eta(\sigma-1)}{\epsilon_1}\right]} f_n^{1-\frac{1}{\eta}}\phi_n^{\frac{\epsilon_1}{\eta}\left[1-\frac{\eta(\sigma-1)}{\epsilon_1}\right]} > 0, \quad \frac{\partial^2 R_\phi}{\partial n^2} > 0,$$

where first inequality holds given there is no negative term, and the second inequality holds given that fixed cost f_r and ϕ_r increase in r, $\eta > 1$ and $1 - \frac{\eta(\sigma-1)}{\epsilon_1} > 0$ (implied by $k > \frac{\eta(\sigma-1)^2-\sigma\eta-1}{\sigma-\eta}$).

When $k = \frac{\eta(\sigma-1)^2-\sigma\eta-1}{\sigma-\eta}$, we can rewrite (13) as $\mu_r\phi_r^{1-\sigma} = A_1^{-1/\eta}f_r^{-1/\eta}$ (as an intermediate step, one can show that the equality $\epsilon_1 = \eta(\sigma-1)$ holds). We substitute the new term into $R_{r\phi} = \sigma\phi^{\sigma-1}f_r\mu_r\phi_r^{1-\sigma}$ to obtain $R_{r\phi} = \sigma A_1^{-1/\eta}\phi^{\sigma-1}f_r^{1-1/\eta}$. Product sales of ϕ will be:

$$R_\phi = \sigma A_1^{-1/\eta}\phi^{\sigma-1}\int_0^n f_r^{1-1/\eta}dr$$

The first and second derivative of R_ϕ with respect to n satisfy:

$$\frac{\partial R_\phi}{\partial n} = \sigma A_1^{-1/\eta} \phi^{\sigma-1} f_n^{1-1/\eta} > 0, \quad \frac{\partial^2 R_\phi}{\partial n^2} = \sigma A_1^{-1/\eta} \phi^{\sigma-1} \left(1 - \frac{1}{\eta}\right) f_n^{-1/\eta} f_n' > 0$$

where the first inequality holds given that there is no negative term, and the second inequality holds given that fixed cost f_r increase in r.

Under the example $k = \frac{\eta(\sigma-1)^2 - \sigma\eta - 1}{\sigma - \eta}$, when f_r is exponential, i.e., $f_r = \gamma e^{\theta r}$ ($\gamma > 0$ and $\theta > 1$), the sales of product ϕ become

$$\begin{aligned} R_\phi &= \sigma A_1^{-\frac{1}{\eta}} \phi^{\sigma-1} \gamma^{1-\frac{1}{\eta}} \int_0^n e^{\theta\left(1-\frac{1}{\eta}\right)r} dr \\ &= \frac{\sigma A_1^{-\frac{1}{\eta}} \phi^{\sigma-1} \gamma^{1-\frac{1}{\eta}}}{\theta\left(1-\frac{1}{\eta}\right)} \left[e^{\theta\left(1-\frac{1}{\eta}\right)n} - 1\right]. \end{aligned}$$

As long as $\theta > 0$, product sales are a convex function of *numeric distribution n.*

A3. Proof of Proposition 4

In case of $1 - \frac{\eta(\sigma-1)}{\epsilon_1} \geq 0$ (which implies $k \geq \frac{\eta(\sigma-1)^2 - \sigma\eta - 1}{\sigma-\eta}$), the proof follows the same steps as Proposition 2. So consider the case in which $1 - \frac{\eta(\sigma-1)}{\epsilon_1} < 0$. Given the observed sales of product ϕ in (15), the first derivative of $R_\phi$with respect to n is:

$$\frac{\partial R_\phi}{\partial n} = \sigma \phi^{\sigma-1} A_1^{\frac{1-\sigma}{\epsilon_1}} f_n^{\epsilon_2} \mu_n^{1-\frac{\eta(\sigma-1)}{\epsilon_1}} > 0.$$

Given that sales decrease in retailer index r in the equilibrium studied, retailer markups also decrease in retailer index r where retailer markup is given in (17) or (19). This implies that both $f_n^{\epsilon_2}$ and $\mu_n^{1-\frac{\eta(\sigma-1)}{\epsilon_1}}$ increase in n, which confirms convexity:

$$\frac{\partial^2 R_\phi}{\partial n^2} > 0$$

A4. Model Simulation Procedures

Table A.1 displays the parameters used in the simulation. Given the parameters, we simulate the economy in which consumers, manufacturers, and retailers are specified by (6), (9), and (16). The fixed cost is specified as $f_r = \gamma e^{\theta r}$ ($\gamma > 0$ and $\theta > 1$). We simulate 10,000 draws u from a uniform distribution from 0 to 1. The corresponding Pareto productivity draws are $\phi = (1-u)^{-\frac{1}{k}} \bar{\phi}$. Then we solve the model by solving for the equilibrium retailer markups by the following procedures (i denotes the i-th loop):

Step 1: Set the initial value of retailers' markups as $\eta_r^{(1)} = \frac{\eta}{\eta-1}$ if it is the start of loop ($i = 1$); otherwise set $\eta_r^{(i)} = \eta_r^{(i-1)}$, where $\eta_r^{(i-1)}$ is obtained from **Step 4** of the last loop ($i \geq 2$).

Step 2: Solve the productivity cutoff ϕ_r using (13) and the $\eta_r^{(i)}$ obtained from ***Step 1***.

Step 3: Given the productivity cutoff for each retailer (obtained from ***Step 2***), calculate the sales of each product in each retailer $R_{r\phi}$, using equation (14) (set $R_{r\phi} = 0$ if $\phi < \phi_r$). With manufacturers' sales in each market, we add them up to get total market sales and the corresponding market shares s_r for each retailer r.

Step 4: Calculate retailers' markups using market shares s_r (obtained from ***Step 3***) and equation (17). Denote the derived markup as $\eta_r^{(i)}$.

Step 5: If the difference between $\eta_r^{(i)}$ and $\eta_r^{(i-1)}$ is smaller than the tolerance, we stop the loop. Otherwise, we loop over ***Step 1*** through ***Step 5*** until markups converge.

Figure 3 in the main text displays the relationship between product shares and the *numeric distribution*. Through all different values of k, the convexity remains robust.

Table A.1: Simulation Parameters

Parameter	Description	Value
σ	Elasticity of substitution (varieties)	4.5
η	Elasticity of substitution (retailers)	3
k	Shape parameter of productivity distribution	[4,8,16]
$\bar{\phi}$	Shift parameter of productivity distribution	1
M	Number of manufacturers	10,000
γ	Shift parameter of fixed cost	100
θ	Elasticity of fixed cost with distance from the cheapest retailers	4
N	Number of retailers	10
P	Aggregate price index	10
w	labor cost	1
Y	GDP	1,000
Tol	Tolerance for markup convergence	1e-6

Notes: $k = 16$ corresponds to the example case (i.e., the sufficient condition to guarantee the convexity between product shares and the numeric distribution).

We also simulate the model with different functional forms for the fixed cost f_r, with all other parameters fixed as displayed in Table A.1. In Figure A.1, we specify f_r in the form of power function, i.e., $f_r = \gamma r^\theta$ where we choose $\gamma = 100$ and $\theta = 2$. In Figure A.2, we instead specify f_r as a concave function of r, i.e., we choose $\gamma = 100$ and $\theta = 0.2$ in the simulation. The relationship between product share and the *numeric distribution* remains convex.

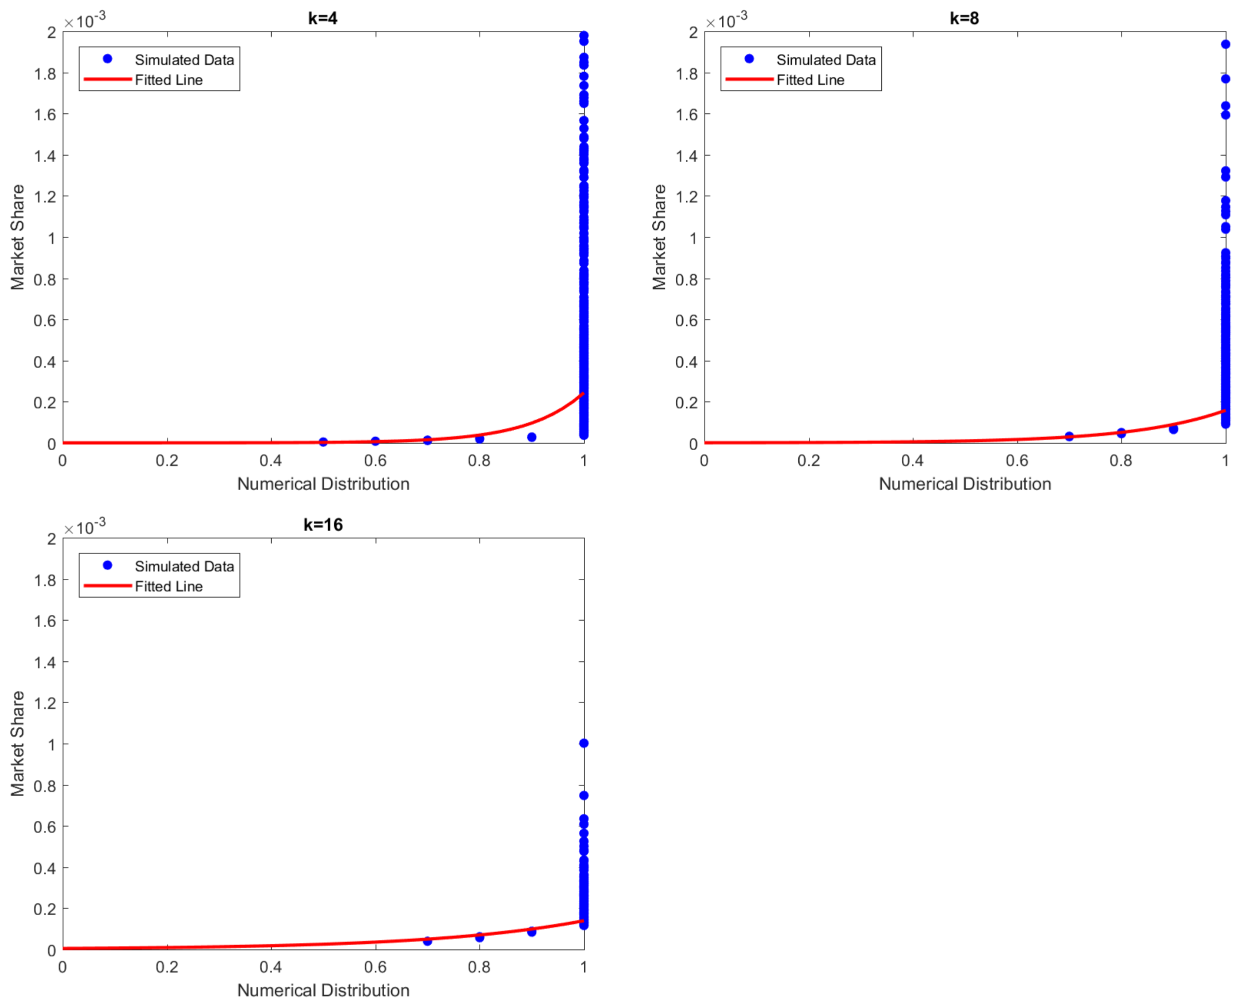

Figure A.1: Convexity between Sales and Numeric Distribution ($f_r = \gamma r^{\theta}$, $\theta = 2$)

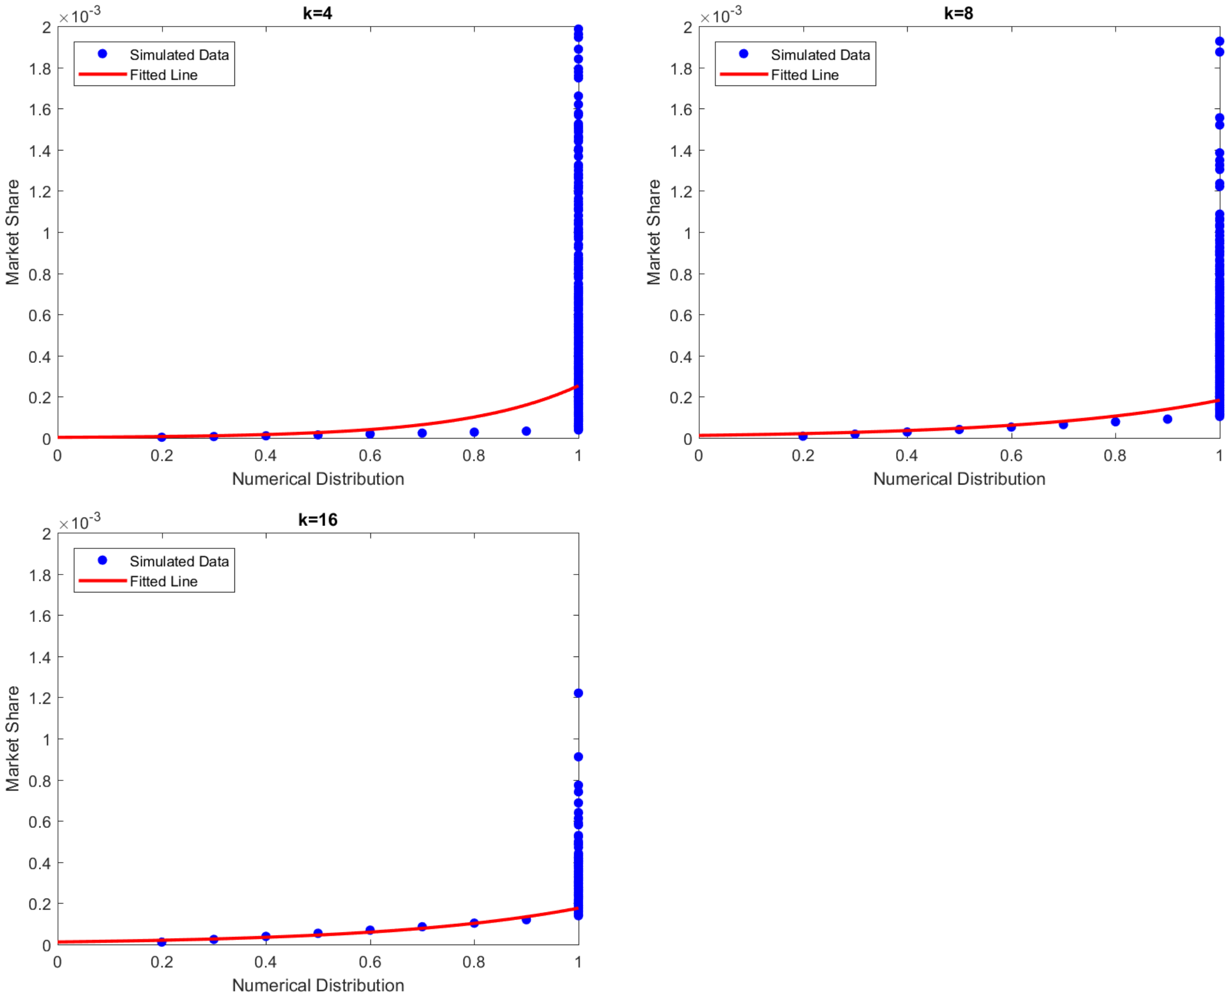

Figure A.2: Convexity between Sales and Numeric Distribution ($f_r = \gamma r^{\theta}$, $\theta = 0.2$)

B. Table Appendix

B1. City coverage of scraped data for Chinese products

Table A.2: Chinese Product Prices Available in the Mobile Application for 2014

	Laundry Detergent		Personal Wash Items		Shampoo		Toothpaste	
City	EANs	Retailers	EANs	Retailers	EANs	Retailers	EANs	Retailers
Beijing	929	11	1,273	11	1,041	10	1,024	11
Changsha	874	10	1,471	11	1,063	9	960	10
Chengdu	778	8	1,214	7	957	8	560	7
Chongqing	870	10	1,419	10	998	11	880	9
Dalian	661	6	986	4	775	5	655	3
Guangzhou	902	14	1,524	16	1,071	12	826	13
Hangzhou	805	8	1,210	8	975	8	788	8
Harbin	729	6	1,063	5	902	6	555	6
Hefei	968	10	1,325	10	1,090	9	1,069	8
Jinan	731	8	1,092	8	901	8	621	7
Kunming	579	5	978	5	773	5	422	5
Ningbo	676	7	1,074	8	842	7	569	7
Shanghai	999	12	1,456	12	1,226	10	1,032	12
Shenyang	929	10	1,383	10	1,084	11	847	10
Shenzhen	966	9	1,674	9	1,195	9	868	9
Suzhou	754	7	1,159	7	956	8	581	7
Tianjin	873	7	1,298	7	1,076	7	900	7
Wuhan	933	11	1,270	12	1,030	10	992	12
Wuxi	798	7	1,164	7	908	7	932	7
Xiamen	896	9	1,551	9	1,067	9	873	9
Xi'an	946	8	1,334	7	1,075	7	-	-

B2. Examples of products that fit the same ICP PPP product description

Table A.3: Examples of Products that Fit the Same ICP PPP Product Description

	Cornflakes Kellogg's 500 gram, range 250-600 gram, milled corn (maize) pre-packed, ready to eat cereals, sugar and(or) other ingredients		Tooth paste, tube, 80 mL, range 50-100 mL, Colgate, Classic Total, exclude whitening
1	KELLOGG'S CORNFLAKES 375GR (F)(ARABIC)	1	COLGATE 100ml TOTAL
2	KELLOGG'S CORNFLAKES 500GR (F) (ARABIC)	2	COLGATE 100ml TOTAL PUMP
3	KELLOGG'S CRUNCHY NUT CORNFLAKES 500GR(F	3	COLGATE 50ML TOTAL 12 CLEAN MINT (FAC)
4	KELLOGG'S HONEYNUT CORNFLKE.375GR(F)(ARA	4	COLGATE 50ml TOTAL
5	KELLOGGS 375g CORN FLAKES	5	COLGATE 50ml TOTAL 12 CLEAN MINT
6	KELLOGGS 375g CRUNCHY NUT CORN FLAKES	6	COLGATE TOTAL 100 ML
7	KELLOGGS 375g HONEY NUT CORN FLAKES	7	COLGATE TOTAL 100ML
8	KELLOGGS 500g CORN FLAKES	8	COLGATE TOTAL 100ML PD
9	KELLOGGS 500g HEALTH WISE BRAN FLAKES	9	COLGATE TOTAL 100ML PD(M.BEN/FL)
10	KELLOGGS ALL BRAN FLAKES 375 GM PKT	10	COLGATE TOTAL 100ML PUMP
11	KELLOGGS C/F 250G (F)	11	COLGATE TOTAL 100ml PD
12	KELLOGGS C/F 375G (F)	12	COLGATE TOTAL 12 100ML PUMP
13	KELLOGGS C/F 500G (F)	13	COLGATE TOTAL 12 50ML
14	KELLOGGS CHOCO CF 375g (ARABIC)	14	COLGATE TOTAL 12 50ml
15	KELLOGGS CORN FLAKES 250GR PKT	15	COLGATE TOTAL 12 CLEAN MINT 50ML GUM
16	KELLOGGS CORN FLAKES 375GR PKT	16	COLGATE TOTAL 12 CLEAN MINT 50ML(FAC)
17	KELLOGGS CORN FLAKES 500 GR PKT	17	COLGATE TOTAL 12 CLEANMINT 50ML (COS)
18	KELLOGGS CORNFLAKES 375g ARABIC	18	COLGATE TOTAL 50ML
19	KELLOGGS CORNFLAKES 500g BOX ARABIC	19	COLGATE TOTAL 50ML (GUM)
20	KELLOGGS CRUMBS CORN FLAKES 595GR(A)ENG	20	COLGATE TOTAL 50ML CLEAN MINT PROT. GUM
21	KELLOGGS CRUNCHYNUT CORNFLAKES 375g ARAB	21	COLGATE TOTAL 50ML(GUM)
22	KELLOGGS FROSTED FLAKES 496GR (ENG)(C)	22	COLGATE TOTAL 50ml
23	KELLOGGS FROSTED FLAKES CORN 397GR(CRT)C	23	COLGATE TOTAL CLEAN MINT 50ml
24	KELLOGGS HONEY NUT C/F 375GR (A)	24	COLGATE TOTAL FRESH STRIPE 100ML
25	KELLOGGS HONEY NUT CORN FLAKES 375GR		
26	KELLOGGS HONEY NUT CORN FLAKES 375g BOX		
27	KELLOGGS M.GRAIN CORNFLAKES 375G(A)CRT(E		
28	KELLOGGS MULTIGRAIN C/FLAKES 375GR PKT		
29	KELLOGS C.F 250GM		
30	KELLOGS C.F 375GM		
31	KELLOGS C.F 500GM		
32	KELLOGS C.F ARABIC 250GM		
33	KELLOGS C.F ARABIC NEW 375GM		
34	KELLOGS C.F. ARABIC 375GM		
35	KELLOGS C.F. ARABIC 500GM		
36	KELLOGS CRUNCHY NUT C.F.500GM		
37	KELLOGS HONEY NUT C.F.375GM		

B3. Robustness: by country and category regressions

Table A.4: Regression Summary by Product Category : GCC Region and the U.S.

(i) Distribution Measure: NUM Distribution					(ii) Distribution Measure: PCV Distribution				
	GCC Average		U.S.			GCC Average		U.S.	
	b_0	b_1	b_0	b_1		b_0	b_1	b_0	b_1
Pooled data	-7.3	4.9	-10.4	5.2	Pooled data	-8.1	5.0	-10.6	5.3
By Category					By Category				
Beans	-6.9	6.7	-9.8	5.4	Beans	-7.8	5.6	-9.8	5.3
Blades	-6.6	5.2	-8.4	5.0	Blades	-7.8	7.2	-8.7	4.9
Bouillon	-4.8	3.7	-7.9	5.7	Bouillon	-6.0	4.7	-8.0	5.6
Cereals	-7.0	6.8	-11.6	5.8	Cereals	-8.2	5.2	-11.7	5.7
Cheese	-7.4	5.0	-11.5	5.7	Cheese	-8.5	4.6	-11.6	5.6
Chewinggum	-6.9	5.1	-8.5	5.8	Chewinggum	-7.6	5.0	-8.7	5.9
Chocolate	-7.8	4.9	-11.2	5.6	Chocolate	-8.9	5.3	-11.3	5.4
Cigarette	-7.4	4.5			Cigarette	-8.0	4.8		
Cookingoil	-7.4	5.8	-9.3	5.4	Cookingoil	-8.2	5.2	-9.3	5.3
Csd	-7.6	4.1	-12.2	6.3	Csd	-8.1	4.4	-12.2	6.2
Deodorant	-8.1	9.6	-10.2	4.9	Deodorant	-8.7	4.9	-10.3	4.8
Detergents	-7.0	4.8	-10.1	4.9	Detergents	-8.0	5.2	-10.1	4.8
Dishwash	-7.1	7.3	-8.5	5.0	Dishwash	-8.1	6.3	-8.6	4.9
Energydrinks	-6.0	5.0	-11.7	6.4	Energydrinks	-6.7	5.2	-11.7	6.3
Fabricconditioner	-6.5	5.8	-9.0	4.8	Fabricconditioner	-7.6	4.8	-9.0	4.7
Insecticides	-5.6	4.9	-9.0	5.3	Insecticides	-6.4	4.6	-9.2	5.3
Juices	-8.4	5.0	-8.0	4.3	Juices	-9.0	4.8	-8.2	4.4
Liquidcordials	-6.2	6.6	-9.2	5.2	Liquidcordials	-7.4	6.6	-9.7	4.8
Malegrooming	-6.3	5.5	-9.9	5.4	Malegrooming	-7.4	5.2	-10.2	5.1
Milk	-7.6	4.9	-8.6	5.3	Milk	-8.4	4.9	-8.7	5.3
Milkpowder	-6.1	4.6	-6.3	5.0	Milkpowder	-7.2	4.8	-6.4	4.9
Powdersoftdrink	-6.5	6.5	-9.3	4.9	Powdersoftdrink	-7.7	6.0	-9.4	4.8
Shampoo	-7.8	6.6	-10.1	4.7	Shampoo	-8.8	5.1	-10.3	4.7
Skincare	-7.8	5.5	-7.6	5.6	Skincare	-9.0	4.7	-8.0	5.3
Skincleansing	-8.2	6.0	-9.1	4.7	Skincleansing	-9.2	5.3	-9.4	4.6
Suncare	-5.6	5.2	-8.5	5.0	Suncare	-6.1	3.9	-8.8	4.8
Tea	-7.5	6.0	-11.1	5.5	Tea	-8.6	5.9	-11.2	5.4
Toothbrush	-7.1	8.1	-8.7	4.5	Toothbrush	-7.9	5.5	-8.9	4.5
Toothpaste	-7.0	4.9	-9.6	4.8	Toothpaste	-8.1	5.0	-9.8	4.9
Water	-6.7	6.4	-11.2	5.6	Water	-7.4	5.6	-11.3	5.4
Summary Statistics					Summary Statistics				
Min	-8.4	3.7	-12.2	4.3	Min	-9.2	3.9	-12.2	4.4
Max	-4.8	9.6	-6.3	6.4	Max	-6.0	7.2	-6.4	6.3
Mean	-7.0	5.7	-9.5	5.3	Mean	-7.9	5.2	-9.7	5.2
Median	-7.0	5.4	-9.3	5.3	Median	-8.0	5.1	-9.4	5.1

Table A.5: By Country and Category Regressions: Numeric Distribution

	Distribution Measure: Numeric Distribution															
	b_0 (constant)							b_1 (slope)							Average	
	U.S.	KUW	QTR	BAH	OMN	UAE	KSA	U.S.	KUW	QTR	BAH	OMN	UAE	KSA	b_0	b_1
Pooled data	-10.4	-7.5	-6.9	-7.1	-7.4	-7.8	-7.4	5.2	4.7	3.8	4.5	5.2	4.8	5.0	-7.4	4.6
By Category																
Beans	-9.8	-7.6	-6.2	-6.4	-6.6	-7.7	-7.2	5.4	7.0	4.7	8.7	7.4	7.0	5.5	-7.3	6.5
Blades	-8.4	-7.6	-6.3	-6.2	-6.4	-7.0	-6.0	5.0	9.9	3.0	4.0	4.4	5.9	4.2	-6.9	5.2
Bouillon	-7.9	-5.1	-4.4	-4.0	-4.9	-5.6	-5.0	5.7	5.1	3.2	2.3	3.3	5.5	2.6	-5.3	3.9
Cereals	-11.6	-7.5	-6.2	-6.7	-7.1	-7.3	-7.5	5.8	7.3	4.4	6.3	7.2	6.7	9.1	-7.7	6.7
Cheese	-11.5	-7.7	-7.0	-7.1	-7.2	-7.5	-7.7	5.7	6.3	3.5	5.1	5.0	5.1	5.1	-8.0	5.1
Chewinggum	-8.5	-6.4	-6.1	-6.9	-7.8	-7.9	-6.5	5.8	4.8	3.7	4.5	8.1	5.0	4.3	-7.2	5.2
Chocolate	-11.2	-8.1	-7.1	-7.1	-8.2	-8.3	-7.9	5.6	5.6	3.3	4.6	6.4	4.5	5.1	-8.3	5.0
Cigarette		-7.5	-6.8	-7.3	-7.4	-8.4	-7.1		4.1	4.1	4.5	4.8	4.8	4.5	-7.4	4.5
Cookingoil	-9.3	-8.1	-6.7	-7.0	-7.1	-7.6	-7.8	5.4	8.6	3.9	4.4	6.2	4.6	7.2	-7.6	5.8
Csd	-12.2	-8.1	-7.2	-7.5	-7.8	-7.3	-7.9	6.3	4.2	3.5	4.7	4.3	3.6	4.1	-8.3	4.4
Deodorant	-10.2	-7.9	-7.4	-7.6	-8.3	-8.7	-8.6	4.9	7.2	8.2	7.6	12.1	10.1	12.2	-8.4	8.9
Detergents	-10.1	-7.1	-6.6	-6.5	-7.2	-7.3	-7.5	4.9	4.3	3.2	4.1	6.6	4.5	5.8	-7.5	4.8
Dishwash	-8.5	-7.8	-6.3	-6.4	-7.4	-7.3	-7.5	5.0	8.6	4.2	4.9	11.5	6.1	8.3	-7.3	6.9
Energydrinks	-11.7	-6.0	-5.5	-5.7	-6.1	-7.3	-5.5	6.4	5.4	4.2	4.7	5.5	5.7	4.6	-6.8	5.2
Fabricconditioner	-9.0	-7.1	-5.9	-5.7	-6.9	-6.4	-7.1	4.8	6.2	4.5	3.7	8.2	4.4	7.8	-6.9	5.6
Insecticides	-9.0	-5.7	-4.9	-5.3	-5.9	-5.9	-5.8	5.3	5.7	3.1	4.4	6.2	4.8	5.1	-6.1	4.9
Juices	-8.0	-8.8	-7.9	-8.0	-8.4	-9.0	-8.6	4.3	4.6	3.8	4.5	5.5	5.5	5.9	-8.4	4.9
Liquidcordials	-9.2	-6.7	-5.3	-5.9	-6.3	-6.2	-6.6	5.2	7.1	4.9	6.3	6.4	6.5	8.4	-6.6	6.4
Malegrooming	-9.9	-6.6	-5.8	-6.1	-6.1	-7.3	-5.9	5.4	6.0	3.3	4.4	5.0	8.7	5.5	-6.8	5.5
Milk	-8.6	-8.2	-6.9	-7.5	-7.8	-7.9	-7.5	5.3	5.7	4.0	4.8	5.2	5.3	4.3	-7.8	4.9
Milkpowder	-6.3	-6.1	-5.6	-5.5	-6.4	-6.7	-6.5	5.0	4.7	3.5	3.4	6.1	5.0	5.1	-6.1	4.7
Powdersoftdrink	-9.3	-7.4	-6.0	-5.8	-6.6	-6.4	-6.9	4.9	7.4	6.8	5.0	6.2	6.2	7.4	-6.9	6.3
Shampoo	-10.1	-7.6	-7.1	-7.5	-7.6	-8.1	-9.0	4.7	6.3	4.5	5.6	6.3	5.7	11.0	-8.1	6.3
Skincare	-7.6	-7.9	-7.3	-7.5	-8.0	-8.3	-7.7	5.6	5.1	3.5	4.5	6.4	6.0	7.8	-7.8	5.5
Skincleansing	-9.1	-8.3	-7.7	-8.0	-8.3	-8.6	-8.3	4.7	7.5	4.5	5.5	6.6	5.7	6.3	-8.3	5.8
Suncare	-8.5	-6.4	-5.3	-5.2	-5.3	-6.6	-4.7	5.0	7.8	4.4	4.4	3.9	6.4	4.6	-6.0	5.2
Tea	-11.1	-7.9	-7.2	-7.1	-7.4	-8.0	-7.6	5.5	6.1	4.9	5.2	7.4	6.6	5.9	-8.0	5.9
Toothbrush	-8.7	-7.9	-6.1	-6.5	-6.7	-8.1	-7.0	4.5	15.5	2.8	4.1	5.0	10.6	11.0	-7.3	7.6
Toothpaste	-9.6	-7.1	-6.3	-6.7	-7.2	-7.5	-7.0	4.8	6.1	3.1	4.2	5.5	5.1	5.6	-7.3	4.9
Water	-11.2	-6.9	-6.1	-6.2	-6.2	-7.3	-7.7	5.6	6.5	5.0	6.3	5.1	5.2	10.3	-7.4	6.3
Summary Statistics																
Min	-12.2	-8.8	-7.9	-8.0	-8.4	-9.0	-9.0	4.3	4.1	2.8	2.3	3.3	3.6	2.6	-8.4	3.9
Max	-6.3	-5.1	-4.4	-4.0	-4.9	-5.6	-4.7	6.4	15.5	8.2	8.7	12.1	10.6	12.2	-5.3	8.9
Mean	-9.5	-7.3	-6.4	-6.6	-7.0	-7.4	-7.1	5.3	6.5	4.1	4.9	6.3	5.9	6.5	-7.3	5.6
Median	-9.3	-7.5	-6.3	-6.6	-7.1	-7.4	-7.3	5.3	6.1	4.0	4.6	6.2	5.6	5.7	-7.4	5.3

Table A.6: By Country and Category Regressions: Product Category Volume

	Distribution Measure: Product Category Volume															
	b_0 (constant)							b_1 (slope)							Average	
	U.S.	KUW	QTR	BAH	OMN	UAE	KSA	U.S.	KUW	QTR	BAH	OMN	UAE	KSA	b_0	b_1
Pooled data	-10.6	-8.7	-7.9	-8.0	-8.4	-8.9	-8.7	5.3	4.8	4.7	5.0	5.1	5.0	5.0	-8.4	4.9
By Category																
Beans	-9.8	-8.6	-7.1	-7.1	-7.3	-8.4	-8.3	5.3	5.9	5.5	5.6	5.6	5.8	5.1	-8.1	5.5
Blades	-8.7	-8.4	-7.6	-7.4	-7.3	-8.3	-7.7	4.9	7.2	6.1	7.1	6.3	7.3	9.0	-7.9	6.8
Bouillon	-8.0	-6.3	-5.3	-4.8	-6.1	-6.8	-6.4	5.6	6.3	4.3	3.3	4.3	5.9	4.1	-6.3	4.8
Cereals	-11.7	-8.2	-8.3	-7.8	-8.5	-8.3	-8.1	5.7	4.8	5.6	5.3	5.5	4.7	5.2	-8.7	5.3
Cheese	-11.6	-8.8	-8.3	-8.4	-8.1	-8.6	-9.0	5.6	4.6	4.6	4.8	4.4	4.6	4.9	-9.0	4.8
Chewinggum	-8.7	-7.5	-6.6	-7.4	-8.3	-8.2	-7.4	5.9	5.1	4.1	4.9	6.6	4.8	4.6	-7.7	5.1
Chocolate	-11.3	-9.5	-8.4	-8.0	-9.2	-9.0	-9.3	5.4	5.6	4.8	5.1	6.1	4.5	5.5	-9.2	5.3
Cigarette		-8.2	-7.5	-8.0	-7.9	-8.6	-7.9		4.7	4.5	5.0	4.8	4.6	5.0	-8.0	4.8
Cookingoil	-9.3	-8.8	-7.5	-7.8	-7.9	-8.5	-8.8	5.3	5.3	4.6	5.2	5.1	4.9	5.9	-8.3	5.2
Csd	-12.2	-8.9	-7.4	-8.1	-8.2	-7.9	-8.3	6.2	4.9	3.7	4.7	4.6	4.2	4.3	-8.7	4.6
Deodorant	-10.3	-8.1	-8.4	-8.6	-9.0	-9.4	-9.0	4.8	3.4	4.9	5.8	5.5	5.1	5.0	-9.0	4.9
Detergents	-10.1	-8.3	-7.3	-7.6	-8.0	-8.4	-8.7	4.8	5.3	4.1	5.3	5.8	5.3	5.6	-8.3	5.2
Dishwash	-8.6	-8.8	-7.2	-7.8	-7.8	-8.3	-8.9	4.9	6.4	5.4	7.6	6.3	5.6	6.4	-8.2	6.1
Energydrinks	-11.7	-6.7	-6.1	-6.3	-6.7	-8.1	-6.0	6.3	5.1	4.6	5.1	5.1	6.1	5.0	-7.4	5.4
Fabricconditioner	-9.0	-8.3	-7.0	-7.1	-7.7	-7.6	-8.0	4.7	5.2	5.2	4.4	5.1	4.4	4.7	-7.8	4.8
Insecticides	-9.2	-6.2	-5.5	-6.0	-6.7	-6.9	-7.1	5.3	4.1	3.7	4.6	5.2	4.9	5.3	-6.8	4.7
Juices	-8.2	-9.6	-8.6	-8.4	-8.8	-9.5	-9.2	4.4	4.8	4.5	4.3	5.0	5.0	5.1	-8.9	4.7
Liquidcordials	-9.7	-8.0	-6.4	-7.2	-8.0	-7.5	-7.5	4.8	6.5	6.0	6.0	7.4	6.3	7.5	-7.8	6.4
Malegrooming	-10.2	-7.6	-6.7	-7.3	-7.6	-8.3	-6.9	5.1	5.1	4.4	5.8	5.6	5.8	4.5	-7.8	5.2
Milk	-8.7	-9.3	-7.6	-7.9	-8.5	-8.8	-8.6	5.3	5.5	4.7	4.6	5.0	5.2	4.5	-8.5	5.0
Milkpowder	-6.4	-7.4	-6.5	-6.5	-7.3	-7.9	-7.6	4.9	4.8	4.5	4.4	5.6	5.0	4.5	-7.1	4.8
Powdersoftdrink	-9.4	-8.3	-7.3	-7.5	-7.4	-7.6	-8.0	4.8	5.8	7.1	5.8	6.2	5.7	5.6	-7.9	5.9
Shampoo	-10.3	-8.8	-8.1	-8.5	-8.4	-9.2	-10.0	4.7	4.8	5.1	5.3	4.8	4.7	5.8	-9.0	5.0
Skincare	-8.0	-9.0	-8.6	-8.8	-9.3	-9.7	-9.0	5.3	4.2	4.4	5.1	5.3	4.8	4.6	-8.9	4.8
Skincleansing	-9.4	-9.2	-8.9	-9.1	-9.0	-9.7	-9.5	4.6	4.8	5.6	6.0	5.1	5.3	5.3	-9.3	5.2
Suncare	-8.8	-6.5	-5.6	-5.4	-6.7	-7.1	-5.3	4.8	4.4	3.7	2.5	4.2	4.2	4.4	-6.5	4.0
Tea	-11.2	-9.0	-8.5	-8.0	-8.4	-9.0	-8.8	5.4	5.8	5.9	5.7	6.2	6.0	5.7	-9.0	5.8
Toothbrush	-8.9	-8.2	-7.3	-7.7	-7.9	-8.7	-7.7	4.5	5.1	5.8	5.7	5.9	5.7	5.0	-8.1	5.4
Toothpaste	-9.8	-8.0	-7.6	-7.7	-8.2	-8.9	-8.4	4.9	4.4	5.0	5.2	5.1	5.7	4.9	-8.4	5.0
Water	-11.3	-8.0	-7.0	-7.0	-6.7	-7.8	-7.8	5.4	5.7	6.0	6.0	4.9	4.7	6.6	-7.9	5.6
Summary Statistics																
Min	-12.2	-9.6	-8.9	-9.1	-9.3	-9.7	-10.0	4.4	3.4	3.7	2.5	4.2	4.2	4.1	-9.3	4.0
Max	-6.4	-6.2	-5.3	-4.8	-6.1	-6.8	-5.3	6.3	7.2	7.1	7.6	7.4	7.3	9.0	-6.3	6.8
Mean	-9.7	-8.2	-7.3	-7.5	-7.9	-8.4	-8.1	5.2	5.2	4.9	5.2	5.4	5.2	5.3	-8.1	5.2
Median	-9.4	-8.3	-7.3	-7.7	-8.0	-8.4	-8.2	5.1	5.1	4.8	5.2	5.2	5.1	5.1	-8.1	5.1

C. Figure Appendix

C1. Robustness: facts on retailers' heterogeneity (UAE and four main products only)

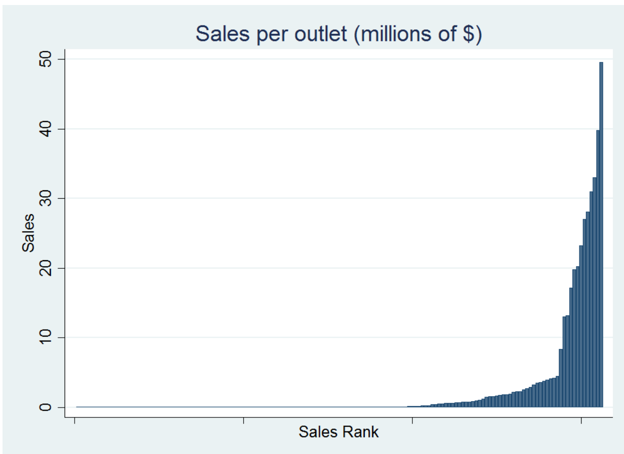

(a)

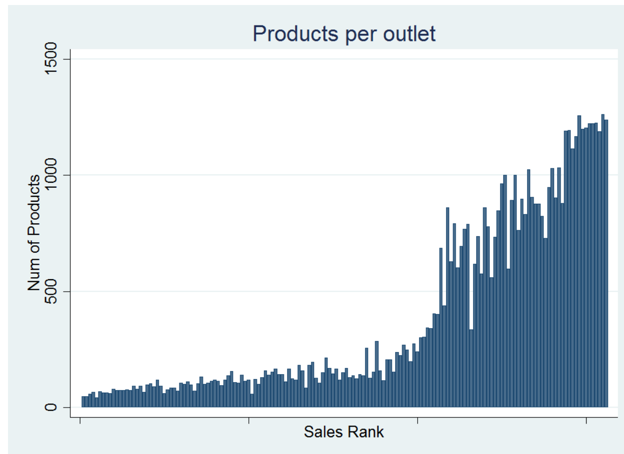

(b)

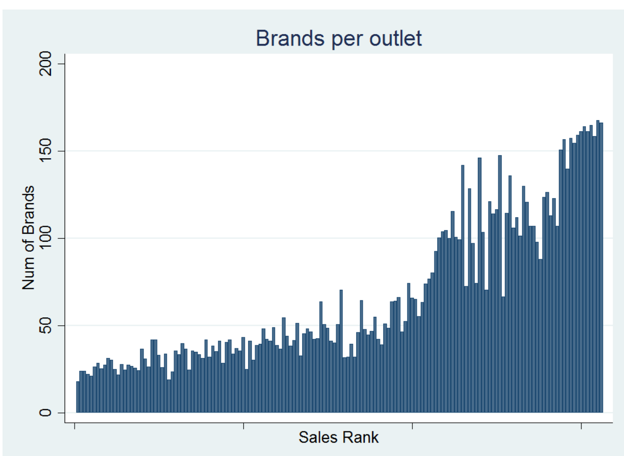

(c)

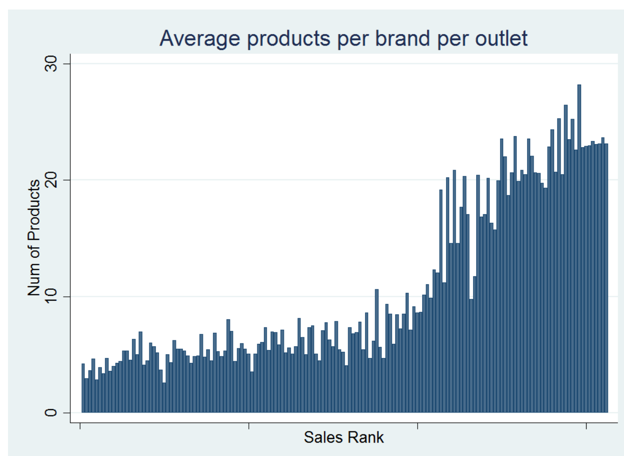

(d)

Note: See Figure 2.

Figure A.3: Facts on retailers' heterogeneity (four main products only), UAE

C2. Estimated inflation series based on alternative product selection to ICP goods

In case two or more items at a store fit the same product definition provided in the ICP product list, we consider alternative practical rules to select the product to represent the ICP product based on prices. For example, we take the minimum, maximum, average, median, or a random price among all products that fit the ICP definition. The variable "mii" indicates that the price of the most-important-item is used, where the most-important-item refers to the product with the largest expenditure share.

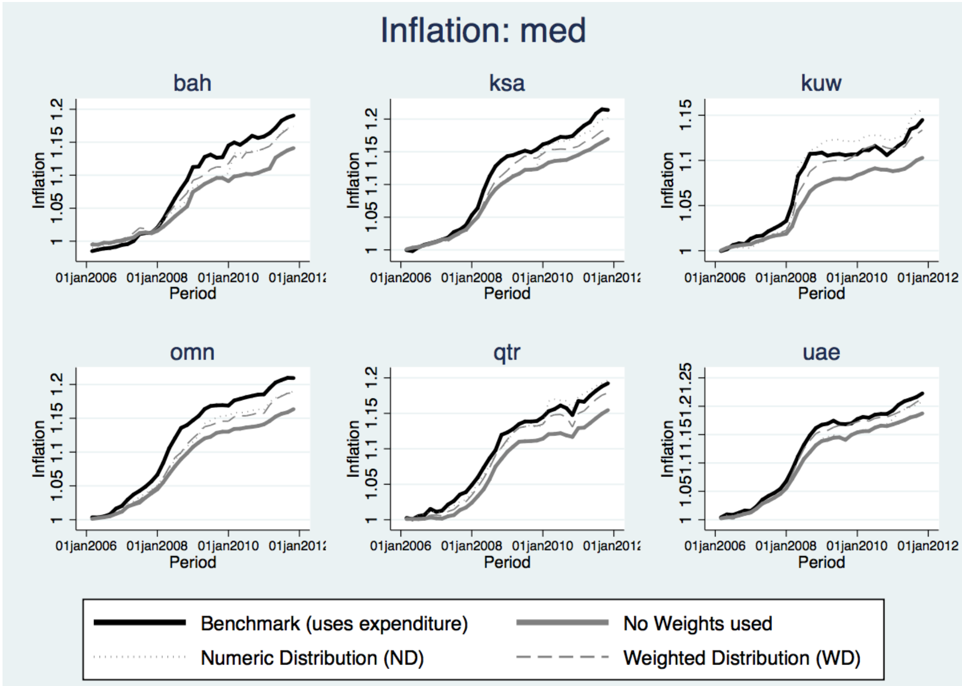

Figure A.5: Criteria of Median Price

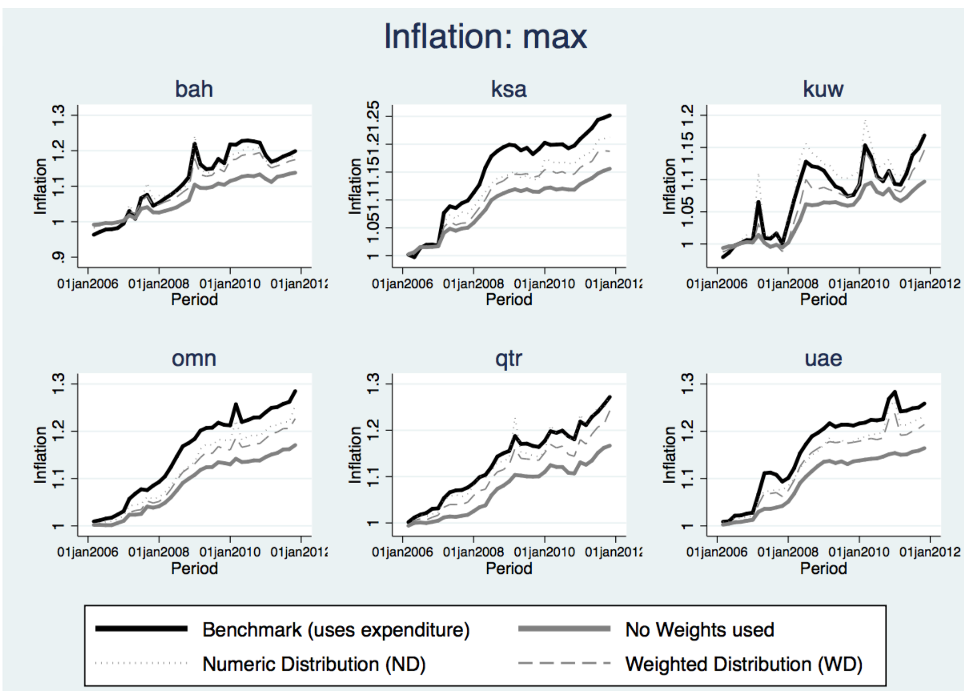

Figure A.4: Criteria of Maximum Price

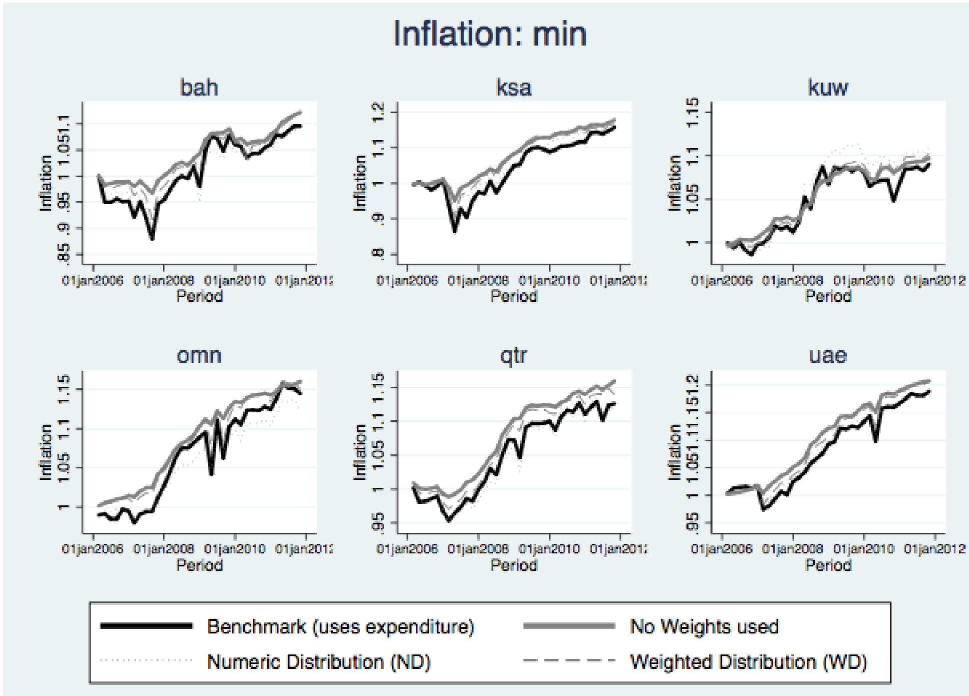

Figure A.6: Criteria of Minimum Price

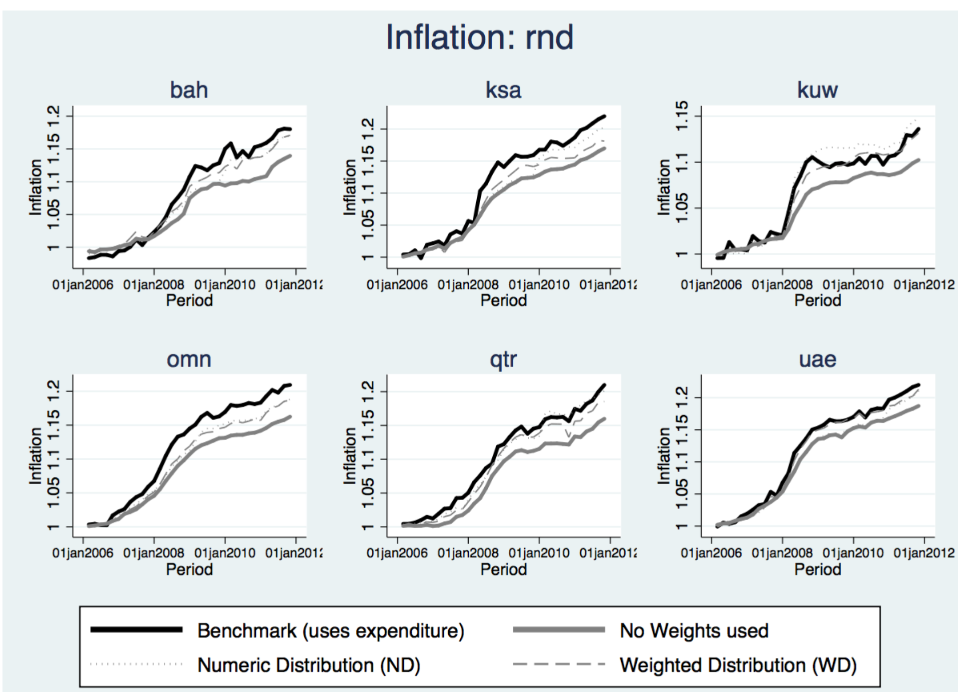

Figure A.7: Criteria of Price for a Random Item

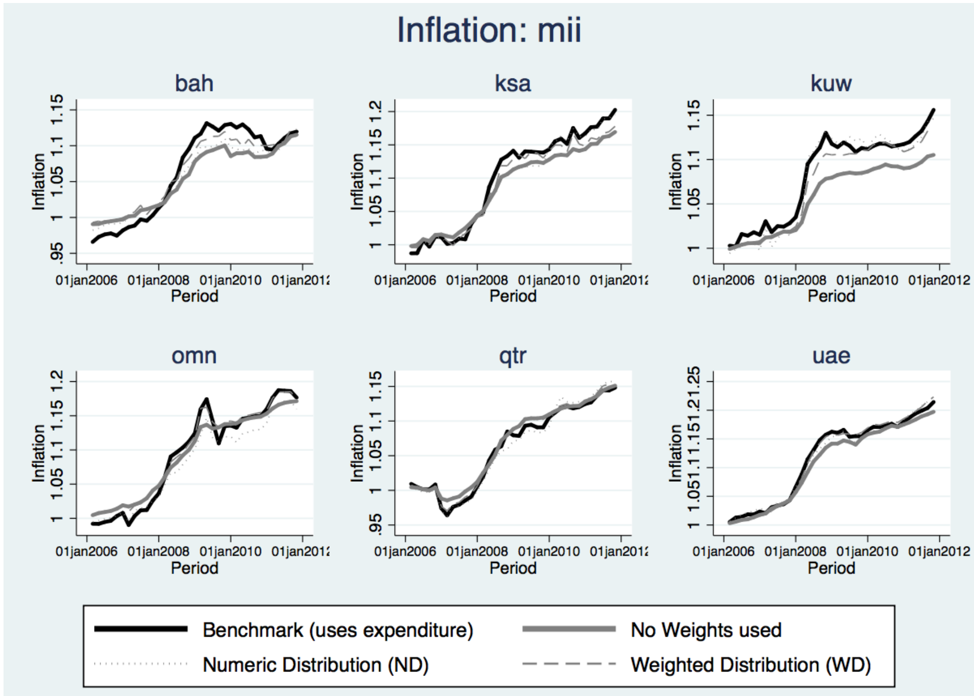

Figure A.8: Criteria of Price for the Most Important Item