NBER WORKING PAPER SERIES

USING THE RETAIL DISTRIBUTION OF SELLERS TO IMPUTE EXPENDITURES SHARES

Alexis Antoniades
Robert C. Feenstra
Mingzhi Xu

Working Paper 28902
http://www.nber.org/papers/w28902

NATIONAL BUREAU OF ECONOMIC RESEARCH
1050 Massachusetts Avenue
Cambridge, MA 02138
June 2021

We thank seminar participants at the American Economic Association, ICP/PPP, ESCOE, the University of Edinburgh, and Peking University for comments. This research was made possible through the support of an NPRP grant from the Qatar National Research Fund. Researchers own analyses calculated based in part on data from Nielsen Consumer LLC and marketing databases provided through the NielsenIQ Datasets at the Kilts Center for Marketing Data Center at The University of Chicago Booth School of Business. The conclusions drawn from the NielsenIQ data are those of the researcher(s) and do not re ect the views of Nielsen. Nielsen is not responsible for, had no role in, and was not involved in analyzing and preparing the results reported herein. Financial support from the China Government Scholarship (CSC) for a doctoral study and a NBER post-doctoral fellowship are gratefully acknowledged. The views expressed herein are those of the authors and do not necessarily reflect the views of the National Bureau of Economic Research.

Using the Retail Distribution of Sellers to Impute Expenditures Shares
Alexis Antoniades, Robert C. Feenstra, and Mingzhi Xu
NBER Working Paper No. 28902
June 2021
JEL No. C81,E01,E31

ABSTRACT

Many price indices must be constructed without quantity data at the elementary level. We show that for some consumer goods in the United States and other countries, one can approximate expenditure shares using weights derived from the retail distribution of sellers. These weights are based on the share of outlets selling an item, or the share of outlets adjusted by the total number of items sold in each. Relative to using no weights, we find that using such imputed weights substantially reduces bias in the frequency of price changes, in annual inflation, and in price comparisons across countries.

Alexis Antoniades
School of Foreign Service
Georgetown University - Qatar
Education City
PO Box 23689
Doha
Qatar
aa658@georgetown.edu

Robert C. Feenstra
Department of Economics
University of California, Davis
One Shields Avenue
Davis, CA 95616
and NBER
rcfeenstra@ucdavis.edu

Mingzhi Xu
Science Building Building No. 1 Room 423N
Peking University
5th Yiheyuan Road
Haidian District
China
mingzhixu@nsd.pku.edu.cn

1 Introduction

Many price indices, whether constructed using data collected from visiting stores or from scraping the web, lack item-specific quantity data for the construction of elementary expenditure aggregates. In the absence of such explicit expenditure weights, elementary aggregates are calculated by taking an unweighted average of the prices that comprise them. While in some cases implicit weights can be introduced through the sampling of the items to be included in the elementary aggregates, this is not the case with indices calculated using data from the web that lack both implicit and explicit weights.[1]

Treating however all items equally is highly problematic and can introduce biases, as we shall demonstrate. Consumers purchase a variety of products, but they exhibit strong preference 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 using scanner data on some consumer goods, the total sales of the top 2% best sellers per product category can account for as much as the total sales of the bottom 96%.[2]

The unavailability of expenditure data at the elementary (or "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 parity parities (PPPs).[3] In the 2011 round, the ICP collected price data from across 199 countries and regions. These data are used to provide direct comparisons of well-being, to compare growth rates by sector, to report price levels, and to assess poverty rates. Moreover, PPP-based GDP is used by the IMF to determine voting rights, quota subscriptions, and financing amounts for its county members. It is also used by the IMF to produce the *World Economic Outlook*. 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,

[1]See Cavallo and Rigobon (2016) in The Billion Prices Project (BPP) for an overview of how online price data offer a promising and rich new source of information for informing economic studies. Because price data collected online are up to date, easy to obtain, and they cover a very large number of retailers, locations, and products, their use in empirical work can be found across several economics disciplines. Cavallo and Rigobon (2011) and Cavallo (2018) use the data to study the distribution of price changes. Gorodnichenko and Weber (2016) 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. An exception to the studies that use online price data with no information on expenditure or quantities is Goolsbee and Klenow (2018). The authors use online transaction data that include both price and quantity to document a lower inflation in online prices than in the CPI. But their data is proprietary and not available for scraping.

[2]Computations are based on sales of Fast Moving Consumer Goods (FMCGs) – defined in note 7 – 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, and results for the United States are similar. A detailed description of the data follows in section 3.

[3]http://blogs.worldbank.org/opendata/world-bank-publish-purchasing-power-parities-december013.

no weight information is collected in the ICP to reflect the quantities of items sold, so all items within the same basic heading are treated equally.[4]

In this paper we show that one can approximate item-specific expenditure weights using weights that can be easily constructed from the retail distribution of sellers. These weights are based on the shares of stores selling an item, or on the share of stores adjusted by the number of all items in those stores. 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. Relative to using no weights, we finds that using using weights imputed from the retail distribution substantially reduces bias in the frequency of price changes, in annual inflation, and in ICP price comparisons across countries.

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 et al. (2020)), or by scraping data from various retailers' web sites (e.g. Cavallo and Rigobon (2011, 2016); Cavallo (2013, 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 (i.e. selling outlets) per item. Because these datasets contain multiple prices per item – reflecting the many outlets an item is available at -- but are not balanced as some items are sold in some outlets while others are not, we can easily obtain measures of the retail distribution of sellers that can then be used to impute expenditure shares. Specifically, using only price information we can construct a metric for retailer distribution by dividing the number of outlets carrying a particular items over the total number of outlets in our sample. We call this metric the *numeric distribution* (*ND*). We can also construct a measure of distribution that takes into account the size of each outlet, where the count of price observations (i.e. items) per retailer at given a point in time 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 item by assuming – as supported by the data – a convex relation between market share and item distribution. That is, we assume that items available in more retailers have higher sales, and that as the number of retail sellers rises then the market share rises at an increasing rate. We use an exponential function to capture the relation between the number of sellers and the the market share, which we argue is a parsimonious choice among possible convex functions.[6] These imputed market shares can then be used to weight price observations and reduce measurement errors in calculations and/or

[4] In both the ICP and in studies that use online data, expenditure data from surveys and from 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.

[6] See note 18. Parameter selection is discussed later in the paper.

estimation.

We illustrate the approach and evaluate its performance in terms of reducing measurement error by considering three important applications: (i) measuring the frequency of price changes, (ii) measuring inflation, and (iii) measuring international price differences. These applications are chosen because they cover a large, important, diverse, and active body of work that is of great interest to academics, policy makers, and practitioners.

For these exercises we use scanner data for a sample of Fast Moving Consumer Goods (FMCGs) sold in the United States and the Gulf Cooperation Council (GCC) countries between 2006 and 2011.[7] First, we compute or estimate the measure of interest using actual prices and quantities. We set the outcome of this estimation to be 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 items are treated (and weighted) equally. This approach mimics the approach taken by researchers working with online and with 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 items within an item 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 again but this time each observation is weighted by an imputed market share derived from estimating the retail distribution from the price data.

In the first application we show that using only prices understates the true frequency of price changes both in the US and in the GCC. This happens because as we document here, items with higher sales experience more frequent price changes. When only price data is used to compute the frequency of price changes, prices of a handful of items that account for the majority of sales and experience frequent price changes are marginalized by the vast number of all other items with infrequent price changes that account for very little sales.

To correct for this measurement error we repeat the exercise by ignoring again expenditure information (which would not have been known to researchers working with online price data), but instead we use information on the 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 the US and each of the GCC countries between January 2006 and December 2011 for FMCGs. By not using expenditure information we *understate* inflation in each country. Specifically, the democratic measure of inflation yields lower inflation by 0.5 percentage points annually in the US and by 0.3 percentage points annually in the GCC as compared to the plutocratic measure that takes into account actual expenditure by

[7] Fast Moving Consumer Goods (FMCGs) is a term from the marketing literature and also used by analytics firms such as AC Nielsen to refer to relatively inexpensive consumer products that are sold quickly. Nearly all goods sold in supermarkets and convenience stores of all sizes are included within FMCGs. Our sample of goods in this paper will focus on 30 product groups of FMCGs within the GCC countries and a set of comparable goods for the United States.

consumers. However, when we compute inflation using prices weighted by an imputed measure of expenditure share derived from the retail distribution metrics, inflation is no longer understated by as much: the downward bias/measurement error is reduced by at least one-half for the United States and two-thirds or more for the GCC countries.

In the third and final application we compute purchasing power parities (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 an outlet fit the same item definition provided in the ICP item list (for example, take the minimum, maximum, average, median, or a random price among all items that fit the definition).

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. Specifically, while we estimate prices to differ by 6% on average among the GCC countries when both prices and expenditure information is 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 the weighted distribution are used, we find prices to differ by 9% and 7%, respectively, hence, reducing measurement error by 75% over the case when only prices are used.

To summarize, in the absence of any information on quantities or on expenditure, using data on the retail distribution provides a very good proxy for expenditure. Even with data from a limited sample of outlets, the convexity between distribution and market share allows us to successfully separate the most important items 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 used in the applications and in section 4 we present the applications. Section 5 concludes and provides directions for further research.

2 The 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):

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

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

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

Numeric distribution (*ND*), also known as physical distribution, reports the share of outlets carrying a particular item. It is the least data-intense measure of the three metrics, but it does not distinguish between outlets with high sales and low sales. *All commodity volume* (ACV) and *product category volume* (PCV) take into account variation in outlet 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 et al. (1989) tortilla chips and instant coffee in the US; Borin et al. (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:

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

For all categories but frozen pizza, they confirmed that a convex relation characterizes market shares, defined as item dollar sales over total category sales, and the retail distribution 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 the cases tested. Similarly to the studies above, we also tested for and confirmed the presence of convexity between market share and the retail distribution on 30 product categories of FMCGs for each of the six

[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.

GCC countries used in our sample.

The evidence suggests that in the absence of any information on item expenditure, the retail distribution can be used to impute market share. 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 the retail distribution solely from price data. In the absence of expenditure data, the *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 item variety as an indicator for outlet size. Specifically, we define the *weighted distribution* (WD) as:

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

Counting the number of products offered by an outlet is a good indicator of its sales. Even if the number of products sold in an outlet is not available, some simpler measures of outlet size might be used.[9] More generally, if universal barcode information is not available to measure the retail distribution, we believe that other information can be used. For example, consider the case of a retail chain that uses its own stock-keeping units rather than universal barcodes for tracking its product items. We have confirmed through examples that: (a) the number of outlets within a chain carrying a particular item vary, with more well-known product brands carried in more locations;[10] (b) there is a convex relationship between an item's *overall* market share and the number of outlets carrying it within the chain.[11]

2.1 An Illustration

Does the availability of products in online data exhibit the same convex relation between market share and the retail distribution as documented in the studies using scanner data? In theory, they should when the online data are based on the items that are available in physical outlets. However, as far as we know, no study has formally put this hypothesis to the test.[12] 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 outlet data with actual expenditure data. The online

[9]See section 4.4 where we discuss this issue in the context of ICP price surveys.

[10]In this example, we collected data for a grocery store chain in the United States with about 350 outlets across seven states.

[11]In this example, we analyzed data from two retail chains in the United Arab Emirates, one with only 10 outlets and another with 43 outlets. Using the Nielsen barcode data to identify items and to measure *overall* market share, we confirm a convex relationship between market share and the retail distribution (ND or WD) computed over *all* retail chains (as in this paper) or over only the chain in question. These results are available on request.

[12]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 the retail distribution and market share, which is what we need to confirm here.

data allow us to construct the two measures of the 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 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 items from a small sample of reported outlets in each city.[13]

The data contain no information on expenditure. The number of outlets included in the dataset varies by city and ranges from 3 to 12. No information on the inclusion or exclusion of outlets exists. Furthermore, we do not know if we can proxy for outlet size with the number of listed items (barcodes) per outlet in the online application. If no price for a given barcode is reported at a particular outlet, this means either that the item is not sold there or that it is sold but the price is not uploaded. Therefore, counting the number of items per outlet may not be a good proxy for actual outlet 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 the 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 the retail distribution.

We use the mobile app dataset to compute the two measures of the retail distribution, and we use the Nielsen dataset to compute market share for each item. We then merge the two datasets, and for the majority of the barcodes found in the mobile application, we now observe both the retail distribution and market share . Finally, we allocate items into bins based on market share and take the median retail distribution (ND or WD) across all items in each market share bin. Scatter plots of the retail distribution and market share for each product category are reported, first using the the numeric distribution (ND) in Figure 1, panel (a), and then using the weighted distribution (WD) in Figure 1, panel (b). The figures confirm that the relation between market share and the computed retail distribution is convex, despite the data issues discussed above that may have compromised our ability to measure the retail distribution accurately.

The implications are significant. First, scholars can exploit the convex relation between market share and the retail distribution to impute market shares from information obtained only from prices, and they can then use imputed market shares to assign item-specific weights. We elaborate more on this next as we consider three important applications. Second, even when data are available from a small sample of outlets, as in our example above using the Chinese data, measures of the retail distribution still provide a very good proxy for expenditure shares.

[13]More information on the dataset is provided in Feenstra et al. (2020).

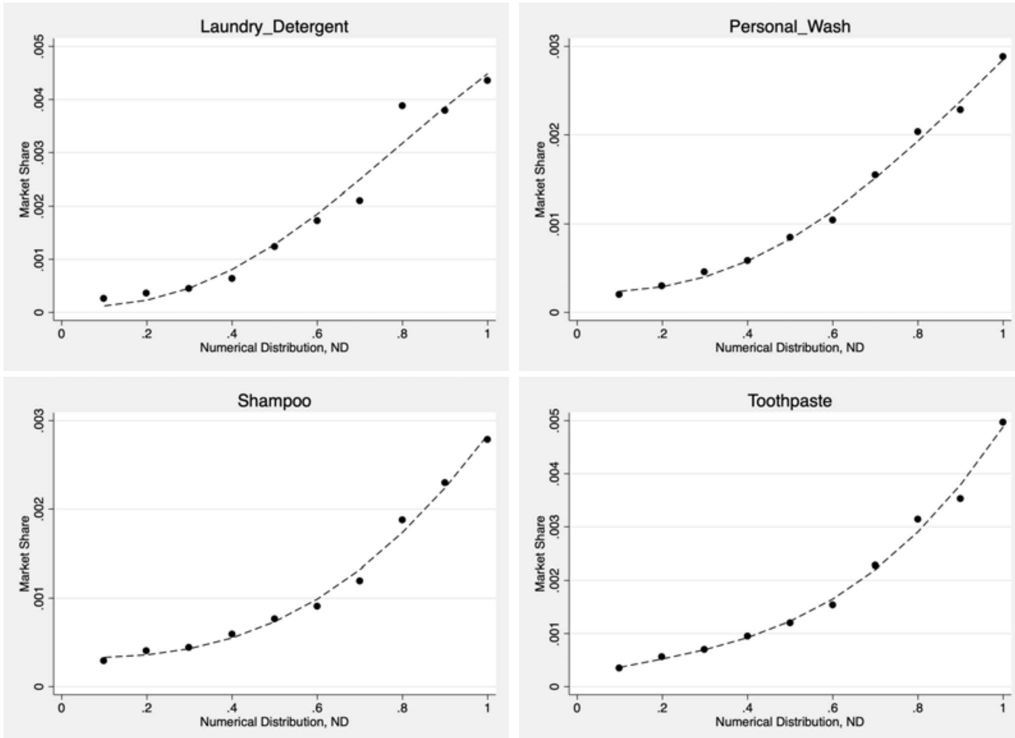

(a) Numeric Distribution

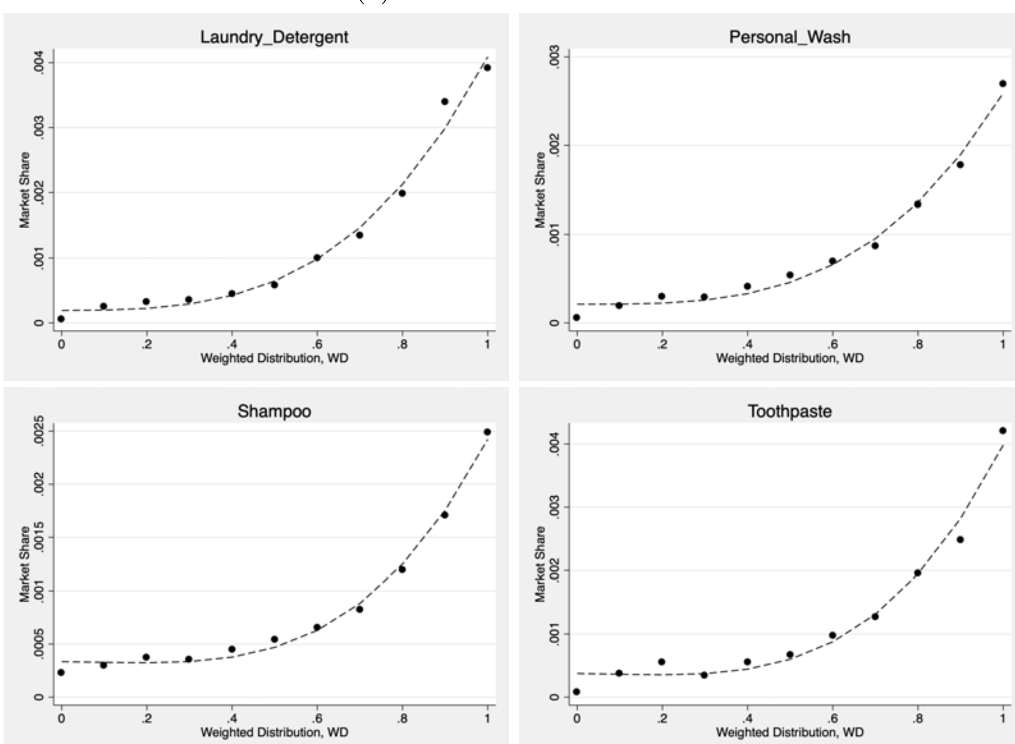

(b) Weighted Distribution

Notes: We use price data from an online app in China to compute the retail distribution (*ND* and *WD*) and Nielsen scanner data to compute market shares for items (barcodes) in the laundry detergent, shampoo, personal wash items, and toothpaste categories. Items are allocated into bins based on the retail distribution (x-axis) and the average market share of all items within a bin is plotted on the graph, against measures of the the numeric distribution (*ND*; panel (a)), and the weighted distribution (*WD*; panel (b)). For more information on the data, see Feenstra et al. (2020).

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

3 Data Description

So far we have (i) drawn attention to the convex relation that exists between the 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 also 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.

Our sample comes from AC Nielsen store scanner data for FMCGs in the United States and the six GCC countries: Bahrain (BAH), Kuwait (KUW), Oman (OMN), Qatar (QTR), Saudi Arabia (KSA) and United Arab Emirates (UAE). The price and quantity information for thousands of items (barcodes) across 30 product groups between January 2006 and December 2011 are used.[14] The frequency of data for the GCC countries is monthly or bi-monthly, and according to Nielsen, these data cover about 85% of all the FMCGs consumed in those countries.

Three important characteristics of the GCC dataset are worth highlighting. First, data are provided for each outlet, across thousands of outlets. By analyzing the data at the outlet 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 of price changes across periods without measurement error.[15] Third, most of the items 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. We confirm that the characteristics of the data from these countries is similar to that from the United States, and which suggests that the findings we present below can be generalized with some confidence to other economies, too.[16]

For the United States, the Nielsen data include observations from five channels: (i) Convenience, (ii) Drug, (iii) Food, (iii) Mass Merchandiser, and (v) Liquor outlets.[17] To make the analysis between US and GCC data more comparable we primarily focus on Food observations, unless otherwise stated. In addition, we pick 129 product modules that closely match the 30 product groups available in the GCC data (the definitions of product groups between the US and the GCC

[14] 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.

[15] In some 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).

[16] 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.

[17] Data at the barcode level for the United States is available by subscription from Nielsen at the Kilts Center for Marketing, https://research.chicagobooth.edu/nielsen.

datasets vary). These modules belong to 21 US product groups of five department groups. Because the frequency of US data is weekly while that of the GCC is monthly or bimonthly, we convert US data to monthly by aggregating up expenditure within a month and considering the price during the first week of the month for each barcode at each outlet.

Descriptive statistics for the GCC and US datasets are provided in Table 1. In total, the dataset provides price and quantity information on 203,719 (non-unique) items sold across 5,851 outlets in the GCC and 111,733 items across 11,046 stores in the US over a six-year period. In the GCC, Qatar and Bahrain are the smallest economies in terms of population, and Saudi Arabia the largest. A closer examination of the data reveals that the majority of items do not exist across all periods and all outlets and that in each product category or product module sales of FMCGs are highly concentrated in a handful of items.

Table 1: Brand Descriptive Statistics

Country	Categories	Products	Outlets	Start Date	End Date
Bahrain	30	24,259	311	Jan 2006	Dec 2011
Kuwait	30	37,660	285	Jan 2006	Dec 2011
Oman	30	40,165	614	Jan 2006	Dec 2011
Qatar	30	24,150	267	Jan 2006	Dec 2011
Saudi Arabia	30	34,447	3,398	Jan 2006	Dec 2011
United Arab Emirates	30	43,038	976	Jan 2006	Dec 2011
United States	21*	111,373	11,046	Jan 2006	Dec 2011

Note: Scanner data for the GCC countries provided by AC Nielsen while for the U.S. by the Kilts Center at the University of Chicago and span sales of FMCGs between 2006 and 2011. 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. * To make the results comparable across the U.S. and the GCC we match the 30 product categories in the GCC with the corresponding US data. The matching yields 129 product modules that belong to 21 product groups and five department groups.

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

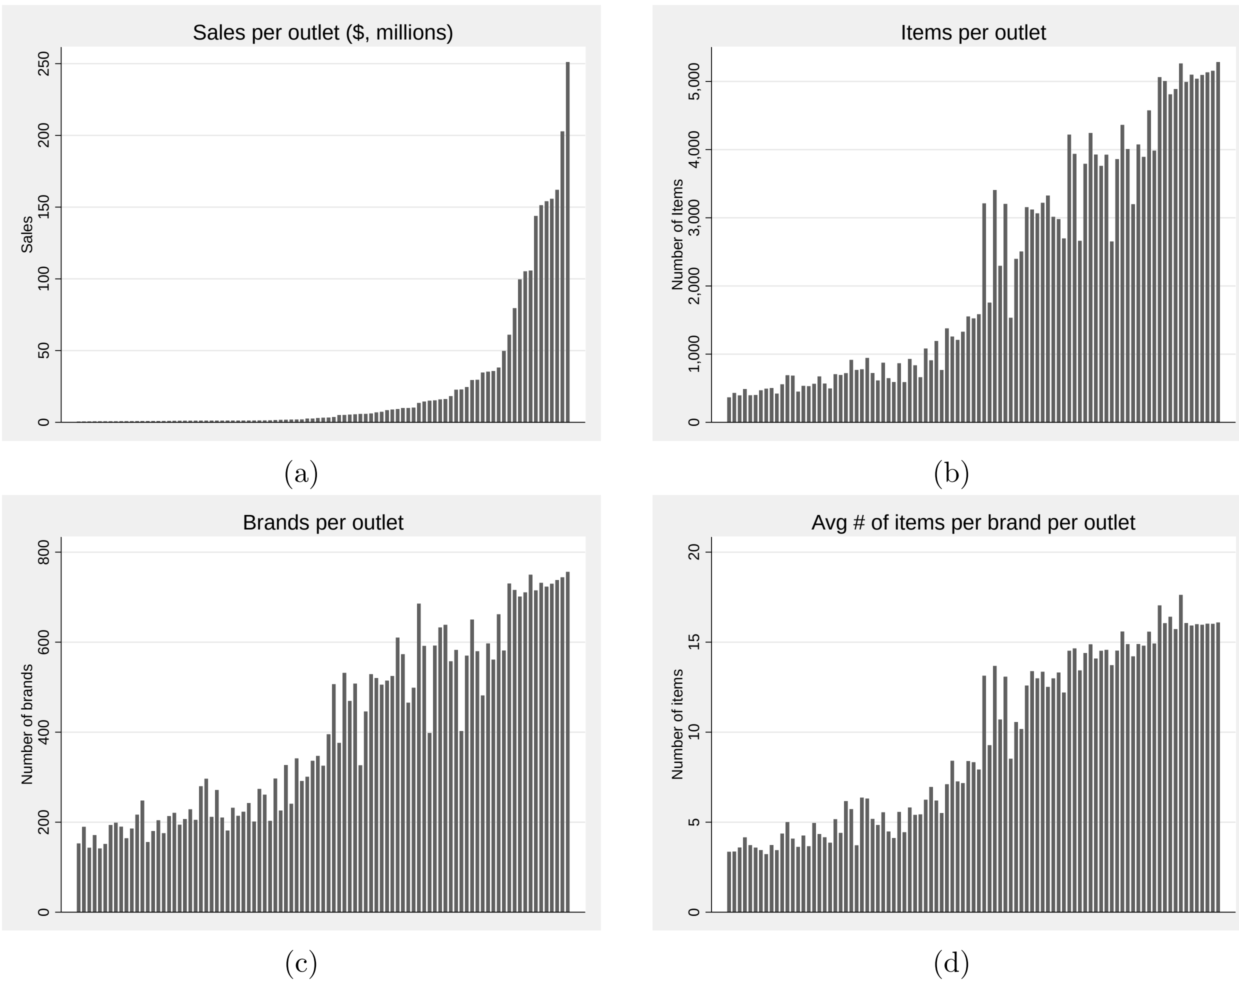

Notes: Retail outlets in UAE are ranked based on average monthly sales (panel (a)). Monthly sales are computed from Nielsen scanner data across 30 product groups between January 2006 and December 2011. Items per outlet (panel (b)), brands per outlet (panel (c)), and average items per brand per outlet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

Figure 2: Facts on Retailers' Heterogeneity, United Arab Emirates

4 Applications

4.1 Imputing Expenditure Shares

We proceed to illustrate how one can exploit the relation between the retail distribution and market size to obtain a proxy for expenditure when expenditure is not observed but prices are. As in the evidence from the marketing literature, we assume a convex relation between market share and the retail distribution. In Online Appendix C 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. As featured in the model, it is the interaction between heterogeneous firms and varying "slotting fee" that yields the convex relation that is observed in the data. We show that assuming heterogeneity in the "slotting fee" incurred by manufacturers is sufficient to generate the convex relation between sales and the distribution measure, which is robust to alternative market structures.

We compare results with no item-specific weights assigned and with (actual based on expenditure shares or imputed from the retail distribution metrics *ND* and *WD*) apply this to three different applications: (i) measuring the frequency 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 item 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 thus, do not assign item-specific weights. 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 the retail distribution extracted solely from price data to impute market shares. These imputed market shares are then used to weight items in the estimations. The outcome of this estimation allows us to test whether our proposed approach, namely, using the retail distribution weights as a proxy for expenditure, reduces measurement error and by how much. We first impute shares using *ND* and then using *WD*.

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

$$\text{Market Share} = \exp\left(a + bND\right) \tag{6}$$

or

$$\text{Market Share} = \exp\left(a + bWD\right) \tag{7}$$

where *ND* and *WD* are item-specific measures of the the numeric and weighted distributions, respectively, and the two regressions need not have the same coefficients.

Table 2: Market Share and Distribution Regression Results

	Dependent variable: ln(market share)					
	GCC		U.S. (Food outlets only)		U.S. (All Outlets)	
	(1)	(2)	(3)	(4)	(5)	(6)
Numeric Distribution (ND)	4.861 (0.001)		4.213 (0.002)		6.215 (0.002)	
Weighted Distribution (WD)		4.984 (0.000)		4.325 (0.002)		5.242 (0.002)
Channel	FOOD	FOOD	FOOD	FOOD	ALL	ALL
Observations	105,383,354	105,383,354	13,143,965	13,143,965	16,165,992	16,165,992
R-squared	0.398	0.549	0.319	0.356	0.323	0.387
Period FE	YES	YES	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. All Outlets include (i) Convenience, (ii) Drug, (iii) Food, (iii) Mass Merchandiser, and (v) Liquor outlets. Robust standard errors are in parentheses. All coefficients are statistically significant at 0.01 level.

In general, the coefficients will not be estimated by researchers who are using this exponential function to predict market shares; rather, the coefficients will be calibrated based on the evidence across other products and/or countries. Here, we employ the Nielsen data to estimate these coefficients for the GCC countries and the United States. Specifically, we first measure the numeric distribution, weighted distribution, and market share for each item across all GCC product groups, countries, and time periods. Then we pool the data and regress item 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 groups are provided in columns 3 and 4. We also report estimates across all outlets in the United States in columns 5 and 6. We observe that regardless of whether GCC or US data are used, the convexity coefficients b are very similar and lie between 4 to 6. In Table 2 we do not report the constant term a (which also depends on the fixed effects), but in general that term is not used for prediction of the market shares because the shares are normalized to sum to unity, which effectively determines the constant a in each application. So a major advantage of the exponential functional form over other alternatives is that there is only a single convexity coefficient b that is needed for the calibration, and we find that this coefficient is quite stable across products and countries.[18] In Online Appendix Tables A.3 to A.5, we provide further regression estimates for each country-group pair across the 30 product groups and seven countries (six GCC countries and the US), where we find that the convexity coefficients b lie between 4 and 7. We will use the pooled GCC estimates of b from Table 2 in all our calculations below, for both the GCC countries and for the United States.

4.2 Frequency of Price Changes

We first compute the frequency 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).

Before presenting the results we make a very important observation: the frequency of price change is proportional to item market share. This is illustrated in Figure 3 where we allocate items into market share bins and report the median frequency of price change for each bin in UAE and in the US.[19] As the figure clearly illustrates, items with higher market share experience more frequent price changes. Goldberg and Hellerstein (2013), while analyzing micro-level firm data in

[18] To ensure that our results below are not sensitive to the specific coefficient of convexity b employed, we replicate the estimations in each of the three applications by allowing the convexity coefficient to vary from 3 to 8 when we impute expenditure shares from the retail distribution metrics. We find that the results are robust to the alternative specifications. We conclude that including measures of the 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 specification 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 (6) and (7) 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. Finally, we also tested whether a quadratic polynomial could provide a better fit. However, as in the case of the logistic function, it requires knowledge on two parameters. Most importantly, the quadratic approximation does not guarantee that imputed expenditure shares will be positive. In such case, one can force the imputed negative expenditure coefficients to be zero but the results obtained in the subsequent exercises were inferior to those presented in the paper.

[19] Scatter plots are reported for the no-sales, no carry-forward case. The results are similar for other cases (e.g. accounting for sales and/or carrying forward prices when gaps are observed) and countries, and when the mean price frequency is reported instead of the median.

the US, document that large firms change prices 2-3 times more frequently than do small firms. Interestingly, we reach the same conclusion by analyzing consumer prices.[20]

We investigate this important finding further by ruling out the possibility that the relation is driven by outliers. To do that, for each of the 129 narrowly defined product modules, we regress the log market share of each barcode at an outlet with the frequency of its price changes. We then plot the estimated coefficients for each of the 129 regressions in Figure 4 below for each module regression, ordered by size. For reference, we also report the coefficient from running the same regression in the UAE for all categories combined. In all but six cases, the regression coefficient is positive and statistically significant. Therefore, even within the narrowly defined modules, retailers change the prices of their best sellers more frequently.

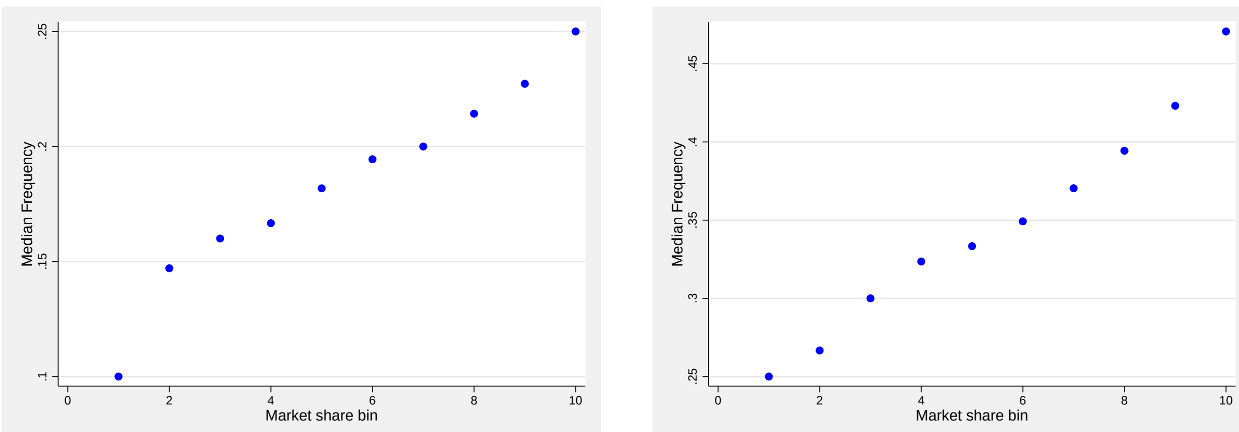

(a) The United Arab Emirates (UAE) (b) The United States of America (USA)

Notes: We use Nielsen scanner data from UAE and the US between January 2006 and December 2011 to calculate the market share of each item and the frequency of its price change. Items are then allocated into bins based on market share, and the median frequency of price change in each bin is plotted, revealing that items with higher sales experience more frequent price changes.

Figure 3: Frequency of Price Changes

The finding that stores adjust prices more frequently for important items suggests that in the absence of any information on quantities or on expenditure that would allow us to distinguish between important and non-important items, averaging the frequency of price changes across all items overstates the degree of price stickiness in the economy.

Table 3 confirms that hypothesis. The frequency of price changes in the UAE and in the US is computed under four alternative specifications: (i) using item-specific weights from expenditure shares (column (1), benchmark); (ii) ignoring expenditure information and treating all items

[20] Similarly, Gorodnichenko et al. (2018) find the online price items that experience a greater number of clicks from buyers – which is used as a proxy for expenditure – have a greater frequency of sales and other price changes.

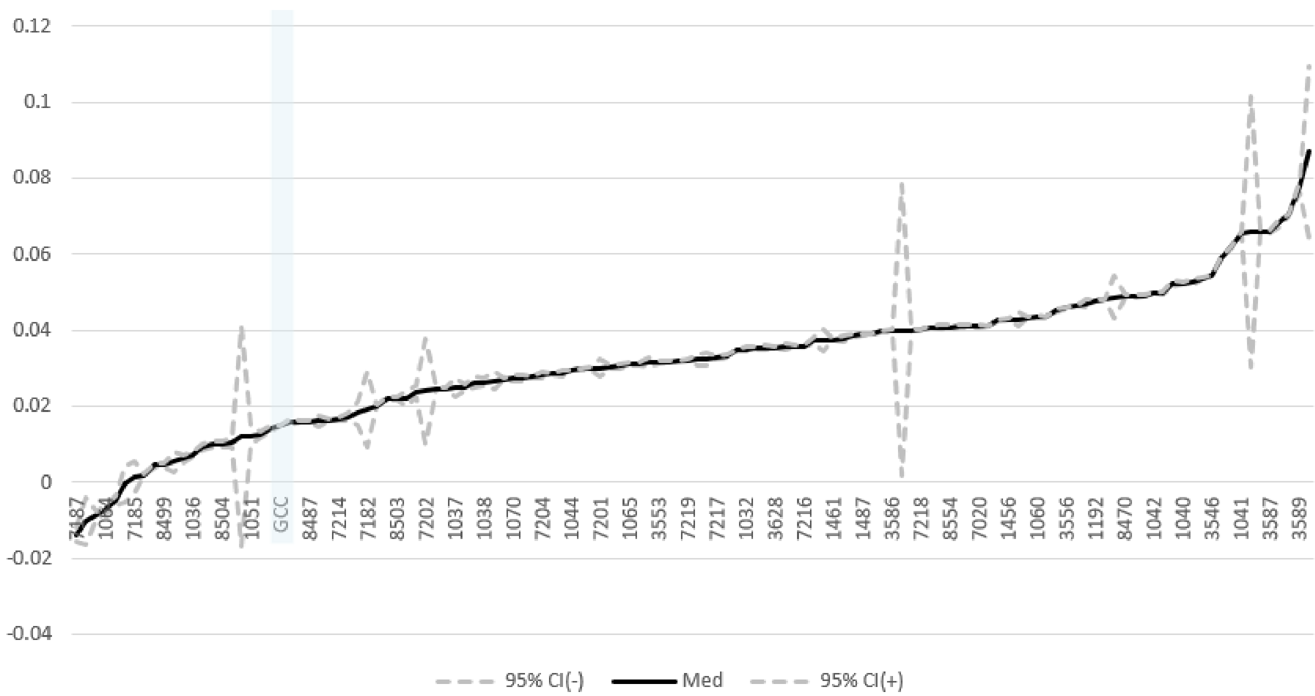

Notes: We report regression coefficients of ln(market share) on the frequency of price changes for each of the 129 US product modules. The coefficients are ranked by order, and standard errors are shown with the dotted lines. Spikes in standard errors occur in modules that contain very few observations. For comparison, we also report the regression coefficient between ln(market share) and the frequency of price changes for all items in the UAE. The results overwhelmingly point to a positive relation between market share and the frequency of price changes. That is, retailers tend to change prices of best sellers more frequently.

Figure 4: Frequency of Price Changes By Product Module

equally (column (2)); (iii) using imputed expenditure shares from the numeric distribution (ND) to assign item-specific weights (column (3)); and (iv) using imputed expenditure shares from the weighted distribution (WD) to assign item-specific weights (column (4)). Panel A reports computations for the frequency of price changes for UAE, and panel B for the US

In the UAE, when both prices and quantities (expenditure shares) are taken into consideration (column 1), we find that the median frequency of price change in UAE is 22% when sales are not taken into account and 17% when they are. By contrast, when information on expenditure is not known and all items are treated equally (column 2), the frequency of price change drops to 13% when sales are not taken into account and to 12% when they are not. That is, not accounting for item importance grossly understates the degree of price stickiness in the economy (by more than 40% in our example). The results are similar using US data (see panel b) and confirm the finding that not assigning item-specific weights in the computation of the frequency of price changes understates its magnitude.[21]

However, when we use measures of the the retail distribution (*numeric distribution* in column 3 and *weighted distribution* in column 4) to impute expenditure shares and then use these imputed shares to weight items, we find that the bias shrinks substantially or even disappears. This happens both for the UAE and the US providing further evidence that imputed shares from the retail distribution metrics are good proxies for actual expenditure shares.

4.3 Inflation

We next turn our attention to measures of inflation and ask whether elementary price indices at the lowest level of aggregation, calculated either without weights, with expenditure weights, or with the imputed weights from the retail distribution, differ from each other. We then check whether these potential discrepancies carry forward in the calculation of higher-level indices.

To match how statistical agencies worldwide and the BLS calculate inflation, we calculate elementary price indexes for each product category using a geometric Laspeyres price index and we then aggregate these up to higher-level indices.[22] The geometric Laspeyers indices for each product group are computed as follows:

$$P_{GL}^{0:t} = \prod_i \left(\frac{p_i^t}{p_i^0}\right)^{w_0^i}, \quad \sum w_i^0 = 1 \tag{8}$$

[21] Nakamura and Steinsson (2008) find the median frequency of price change for consumer goods in the US to be 19% to 20%. But their results are not directly comparable to ours as they examine the entire set of products that span the CPI basket (and use data from earlier years). In their Appendix, however, they do report estimates by Entry Level Items (ELI) that look not much different from ours. For example, they report a 37% frequency of price change for milk, 39% for carbonates soft drinks, and 32% for cheese products.

[22] For more information on how to compute price indices and more details on the approach we take, see Chapter 9 of International Monetary Fund (2020).

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

		(1) Benchmark P & Q	(2) P	(3) P & NUM	(4) P & ACP
	A. United Arab Emirates				
(i) With sales					
	Contiguous observations	0.22	0.13	0.20	0.19
	Carrying regular price forward during sales and stockout	0.20	0.12	0.20	0.18
(ii) Without sales					
	Contiguous observations	0.17	0.10	0.17	0.15
	Carrying regular price forward during sales and stockout	0.16	0.09	0.16	0.14
	B. United States				
(i) With sales					
	Contiguous observations	0.40	0.38	0.40	0.40
	Carrying regular price forward during sales and stockout	0.38	0.33	0.36	0.60
(ii) Without sales					
	Contiguous observations	0.32	0.31	0.33	0.33
	Carrying regular price forward during sales and stockout	0.29	0.26	0.29	0.29

Note: The table reports the median frequency of price changes. Data are provided by Nielsen and cover sales of FMCGs between 2006 and 2011 in the GCC countries and in the U.S. 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.

where w_0^i is the expenditure share of each item within a product group. When equal weights are applied (as in the case when no information on expenditure shares is available), the geometric Laspeyeres index reduces to the Jevons index. As before, we calculate four alternative elementary price indices : (1) using actual expenditure shares as weights (our benchmark case), (2) using no weights and treating all products equally, (3) using imputed shares from *ND* as weights, and (4) using imputed shares from *WD* as weights. When weights are used, these are normalized to 1 in each period. Table 4 reports the *difference* in the annualized inflation measures obtained when using no weights, or using *ND* or *WD*, versus the benchmark case of actual expenditure weights.

We aggregate the elementary price indices to higher-level aggregates with:

$$P_{GL}^{0:t} = \sum w_j^b P_j^{0:t}, \quad \sum w_j^b = 1 \tag{9}$$

where $P_j^{0:t}$ denotes the elementary price index of product group j till period t and w_j^b is a predetermined, period b weight assigned to each elementary group. In the case of the GCC countries, we either assume that expenditure shares w^b for each elementary aggregate are not known at such a disaggregate level so that all elementary aggregate indices are treated equally (Table 4, panel a), or that these are known so that heterogeneous weights are used (Table 4, panel b). In the latter case, we use overall sales in 2006 by product group to create the expenditure shares w_j^b.

For the United States, we also consider the case where all aggregates are treated equally when aggregating (Table 4, panel a, and Figure 5a). But we also exploit the structure of the Nielsen dataset to aggregate (without weights) the product group elementary price indices to five department groups, and then aggregate (with weights) the department groups again to produce a higher-level index. We believe that this is a more meaningful exercise as the BLS is likely to have expenditure shares at the department group level but not at the product group. We use again overall 2006 sales to obtain the weights for each department. These, along with their corresponding market shares, are: (i) Dairy, 28%, (ii) Dry Grocery, 61%, (iii) General Merchandise, 1%, (iv) Health and Beauty Care, 3%; and (v) Non-Food Grocery, 7%. The results are available in Table 4, panel b, and Figure 5b. Inflation figures for each of the 21 product groups and for each department group are available in the Online Appendix for each of the four alternative versions.

Regardless of what assumptions we make on the availability of weight information w_j^b, we find strong evidence that excluding weights at the calculation of elementary price indices results in under-reporting of inflation: annual inflation computed without any weights is lower than when using expenditure weights by 0.3 percentage points on average for the GCC countries and 0.4-0.5 percentage points for the US. But we find that using imputed weights from ND and WD in the elementary price aggregates reduces the bias in the higher-level aggregates: the downward bias is reduced by two-thirds or more for the GCC countries and at least one-half for the United States. Therefore, we conclude that when calculating inflation, using imputed weights from ND and WD provide a meaningful proxy for actual expenditure shares.

Table 4: Summary of Inflation Measures

County / Region	Total Gap (1) No wgt	(2) ND	(3) WD	RMSE ratio (4) ND	(5) WD
(a) Higher level aggregate - no weights					
Bahrain	-0.2%	-0.1%	0.0%	0.29	0.48
Kuwait	0.0%	0.2%	0.3%	0.33	0.21
Oman	0.0%	-0.1%	0.0%	0.48	0.65
Qatar	-0.1%	0.0%	0.1%	0.80	0.81
Saudi Arabia	-0.6%	-0.3%	-0.1%	0.34	0.75
United Arab Emirates	-0.9%	-0.6%	-0.4%	0.40	0.67
Average GCC	***-0.3%***	***-0.1%***	***0.0%***	**0.44**	**0.60**
United States	**-0.5%**	**-0.2%**	**-0.1%**	**0.65**	**0.71**
(b) Higher level aggregate - expenditure weights based on total 2006 sales					
Bahrain	-0.3%	-0.1%	0.0%	0.11	0.61
Kuwait	-0.3%	0.1%	0.2%	0.81	0.76
Oman	0.1%	0.1%	0.0%	0.72	0.74
Qatar	-0.2%	0.3%	0.3%	0.51	0.34
Saudi Arabia	-0.2%	-0.2%	0.0%	0.54	0.65
United Arab Emirates	-1.0%	-0.8%	-0.6%	0.29	0.50
Average GCC	***-0.3%***	***-0.1%***	***0.0%***	**0.50**	**0.60**
United States	**-0.4%**	**-0.2%**	**-0.2%**	**0.51**	**0.52**

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.

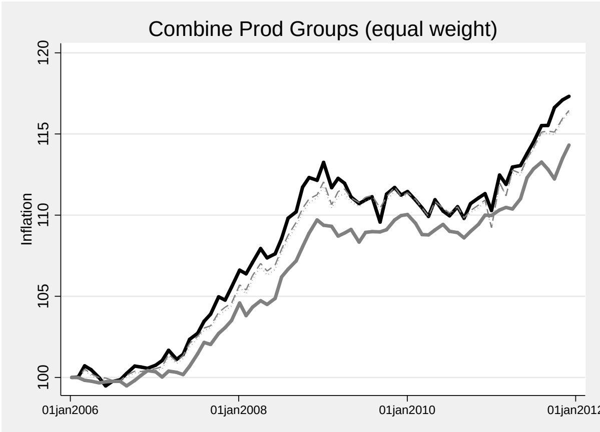

(a) US Inflation Aggregate - No Department Weights

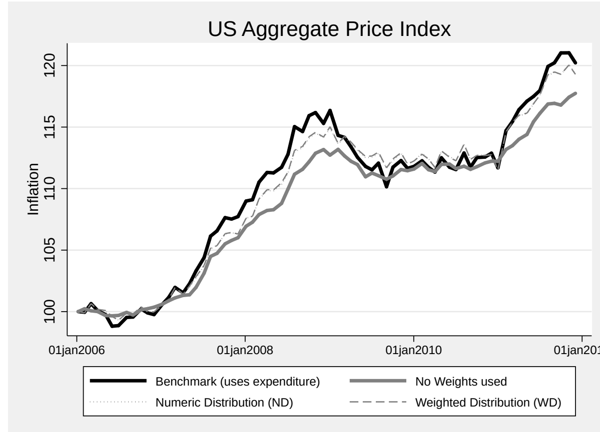

(b) US Inflation Aggregate - with Department Weights

Figure 5: US Inflation Measures

4.4 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. [23]

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

[23]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 in World Bank (2014) and at the World Bank, International Comparison Program (ICP), http://siteresources.worldbank.org/ICPEXT/Resources/ICP_2011.html.

aggregate) stages. Rightly, the World Bank characterizes the ICP as the largest and most complex statistical exercise in the world.[24]

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" 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 items are selected. This happens because multiple varieties (barcodes) of the same product match the same ICP product description. For example, 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".[25] 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 outlets to be audited (n=10, 20, or 50). We then ask the simulation to pick those n outlets out of the universe of stores in our sample by selecting large outlets 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 outlets, with the largest outlets 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. We give guidance as to which price must be quoted if multiple varieties of a product at

[24]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.

[25]For a full list of products that fit the same definition in the case of toothpaste and cereals, see Table A.2 in the Online Appendix.

an outlet 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 judged from shelf space). Finally, and most important, we suppose that the price auditors keep track of whether a product matching the definition is available in the chosen outlet or note. That information is used to construct the numeric and weighted retail distributions. We recognize that the number of items available in each outlet – as needed to compute the weighted retail distribution – might not be available to the ICP. In that case, simpler measures of size might be collected: for example, price auditors can still report information on outlet type (hypermarket, grocery, self-service) or the number of checkout counters.

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*) 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.

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.[26] However, when information on the numeric and weighted distribution

[26]For robustness, we also report estimation results when the average (instead of the median) across the 50

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 the retail distribution reduces measurement bias substantially.[27]

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	0.17	0.09	0.07
avg	20	0.05	0.18	0.06	0.07
avg	50	0.08	0.20	0.06	0.09
max	10	0.07	0.19	0.11	0.08
max	20	0.06	0.22	0.07	0.08
max	50	0.14	0.27	0.13	0.16
med	10	0.03	0.10	0.05	0.01
med	20	0.02	0.13	0.04	0.03
med	50	0.07	0.16	0.05	0.08
mii	10	0.13	0.16	0.07	0.10
mii	20	0.14	0.20	0.07	0.12
mii	50	0.10	0.17	0.07	0.11
min	10	0.04	0.10	0.05	0.04
min	20	0.04	0.11	0.04	0.04
min	50	0.02	0.07	0.05	0.03
rnd	10	0.02	0.12	0.05	0.06
rnd	20	0.06	0.14	0.05	0.09
rnd	50	0.08	0.16	0.03	0.10

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 "rnd"), (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").

iterations is computed. The results, available in the Online Appendix, are similar to those reported in Table 5.

[27] 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 yield similar results.

5 Conclusions

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 overstate price stickiness, understate inflation, and overstate price differences.

To overcome the challenge imposed by the lack of expenditure data, we propose that researchers use information on the retail distribution of sellers to impute expenditure shares. By exploiting the convexity that characterizes the relation between the number of sellers and a product's market share, one can build measures of imputed expenditure shares. Our approach, which we motivate through evidence in the literature but also through a micro-founded framework (see the Online Appendix), 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 demonstrate that using weights imputed from the retail distribution substantially reduces bias in the frequency of price changes, in annual inflation, and in ICP price comparisons across countries.

Our results are based on scanner data, so applicability to the entire set of entry level items that comprise the CPI needs to be studied further as most of the items are not barcoded. We believe, however, that our methodology could potentially be applied to the entire set of entry level items that comprise the CPI although further investigation is needed to assess the extent to which this is possible. There are several challenges, but also potential ways around them. First, not all items are barcoded or the coding standard applied may vary across retailers. For that, one may use product definition to check availability. In the ICP exercise we conducted, we used the product definitions provided by the World Bank, and not barcodes, to identify the availability of these items across each of the GCC countries. Second, data may come from a limited number of outlets. This does not appear to be a problem based on our analysis. Specifically, in the exercise above using Chinese data we showed that the measures of the retail distribution based on less than a dozen outlets from each city were good proxies for actual expenditure shares (Figure 1). Furthermore, in the ICP exercise we showed that the improvements gained from imputing shares from retail distribution metrics obtained from auditing 10 outlets was not much different from that obtained when 50 outlets were audited (Table 5). Third, there may be items such as services that our way of measuring availability using retail distribution may not be applicable. In such a case, availability could potentially be captured in more creative ways, such as the number of results on Google searches, the number of reviews on online marketplaces, or the number of mentions on social media. Finally, note that availability can be used either to impute weights (as we have done here) or simply as an additional way for the BLS to perform sampling selection. How well this approach may work, and for what items, is an open question but it could potentially complement and extend the methodology that we have presented here.

References

Antoniades, Alexis, "The future of price statistics: Innovations in data, technology, and methods," 2016.

_ **and Nicola Zaniboni**, "Exchange rate pass-through into retail prices," *International Economic Review*, 2016, *57* (4), 1425–1447.

Borin, Norm A, Cynthia Van Vranken, and Paul W Farris, "A pilot test of discrimination in the Japanese distribution system," *Journal of Retailing*, 1991, *67* (1), 93.

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

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

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

_ **and Roberto Rigobon**, "The distribution of the Size of Price Changes," Technical Report, National Bureau of Economic Research 2011.

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

Deaton, Angus and Bettina Aten, "Trying to understand the PPPs in ICP 2011: Why are the results so different?," *American Economic Journal: Macroeconomics*, 2017, *9* (1), 243–64.

Farris, Paul, James Olver, and Cornelis De Kluyver, "The relationship between distribution and market share," *Marketing Science*, 1989, *8* (2), 107–128.

Feenstra, Robert C, Mingzhi Xu, and Alexis Antoniades, "What is the Price of Tea in China? Goods Prices and Availability in Chinese Cities," *The Economic Journal*, 2020, *130* (632), 2438–2467.

Goldberg, Pinelopi Koujianou and Rebecca Hellerstein, "A structural approach to identifying the sources of local currency price stability," *Review of Economic Studies*, 2013, *80* (1), 175–210.

Goolsbee, Austan D and Peter J Klenow, "Internet rising, prices falling: Measuring inflation in a world of e-commerce," *American Economic Review*, 2018, *108*, 488–92.

Gorodnichenko, Yuriy and Michael Weber, "Are sticky prices costly? Evidence from the stock market," *American Economic Review*, 2016, *106* (1), 165–99.

_ **, Slavik Sheremirov, and Oleksandr Talavera**, "Price setting in online markets: Does IT click?," *Journal of European Economic Association*, 2018, *16*, 1764–1811.

Hottman, Colin J, Stephen J Redding, and David E Weinstein, "Quantifying the sources of firm heterogeneity," *The Quarterly Journal of Economics*, 2016, *131* (3), 1291–1364.

Inklaar, Robert and DS Rao, "Cross-country income levels over time: did the developing world suddenly become much richer?," *American Economic Journal: Macroeconomics*, 2017, *9* (1), 265–90.

International Monetary Fund, "Calculating consumer price indices in practice." Chapter 8 in Consumer Price Index Manual," Technical Report, International Monetary Fund 2020.

Kruger, Michael and Brian Harper, "Market Share and Product Distribution: Re-Tested and Extended," 2006.

Kuksov, Dmitri and Amit Pazgal, "Research note: the effects of costs and competition on slotting allowances," *Marketing Science*, 2007, *26* (2), 259–267.

Mercer, Alan, *Implementable Marketing Research*, Prentice Hall, 1992.

Nakamura, Emi and Jón Steinsson, "Five facts about prices: A reevaluation of menu cost models," *The Quarterly Journal of Economics*, 2008, *123* (4), 1415–1464.

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

Rao, Akshay R and Humaira Mahi, "The price of launching a new product: Empirical evidence on factors affecting the relative magnitude of slotting allowances," *Marketing Science*, 2003, *22* (2), 246–268.

Reibstein, David J and Paul W Farris, "Market share and distribution: A generalization, a speculation, and some implications," *Marketing Science*, 1995, *14* (3_supplement), G190–G202.

World Bank, *Purchasing power parities and the real size of world economies: a comprehensive report of the 2011 International Comparison Program*, The World Bank, 2014.

Online Appendix

A. Table Appendix

Table A.1: Availability of Chinese Product Prices Scraped from 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	-	-

Table A.2: 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		

Table A.3: Regression Summary by Product Category : GCC Region and the U.S.

(i) Distribution Measure: NUM Distribution			(ii) Distribution Measure: PCV Distribution		
	GCC Average $\hat{b}$	U.S. $\hat{b}$		GCC Average $\hat{b}$	U.S. $\hat{b}$
Pooled data	4.9	5.2	Pooled data	5.0	5.3
By Category			By Category		
Beans	6.7	5.4	Beans	5.6	5.3
Blades	5.2	5.0	Blades	7.2	4.9
Bouillon	3.7	5.7	Bouillon	4.7	5.6
Cereals	6.8	5.8	Cereals	5.2	5.7
Cheese	5.0	5.7	Cheese	4.6	5.6
Chewing gum	5.1	5.8	Chewing gum	5.0	5.9
Chocolate	4.9	5.6	Chocolate	5.3	5.4
Cigarette	4.5	n.a	Cigarette	4.8	n.a
Cooking oil	5.8	5.4	Cooking oil	5.2	5.3
Carbonated soft-drinks	4.1	6.3	Carbonated soft-drinks	4.4	6.2
Deodorant	9.6	4.9	Deodorant	4.9	4.8
Detergents	4.8	4.9	Detergents	5.2	4.8
Dish washer	7.3	5.0	Dish washer	6.3	4.9
Energy drinks	5.0	6.4	Energy drinks	5.2	6.3
Fabric conditioner	5.8	4.8	Fabric conditioner	4.8	4.7
Insecticides	4.9	5.3	Insecticides	4.6	5.3
Juices	5.0	4.3	Juices	4.8	4.4
Liquid cordials	6.6	5.2	Liquid cordials	6.6	4.8
Male grooming	5.5	5.4	Male grooming	5.2	5.1
Milk	4.9	5.3	Milk	4.9	5.3
Milk powder	4.6	5.0	Milk powder	4.8	4.9
Powder soft-drink	6.5	4.9	Powder soft-drink	6.0	4.8
Shampoo	6.6	4.7	Shampoo	5.1	4.7
Skincare	5.5	5.6	Skincare	4.7	5.3
Skin cleansing	6.0	4.7	Skin cleansing	5.3	4.6
Sun-care	5.2	5.0	Sun-care	3.9	4.8
Tea	6.0	5.5	Tea	5.9	5.4
Toothbrush	8.1	4.5	Toothbrush	5.5	4.5
Toothpaste	4.9	4.8	Toothpaste	5.0	4.9
Water	6.4	5.6	Water	5.6	5.4
Summary Statistics			Summary Statistics		
Min	3.7	4.3	Min	3.9	4.4
Max	9.6	6.4	Max	7.2	6.3
Mean	5.7	5.3	Mean	5.2	5.2
Median	5.4	5.3	Median	5.1	5.1

Table A.4: By Country and Category Regressions: Numeric Distribution

	$\hat{b}$ (slope)							
	U.S.	KUW	QTR	BAH	OMN	UAE	KSA	Average $\hat{b}$
Pooled data	5.2	4.7	3.8	4.5	5.2	4.8	5.0	4.6
By Category								
Beans	5.4	7.0	4.7	8.7	7.4	7.0	5.5	6.5
Blades	5.0	9.9	3.0	4.0	4.4	5.9	4.2	5.2
Bouillon	5.7	5.1	3.2	2.3	3.3	5.5	2.6	3.9
Cereals	5.8	7.3	4.4	6.3	7.2	6.7	9.1	6.7
Cheese	5.7	6.3	3.5	5.1	5.0	5.1	5.1	5.1
Chewing gum	5.8	4.8	3.7	4.5	8.1	5.0	4.3	5.2
Chocolate	5.6	5.6	3.3	4.6	6.4	4.5	5.1	5.0
Cigarette	n.a	4.1	4.1	4.5	4.8	4.8	4.5	4.5
Cooking oil	5.4	8.6	3.9	4.4	6.2	4.6	7.2	5.8
Carbonated soft-drinks	6.3	4.2	3.5	4.7	4.3	3.6	4.1	4.4
Deodorant	4.9	7.2	8.2	7.6	12.1	10.1	12.2	8.9
Detergents	4.9	4.3	3.2	4.1	6.6	4.5	5.8	4.8
Dishwasher	5.0	8.6	4.2	4.9	11.5	6.1	8.3	6.9
Energy drinks	6.4	5.4	4.2	4.7	5.5	5.7	4.6	5.2
Fabric conditioner	4.8	6.2	4.5	3.7	8.2	4.4	7.8	5.6
Insecticides	5.3	5.7	3.1	4.4	6.2	4.8	5.1	4.9
Juices	4.3	4.6	3.8	4.5	5.5	5.5	5.9	4.9
Liquid cordials	5.2	7.1	4.9	6.3	6.4	6.5	8.4	6.4
Male grooming	5.4	6.0	3.3	4.4	5.0	8.7	5.5	5.5
Milk	5.3	5.7	4.0	4.8	5.2	5.3	4.3	4.9
Milk powder	5.0	4.7	3.5	3.4	6.1	5.0	5.1	4.7
Powder soft-drink	4.9	7.4	6.8	5.0	6.2	6.2	7.4	6.3
Shampoo	4.7	6.3	4.5	5.6	6.3	5.7	11.0	6.3
Skincare	5.6	5.1	3.5	4.5	6.4	6.0	7.8	5.5
Skin cleansing	4.7	7.5	4.5	5.5	6.6	5.7	6.3	5.8
Sun-care	5.0	7.8	4.4	4.4	3.9	6.4	4.6	5.2
Tea	5.5	6.1	4.9	5.2	7.4	6.6	5.9	5.9
Toothbrush	4.5	15.5	2.8	4.1	5.0	10.6	11.0	7.6
Toothpaste	4.8	6.1	3.1	4.2	5.5	5.1	5.6	4.9
Water	5.6	6.5	5.0	6.3	5.1	5.2	10.3	6.3
Summary Statistics								
Min	4.3	4.1	2.8	2.3	3.3	3.6	2.6	3.9
Max	6.4	15.5	8.2	8.7	12.1	10.6	12.2	8.9
Mean	5.3	6.5	4.1	4.9	6.3	5.9	6.5	5.6
Median	5.3	6.1	4.0	4.6	6.2	5.6	5.7	5.3

Table A.5: By Country and Category Regressions: Product Category Volume

	$\hat{b}$ (slope)							
	U.S.	KUW	QTR	BAH	OMN	UAE	KSA	Average $\hat{b}$
Pooled data	5.3	4.8	4.7	5.0	5.1	5.0	5.0	4.9
By Category								
Beans	5.3	5.9	5.5	5.6	5.6	5.8	5.1	5.5
Blades	4.9	7.2	6.1	7.1	6.3	7.3	9.0	6.8
Bouillon	5.6	6.3	4.3	3.3	4.3	5.9	4.1	4.8
Cereals	5.7	4.8	5.6	5.3	5.5	4.7	5.2	5.3
Cheese	5.6	4.6	4.6	4.8	4.4	4.6	4.9	4.8
Chewing gum	5.9	5.1	4.1	4.9	6.6	4.8	4.6	5.1
Chocolate	5.4	5.6	4.8	5.1	6.1	4.5	5.5	5.3
Cigarette	n.a	4.7	4.5	5.0	4.8	4.6	5.0	4.8
Cooking oil	5.3	5.3	4.6	5.2	5.1	4.9	5.9	5.2
Carbonated soft-drinks	6.2	4.9	3.7	4.7	4.6	4.2	4.3	4.6
Deodorant	4.8	3.4	4.9	5.8	5.5	5.1	5.0	4.9
Detergents	4.8	5.3	4.1	5.3	5.8	5.3	5.6	5.2
Dishwasher	4.9	6.4	5.4	7.6	6.3	5.6	6.4	6.1
Energy drinks	6.3	5.1	4.6	5.1	5.1	6.1	5.0	5.4
Fabric conditioner	4.7	5.2	5.2	4.4	5.1	4.4	4.7	4.8
Insecticides	5.3	4.1	3.7	4.6	5.2	4.9	5.3	4.7
Juices	4.4	4.8	4.5	4.3	5.0	5.0	5.1	4.7
Liquid cordials	4.8	6.5	6.0	6.0	7.4	6.3	7.5	6.4
Male grooming	5.1	5.1	4.4	5.8	5.6	5.8	4.5	5.2
Milk	5.3	5.5	4.7	4.6	5.0	5.2	4.5	5.0
Milk powder	4.9	4.8	4.5	4.4	5.6	5.0	4.5	4.8
Powder soft-drink	4.8	5.8	7.1	5.8	6.2	5.7	5.6	5.9
Shampoo	4.7	4.8	5.1	5.3	4.8	4.7	5.8	5.0
Skincare	5.3	4.2	4.4	5.1	5.3	4.8	4.6	4.8
Skin cleansing	4.6	4.8	5.6	6.0	5.1	5.3	5.3	5.2
Sun-care	4.8	4.4	3.7	2.5	4.2	4.2	4.4	4.0
Tea	5.4	5.8	5.9	5.7	6.2	6.0	5.7	5.8
Toothbrush	4.5	5.1	5.8	5.7	5.9	5.7	5.0	5.4
Toothpaste	4.9	4.4	5.0	5.2	5.1	5.7	4.9	5.0
Water	5.4	5.7	6.0	6.0	4.9	4.7	6.6	5.6
Summary Statistics								
Min	4.4	3.4	3.7	2.5	4.2	4.2	4.1	4.0
Max	6.3	7.2	7.1	7.6	7.4	7.3	9.0	6.8
Mean	5.2	5.2	4.9	5.2	5.4	5.2	5.3	5.2
Median	5.1	5.1	4.8	5.2	5.2	5.1	5.1	5.1

B. Figure Appendix

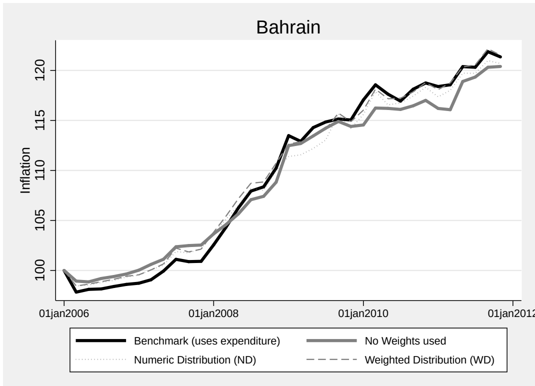

(a)

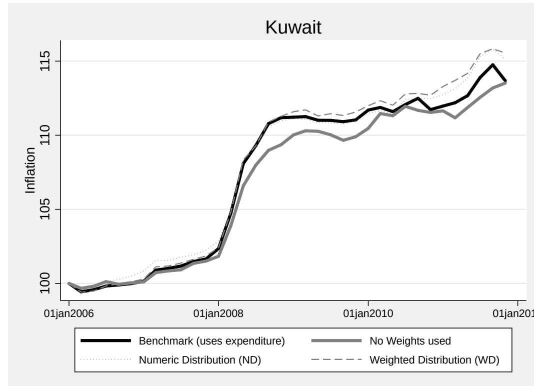

(b)

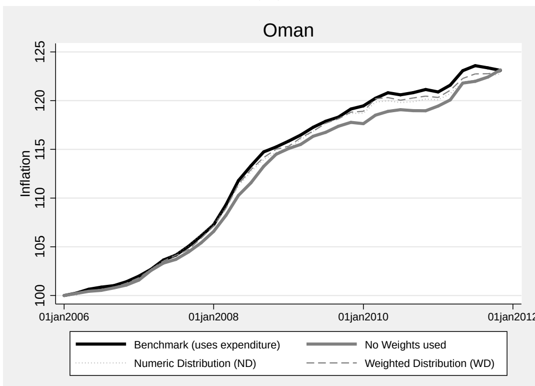

(c)

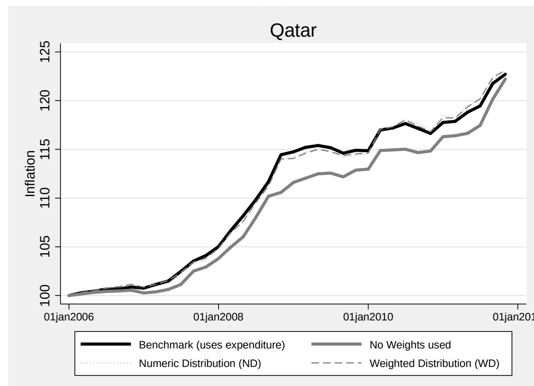

(d)

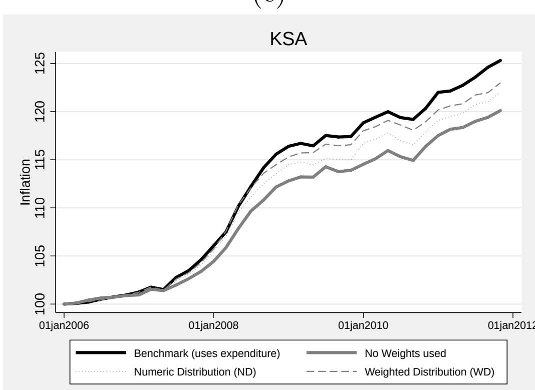

(e)

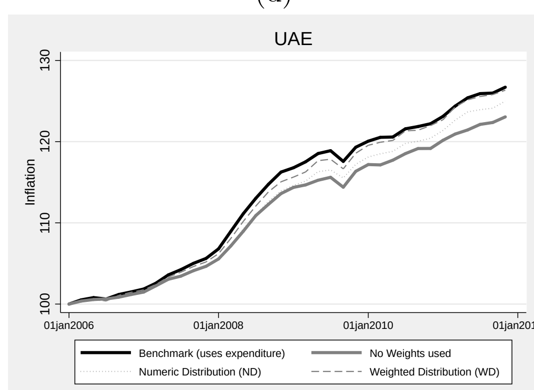

(f)

Figure B.1: Alternative Aggregate GCC Inflation Measures (No Weights)

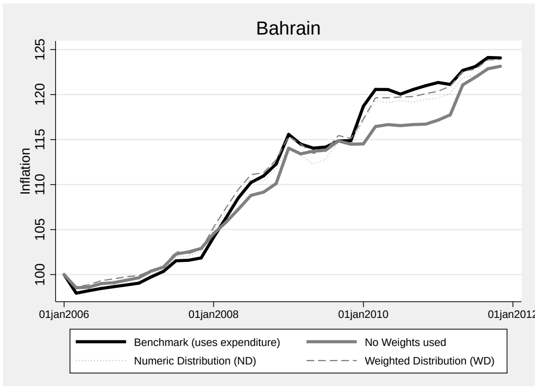

(a)

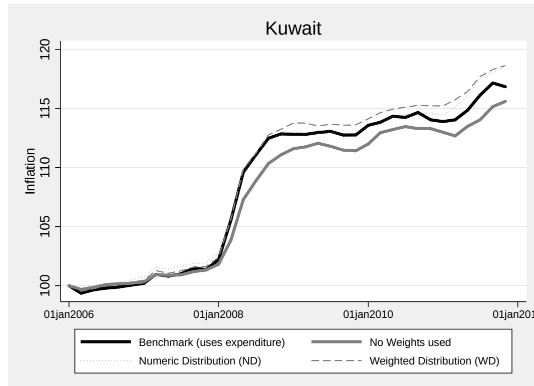

(b)

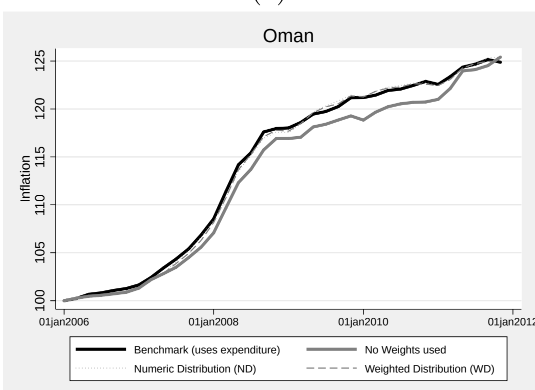

(c)

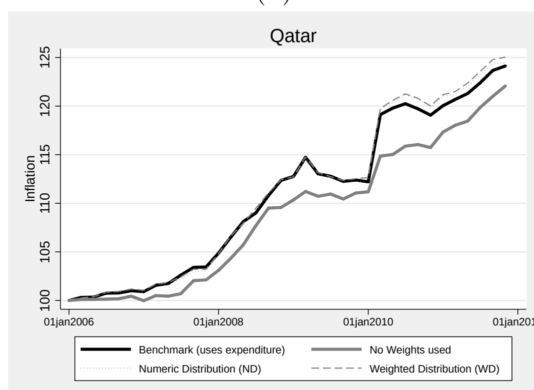

(d)

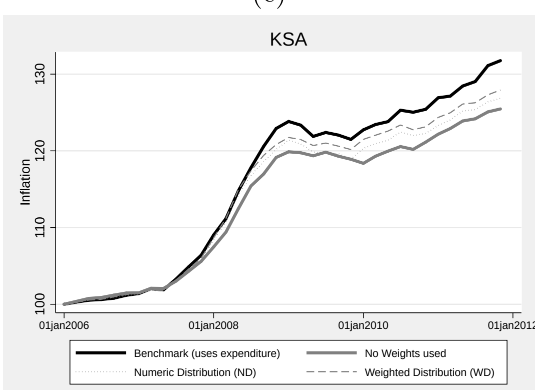

(e)

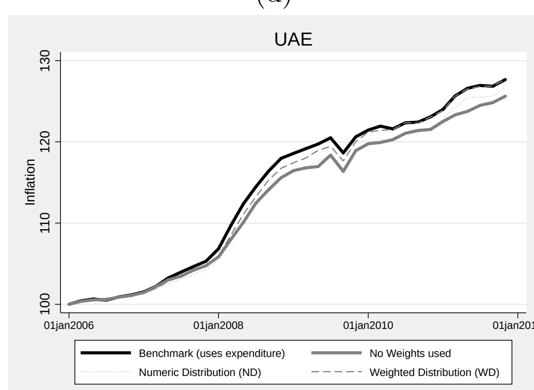

(f)

Figure B.2: Alternative Aggregate GCC Inflation Measures (Expenditure Weights)

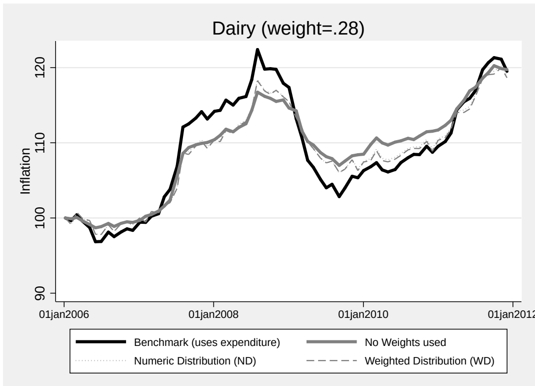

(a)

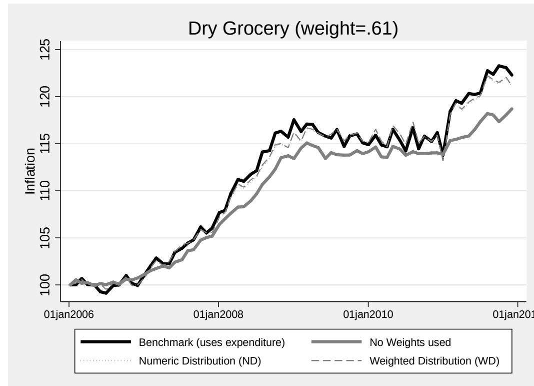

(b)

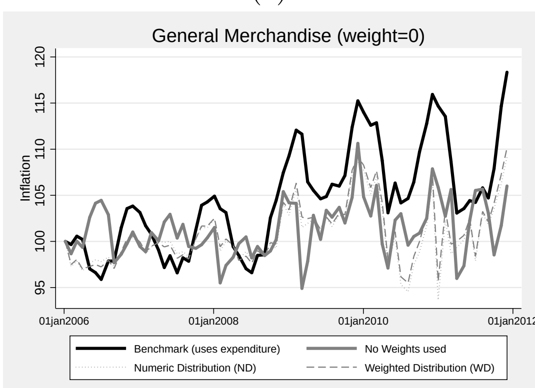

(c)

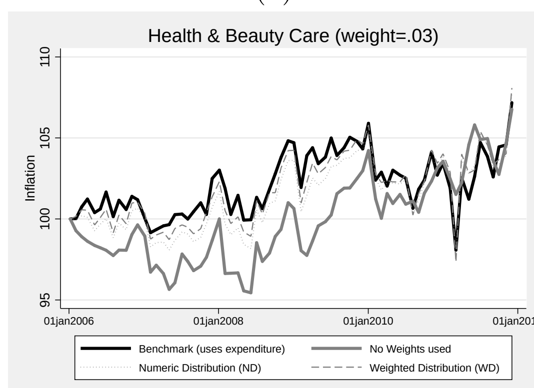

(d)

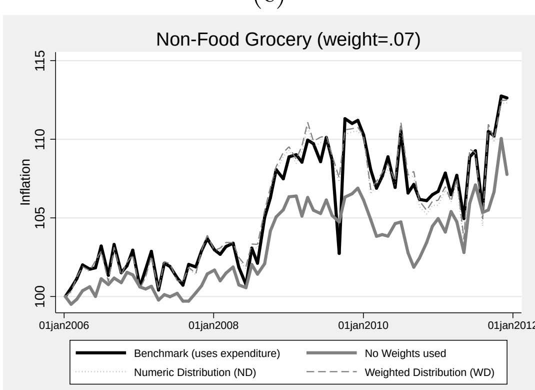

(e)

Figure B.3: US Inflation by Department

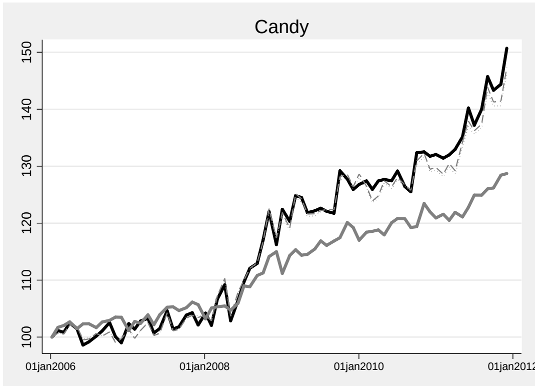

(a)

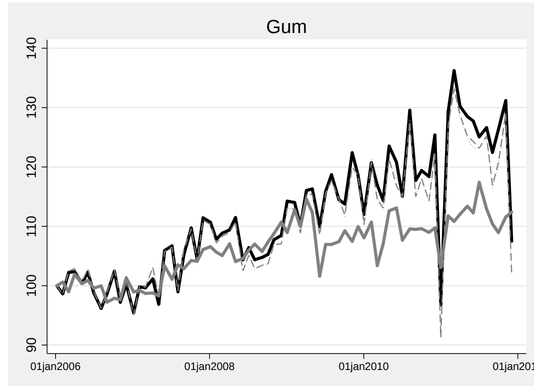

(b)

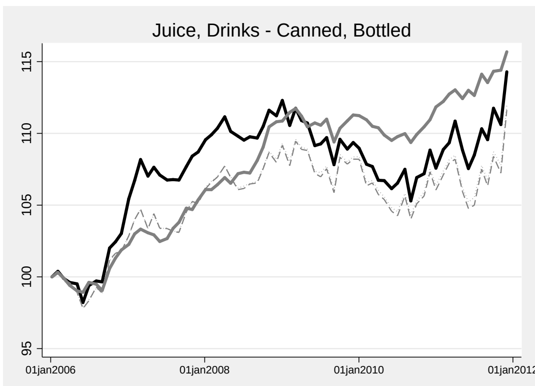

(c)

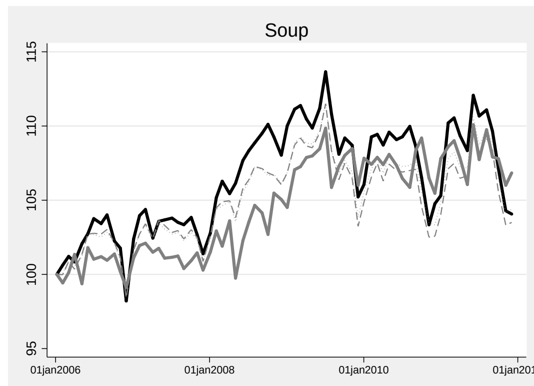

(d)

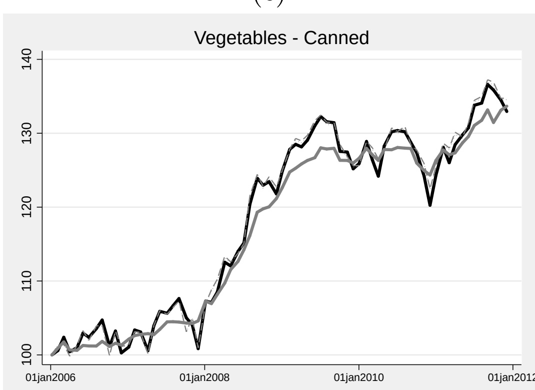

(e)

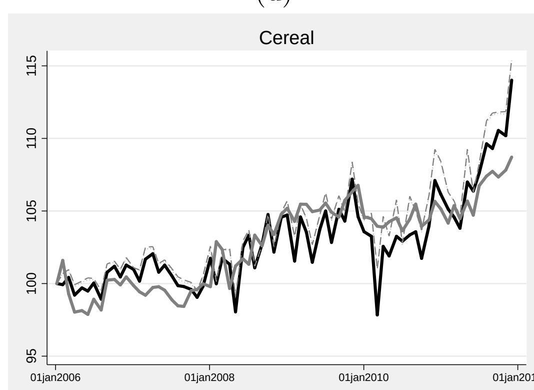

(f)

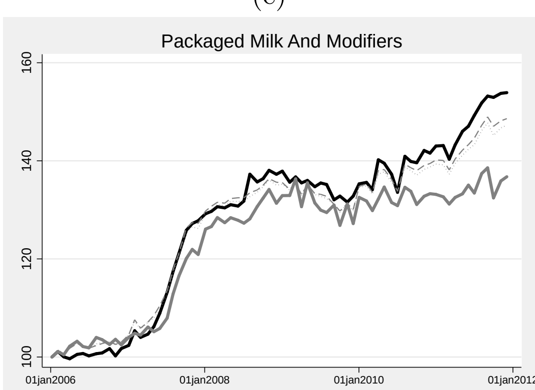

(g)

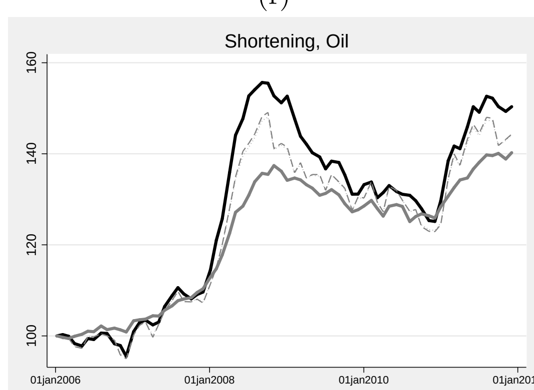

(h)

Figure B.4: US Inflation by Group (Part I)

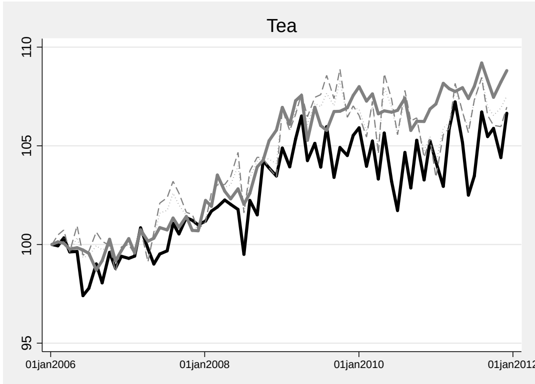

(a)

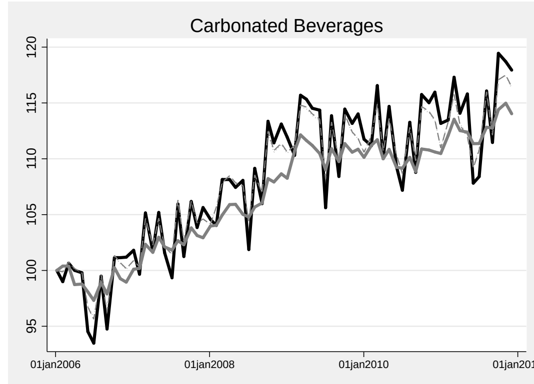

(b)

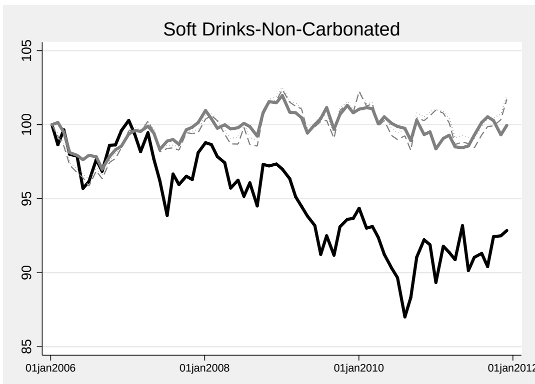

(c)

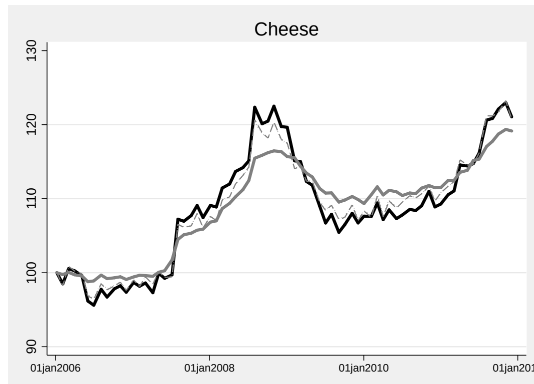

(d)

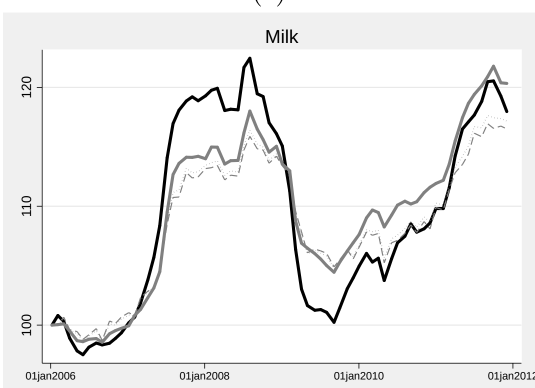

(e)

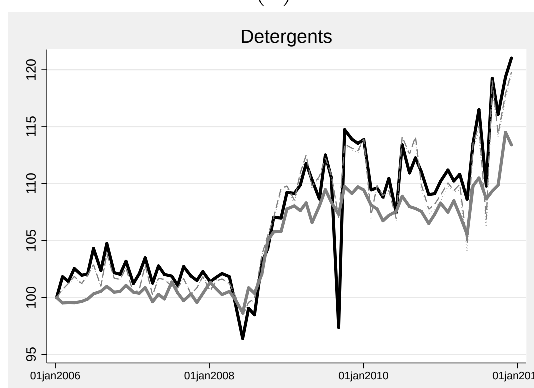

(f)

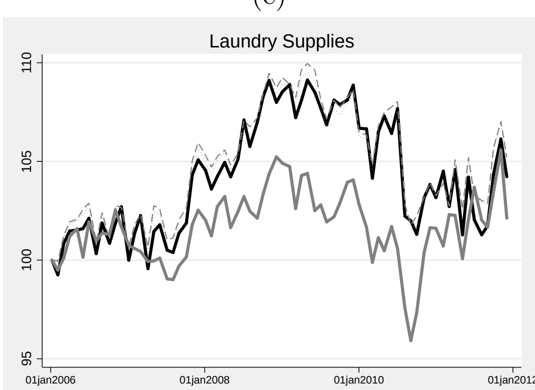

(g)

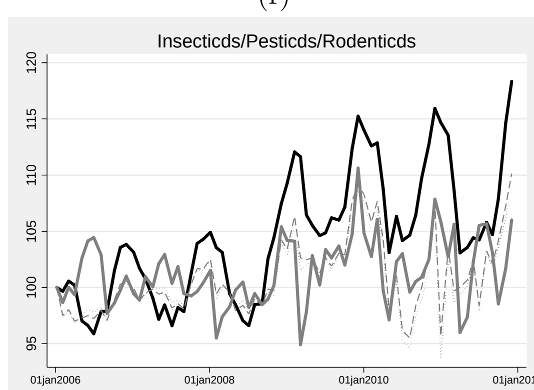

(h)

Figure B.5: US Inflation by Group (Part II)

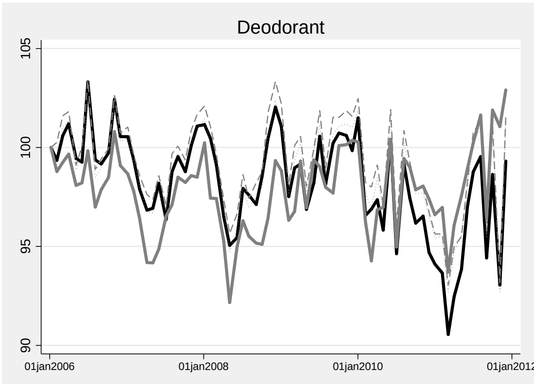

(a)

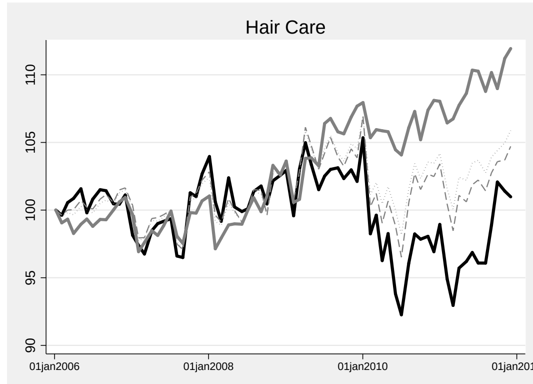

(b)

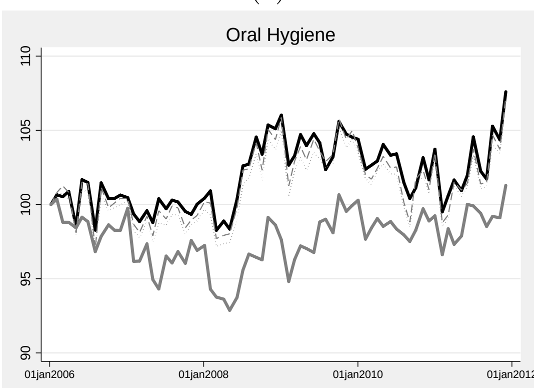

(c)

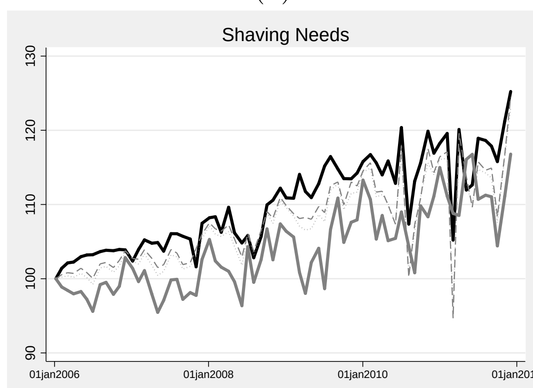

(d)

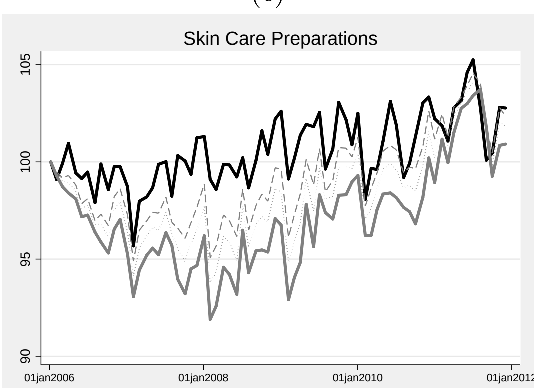

(e)

Figure B.6: US Inflation by Group (Part III)

C. Theory Appendix

C1. Theoretical Foundation of Convexity

Having reviewed evidence from the literature and the data for the convex relation between market share and the 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 building on the Melitz (2003) model, characterizes both manufacturers' and retailers' decisions under alternative market structure settings. As featured in the model, it is the interaction between heterogeneous firms and varying "slotting fee" that yields the convex relation that is observed in the data. We show that assuming heterogeneity in the slotting fee incurred by manufacturers 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 analysis 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 et al. (2016); Feenstra et al. (2020), and is assumed to be nested CES, as follows:

$$U = \left(\int_{r \in \Omega} X_r^{\frac{\eta-1}{\eta}} \right)^{\frac{\eta}{\eta-1}}, \; X_r = \left(\int_{j \in J_r} x_{rj}^{\frac{\sigma-1}{\sigma}} \right)^{\frac{\sigma}{\sigma-1}}, \; \sigma > \eta \tag{10}$$

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,

$$x_{rj} = p_{rj}^{-\sigma} P_r^{\sigma-\eta} P^{\eta-1} Y. \tag{11}$$

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:

$$P_r = \left(\int_{j \in J_r} p_{rj}^{1-\sigma} \right)^{\frac{1}{1-\sigma}}, \; P = \left(\int_{r \in \Omega} P_r^{1-\eta} \right)^{\frac{1}{1-\eta}} \tag{12}$$

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 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 slotting fee, which we call the *slotting fee*. For simplicity, we treat the slotting fee as exogenous (i.e., not chosen by retailers) so that it becomes a slotting fee and paid by manufacturers.[28] 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 slotting fee and a retailer of $r = 1$ has the highest slotting fee ($\partial f/\partial r > 0$). We study the equilibrium in which manufacturers will prefer to sell in retailers with lower slotting fee. That is, we assume that manufacturers go to retailers with the lowest slotting fee first and then to those with increasing higher slotting fee 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.[29]

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:

[28] The marketing literature refers to the slotting fee ($f(r)$) 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.

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

$$\pi_{\phi} \equiv \max_{q_{r\phi}} \int_{0}^{r_{\phi}} \pi_{r\phi} dr = \max_{q_{r\phi}} \int_{0}^{r_{\phi}} \left(q_{r\phi} x_{r\phi} - f\left(r\right)\right) dr, \tag{13}$$

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, $f\left(r\right)$ denotes the slotting 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 (13), $q_{r\phi}$ 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:

$$q_{r\phi} = \frac{\sigma}{\sigma - 1} \frac{w}{\phi}, \ \forall r \in [0, r_{\phi}]. \tag{14}$$

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

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

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

$$\frac{1}{\sigma - 1} \left(\frac{\sigma - 1}{\sigma}\right)^{\sigma} Y P^{\eta - 1} w^{1-\sigma} \times \mu_r^{-\sigma} P_r^{\sigma - \eta} \times \phi_r^{\sigma - 1} = f\left(r\right). \tag{15}$$

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 $f\left(r_1\right) < f\left(r_2\right)$ then $\phi_{r_1} < \phi_{r_2}$).[30] 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 (12):

[30]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 $\left(\phi_r, f\left(r\right)\right)$. 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.

$$
\begin{aligned}
P_r &= \left[M \int_{\phi_r}^{+\infty} p_{r\phi}^{1-\sigma} g\left(\phi\right) d\phi \right]^{1/(1-\sigma)} \\
&= \frac{\sigma}{\sigma-1} \left(\frac{k}{k-\sigma+1} \right)^{\frac{1}{1-\sigma}} \bar{\phi}^{\frac{k}{1-\sigma}} M^{\frac{1}{1-\sigma}} w \times \mu_r \phi_r^{\frac{k-\sigma+1}{\sigma-1}},
\end{aligned} \tag{16}
$$

Substituting (16) back to (15), we could solve the cutoff of productivity ϕ_r:

$$
\phi_r^{\epsilon_1} = A_1 f\left(r\right) \mu_r^{\eta}, \tag{17}
$$

where ϵ_1 and A_1 are defined as:

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

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

$$
R_{r\phi} = \sigma \phi^{\sigma-1} A_1^{\frac{1-\sigma}{\epsilon_1}} f\left(r\right)^{\epsilon_2} \mu_r^{1-\frac{\eta(\sigma-1)}{\epsilon_1}}, \tag{18}
$$

where the equality uses (17). 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:

$$
\begin{aligned}
R_\phi &= \int_0^{r_\phi} R_{r\phi} dr \\
&= \int_0^{n} R_{r\phi} dr \\
&= \sigma \phi^{\sigma-1} A_1^{\frac{1-\sigma}{\epsilon_1}} \int_0^{n} f\left(r\right)^{\epsilon_2} \mu_r^{1-\frac{\eta(\sigma-1)}{\epsilon_1}} dr.
\end{aligned} \tag{19}
$$

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 the *numeric distribution* n (see Appendix C2). 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:

$$\max_{p_{rj},j\in J_r}\left[\sum_{j\in J_r}\left(p_{rj}-q_{rj}\right)x_{rj}\right]\Leftrightarrow\max_{p_{r\phi},\phi>\phi_r}\left[M\int_{\phi_r}^{+\infty}\left(p_{r\phi}-q_{r\phi}\right)x_{r\phi}g(\phi)d\phi\right] \tag{20}$$

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 et al. (2020), and the pricing rule of retailer r is:

$$p_{r\phi}=\mu_r q_{r\phi},\ \ with\ \ \mu_r\equiv 1+\frac{1}{(\eta-1)\left(1-s_r\right)},\ \forall\phi>\phi_r, \tag{21}$$

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* (21)*, and product sales are convex in the numeric distribution if:*

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

The proof of Proposition 2 is in Appendix C3, 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:

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

which could be further simplified given (17) as:

$$\max_{\mu_r} \left[\frac{\sigma k M \bar{\phi}^k A_1^{-\frac{k}{\epsilon_1}}}{k-\sigma+1} f(r)^{1-\frac{k}{\epsilon_1}} (\mu_r - 1)\mu_r^{-\frac{\eta k}{\epsilon_1}} \right]. \tag{22}$$

The first order condition of (22) with respect to μ_r implies that:[31]

$$\mu_r = 1 + \frac{1}{\eta k/\epsilon_1 \left[\eta - (\eta - 1) s_r\right] - 1}, \tag{23}$$

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:[32]

$$k > 1 + \frac{\eta+1}{\sigma(\eta-1)}. \tag{24}$$

Similar to the pricing rule for shortsighted retailers in (21), the markup of a farsighted retailer also positively depends on its market share. Therefore, we derive a proposition similar to Proposition 2.[33]

Proposition 3. *When retailers are farsighted, retailers' markups positively depend on their market shares as in* (23)*, 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 the 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 C4. 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

[31] 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 $\partial \ln P_r/\partial \ln \mu_r = 1$.
[32] In the extreme case in which there is only one retailer, the markup is $\mu_r = \eta k/\epsilon_1/(\eta k/\epsilon_1 - 1)$.
[33] Proposition 3 implicitly assumes that model parameters satisfies (24).

higher slotting fee. 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.

Model Simulation

To provide an overview of how well the model generates the convex relation between product market share and the retail distribution, we perform a simulation exercise for the case in which retailers are shortsighted.[34] In the simulation, we simulate the sales and the *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 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 (10), (13), and (20). In practice, we specify the slotting fee as $f(r) = \gamma e^{\theta r}$ ($\gamma > 0$ and $\theta > 1$) and simulate 10,000 draws u from a uniform distribution from zero to one. 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 C.1 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*.[35]

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 slotting fee incurred by manufacturers to sell in retailers vary across retailers. Our theoretical results further corroborate the robustness of using the retail distribution to approximate product sales when they are absent.

[34] The pattern for farsighted retailers remains similar, as is also discussed in Proposition 3. The detailed procedure for simulation is provided in Appendix C5.

[35] In Figure C.2 and C.3, we also simulate the model with different functional forms for the slotting fee $f(r)$, and the convex relationship between product market share and the *numeric distribution* remains robust. 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.

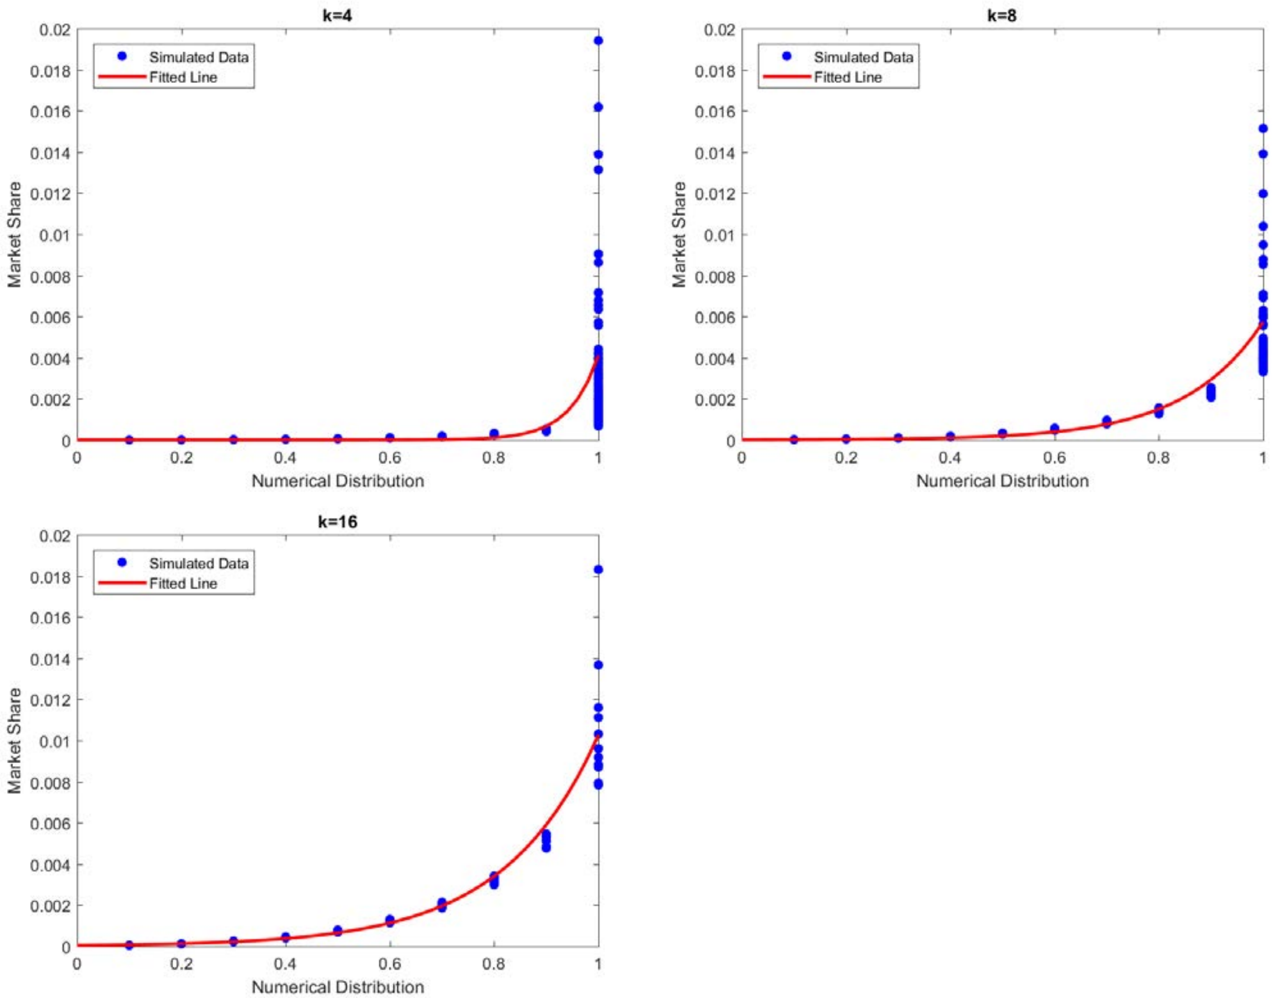

Notes: Parameter value $k = 16$ satisfies parameter restriction in Proposition 2 and 3.

Figure C.1: Convexity between Sales and Numeric Distribution

C2. 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 ($\epsilon_2 > 0$):

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

where the first inequality holds given that there is no negative term, and the second inequality holds given that slotting fee $f(r)$ increase in r.

C3. Proof of Proposition 2

We rewrite (19) 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 (17), 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 slotting fee $f(r)$ 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 (17) 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:

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

where the first inequality holds given that there is no negative term, and the second inequality holds given that slotting fee $f(r)$ 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 the *numeric distribution* n.

C4. 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 (19), 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 (21) or (23). This implies that both $f(n)^{\epsilon_2}$ and $\mu_n^{1-\frac{\eta(\sigma-1)}{\epsilon_1}}$ increase in n, which confirms convexity:

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

C5. Model Simulation Procedures

Table C.1 displays the parameters used in the simulation. Given the parameters, we simulate the economy in which consumers, manufacturers, and retailers are specified by (10), (13), and (20). The slotting fee 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 (17) 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 (18) (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 (21). 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 C.1 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 C.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 slotting fee	100
θ	Elasticity of slotting fee 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 slotting fee $f(r)$, with all other parameters fixed as displayed in Table C.1. In Figure C.2, 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 C.3, 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.

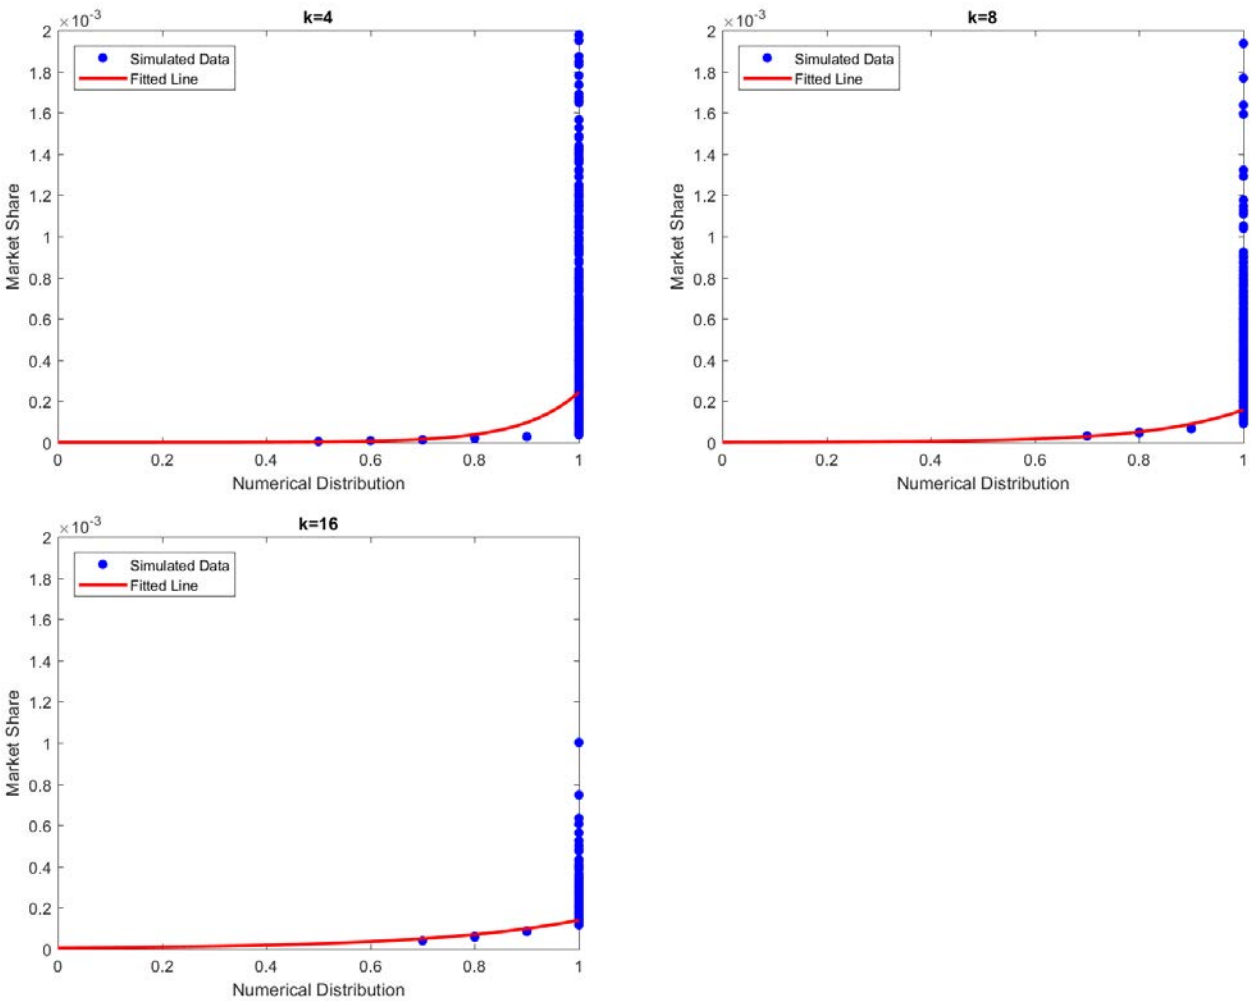

Figure C.2: Convexity between Sales and Numeric Distribution ($f(r) = \gamma r^{\theta}$, $\theta = 2$)

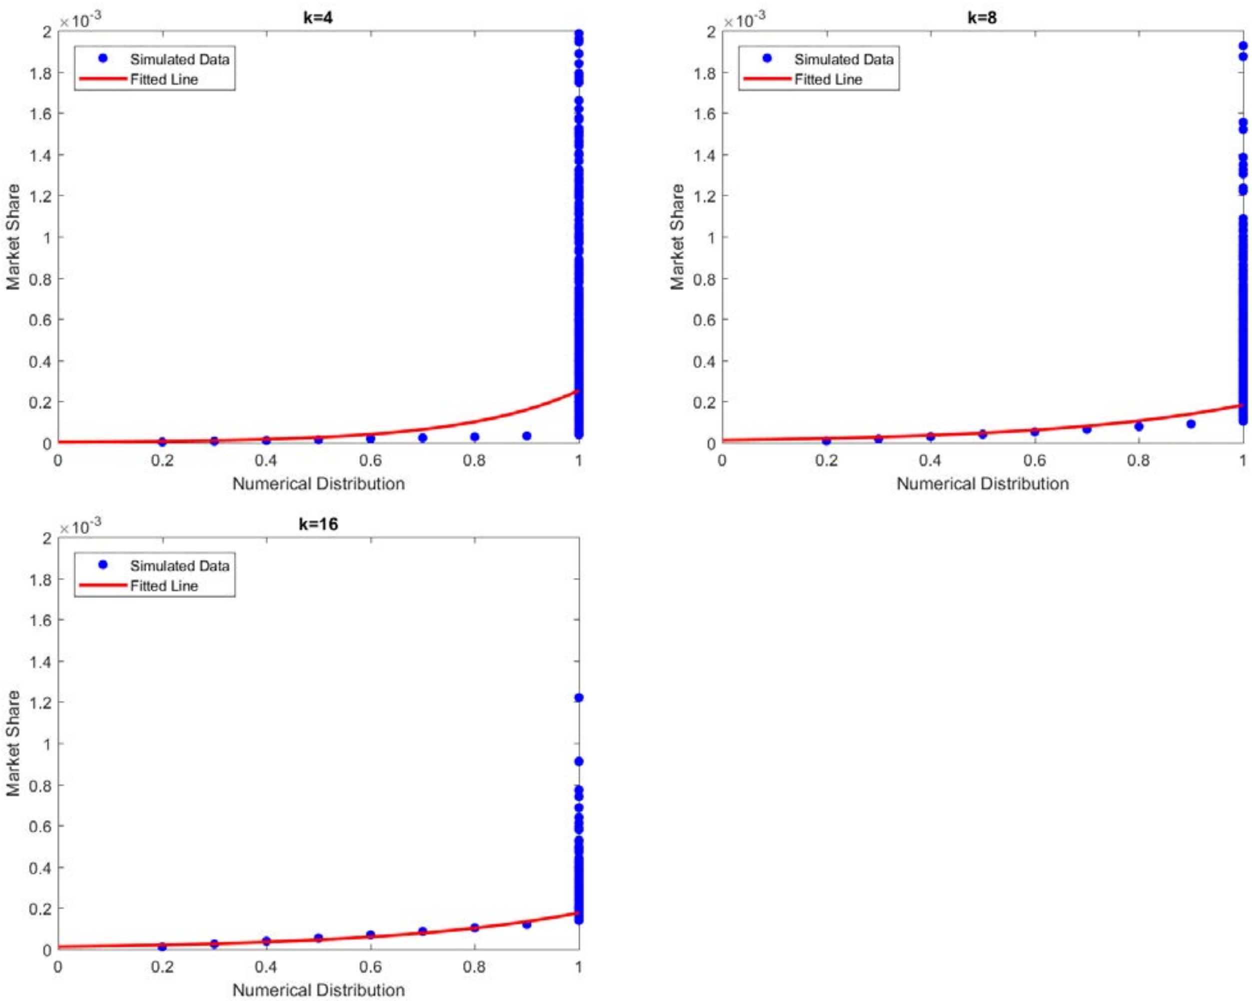

Figure C.3: Convexity between Sales and Numeric Distribution ($f(r) = \gamma r^{\theta}$, $\theta = 0.2$)