Large Firms, Consumer Heterogeneity and the Profit Share*

Robert C. Feenstra† — UC Davis and NBER

Luca Macedoni‡ — Aarhus University

Mingzhi (Jimmy) Xu§ — Peking University

November 30, 2023

Abstract

We examine the relationship between large firms and the profit share in a model that features oligopolistic competition and consumer heterogeneity. Conditional on the sales distribution, the presence of consumer heterogeneity increases the profit share because it increases firm-level markups. Using data on purchases at the household-barcode level from Nielsen, we quantify the role of consumer heterogeneity, finding that the average markup and the aggregate profit share are 20 and 6.4 percentage points larger than those predicted by a model of a representative consumer. Based on our evidence from 2004 to 2018, and extrapolating it to 1990 through 2021, we predict that the profit share in retailing could have increased by over 4 percentage points over this longer period due to rising income inequality across consumers.

Keywords: Firm Heterogeneity; Multiproduct firms; Consumer Heterogeneity; Scanner data.

JEL Code: D12, L11, L25, O51.

*Researchers own analyses calculated (or derived) based in part on data from Nielsen Consumer LLC and marketing databases provided through the NielsenIQ Datasets at the Kilts Center for Marketing Data Center at The University of Chicago Booth School of Business. The conclusions drawn from the NielsenIQ data are those of the researchers and do not reflect the views of NielsenIQ. NielsenIQ is not responsible for, had no role in, and was not involved in analyzing and preparing the results reported herein. We thank Yucheng Yang for sharing the GS1 Global Data. We thank Philipp Schröder, Frederic Warzynski, and Zhiyuan Chen for their suggestions and feedback. We thank seminar participants at Aarhus University, ETOS, SFSU, and Renmin University. Financial support from the Natural Science Foundation of China (Grant No. 72322007) are gratefully acknowledged by Mingzhi Xu.

†rcfeenstra@ucdavis.edu

‡lmacedoni@econ.au.dk

§mingzhixu@nsd.pku.edu.cn

1 Introduction

Since the 1980s, the labor share of income has declined in the United States and globally (Elsby et al., 2013; Karabarbounis and Neiman, 2014) while the profit share has risen (Barkai, 2020). Autor et al. (2020) attribute this pattern to the competition-induced allocation of production in favor of superstar firms, which are large in size and have high markups and profit margins. In a model of monopolistic competition, they show that an increase in competition, due to international trade or technological progress, causes a reduction in markups of all surviving firms and a reallocation of production towards high-markup firms, leading to an increase in the profit share.[1]

Despite the robust empirical evidence confirming the role of large firms in explaining the rising profit share in Autor et al. (2020), there are challenges in constructing a model consistent with these facts. First, while the empirical analysis of Autor et al. (2020) focuses on superstar firms, their theory relies on monopolistic competition between "small" firms, who take market conditions as given. Second, their theory depends on a sufficiently skewed firm productivity distribution (i.e., a log-convex productivity distribution), and thus, the results require a stronger condition on productivity distribution than commonly imposed, such as the Pareto distribution (which is log-linear).[2] Therefore, it is natural to explore the case of large firms in a more general setting that allows for increases in the profit share without strong assumptions on the productivity distribution.

This paper introduces consumer taste heterogeneity into a model with large oligopolistic firms. Taste differences are modeled as heterogeneity in the demand shifters that consumers have across disaggregate products. This dimension is motivated by Hottman et al. (2016), who document that 50-70% of firm size heterogeneity is driven by heterogeneity in product-specific demand shifters. While Hottman et al. (2016) consider a representative consumer, Neiman and Vavra (2023) investigate consumer heterogeneity and argue that this dimension leads to an extra source of consumer gains from variety. They do not find, however, a significant impact of consumer heterogeneity on aggregate market power. This finding may be related to their assumption that consumer preferences across products are distributed as Pareto (but with different preferred products for each consumer). So the Pareto assumption on taste parameters – like the analogous Pareto assumption on firm productivities – combined with their assumption that there is no product that is preferred by a positive mass of

[1]There has been a general consensus that the declining labor share and rising profit share is significant and robust (Autor et al., 2020). The pattern is also broadly consistent with other perspectives associated with the declining labor share, such as the treatment of depreciation and capitalization when measuring capital (Smith et al., 2019; Koh et al., 2020), increased cost of housing (Rognlie, 2016), the underestimated labor income due to self-employment (Elsby et al., 2013) and the increasing human capital among top earners (Smith et al., 2019).

[2]While it has been common to use the Pareto distribution in the Melitz model since Chaney (2008), recent research has shown that the truncated Pareto or the log-normal distribution for productivity is a better fit to the trade data: see Fernandes et al. (2018) and Adão et al. (2020), for example. Neither of these distributions is globally log-convex, as required by Autor et al. (2020).

consumers, appears to rule out any impact on the aggregate markup and profit share.[3]

In contrast, we suppose that discrete groups of consumers – which we refer to as "consumer segments" – have identical preferences (as measured by their demand shifters) within segments but differing preferences across segments. We find that markups are in general higher in this heterogeneous consumer setting than the markups of a firm with identical aggregate market share selling to a representative consumer. Hence, conditional on the same firm-level shares, the implied markups in a model with heterogeneous consumers are larger than in a model with a representative consumer. This is because a firm facing heterogeneous consumers has different demand elasticities across segments so that it can exploit the (relatively) lower demand elasticity in one segment, charging higher prices for all consumers.

To determine the extent to which consumer heterogeneity contributes to high and rising profits of large, multi-product firms, we leverage rich scanner data from Nielsen that contains information on sales, prices, and quantities at the barcode level in US supermarkets. We further combine the data with information on purchases at the household-retailer-barcode level for a sample of households. Using a nested-CES demand system, we derive a segment-specific demand shifter for each barcode associated with the appeal of the barcode for each segment. In this manner, we divide consumers into segments according to their inferred demand shifters.[4]

Allowing for this consumer heterogeneity, we estimate that the implied markups are 20 percentage points higher on average than predicted by the formula for a representative consumer model (such as Hottman et al. (2016)). Furthermore, the difference is significantly greater when considering the largest firms: for the top five and one firms, respectively, markups are 27 and 39 percentage points higher than the counterparts in a model with the representative consumer. A similar result occurs when we consider total profits. The aggregate profit share is 6.4 percentage points larger under consumer heterogeneity, although there are differences across product groups: the effect of consumer heterogeneity on the profit share is small for fresh produce and frozen meat (both close to zero), but large for beer (18.6) and diapers (20.7).

The aggregate profit share predicted by our model increases slightly from 35.2% to 36% from 2004 to 2010 and then it declines to 33% in 2018. The decline in the last eight years of our sample is at odds with the results in the literature (Autor et al., 2020), and it is probably

[3]In a trade model with heterogeneous firms that differ in their productivities, using the Pareto distribution for productivities rules out any impact of import competition on the profit share, even with very general consumer preferences and variable firm markups (Arkolakis et al., 2019); this result does not hold, however, with alternative distributions such as the truncated Pareto (Feenstra, 2018).

[4]We do not consider how product characteristics (such as flavor or functionality) may affect demand shifters due to data limitations. For an application of such a case to the ice cream product group, see Draganska et al. (2007).

driven by our focus on the retail sector using Nielsen Data: these scanner data may not be representative of the entire US economy, and the Nielsen Home Scanner data that we rely on may also lack representation of national consumption expenditures due to potential bias from data collection (Zhen et al., 2009; Leicester, 2013).[5] Despite these limitations, a robust finding of our analysis is that large firms combined with consumer heterogeneity lead to high profit shares: the aggregate profit share in 2004 is 9.7 percentage points higher than what is predicted by a model of small firms, and the difference increases to 10.6 percentage points in 2010 and it drops to 7.9 in 2018. We decompose such differences into the component due to consumer heterogeneity and the component due to firm size. We find that both components play a role in the change in the profit share. The contribution of consumer heterogeneity, which is the difference in profit share predicted by our model and a model with large firms and a representative consumer, increases from 5.6 in 2004 to 6.3 in 2010 and then falls to 4.8 in 2018. The contribution of firm size, which is the difference in prediction from a model with large firms and a representative consumer and a model with small firms, follows the same pattern but it is of lesser magnitude. In particular, the difference in profit shares goes from 4 percentage points in 2004 to 4.3 in 2010 and then it falls to 3 in 2018.

To further investigate the role of consumer heterogeneity in explaining the increase in the aggregate profit share, we focus on three key aspects: changes in consumer heterogeneity, changes in the strategies followed by firms to deal with heterogeneous consumers, and changes in the fixed costs associated with targeting segments. The profit share we measure is significantly correlated with measures of consumer heterogeneity, such as income inequality, the number of different consumer segments generated by our algorithm, and the variance of the segment-specific taste component. In other words, as heterogeneity increases (decreases) in our sample, so does the ability of firms to take advantage of this feature and obtain higher (lower) profits. Focusing on income inequality, we find that an increase in the Gini coefficient leads to a significant increase in the profit share during our sample period. Extending this to the longer period of 1990 to 2021, we predict that the rise in the Gini of 0.06 (from 0.43 in 1990 to 0.49 in 2021)[6] would correspond to an increase in the aggregate profit share of over 4 percentage points between 1990 and 2021.

Related Literature Our paper relates to the growing body of literature on the determinants of firm size heterogeneity. The seminal work by Hottman et al. (2016) showed that

[5]See section 2 for more details on the potential bias. The principal reason that we believe the Nielsen Home Scanner data is not nationally representative is that the measures of income inequality we compute from it do not show the slow and steady rise in US income inequality as found in national data; see section 6.1.

[6]Data on the Gini coefficient is sourced from Statista. For details, see footnote 41.

among the possible determinants of firm size, product appeal and product scope are the most quantitatively relevant.[7] Product appeal has been further examined by Faber and Fally (2021), who find that larger firms tend to cater to the taste of the richer consumers within a country, and Aw-Roberts et al. (2018), who document that the taste component of product appeal could be quantitatively relevant in several industries.[8] Failure to account for the heterogeneity in such demand shifters can lead to substantial bias in estimating firms' productivity heterogeneity (Foster et al., 2008; Blum et al., 2018).[9] This paper examines the contribution of consumer taste heterogeneity in explaining the size heterogeneity of firms. We are agnostic about the sources of taste heterogeneity.[10]

The empirical literature has predominantly assumed that differences in taste manifest clearly when considering consumers from different geographical areas. For instance, Coşar et al. (2018) examines taste differences across countries in the car market and Jäkel (2019) in the chocolate market. Atkin (2013) studies taste heterogeneity for food products across regions in India, and builds a model of habit formation to endogenize the development of taste heterogeneity. Our paper relaxes the assumption that geography or income is the only determinant of taste heterogeneity, by dividing consumers into segments according to their demand shifters. Our approach stands in line with the review by Nevo (2011), who argues that standard consumer attributes (e.g., income, education) are only partially correlated with consumer choices and that unobserved variables (at least to the researcher) are quantitatively more relevant.[11]

As already mentioned, a closely related paper is Neiman and Vavra (2023) who also

[7] Macedoni and Xu (2018) explore the role of supply factors such as productivity and flexibility in determining firm product scope.

[8] The authors assume that differences in terms of demand shifters of one product across countries of destination can be attributed to differences in taste across countries. As they use customs data and a product has a higher level of aggregation, we consider their findings as industry-level evidence.

[9] In the context of global markets, demand heterogeneity has been considered as an important factor for the selection of export destinations by firms and for export performance. Khandelwal (2010), Johnson (2012) and Feenstra and Romalis (2014) estimate demand shifters across countries using country and product-level trade data; using firm-level data, Baldwin and Harrigan (2011), Kugler and Verhoogen (2012), Crozet et al. (2012), Hallak and Sivadasan (2013), Di Comite et al. (2014) and Roberts et al. (2018) document that heterogeneous demand (i.e., the firm-level demand residuals) are important to account for firms' exporting patterns. Several factors are responsible for the heterogeneity in demand shifters. For instance, product appeal can originate from the quality of a product that makes it perceived as intrinsically better (Fajgelbaum et al., 2011), or it can arise from the closeness of the product to the taste of consumers (Coşar et al., 2018). Furthermore, product appeal is influenced by the distribution channels, which are only partially controlled by manufacturers (Guan et al., 2019).

[10] The difference in consumer taste can be traced to genetic reasons according to Drewnowski et al. (1997).

[11] This is apparent when considering demographic characteristics of different segments of consumers. In Appendix A.9.1, in which we provide a detailed analysis for the product group of oral hygiene products, we show that demographic characteristics only partially account for different purchasing patterns. Along these lines, Neiman and Vavra (2023) find that demographic characteristics such as income, race and education cannot explain the household expenditure patterns that they focus on.

assume heterogeneous consumers. They document that over the last 15 years, households have increasingly concentrated their expenditure on specific barcode items with the preferred items differing across them, or what they call "niche" consumption. That pattern is rationalized by increasing product variety along with an increase in the fixed costs of learning about each variety. Brand (2021) finds a reduction in the price sensitivity of consumers in supermarkets during the last decade, which he also attributes to greater "niche" consumption. He argues that improvements in supply-chain management have allowed producers to increasingly focus on newer products that target particular consumers, and the reduced price sensitivity is associated with rising markups. In our paper, we also study how firms exploit consumer heterogeneity to charge higher markups and evaluate the effects of this pattern on the aggregate profit share, finding that consumer heterogeneity contributes to increasing the profit share.[12]

The paper is organized as follows. After briefly reviewing the data in section 2, section 3 describes the model that guides our empirical analysis. Section 4 outlines the strategy we use to divide consumers into segments and identify the segment taste shifter for each product. Section 5 applies the model to evaluate markups and the profit share, while section 6 focuses on how firms adjust their targeting strategies in response to consumer heterogeneity. Section 7 concludes and additional results are gathered in the Online Appendix.

2 Data

We use the Retail Scanner and the Consumer Panel (or the Home Scanner) databases collected by Nielsen (US) and provided by the Kilts Center for Marketing Data Center at the University of Chicago Booth School of Business.[13] The retailer scanner data provides the price and quantity for each transaction of barcode product by a store, and stores are widely distributed across the United States. A barcode is a 12-digit Universal Product Code (UPC). The transaction is available on a weekly basis. We use data from the GS1 Company Database to identify the company owning a given UPC. The GS1 is a not-for-profit information standards organization, which provides barcodes to manufacturers globally. The GS1 Company Database includes information documented by a firm when registering a barcode with GS1, such as company name, location, and the company GS1 prefix information. A

[12]In Appendix A.11, We also find an increase in the fixed marketing costs of targeting a segment. That a change in the fixed cost technology can drive up profits is also a mechanism explored Hsieh and Rossi-Hansberg (2021) for service industries, including retailing.

[13]While we focus on 2012 for cross-section results, we also provide time-series evidence over 2004-2018. For more detailed information about the Retail Scanner and Consumption Panel Databases, see Hottman et al. (2016), Faber and Fally (2021), and Feenstra et al. (2020).

GS1 company prefix is a unique identifier assigned to the beginning digits of every barcode registered by that company. The prefix allows us to assign individual product UPC codes to the owning company, which identifies the company that owns the brand (Hottman et al. (2016)). For instance, a firm in our data is Procter & Gamble, which owns several brands, such as Crest in the oral hygiene products, and Pantene and Head & Shoulders in the hair care products.[14] The unique numeric barcode and firm identifiers allow us to calculate firm sales and product scope, i.e., the number of barcode products produced by a firm. From the Retail Scanner Database, we source our firm-specific variables.

We trace household consumption at the barcode level using the home scanner data, which is the household-level panel for 2004-2018 and contains about 530 million unique transactions on non-service retail spending made by 186,233 households distributed across 2,978 US counties. We observe the associated price and quantity at the barcode level for each purchase by a consumer in different stores from this data.[15] In the analysis, we aggregate the expenditure and quantity (of each barcode item) across stores to focus on the consumers' annual purchases of each item. Along with the detailed budget shares for households, we also have information on some characteristics such as race, age, education, marital status, income level, and presence of children. Our dataset focuses on 64 product groups, such as food, liquor, snacks, household appliances, and personal care products. We should note that the panelists do not comprise a representative sample and

We acknowledge that the Nielsen Consumer Panel data may not be representative of US consumers and that there might be biases in the data collection.[16] In fact, the data is collected by households who scan their purchases with a home hand scanner. These households can participate for as long as they wish and receive redeemable points for consumer goods as an incentive.[17] Zhen et al. (2009) find that higher-income households and households with more members, who have larger expenditures, may have higher opportunity costs of time,

[14] We are unable to obtain a firm identifier for 12% of UPCs when we match GS1 data with the UPC lists. In these cases, we drop the UPCs without a firm identifier. It is important to note that if a grocery chain hires a company to produce a product under the grocery chain's brand, the grocery chain is considered the firm identifier, as mentioned by Hottman et al. (2016).

[15] Nielsen collects these data using scanner devices, with which households scan each transaction they have made after shopping.

[16] Nielsen uses household projection factors to ensure that the aggregate data is representative of the United States as a whole. In our analysis, these projection factors may affect the relative size of segments in an a priori ambiguous way. On the one hand, most demographic characteristics are equally represented across segments, and thus, employing the projection factors may leave our results unchanged. On the other hand, given that individuals with the same observable characteristics tend to purchase different products, using Nielsen's projection factors may lead to even larger effects of consumer heterogeneity by artificially amplifying differences across segments.

[17] See https://mamasmission.com/join-nielsen-home-scan-consumer-panel-limited-space-available/ for an advertisement calling for Nielsen Home Scan Consumer Panel participants.

and, therefore, might skip reporting some of their purchases. Similar evidence is documented by Leicester (2013), who finds that those who agree to participate in scanner data may be a self-selected sample of more price-conscious households.

3 Model

The aim of the model is twofold. First, we use the structure on the demand side of the model to identify consumer heterogeneity in the data. Second, we study the implications of consumer heterogeneity for markups, profits, and marketing costs. Our model follows Hottman et al. (2016) but with one crucial difference: there are segments of different consumers with different preferences. In particular, we assume that the nested-CES preference structure with demand shifters applies to each consumer segment, where demand shifters for each product differ across the segments. The elasticities of substitution are identical across consumer groups, however. Our model features a discrete number of multiproduct firms that compete oligopolistically: both prices and the set of products are decided by taking into account the strategic interactions between firms.

3.1 Demand

Consumers are divided (exogenously) into $n = 1, ..., N$ types or segments. We use the terms "consumer types" and "consumer segments" interchangeably. Each segment is made of $L_n > 0$ consumers with an average per-capita income of $w_n > 0$ and identical tastes within each segment. The total number of consumers is $L = \sum_{n=1}^{N} L_n$. We stress that it is not the average income of consumers that distinguishes segments, but rather, their demand shifters for UPC items, as we describe now.

All consumers have a nested preference structure with three layers. The first layer consists of a continuum of product groups indexed by g, over which consumers have the following Cobb-Douglas utility function:

$$\ln U_n = \int z_g \ln U_{gn} dg, \qquad \int z_g dg = 1, \tag{1}$$

where z_g is a product group shifter that is common across consumers of all segments. The sub-utility U_{gn} is a CES aggregate over the bundle of products produced by $f = 1, ..., F_g$ firms:

$$U_{gn} = \left[\sum_{f=1}^{F_g} \left(z_{fg} Q_{fgn} \right)^{\frac{\sigma^g - 1}{\sigma^g}} \right]^{\frac{\sigma^g}{\sigma^g - 1}}, \tag{2}$$

where $\sigma^g > 1$ is the elasticity of substitution across firms in product group g and z_{fg} is a firm-specific demand shifter. The firm-specific demand shifter is common across segments.

The third and most disaggregate level of the utility function is the consumption Q_{fgn}, which is a CES aggregate over the varieties that the firm f produces and sells to a consumer segment n. We identify a "variety" as a UPC item, denoted indexed by $i \in I_{fg}$, where I_{fg} is the set of all UPC's sold by firm f in product group g. We allow consumer segment n to potentially demand a subset of these UPC's, denoted by $I_{fgn} \subseteq I_{fg}$. The consumption Q_{fgn} is then given by:

$$Q_{fgn} = \left[\sum_{i \in I_{fgn}} \left(z_{ifgn} q_{ifgn} \right)^{\frac{\eta^g - 1}{\eta^g}} \right]^{\frac{\eta^g}{\eta^g - 1}}, \tag{3}$$

where z_{ifgn} is a UPC-firm-good demand shifter that by assumption is *common* to all consumers within the segment n, while q_{ifgn} is the per-capita quantity consumed in that segment, and $\eta^g > 1$ is the elasticity of substitution between UPCs of a firm.

To provide a clearer understanding of the preference structure, let us examine an example using Nielsen data. The first preference layer pertains to product groups, reflecting consumers' preferences for aggregate consumption of categories such as oral hygiene products, baby food, and hair care products. Within each product group, the second layer considers consumers' valuation of aggregate consumption across different companies. For example, in the oral hygiene product group, this layer combines the consumption of products from Procter & Gamble and Colgate-Palmolive (with elasticity of substitution across companies σ^g). Lastly, the third layer focuses on aggregating UPC products within individual firms. In the case of Procter & Gamble's oral hygiene products, this includes a variety of Crest products as well as other associated brands (with elasticity of substitution across products η^g).

If $\eta^g > \sigma^g$, varieties within a firm are more substitutable than varieties across firms, i.e., different varieties of Colgate are more substitutable than Colgate and Crest. By contrast, if $\eta^g < \sigma^g$, varieties within a firm are less substitutable than varieties across firms. In this latter case, varieties within a firm might be complements depending on the firm market share (see Appendix A.1.1). In the quantitative analysis, all our product groups exhibit a higher elasticity within-firm than across-firm and, on average, $\bar{\sigma^g} = 4.2$ and $\bar{\eta^g} = 7.4$.

With this demand structure, the per-capita demand for a variety identified by its UPC-firm-group-segment is given by:

$$q_{ifgn} = z_g w_n P_{gn}^{\sigma^g - 1} z_{fg}^{\sigma^g - 1} P_{fgn}^{\eta^g - \sigma^g} z_{ifgn}^{\eta^g - 1} p_{ifg}^{-\eta^g}, \tag{4}$$

where p_{ifg} is the price of the UPC and it is common across segments. In other words, we assume that firms cannot price discriminate across consumer segments, as we will confirm empirically (see section 4.2.2). The CES price indexes will still depend on the segment n, because they are defined inclusive of the demand shifters:

$$P_{gn} = \left[\sum_{f=1}^{F_g} \left(\frac{P_{fgn}}{z_{fg}} \right)^{1-\sigma^g} \right]^{\frac{1}{1-\sigma^g}}, \tag{5}$$

$$P_{fgn} = \left[\sum_{i \in I_{fgn}} \left(\frac{p_{ifg}}{z_{ifgn}} \right)^{1-\eta^g} \right]^{\frac{1}{1-\eta^g}}. \tag{6}$$

For the remainder of the model section, we drop the product group index g to ease the notation. The elasticity of the price indexes with respect to prices can be expressed as appropriately defined market shares:

$$\frac{\partial P_n}{\partial P_{fn}} \frac{P_{fn}}{P_n} = s_{fn} \equiv \frac{\left(\frac{P_{fn}}{z_f} \right)^{1-\sigma}}{\sum_{f'=1}^{F} \left(\frac{P_{f'n}}{z_{f'}} \right)^{1-\sigma}}, \tag{7}$$

$$\frac{\partial P_{fn}}{\partial p_{if}} \frac{p_{if}}{P_{fn}} = s_{ifn} \equiv \frac{\left(\frac{p_{if}}{z_{ifn}} \right)^{1-\eta}}{\sum_{j \in I_f} \left(\frac{p_{jf}}{z_{jfn}} \right)^{1-\eta}}, \tag{8}$$

where s_{fn} is firm f's revenue share within the expenditures of consumers in segment n, with $\sum_{f=1}^{F} s_{fn} = 1$, and s_{ifn} is the revenue share of barcode product i within the expenditure of consumers in segment n on firm f's UPCs, with $\sum_{i \in I_f} s_{ifn} = 1$.

These preferences allow us to nest the results of a representative consumer, in which case all segments are identical. We formalize the difference between that case and the heterogeneous consumers by reintroducing the goods subscript g and using the definition:

Definition 1 *The representative consumer (RepC) model has* $z_{ifgn} = z_{ifg}$ $\forall$ g *and* n*, while* $\forall$ g *the heterogeneous consumer (HetC) model has* $z_{ifgm}/z_{ifgn} \neq z_{jf'gm}/z_{jf'gn}$ *for at least two firms* $f \neq f'$ *each selling a UPC item* $i \neq j$ *to at least two consumer segments* $m \neq n$*, with* $i \in I_{fgm} \cap I_{fgn} \neq \emptyset$ *and* $j \in I_{f'gm} \cap I_{f'gn} \neq \emptyset$*.*

In other words, the *HetC* model has at least two firms f and f' selling the UPCs i and j, respectively, to consumer segments m and n with relative demand parameters that differ. In contrast, if the demand parameters are equal across all consumer segments, then we are in the *RepC* model, which is equivalent to a single consumer segment.

3.2 Technology

The production of one unit of UPC i by firm f requires c_{if} units of a numeraire resource, which is meant to represent the combinations of factors of production used by firms. UPCs within a firm can vary in their unit costs to capture the idea of core competence, which is a common and verified assumption in the literature (Eckel and Neary, 2010; Arkolakis et al., 2021). Furthermore, we assume that c_{if} is independent of the quantity sold and of the consumer segment that is purchasing.

Our assumption that unit costs per UPC are the same for all consumer segments means that we assume away differences in variable trade costs associated with reaching different segments. Hence, our model best represents the case in which firms face different consumer types within the same geographical area. This assumption is mainly motivated by data availability: since we do not know the geographic origin of UPCs, it is difficult to measure any differences in trade costs by segment. In addition, as price discrimination (or lack thereof) across different geographic locations is a well-studied phenomenon (Atkeson and Burstein, 2008; Cavallo et al., 2014; Simonovska, 2015; DellaVigna and Gentzkow, 2019), we prefer to focus here on optimal pricing across heterogeneous consumers when discrimination is not allowed.

3.3 Pricing

We do not consider any fixed costs of the firm.[18] Firm f's profits within a product group are equal to:

$$\pi_f = \sum_{i \in I_f} \sum_{n=1}^{N} L_n (p_{if} - c_{if}) q_{ifn}. \tag{9}$$

Firms maximize their profits by choosing p_{if} $\forall i \in I_f$, as well as the product set I_f. Firms take as given the vector of UPC-segment-specific demand shifters z_{ifn}, and the marginal costs of production c_{if}.

In a standard multiproduct firm setting with Bertrand competition, a firm internalizes two effects from raising the price of a UPC. First, raising the price increases its own price index P_{fn}, which increases the demand for all its other UPCs. Second, it increases the aggregate price index P_n, further raising the demand for all its other UPCs depending on how large the firm is. In our framework, there is an extra dimension due to segment heterogeneity – firms internalize the fact that raising the price of one variety has different effects across segments. The fact that firms recognize their impact in these three dimensions (or the first

[18] We introduce fixed costs in Appendix A.11.

two with representative consumers) is what we mean by "large" firms. In contrast, if firms treat the price indexes – by acting as if their shares were zero in (7) and (8) – then they are "small", atomistic firms.

The equilibrium markups $\mu_{if} \equiv p_{if}/c_{if}$ for all UPC-firm (if) pairs are implicitly defined by the first-order condition:[19]

$$(\sigma-1)\left[1-\frac{\sum_{n=1}^{N}\alpha_n s_{fn}^2 s_{ifn}}{\sum_{n=1}^{N}\alpha_n s_{fn} s_{ifn}}\right]=\frac{\eta}{\mu_{if}}-\frac{\sum_{n=1}^{N}\alpha_n s_{fn} s_{ifn}\left[(\sigma-1)s_{fn}+(\eta-\sigma)\right]\sum_{i\in I_f}\frac{s_{ifn}}{\mu_{if}}}{\sum_{n=1}^{N}\alpha_n s_{fn} s_{ifn}} \tag{10}$$

where $\alpha_n = L_n w_n / \sum_{n=1}^{N} L_n w_n$ denotes the share of income of consumers in segment n. Several features of these markups should be noted.

First, we have indexed the markups by the UPC item i sold by firm f, because markups can differ across the items sold by a multiproduct firm even if the marginal costs of those items happen to be equal. This feature of *variable markups* with heterogeneous consumers is also found by Neiman and Vavra (2023) and will be demonstrated empirically in our model. It stands in marked contrast to the case of a multiproduct firm in a *RepC* model. The *RepC* model is equivalent to the multiproduct firm selling to one segment, i.e., $N = 1$ in the summations within (10). The markup is then immediately solved from (10) as:

$$\mu_{\mathrm{RepC}}=\frac{\sigma-(\sigma-1)s_f}{(\sigma-1)\left(1-s_f\right)}, \tag{11}$$

where $\sum_{i\in I_f} s_{ifn} = s_{fn} = s_f$ is the *total market share* of the firm since there is only one segment, $N = 1$. The markup of a firm in (11) is increasing in its total market share s_f, and *only* depends on that total share in the *RepC* model. Even though a firm might be selling multiple products with different individual market shares, because it internalizes the cross-demand effects of changing any individual price, it ends up charging the same markup on all its products (regardless of their marginal costs).[20]

We can contrast the *RepC* markup in (11) with the markup charged by a firm with the same total market share, but facing heterogeneous consumers. For simplicity, consider the markup in (10) for single-product firms in the *HetC* model satisfying Definition 1. Then we have the following result:

[19] The derivation of (10) as well as the proof of Proposition 1 are in Appendix A.1.

[20] This result is demonstrated in the CES case for a representative consumer by Hottman et al. (2016) and Feenstra (2004, p. 267).

Proposition 1 *The HetC model with single product firms has markups that are at least as high as the markups for single or multiproduct firms in the RepC model, conditional on having the same firm-level market shares, and strictly higher for at least one firm.*

To establish this result, consider the optimal markup in (10) of a firm f with only one UPC, i.e., $I_f = \{i\}$. As a result, $s_{ifn} = 1$. Firm f's markup is then computed from (10) as:

$$\mu_{\text{HetC}} = \frac{\sigma - (\sigma - 1)\frac{\sum_{n=1}^{N} \alpha_n s_{fn}^2}{\sum_{n=1}^{N} \alpha_n s_{fn}}}{(\sigma - 1)\left[1 - \frac{\sum_{n=1}^{N} \alpha_n s_{fn}^2}{\sum_{n=1}^{N} \alpha_n s_{fn}}\right]}. \tag{12}$$

We see that for a single-product firm in the *HetC* model, the markup in (12) is increasing in $\sum_{n=1}^{N} \alpha_n s_{fn}^2 / \sum_{n=1}^{N} \alpha_n s_{fn}$, which is the sum of squared market shares across segments, weighted by α_n and divided by the firm's total market share $s_f = \sum_{n=1}^{N} \alpha_n s_{fn}$. In contrast, the markup in (11) is increasing in s_f. By Jensen's inequality, $\sum_{n=1}^{N} \alpha_n s_{fn}^2 / \sum_{n=1}^{N} \alpha_n s_{fn} \geq \sum_{n=1}^{N} \alpha_n s_{fn} = s_f$, and in the proof of Proposition 1 we show that this inequality holds strictly for at least one firm given the heterogeneity in demand shifters in Definition 1. For the same firm-level market shares, therefore, at least one single product firm that sells to multiple segments will have a higher markup than a firm selling to a single segment. Since the markup in (12) increases with s_{fn}^2, we can also conclude that the larger the variance in s_{fn} across segments, the higher the firm's markups.

The intuition for this result is as follows. In the presence of consumer taste heterogeneity, firms face different perceived demand elasticity across segments. Rather than charging a markup consistent with an average of the perceived demand elasticities, however, firms optimally use a greater weight on low-elasticity segments. The higher profits obtained by charging higher markups to low-elasticity and all other segments more than offset the loss in demand on the high-elasticity segments.

The effect of the two elasticities of substitution σ and η on markups of the *HetC* model is similar: higher σ or η, all else constant, yield lower markups. Using numerical methods we show that changing σ has a much larger effect on firm-level markups than η. In fact, while σ affects the average markup level of a firm, η affects more the dispersion of markups across the products within a firm (see Figure A.1 in the Appendix). We also note that in the *RepC* model, η does not enter the markup formula, and, thus, it only affects markups when consumers are heterogeneous.

The Aggregate Profit Share. We have shown above that the markup of at least one single-product firm selling to heterogeneous consumers is larger than the markup of a firm, with identical aggregate market share, selling to a representative consumer. Hence, conditional on the firm-level market share, the implied profit share of a model with heterogeneous consumers is larger than that of a model of a representative consumer. In addition, since markups under oligopoly are larger than markups under the assumption of atomistic firms, the implied profit share of the representative consumer model with large firms is larger than that implied by a model of small firms. So we obtain a ranking of markups across these three models, conditional on the firm-level market shares.

We will measure the aggregate share of profits in each product group as follows:

$$\frac{\Pi}{R} = \frac{\sum_{f=1}^{F}\sum_{i\in I_f}\pi_{if}}{\sum_{f=1}^{F}\sum_{i\in I_f} r_{if}} = \sum_{f=1}^{F}\sum_{i\in I_f}\frac{\mu_{if}-1}{\mu_{if}}\frac{r_{if}}{\sum_{f=1}^{F}\sum_{i\in I_f} r_{if}} = \sum_{f=1}^{F}\sum_{i\in I_f}\frac{\mu_{if}-1}{\mu_{if}} s_{if}s_f, \tag{13}$$

where r_{if} denotes the revenues of a firm f in product i and $\frac{\mu_{if}-1}{\mu_{if}}$ is the ratio between profit margin and markups at the product-firm level, which is the profit share at the product-firm level. The aggregate share of profits is the sum of each product-level profit share, weighted by the product revenue share within a firm s_{if} and firm market share s_f.

In our quantitative exercise, we observe the market shares at a high level of disaggregation and apply the model to measure the implied markups, thus estimating the aggregate share of profits as shown above. We stress that in order to estimate the markups from (10) we need to have the UPC-firm share (if) sold within each product group g *to each* consumer segment n, or s_{ifgn}. Measuring these shares will depend on our identification of meaningful consumer segments, which we consider in the next section.

4 Consumer Heterogeneity

As a first step, we describe the strategy we follow to estimate the segment-specific demand shifters. Reintroducing the product group notation g, we take logs of (4) to decompose the demand for a UPC as:

$$\ln q_{ifgn} = \underbrace{\ln\left(z_g w_n P_{gn}^{\sigma^g-1} z_{fg}^{\sigma^g-1} P_{fgn}^{\eta^g-\sigma^g}\right)}_{\text{Firm-segment}\equiv\ln\Phi_{fgn}} + (\eta^g-1)\times \underbrace{\ln z_{ifgn}}_{\text{UPC-Firm-segment}} - \eta^g \ln p_{ifg}. \tag{14}$$

After controlling for the price, the quantity demanded of a UPC by a segment is a function of a firm-segment component, which we label Φ_{fgn}, and a firm-segment-UPC component, denoted by z_{ifgn}. Let us assume that the demand shifter z_{ifgn} can be decomposed into two components: first, $\bar{\phi}_{ifg}$ is a firm-UPC component that captures the quality or appeal of the product and is globally recognized across consumers' segments; and second, ϕ_{ifgn} is a firm-UPC-segment component that captures the taste of consumers in segment n for the UPC i produced by firm f. In particular:

$$\ln z_{ifgn} = \ln \bar{\phi}_{ifg} + \ln \phi_{ifgn}. \tag{15}$$

Because $\bar{\phi}_{ifg}$ captures the demand shifter that is commonly recognized across consumers' segments, without loss of generality, we normalize the taste component (firm-UPC-segment) of the demand shifter so that it has mean zero:

$$\frac{1}{N_g}\sum_{n=1}^{N_g} \ln \phi_{ifgn} = 0. \tag{16}$$

Hence, differences in taste *across segments* are represented by the difference of ϕ_{ifgn} from zero. To empirically identify such demand shifters, we follow a two-step procedure, which we describe in the following sections. First, we divide consumers into segments (sections 4.1 and 4.2). Second, we estimate the segment-specific demand shifters using (14) (section 4.3).

4.1 Assignment of Consumers to Segments

As noted earlier, consumer types are defined as having identical demand shifters for UPC items within a product group. Empirically, we allocate each of the L household into one of the non-overlapping sets H_n, where $\cup_{n=1}^{N} H_n = \{1, ..., L\}$, according to the similarity of their budget shares. Let x_{ifh} denote the budget share of household h on UPC-firm if relative to the entire sample budget share on the same UPC, so x_{ifh} is defined as:

$$x_{ifh} = \frac{q_{ifh}p_{if}}{w_h} \Big/ \frac{\sum_{h' \in H_n} q_{ifh'}p_{if}}{\sum_{h' \in H_n} w_{h'}}.$$

where w_h is the total expenditures of consumer h in a given product group. We apply this definition to each product group separately, and thus, x_{ifh} is defined for all consumers and products within a product group. The relative budget share of a consumer captures a simple measure of its demand shifter, or tastes, relative to the aggregate demand shifter.[21]

Empirically, we use the k-means clustering algorithm to assign each consumer h in our data to a segment, or cluster H_n. The k-means clustering is a statistical algorithm for classifying objects in clusters based on their attributes in a certain number of categories (Everitt,

[21] We have verified that the results are robust to using the absolute purchases on goods rather than the relative budget shares.

1993).[22] The method employs the distance between objects as a dissimilarity measure when forming the clusters such that each object belongs to the cluster with the nearest mean.

The assignment of consumers into a segment is done by minimizing the within-cluster dissimilarity between individuals in terms of the relative budget shares x_{ifh}.[23] The objective function of the algorithm is:

$$\min \sum_{n=1}^{N} \sum_{h \in H_n} \sum_{f=1}^{F} \sum_{i \in I_f} ||x_{ifh} - m_{ifn}||, \tag{17}$$

where m_{ifn} is the segment-specific mean of the relative budget share x_{ifh}, and $||x_{ifh} - m_{ifn}||$ is, therefore, a measure of the dissimilarity between the relative budget shares of individual h and the mean for the segment n.[24]

The k-means clustering algorithm requires two choices: i) the number of segments N and ii) the measure of the dissimilarity defined by the operator $||\ ||$. The number of segments N is chosen in order to maximize the Calinski-Harabasz pseudo-F statistics. A higher value of the F statistics is associated with a larger dissimilarity across clusters. In practice, our algorithm requires two loops. In the inner loop, given N, the algorithm divides consumers into segments and computes the Calinski-Harabasz pseudo-F statistics. We run the inner loop for $N = 2, ..., 15$, and choose the number of segments N that maximizes the Calinski-Harabasz pseudo-F statistics. Finally, we use the Minkowski distance of order one, i.e., the absolute value, as the measure of dissimilarity.[25]

[22]K-means clustering finds similarity between the points and it is then able to group them into clusters, based on similarity. It's an unsupervised machine-learning algorithm, which has been widely used in economic analysis to classify states into regions with similar cycles (Crone and Clayton-Matthews, 2005), to group countries based on economic performance (Levy-Yeyati and Sturzenegger, 2005; Caballero, 2014; Ghosh et al., 2014), and to classify firms into technological clusters (Chyi et al., 2012).

[23]As a robustness check, we also consider two additional criteria to divide consumers into segments based on the geographical location of consumers, and their preferred distribution channel, i.e., type of store: see Appendix A.7. In the main text, we focus on the consumer type as defined by the relative budget shares because, in our model, firms cannot price discriminate across consumers, and there are no differences in costs to reach different segments. Empirically, there is little evidence of price discrimination within retail chains across geographic locations, though different chains charge different prices (DellaVigna and Gentzkow, 2019).

[24]We use simple weights in dividing consumers into segments. Using different weights (as the projection factors provided by Nielsen) would not affect qualitatively the results. In addition, as discussed in section 4.2, though the demographic characteristics of segments may vary across years to reflect changes in the sample of households, of goods, and of tastes, we do not find significant trends across years for most characteristics. Hence, modifying the weights used in the clustering algorithm based on the demographic characteristics of households may lead to an overestimation of the differences between segments.

[25]In Stata, this is achieved with the command cluster k-means **x**, start(firstk) measure(L(1)) k(5). Our results remain robust to using Minkowski distance of different order to cluster consumers.

4.2 Segmentation Results

Representation of Consumer Characteristics

To demonstrate the above method, let us consider the case of oral hygiene products. The k-means clustering algorithm divides households into seven types (where the Calinski-Harabasz pseudo-F statistics is maximized). Table 1 summarizes the main characteristics of each segment. For exposition purposes, we rank segments by their aggregate expenditure shares in the toothpaste market: Segment 1 is the largest segment, as it captures 29% of the market, while Segment 7 is the smallest, as it captures 4% of the market. Segment 3 exhibits the largest expenditure per household on toothpaste ($56), while Segment 6 is the smallest ($35). While segments 1 and 2 have average demographic characteristics, segments 3, 4, and 7 are over-represented by high-income households, and segments 3 and 5 by young households. Segment 6 is over-represented by low-income and older consumers. Race and ethnicity are fairly similar across segments (see Appendix A.9). The segments also differ in their preferred company: Procter & Gamble is the preferred company by segments 1 and 2, with its Crest brand being the most popular. Warner-Lambert (Listerine) is the preferred company for Segment 2, GlaxoSmithKline (Aquafresh) for Segment 3, Block Drug (Sensodyne) for Segment 5, Colgate-Palmolive for Segment 6, and Church & Dwight (AIM) for segment 7.

Table 1: Segments characteristics - Oral Hygiene

Segment #	Exp per Household ($)	Total Share	Main Brand	Distinctive Features
1	44	0.29	Procter & Gamble	Average
2	44	0.27	Warner-Lambert	Average
3	56	0.14	Glaxosmithkline	High-Income, Young
4	49	0.12	Procter & Gamble	High-Income
5	51	0.08	Block Drug	Young
6	35	0.05	Colgate-Palmolive	Low-Income, Older
7	48	0.04	Church & Dwight	High-Income

In applying the method to all product categories, we use the relative expenditure shares of UPCs purchased by at least 1% of the households in the data so that we have heterogeneous consumption purchases within a segment. Table A.1 in the Appendix reports the number of segments by consumer type per product category. The number of segments according to this criterion ranges from two (e.g., breakfast food) to seven (e.g., oral hygiene and cosmetics).

We analyze the differences across segments in terms of their demographic characteristics. Some demographic characteristics such as age, income, marital status, and education tend to be more important than race and the presence of children, though overall, demographic characteristics play a very limited role in defining segments.[26] The importance of some demographic characteristics increases when we examine segment characteristics within product groups.[27] Though the demographic characteristics of segments may vary across years to reflect changes in the sample of households, of goods, and of tastes, we do not find significant trends across years for most characteristics.[28]

4.2.1 Representation of Geographic Space

We also provide evidence that the geographic distribution of consumer segments is widely spread across space in the United States. If some types of consumers are concentrated in particular regions, firms may practice price discrimination across segments by setting different prices in different locations, which is inconsistent with the conceptual framework. We test whether the conditional population share of one type of consumer in a region equals its unconditional population share at the national level. As shown in the Appendix, there are low rejection rates (i.e., less than 19%) through all years, suggesting that the probability that certain types of consumers are clustered in some states is low.[29]

4.2.2 Potential Threat of Price Discrimination

To ensure there is no (potential) price discrimination across segments, we follow DellaVigna and Gentzkow (2019) and use the following price equation, where we regress household-UPC-location prices on UPC-, segment-, retailer-, and location-fixed effects:

$$\ln p_{ifdrh} = \delta_{if} + \delta_d + \delta_r + \delta_n + \epsilon_{ifdrh} \tag{18}$$

where $\ln p_{ifdrh}$ denotes the average price of UPC product i supplied by firm f, which household h in region d purchased from retailer r. We run equation (18) by product group and

[26] In Table A.3, we provide measures of the variability of the share of households with different demographic characteristics across segments.

[27] In Figures A.3 and A.4, we show measures of the difference in demographic characteristics across segments per product group, finding that for some product groups, certain demographic characteristics become more important.

[28] See Figures A.5 and A.6. There is one notable difference in household incomes. In fact, the difference across segments tends to decrease after 2010.

[29] See Table A.4. Table A.5 lists product groups where our null hypothesis has been rejected in the test. To rule out the possibility that results are not driven by some outlier states, especially those with the least represented states in the consumer survey, we repeat the exercise and drop the lowest represented states with the lowest number of consumers. This set of results is provided in panel B of Table A.4, and product groups where our null hypothesis has been rejected in the test are provided in A.6. More detailed discussion is provided in appendix A.3.

year and conduct a joint test of whether the coefficients of segment dummies are statistically significant. The rejection of the null hypothesis suggests that the price of an identical product may be significantly higher or lower for a certain type of consumer, suggesting potential price discrimination across consumer segments.

Table 2 summarizes the results of various statistical tests based on (18). The second column reports the joint test on whether segment dummies are significantly different from each other. The results show that in only 2 to 13 product categories (out of 64) between 2004 and 2018, we reject the null hypothesis. In columns (3) to (5), we report the contributions of different fixed effects to the model fit in terms of R^2. In panels A and B, a region is defined by a county and state, respectively. In both panels, product heterogeneity alone explains about 64% of the price variation for all product categories across space. The region, retailer, and segment fixed effects additionally explain about 3.0% of price variations in panel A, and 2.4% in panel B. The last column reports the R^2 explained by segment fixed effects. The small number suggests that consumer segments play a minimal role in explaining price variations. The pattern remains robust to analyzing different years. Overall, the analysis of the explanatory power of segment fixed effects shows that the extent of price discrimination across consumer segments is not economically significant.

4.2.3 Robustness Checks of Clustering Algorithm

We additionally provide robustness checks using the K-mean clustering algorithm to categorize consumers based on (i) the number of UPC products by brand as a share of the total number of products purchased by the household, and (ii) a dummy variable indicating whether the households purchase a brand. In each case, we repeat the clustering process by year and product category, and we compute the probability that a consumer is classified into a segment where the majority of consumers had also appeared in a given segment in the baseline case.[30] Overall, we observe that consumers are still put together in more or less the same segments generated in the baseline under the two alternative criteria in clustering.

4.3 Estimation of the Taste Parameter

Having divided consumers into segments, we estimate product and product-segment-specific demand shifters. To achieve this, suppose that (14) represents the demand by consumer h in segment n. That is, we replace segment n in (14) by consumer h, and then for each UPC

[30]See Appendix A.4 for an example of this procedure, and the distribution of this probability is summarized in Figure A.7. Figure A.8 compares the average number of segments per product group (2004-2018) generated by the clustering algorithm based on the relative budget share (i.e., baseline), (i), and (ii), respectively.

Table 2: Test for Potential Price Discrimination

Year	Number of Rejections (F-stat) H_0: $\delta_1 = \delta_2 = ... = \delta_n$	R^2 Decomposition: Product only (δ_{if})	Region-retailer-segment ($\delta_d + \delta_r + \delta_n$)	Segment only (δ_n)
Panel A: County as Region				
2004	3	0.640	0.029	0.0002
2005	3	0.643	0.029	0.0002
2006	4	0.645	0.029	0.0002
2007	4	0.640	0.027	0.0003
2008	2	0.637	0.027	0.0003
2009	4	0.629	0.028	0.0003
2010	7	0.621	0.027	0.0003
2011	5	0.623	0.026	0.0003
2012	7	0.613	0.028	0.0003
2013	8	0.615	0.028	0.0003
2014	6	0.625	0.029	0.0002
2015	13	0.629	0.028	0.0003
2016	10	0.634	0.028	0.0002
2017	7	0.637	0.029	0.0002
2018	7	0.637	0.029	0.0003
Panel B: State as Region				
2004	4	0.640	0.023	0.0003
2005	3	0.643	0.023	0.0003
2006	3	0.645	0.023	0.0002
2007	4	0.640	0.022	0.0003
2008	3	0.637	0.022	0.0003
2009	6	0.629	0.022	0.0003
2010	5	0.621	0.021	0.0003
2011	6	0.623	0.021	0.0003
2012	7	0.613	0.022	0.0003
2013	10	0.615	0.022	0.0003
2014	6	0.625	0.023	0.0003
2015	14	0.629	0.023	0.0004
2016	11	0.634	0.023	0.0003
2017	7	0.637	0.024	0.0003
2018	9	0.637	0.024	0.0003

Note: The table reports statistics tests for the potential price discrimination across consumer segments based on equation (18). We run the regression for each of the 64 product groups by year independently, and the second column reports the number of rejections of the null (out of 64).

we take the average across all consumers in the same segment n to obtain:

$$\overline{\ln q_{ifgh}} = \ln \Phi_{fgn} + (\eta^g - 1) \ln z_{ifgn} - \eta^g \ln p_{ifg}, \tag{19}$$

where $\overline{\ln q_{ifgh}} \equiv \sum_{h \in H_n} \ln q_{ifgh}/L_n$ is the average quantity demand by consumers in segment n. By assumption, $z_{ifgh} = z_{ifgn}\ \forall h \in \{H_n\}$, hence, $\sum_{h \in n} \ln z_{ifgh}/L_n = \ln z_{ifgn}$. We estimate the firm-segment demand shifter $\ln \Phi_{fgn}$ in (19) by running this regression separately for each firm-segment pair fn in each product group g. In other words, for each sample consisting of UPC products supplied and consumed by the firm-segment pair, $(\eta^g - 1) \ln z_{ifgn}$ is obtained as the residual of regression of $\overline{\ln q_{ifgn}} + \eta^g \ln p_{ifg}$ on a constant that captures the value $\ln \Phi_{fgn}$. We use the elasticity of substitution η^g from Hottman et al. (2016), so from this procedure we obtain the estimates $\hat{z}_{ifgn}$.[31]

With the demand shifters $\hat{z}_{ifgn}$ for each UPC item, we apply the normalization outlined in (15) and (16) to obtain the common component of the demand shifter $\ln \hat{\bar{\phi}}_{ifg}$ and the segment-specific taste parameter $\ln \hat{\phi}_{ifgn}$. The segment's taste for a particular UPC item is captured by the deviation of that demand shifter from the average demand shifter across consumer segments for that UPC. The identifying assumption is that, conditional on prices and on the firm-segment component of demand, differences in consumption patterns across segments reflect different preferences. In other words, we condition on a given firm-UPC item, which is sold at the same price and with the same characteristics across different segments, and then demand differences reflect the heterogeneity of tastes.

5 Markups and the Profit Share

Combining our data on UPC-segment-specific sales, we evaluate the implications of consumer heterogeneity for firm markups and profits. We first examine the cross-section across firms in 2012 and then examine the change in average profits and markups over time.

[31] Hottman et al. (2016) estimate the elasticities of substitution following an approach introduced in the literature by Feenstra (1994) and subsequently used by Broda and Weinstein (2006). Since the demand elasticity is the same across consumer segments, the aggregate demand for a product predicted by our model is isomorphic to that of Hottman et al. (2016). Hence, our model with heterogeneous consumers nests theirs. Since the estimation of the elasticities of substitution is based on the aggregate demand for a product variety, it is appropriate for us to use the demand elasticities η^g and σ^g from Hottman et al. (2016) in our quantitative exercise.

5.1 Markups and Profits Across Firms

For each product group, we compute segment, firm-segment, and UPC-firm-segment-specific market shares. We use the elasticities of substitution from Hottman et al. (2016) to solve for UPC-specific markups using equation (10). Panel (a) of Figure 1 shows the relationship between the implied UPC markups in our model and the markups that would be implied by the *RepC* formula (11). There is significant heterogeneity of markups within a firm. This is in contrast to the *RepC* result that markups are constant within a firm. For instance, for the top-selling firm in shaving needs, markups $\mu_{if} = p_{if}/c_{if}$ vary from 1.69 to 2.81.

While most UPCs' markups are larger than the *RepC* formula, there are some UPCs where markups are lower. For the top-selling firm in shaving needs, as mentioned just above, the *RepC* formula predicts markups of 1.81 and there are a few products sold by that firm with markups below that value.[32] Across product groups, UPCs with markups below the *RepC* prediction account for 6.1% of UPCs, and they are UPCs with low market share and low dispersion of market share across consumer segments. In panel (b) of Figure 1, we report the weighted average markup by a firm, where the weights are the UPC's market share within a firm. The presence of consumer heterogeneity generates larger markups than those implied by the *RepC* formula. For the top firm in shaving needs, average markups are 2.67 against the *RepC* formula of 1.81.

Figure 1: Implied Markups

(a) UPC-specific Markups (b) Average Firm Level Markups

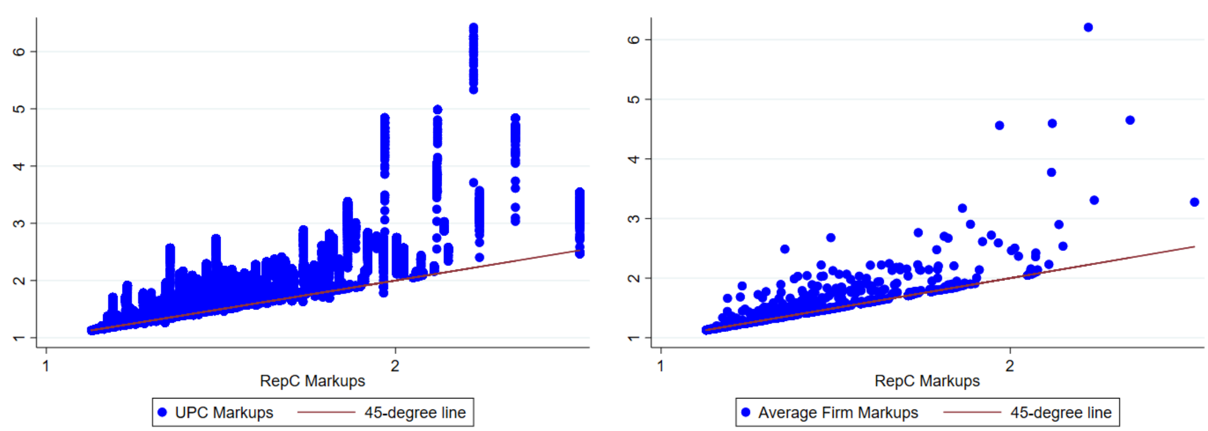

The difference between our markup formula and the *RepC* result differs considerably across product groups and across firms within a product group. On average, the presence

[32] Notice that having a lower markup than the *RepC* formula for some items sold by a multiproduct firm in the *HetC* model does not contradict Proposition 1, because that result assumes that firms in the *HetC* model are selling only a single product.

of taste heterogeneity across consumer types implies that markups are 20 percentage points larger on average across product groups. For the top five and top one firm within each product group, average markups across groups are 27 percentage points and 39 percentage points larger. Therefore, the presence of consumer heterogeneity generates markups that are higher relative to the case of consumer homogeneity, and these differences are magnified for large firms.

Figure 2 shows the average percentage point markup difference between our *HetC* model and the *RepC* formula by product group, as well as the difference for the top five and top one firms. The most significant differences are found in dairy product groups (e.g., eggs, butter and personal care (e.g., shaving needs, diapers). These are product groups with a relatively larger number of consumer segments and relatively higher concentrations. For shaving needs, the average markup in the *RepC* model is 1.67 and it rises to 2.2 in the *HetC* model, while for the top firm these markups are 1.81 and 2.67. The smallest markup differences are found in wine, health aids, and fresh produce.[33]

Figure 2: Difference in Markups by Product Group: *HetC* versus *RepC*

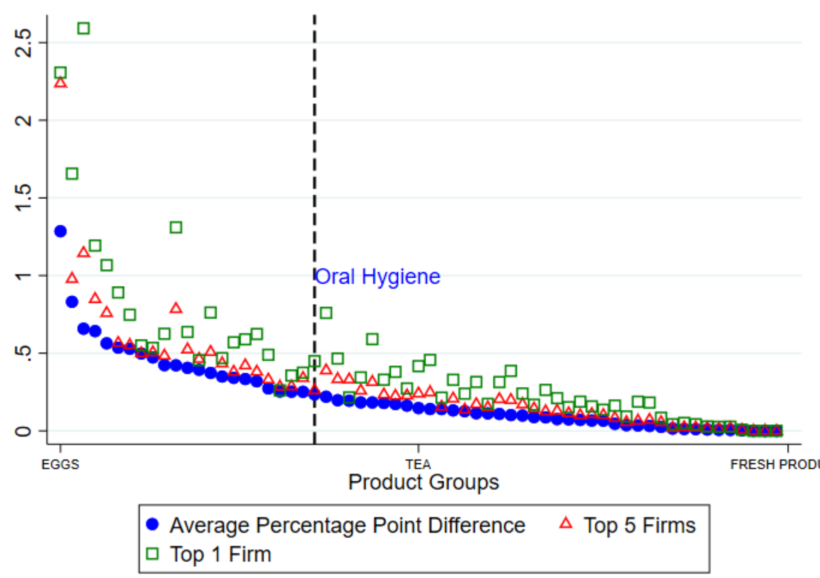

Similar results emerge when we compute the profit share as the ratio between variable profits and revenues, as in (13). The aggregate profit share in 2012 is 35.5% and this value is in line with the results of Karabarbounis and Neiman (2014), who consider a wider sample of firms than just retailers. In Figure 3, we compare the profits implied by our model to those that emerge in the *RepC* model of a representative consumer. Panel (a) shows that profits are always larger in our *HetC* model, regardless of firm size. The aggregate profit

[33]In Table A.2, we report the differences by product group.

share (35.5%) is 6.4 percentage points larger than with a representative consumer (29.1%) in 2012.[34] There is considerable heterogeneity across product groups, as the largest difference is in the disposable diaper category, and the smallest in fresh produce. In a model with small firms, and constant markups, the implied profit share is 25.4%, which is 10 percentage points smaller than our *HetC* case.

Figure 3: Total Profits: *HetC* versus *RepC*

(a) Firm-Level (b) Product Group-Level

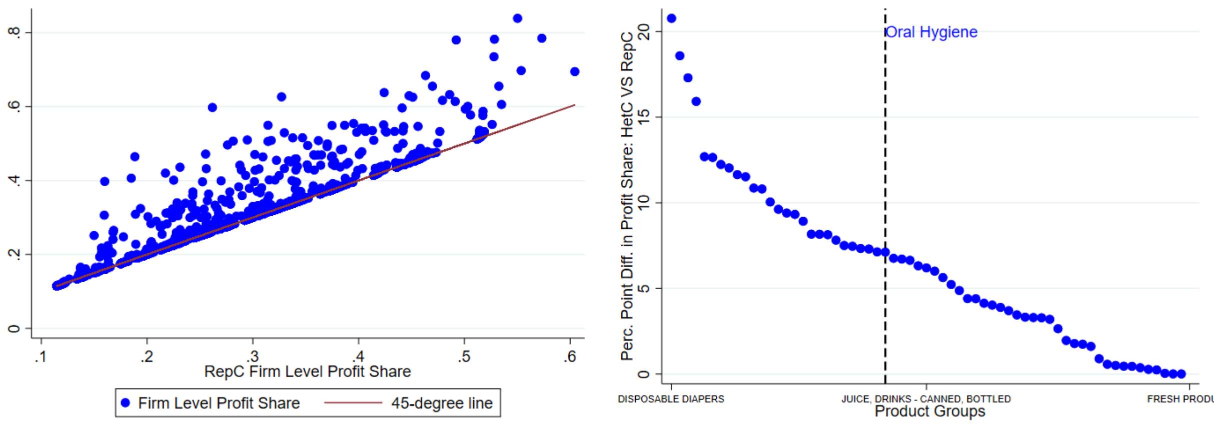

We also examine how firm size and its interaction with the *HetC* model affects the firm-level profit share. First, note that in each product group, a few large firms account for much of the market share. On average across product groups, the largest firm accounts for 29% of the market sales, the largest 5 firms for 66%, and the largest 10 for 78%.[35] Within each product group too, a small number of large firms hold a significant portion of the market share. On average, the largest firm in a product group represents 29% of the market sales, while the top 5 firms account for 66%, and the top 10 firms contribute to 78%.[36]

Figure 4 displays the percentage point difference between the firm-level profit share predicted by the *HetC* model and the *RepC* model, in comparison to the profit share predicted

[34] In computing the profit share and the aggregate markup for each product group, we use the same weights across UPC items and firms as shown in (13). The only difference between these two calculations, therefore, is that the profit share is averaging the Lerner index $(\mu_{if}-1)/\mu_{if}$, whereas the aggregate markup is averaging μ_{if}. To get a rough idea of the difference in these calculations, consider the average markups for the *HetC* model (1.64) and *RepC* model (1.44), which differ by 20 percentage points. The percentage point difference between the corresponding Lerner indexes is 8.5 (i.e., $(1-1.64^{-1})-(1-1.44^{-1})=0.085$).

[35] Figure A.9 in the Appendix provides the breakdown across product groups highlighting the heterogeneity of the role of large firms across groups: in disposable diapers, the largest firm accounts for more than half of the market sales, while in medications and remedies, the largest firm accounts for 10% of the market sales.

[36] In the Appendix, Figure A.9 offers a detailed breakdown across product groups, emphasizing the varying importance of large firms in different groups. For instance, the largest firm in the disposable diaper group dominates with over 50% of the market sales, whereas in the medications and remedies group, the leading firm only represents 10% of the market sales.

by the small firms model. This difference is plotted against each firm's market share. For firms with a market share close to zero, the difference between the *RepC* model and the small firms model is minimal. However, as the firm's market share grows, this difference also increases. The gap between the *HetC* model and the small firms model is larger at any given level of market share and becomes even more pronounced at higher market share levels.

Figure 4: Firm Profit Share Relative to Small Firms Model (Percentage Point Difference)

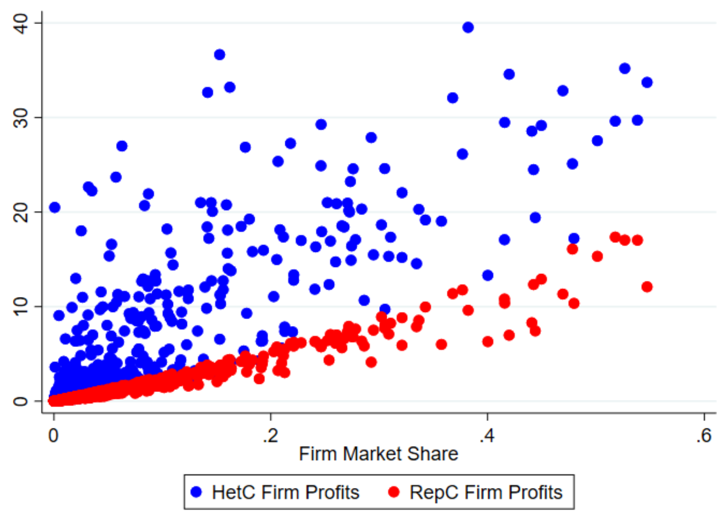

The figure plots the percentage point difference between the firm-level profit share under *HetC* and *RepC* model relative to the profit share of a small firms model. In particular, for each firm, we report $100 * ((\pi/r)_{HetC} - (\pi/r)_{MC})$ and $100 * ((\pi/r)_{RepC} - (\pi/r)_{MC})$.

5.2 Profits Over Time

In this section, we study the change in the aggregate and product group-specific profit share from 2004 to 2018. We apply the clustering algorithm to the households in the sample of each year. Note that the Consumer Data is more akin to a repeated cross-section than to panel data. We then apply our model to infer markups.

Panel (a) of Figure 5 reports the aggregate profit share across models. Our *HetC* model of large firms and consumer heterogeneity predicts the largest profit share across years. From 2004 to 2010, the profit share predicted by the *HetC* increases by 0.8 percentage points from 35.2% to 36%. In the following years, the profit share declines until the 2018 value of 33%. A similar pattern arises in the *RepC* model, where the predicted profit share exhibits a slight increase till 2010 and then it begins to decline. In the model with small firms, the profit share is relatively constant and slightly changes due to changes in the expenditure composition across product groups.

Relative to a model of small firms, which features constant markups, our baseline model has two additional margins: large firms and consumer heterogeneity. We quantify the importance of each channel in panel (b) of Figure 5, where we compute the difference in the profit share between the *HetC* model and the *RepC* model and between the *RepC* model and the small firms model. The difference in the profit share predicted by *HetC* and the *RepC* models grows from 5.6 percentage points in 2004 to 6.4 percentage points in 2012 and then declines until the 2018 value of 4.8. Hence, the role of consumer heterogeneity increases in the first part of the sample and declines afterward. The difference between *RepC* and small firms slightly grows from 4 percentage points in 2004 to 4.3 percentage points in 2010. In line with the change in the profit share, there is a subsequent decline in the difference which attains the lowest value of 3 percentage points in 2018. Thus, the role of large firms follows the same pattern as the role of consumer heterogeneity, with an earlier peak.[37]

The qualitative pattern of an earlier increase and subsequent decline in the profit share is robust to 1) holding constant the expenditure shares across product groups, 2) dropping a few product groups that experience a sharp decline in the profit share in a single year, and 3) using alternative clustering criteria.[38]

Figure 5: Profit Share Over Time

(a) Aggregate Profit Share Across Models (b) Percentage Point Difference Across Models

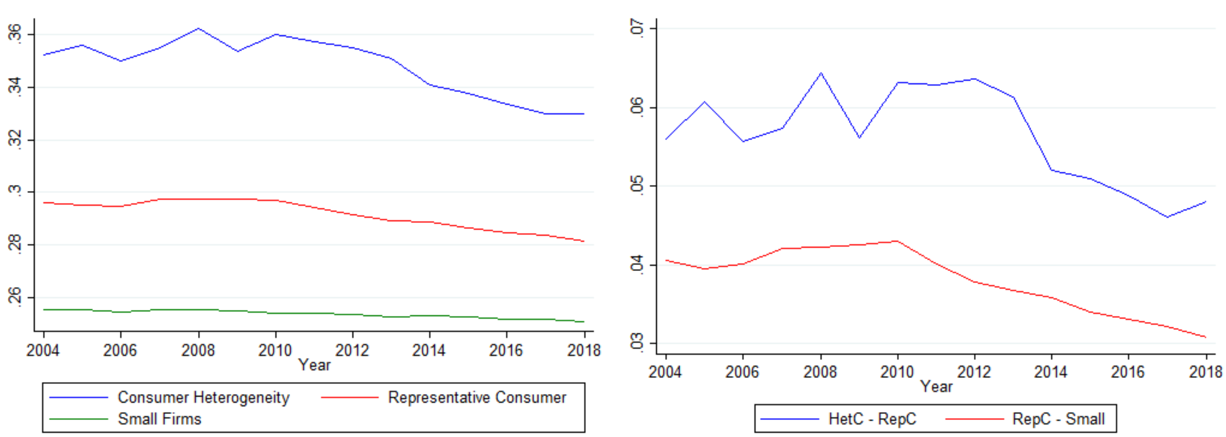

As shown in the cross-section results, there is considerable heterogeneity across sectors, which we illustrate in Figure 6. In some product groups, the profit share has increased, as in crackers, while in others, such as in prepared foods and baby food, it has decreased.

[37]In Figure A.13 of the appendix, we show that the decline in the profit share after 2010 is associated with an increase in the number of firms, which reduces the average market power, and a reduction in the average scope. These changes in the number of firms and in their scope can contribute to the decline in profits after 2010 for the *RepC* model, in particular, where consumer heterogeneity does not play a role.

[38]Results of these robustness checks are provided in Appendix A.5.

The effects of consumer heterogeneity also vary across groups. In panel (b) of Figure 6, we show that there is a positive correlation between the change in profit share and the change in the difference between our prediction and the prediction of the representative consumer model. This result further highlights the importance of consumer heterogeneity in driving the changes in the profit share.[39]

Figure 6: Percentage Point Change in Profit Share Across Product Groups

(a) Change in Profit Share　　(b) Contribution of Consumer Heterogeneity

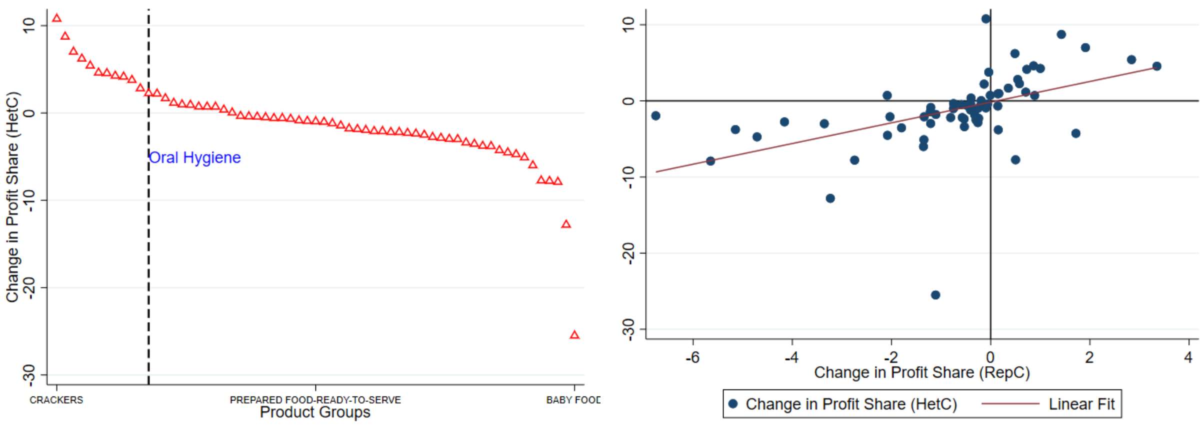

6 Explaining the Changes in the Profit Share

This section investigates possible determinants of the change in the profit share documented above. In particular, we consider changes in consumer heterogeneity, consumer preferences, and the strategies with which firms deal with consumer heterogeneity. The rise and fall in the profit share we document is correlated with a rise and then fall in measures of consumer heterogeneity, such as income inequality, the number of different consumer segments, and the variance of the segment-specific taste component. While until 2012, firms were better able to introduce varieties to target different segments, this trend reverses and firms target fewer and fewer segments. Finally, we document a small rise in the fixed marketing costs in the last three years of our sample.

[39] In Figure A.14, we report the 95% confidence intervals for the profit share computed in the three models and the difference across models holding constant the expenditure shares for each product group at its 2004 level. The detailed information for constructing confidence intervals is provided in section A.6.

6.1 Changes in Consumer Heterogeneity

We first examine how changes in the degree of consumer heterogeneity correlate with changes in the profit share. To measure the extent of consumer heterogeneity, we focus on two key measures: income inequality and the number of consumer segments. Intuitively, changes in the level of income inequality are a proxy for the overall degree of consumer heterogeneity. Using the number of segments as a proxy for the degree of consumer heterogeneity requires some additional explanation. Recall that we choose the number of consumer segments to maximize the Calinski-Harabasz pseudo-F statistics. In other words, we adopted an empirical criterion to choose the number of segments in each product group. For this reason, a reduction in the number of segments for a product group implies that we obtain better F statistics by dividing consumers into fewer segments. This means that consumers have become less heterogeneous in their purchases.

Figure 7: Change in Profit Share and Inequality of the United States (2004-2018)

(a) Change in Inequality Measure (b) Profit Share *HetC*: Gini

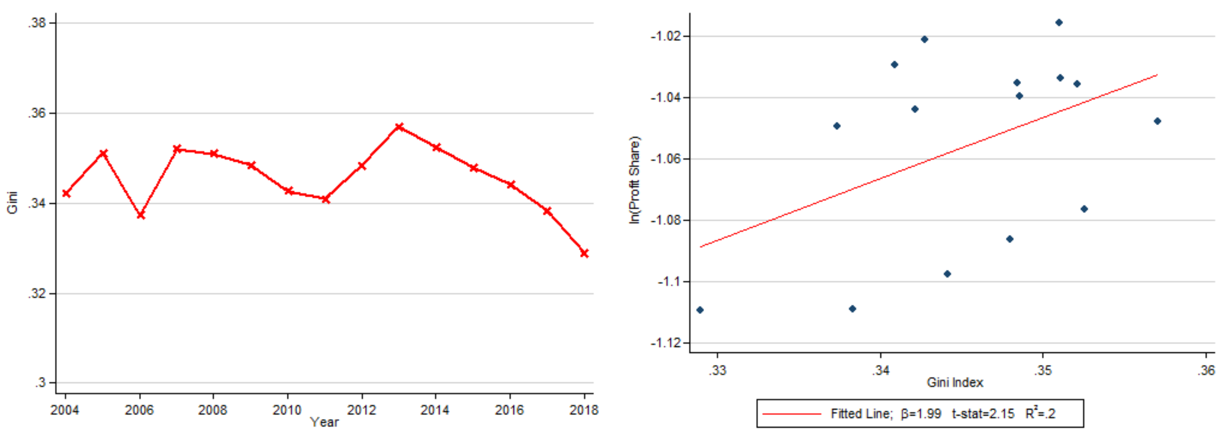

Using the household income information provided by Nielsen Home Scanner Data, we compute income inequality measure (i.e., the Gini index). The change in inequality measures from 2004 to 2018 for the United States based on Nielsen Home Scanner Data are displayed in panel (a) of Figure 7. There was an increase in the first period from 2004 to 2013, and all inequality measures declined since 2013. Although the peak income inequality does not coincide with the peak in the profit share, the broad pattern of the two variables coincide. We report in panel (b) the positive correlation between profit share and various income inequality, and both figures suggest a significantly positive correlation between income inequality measures and the change in the profit share over time.[40]

[40]In Figure A.24, we report the positive correlation between the level of income inequality and the gap

To provide an example of the economic significance of our findings, we extend our results over a more prolonged span from 1990 to 2021. Based on data from Statista, the U.S. Household Income Gini Index rose from 0.43 in 1990 to 0.49 in 2021, which marks an increase of 0.06 over three decades.[41] This slow increase in the Gini index at the national level differs from the pattern in panel (a) of Figure 7, where the Gini index is computed using the Nielsen Home Scanner Data. This is the principal reason we believe that the Nielsen sample is not nationally representative. Nevertheless, we can exploit the positive correlation between the profit share and income inequality to make our prediction about how the profit share increased over 1990-2011. Considering that the semi-elasticity of profit share relative to the Gini index stands at 1.99, as shown in panel (b) of Figure 7, we infer that the relative change in profit share induced by inequality equals $1.99 \times 0.06 = 0.12$, or 13% (since $\exp(0.12) = 1.13$). Considering that the aggregate profit share (with heterogeneous consumers) in the retail sector was about 33 percent at the end of our sample in Figure 5, then we can attribute 13% of that amount – or over 4 percentage points – to the growth in income inequality amongst consumers between 1990 and 2021.

A similar pattern arises when we examine the number of segments that are optimally selected by our clustering algorithm. Figure 8 panel (a) shows the average number of segments per product group over time. There is an increase in the first period and then a decline. This means that consumers have been purchasing more similar sets of products in 2018 than in 2004. In panel (b) of Figure 8, we plot the log change in the *HetC* profit share computed at the product group level between 2018 and 2004 against the log difference in the number of segments within each product group. In panel (c) of Figure 8, we plot the log change in the percentage point difference between *HetC* and *RepC* profit share computed at the product group level between 2018 and 2004 against the log difference in the number of segments within each product group. The growth in the number of segments is associated with a larger increase in the profit share and in the difference between *HetC* and *RepC*.

6.1.1 Evidence from Variance Decomposition

To understand the quantitative contribution of appeal heterogeneity across segments to the UPC-specific demand shifters, we follow the strategy used in the Hottman et al. (2016) and Bernard et al. (2021) and decompose the variance of the UPC demand shifter $\ln z_{ifgn}$ into the common component $\ln \bar{\phi}_{ifg}$ and the segment-specific component $\ln \phi_{ifgn}$. Since the log of the demand shifter $\ln z_{ifgn}$ is the sum of the two components $\ln \bar{\phi}_{ifg} + \ln \phi_{ifgn}$, the variance

between the *HetC* and *RepC* models.

[41] Additional data on the Gini index trends in the U.S. can be sourced from https://www.statista.com/statistics/219643/gini-coefficient-for-us-individuals-families-and-households.

Figure 8: Change in Profit Share and Number Segments

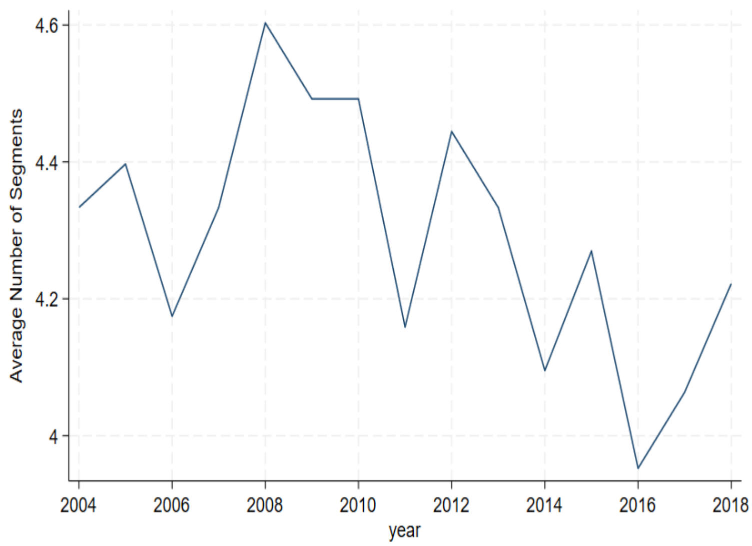

(a) Change in Number of Segments (2004-2018)

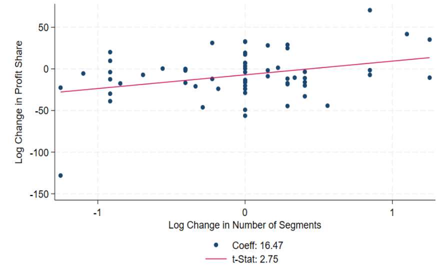

(b) Profit Share *HetC*

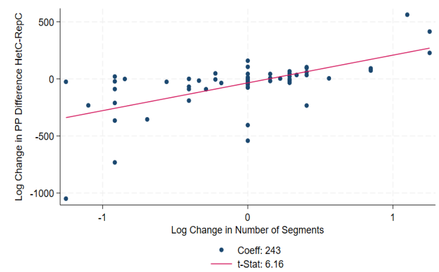

(c) ΔProfit Share

decomposition can be done exploiting the OLS properties and separately regress $\ln \bar{\phi}_{ifg}$ and $\ln \phi_{ifgn}$ on the demand shifter $\ln z_{ifgn}$. The coefficients from such regressions represent the contribution of each of the two components in explaining the variance of the demand shifters, and the two coefficients sum to one.

Figure 9: Contribution of Segments $\ln \phi_{ifgn}$ to Explain Product Appeal

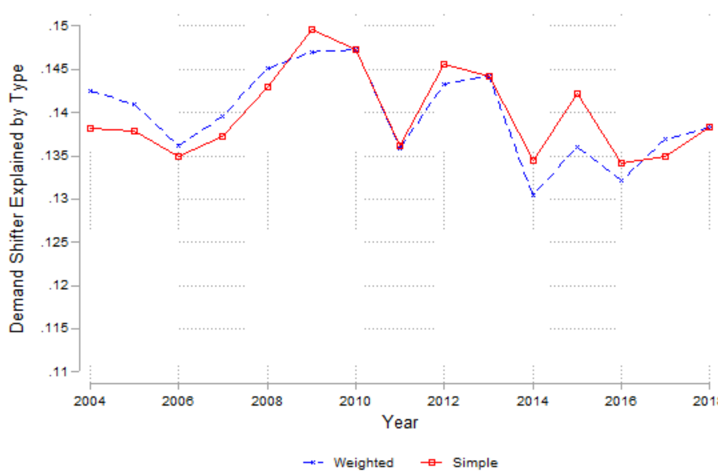

We run the variance decomposition by product group and year and report the average result across product groups in Figure 9, using a simple average or the weighted average, where the weights are the expenditure shares on product groups. The average explanatory power of the segment-specific component $\ln \phi_{ifgn}$ ranges between 13% to 15%.[42] There is significant heterogeneity across product groups: for ice creams, the explanatory power of the segment demand shifter is almost 52%, while for electronics, records, and tapes is around 4% (see Figure A.19).

On average, the explanatory power of the segment-specific component of the demand shifter increases by approximately one percentage point from 2004 to 2010 and then declines to its initial value in 2018. This indicates that the dispersion of demand shifters around the mean across segments has been first increasing and then decreasing. In other words, segments have grown more heterogeneous from 2004 to 2010 and then less heterogeneous in the subsequent years. This pattern is approximately in line with the rise and then decline in the measures of consumer heterogeneity we presented in the previous section.

6.2 Firm Personalization Strategies

In the presence of consumer heterogeneity, firms of different sizes may adopt different strategies in terms of targeting, or "personalization", strategies. There are two key margins to capture. First, how differently do consumer segments perceive the products of a firm? This is the *intensiveness of personalization* measured by how different the demand shifters across segments for a particular UPC are. It reflects the ability of a firm to produce varieties that are perceived differently across consumer segments. Let $\text{var}\phi_{ifg}$ be the variance of the segment component of the demand shifter of a product across consumer segments. We measure the intensiveness of personalization as:

$$IP_{fgt} = \text{median}_i\{\text{var}\phi_{ifg}\},$$

which is the median level of the within-product variance of ϕ_{ifg}. In the extreme case in which all firm products have the same demand shifters across segments as in the *RepC* model, IP_{fg} assumes the minimum value of 0. Larger values of IP_{fg} imply that firm products are perceived differently by different segments.

How are products targeted to different consumer segments? This is the *extensiveness of personalization*, which captures the ability of a firm to design products that appeal to the

[42] In a robustness exercise, in which we consider a division into segments by state, retail chain, and consumer type (see Appendix A.7), we find that the segment-specific demand shifter accounts for approximately 37% of the variance of demand shifters. About 22% of the variance is due to heterogeneity in demand shifters across retail chains, while heterogeneity across states and consumer types accounts for 7-8%. Hence, focusing only on the division across segments by consumer type is likely to yield conservative estimates for the effects of consumer heterogeneity on markups, profits, and other firm performance variables.

taste of different consumer segments. In particular, for each UPC item i, we let $n(i)$ denote the consumer segment that a firm is targeting, by which we mean that $z_{ifn(i)} > z_{ifm}$ for all $m = 1, ..., N$ with $m \neq n(i)$.[43] In other words, in terms of appeal, such a UPC i ranks first in the segment $n(i)$. We measure the extensiveness of personalization as:

$$EP_{fgt} = \frac{\#\text{ of segments with at least one product that ranks first}}{\#\text{ of consumer segments}},$$

where EP_{fg} ranges from $1/N$ to 1. If all firm products rank first in the same segment, the firm is only targeting one segment, and thus $EP_{fg} = 1/N$. If a firm has at least one product that ranks first in all segments, such a firm reaches the maximum level of $EP_{fg} = 1$. Small values of EP_{fg} do not imply that some segments do not purchase any products from firm f, but just that most products' demand shifters are the highest for only a few segments.[44]

In order to study how the two measures have evolved over time, in Figure 10, we plot the average IP_{fg} and EP_{fg} across product groups over time. We consider both a simple average (panel (a)) and a weighted average (panel (b)). While IP_{fg} exhibits an upward trend, EP_{fg} follows the pattern identified for the profit share, with an increase in the initial years, followed by a decline in 2012. The results on the intensiveness of personalization indicate that, on average, products of the same firm have in later years been perceived more differently than in the early part of the sample. In other words, firms have started to offer more niche products than mainstream products. By contrast, the results on the extensiveness of personalization indicate that, on average, firms have been targeting more and more segments up until 2012. In these years, firms were increasingly able to exploit the heterogeneity of consumers by targeting more segments and, thus, limiting the within-segment cannibalization of products. However, EP_{fg} declines after 2012 as firms targeted fewer and fewer segments.

[43] We implicitly assume that the demand shifter of the targeted segment is different from the demand shifter for the same UPC for all other segments. Hence, one UPC can only target one segment only. Given our empirical methodology for measuring the demand shifters (described in section 4), this assumption is always satisfied.

[44] The two concepts of firm personalization are related to firm size. In particular, in Figures A.21 and A.22 of the Appendix, we plot the coefficients obtained from a regression of IP_{fg} and EP_{fg} on firm sales and scope within a product group. In terms of the intensiveness of personalization, we find a statistically insignificant relationship between firm size and IP_{fg} for most product groups. In contrast, we find a positive and statistically significant coefficient for all product groups between the extensiveness of personalization and firm size. The robust pattern indicates that larger firms target more segments than smaller firms. In Appendix A.8, we report the relationship between the intensiveness and extensiveness of personalization and firm using alternative segmentation criteria. We also document a positive relationship between the share of products firms use to target a segment and the segment's size. For details, see Appendix A.8.1.

Figure 10: Firm Personalization Strategies Over Time

(a) Simple Average (b) Weighted Average

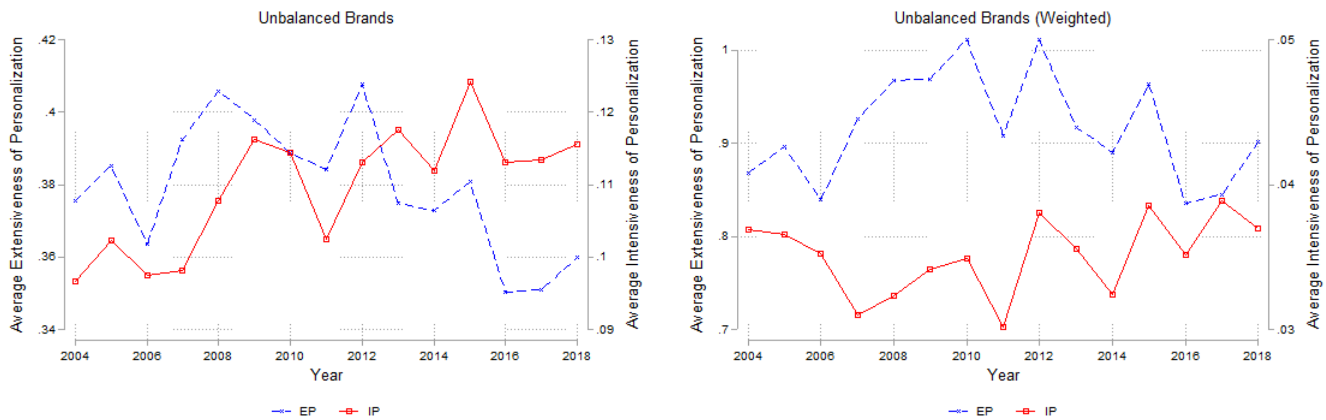

7 Conclusions

This paper studies how the interaction between large firms and consumer heterogeneity can contribute to the rising profit share. We build a model with oligopolistic multi-product firms, which sell to segments that exhibit different demand shifters. Our model predicts that the presence of consumer heterogeneity increases firm-level markups and profits.

We apply our model to the Nielsen dataset on household purchases at the barcode level and estimate product-segment-specific demand shifters for a large number of product groups. This difference in these demand shifters across consumer segments leads oligopolistic firms to charge higher prices than would occur with a representative consumer. We find that the presence of taste heterogeneity across consumer types implies that markups are 20 percentage points larger on average than with a representative consumer, but even higher for the largest firms in each product group. The aggregate profit share is 6.4 percentage points higher with heterogeneous consumers. Furthermore, as consumer heterogeneity expands – as measured by the Gini index of income inequality – then the aggregate profit share expands. Using the impact of income inequality on the log profit share over 2004-2018 and applying this to the US national change in income inequality over 1990-2021, we predict that increased inequality can account for an increase in the profit share of over 4 percent.

We also consider other measures of consumer heterogeneity, such as the number of consumer segments generated by our algorithm, and the contribution of segment-specific product appeal in explaining the total variation in product appeal. Each of these measures is positively associated with the profit share. We further argue that changes in segment-specific demand shifters can be driven by firms' strategies in targeting, or personalizing, their products to a given segment. A product targets a segment if it has the largest demand shifter in that segment relative to all other segments. We examine how intensively firms personalize their products for the taste of given segments, i.e., how differently are the products of a firm perceived by different segments. The ability of a firm to target segments also has an extensive component, i.e., the number of segments targeted. We find that the first measure increases over time, and the second increases in the early part of our sample, when the profit share also rises. In sum, our results point to a stronger role for consumer heterogeneity in explaining changes in the share of profits for the United States than previously found.

References

R. Adão, C. Arkolakis, and S. Ganapati. Aggregate Implications of Firm Heterogeneity: A Nonparametric Analysis of Monopolistic Competition Trade Models. *NBER Working Paper 28081*, 2020.

C. Arkolakis, A. Costinot, D. Donaldson, and A. Rodríguez-Clare. The Elusive Pro-Competitive Effects of Trade. *The Review of Economic Studies*, 86(1):46–80, 2019.

C. Arkolakis, S. Ganapati, and M.-A. Muendler. The Extensive Margin of Exporting Products: A Firm-level Analysis. *forthcoming in American Economic Journal: Macroeconomics*, 2021.

A. Atkeson and A. Burstein. Pricing-to-Market, Trade Costs, and International Relative Prices. *American Economic Review*, 98(5):1998–2031, 2008.

D. Atkin. Trade, Tastes, and Nutrition in India. *American Economic Review*, 103(5):1629–63, 2013.

D. Autor, D. Dorn, L. F. Katz, C. Patterson, and J. Van Reenen. The Fall of the Labor Share and the Rise of Superstar Firms. *The Quarterly Journal of Economics*, 135(2): 645–709, 2020.

B. Y. Aw-Roberts, Y. Lee, and H. Vandenbussche. Decomposing Firm-Product Appeal: How important is Consumer Taste? CEPR Discussion Papers 12707, C.E.P.R. Discussion Papers, 2018.

R. Baldwin and J. Harrigan. Zeros, Quality, and Space: Trade Theory and Trade Evidence. *American Economic Journal: Microeconomics*, 3(2):60–88, 2011.

S. Barkai. Declining Labor and Capital Shares. *The Journal of Finance*, 2020.

A. B. Bernard, E. Dhyne, G. Magerman, K. Manova, and A. Moxnes. The Origins of Firm Heterogeneity: A Production Network Approach. *forthcoming in Journal of Political Economy*, 2021.

B. S. Blum, S. Claro, I. Horstmann, and D. A. Rivers. The ABCs of Firm Heterogeneity: The Effects of Demand and Cost Differences on Exporting. Technical report, Working paper, University of Toronto, Toronto, 2018.

J. Brand. Differences in Differentiation: Rising Variety and Markups in Retail Food Stores. *University of Texas at Austin Working Paper*, 2021.

C. Broda and D. Weinstein. Globalization and the gains from variety. *The Quarterly Journal of Economics*, 121(2):541–585, 2006.

J. A. Caballero. Do Surges in International Capital Inflows Influence the Likelihood of Banking Crises? *The Economic Journal*, 126(591):281–316, 2014.

A. Cavallo, B. Neiman, and R. Rigobon. Currency unions, product introductions, and the real exchange rate. *The Quarterly Journal of Economics*, 129(2):529–595, 2014.

T. Chaney. The Intensive and Extensive Margins of International Trade. *American Economic Review*, 98(4):1707–1721, 2008.

Y.-L. Chyi, Y.-M. Lai, and W.-H. Liu. Knowledge Spillovers and Firm Performance in the High-technology Industrial Cluster. *Research Policy*, 41(3):556–564, 2012.

A. K. Coşar, P. L. Grieco, S. Li, and F. Tintelnot. What Drives Home Market Advantage? *Journal of International Economics*, 110(Supplement C):135 – 150, 2018.

T. M. Crone and A. Clayton-Matthews. Consistent Economic Indexes for the 50 States. *Review of Economics and Statistics*, 87(4):593–603, 2005.

M. Crozet, K. Head, and T. Mayer. Quality Sorting and Trade: Firm-level Evidence for French Wine. *The Review of Economic Studies*, 79(2):609–644, 2012.

S. DellaVigna and M. Gentzkow. Uniform Pricing in U.S. Retail Chains*. *The Quarterly Journal of Economics*, 134(4):2011–2084, 2019.

F. Di Comite, J.-F. Thisse, and H. Vandenbussche. Verti-zontal differentiation in export markets. *Journal of International Economics*, 93(1):50–66, 2014.

M. Draganska, K. Seim, and M. Mazzeo. Beyond Plain Vanilla: Modeling Joint Product Assortment and Pricing Decisions. *Quantitative Marketing and Economics*, 7, 11 2007.

A. Drewnowski, S. A. Henderson, and A. B. Shore. Taste responses to naringin, a flavonoid, and the acceptance of grapefruit juice are related to genetic sensitivity to 6-n-propylthiouracil. *The American Journal of Clinical Nutrition*, 66(2):391–397, 1997.

C. Eckel and J. P. Neary. Multi-Product Firms and Flexible Manufacturing in the Global Economy. *Review of Economic Studies*, 77(1):188–217, 2010.

M. Elsby, B. Hobijn, and A. Sahin. The Decline of the U.S. Labor Share. *Brookings Papers on Economic Activity*, 44(2 (Fall)):1–63, 2013.

B. Everitt. *Cluster Analysis*. London: Edward Arnold, 1993.

B. Faber and T. Fally. Firm Heterogeneity in Consumption Baskets: Evidence from Home and Store Scanner Data. *The Review of Economic Studies*, 89(3):1420–1459, 2021.

P. Fajgelbaum, G. M. Grossman, and E. Helpman. Income Distribution, Product Quality, and International Trade. *Journal of Political Economy*, 119(4):721 – 765, 2011.

R. Feenstra. New product varieties and the measurement of international prices. *American Economic Review*, 84(1):157–77, 1994.

R. Feenstra. *Advanced International Trade*. Princeton: Princeton University Press, 2004.

R. C. Feenstra. Restoring the Product Variety and Pro-competitive Gains from Trade with Heterogeneous Firms and Bounded Productivity. *Journal of International Economics*, 110:16–27, 2018.

R. C. Feenstra and J. Romalis. International Prices and Endogenous Quality. *The Quarterly Journal of Economics*, 129(2):477–527, 2014.

R. C. Feenstra, M. Xu, and A. Antoniades. What is the Price of Tea in China? Towards the Relative Cost of Living in Chinese and US Cities. *The Economic Journal*, 2020.

A. M. Fernandes, P. J. Klenow, S. Meleshchuk, M. D. Pierola, and A. Rodríguez-Clare. The Intensive Margin in Trade: How Big and How Important. *NBER Working Paper 25195*, 2018.

L. Foster, J. Haltiwanger, and C. Syverson. Reallocation, Firm Turnover, and Efficiency: Selection on Productivity or Profitability? *American Economic Review*, 98(1):394–425, 2008.

A. R. Ghosh, M. S. Qureshi, J. I. Kim, and J. Zalduendo. Surges. *Journal of International Economics*, 92(2):266–285, 2014.

X. Guan, M. Mantrala, and Y. Bian. Strategic Information Management in A Distribution Channel. *Journal of Retailing*, 95(1):42–56, 2019.

J. C. Hallak and J. Sivadasan. Product and Process Productivity: Implications for Quality Choice and Conditional Exporter Premia. *Journal of International Economics*, 91(1): 53–67, 2013.

C. Hottman, S. J. Redding, and D. E. Weinstein. Quantifying the Sources of Firm Heterogeneity. *The Quarterly Journal of Economics*, 131:1291—1364, 2016.

C.-T. Hsieh and E. Rossi-Hansberg. The Industrial Revolution in Services. *NBER Working Paper 25968*, 2021.

I. C. Jäkel. Product appeal, Differences in Tastes, and Export Performance: Evidence for Danish Chocolate and Confectionery. *International Journal of Industrial Organization*, 63:417 – 459, 2019.

R. C. Johnson. Trade and Prices with Heterogeneous Firms. *Journal of International Economics*, 86(1):43–56, 2012.

L. Karabarbounis and B. Neiman. The Global Decline of the Labor Share. *The Quarterly Journal of Economics*, 129(1):61–103, 2014.

A. Khandelwal. The Long and Short (of) Quality Ladders. *The Review of Economic Studies*, 77(4):1450–1476, 2010.

D. Koh, R. Santaeulàlia-Llopis, and Y. Zheng. Labor Share Decline and Intellectual Property Products Capital. *Econometrica*, 88(6):2609–2628, 2020.

M. Kugler and E. Verhoogen. Prices, Plant Size, and Product Quality. *The Review of Economic Studies*, 79(1):307–339, 2012.

A. Leicester. The potential use of in-home scanner technology for budget surveys. *NBER Working Paper*, (w19536), 2013.

E. Levy-Yeyati and F. Sturzenegger. Classifying Exchange Rate Regimes: Deeds vs. Words. *European Economic Review*, 49(6):1603–1635, 2005.

L. Macedoni and M. Xu. Flexibility and Productivity: Towards the Understanding for Multi-product Exporters. *Aarhus University Working Paper*, 2018.

B. Neiman and J. Vavra. The rise of niche consumption. *American Economic Journal: Macroeconomics*, 15(3):224–64, 2023.

A. Nevo. Empirical Models of Consumer Behavior. *Annual Review of Economics*, 3(1):51–75, 2011.

M. J. Roberts, D. Yi Xu, X. Fan, and S. Zhang. The Role of Firm Factors in Demand, Cost, and Export Market Selection for Chinese Footwear Producers. *The Review of Economic Studies*, 85(4):2429–2461, 2018.

M. Rognlie. Deciphering the Fall and Rise in the Net Capital Share: Accumulation or Scarcity? *Brookings Papers on Economic Activity*, 2015(1):1–69, 2016.

I. Simonovska. Income Differences and Prices of Tradables: Insights from an Online Retailer. *The Review of Economic Studies*, 82(4):1612–1656, 2015.

M. Smith, D. Yagan, O. Zidar, and E. Zwick. Capitalists in the Twenty-first Century. *The Quarterly Journal of Economics*, 134(4):1675–1745, 2019.

C. Zhen, J. L. Taylor, M. K. Muth, and E. Leibtag. Understanding differences in self-reported expenditures between household scanner data and diary survey data: a comparison of homescan and consumer expenditure survey. *Applied Economic Perspectives and Policy*, 31(3):470–492, 2009.

A Appendix

A.1 Model Derivations

A.1.1 Complements and Substitutes.

Consider the demand for a product i by firm f in a segment n and product group g. For ease of notation, let us focus on a single product group and a single segment, so that the only subscripts left are product i and firm f:

$$q_{if} = zwP^{\sigma-1}z_f^{\sigma-1}P_f^{\eta-\sigma}z_{if}^{\eta-1}p_{if}^{-\eta}, \tag{20}$$

Let us now consider the cross-price elasticity of demand for product i with respect to the price of product j produced by the same firm:

$$\begin{aligned}
\frac{d\ln q_{if}}{d\ln p_{jf}} &= (\sigma-1)\frac{d\ln P}{d\ln P_f}\frac{d\ln P_f}{d\ln p_{jf}} + (\eta-\sigma)\frac{d\ln P_f}{d\ln p_{jf}} = \\
&= \frac{d\ln P_f}{d\ln p_{jf}}\left((\sigma-1)\frac{d\ln P}{d\ln P_f} + \eta - \sigma\right) = \\
&= s_{jf}\left((\sigma-1)s_f + \eta - \sigma\right)
\end{aligned}$$

The elasticity is positive and the two products are substitutes if:

$$\eta > \sigma - (\sigma-1)s_f \tag{21}$$

If $\eta > \sigma$ then the products within all firms are positive. If $\eta < \sigma$, then whether or not products within a firm are complements or substitutes depends on the market share of the firm. If $s_f > \frac{\sigma-\eta}{\sigma-1}$, then products are substitutes and, vice versa, products are complements if $s_f > \frac{\sigma-\eta}{\sigma-1}$.

A.1.2 Profits

Revenues for a firm f from UPC i in segment n are given by:

$$L_n p_{ifn} q_{ifn} = z_g L_n w_n P_n^{\sigma-1} z_f^{\sigma-1} P_{fn}^{\eta-\sigma} z_{ifn}^{\eta-1} p_{if}^{1-\eta} = z_g L_n w_n \left(\frac{\frac{P_{fn}}{z_f}}{P_n}\right)^{1-\sigma}\left(\frac{\frac{p_{if}}{z_{ifn}}}{P_{fn}}\right)^{1-\eta},$$

where z_g denotes the demand shifter for good g, but for convenience, we omit the subscript g on all other variables and parameters. Summing over the varieties sold to segment n yields the following segment-specific revenues:

$$\sum_{i\in I_f} L_n p_{ifn} q_{ifn} = z_g L_n w_n \left(\frac{\frac{P_{fn}}{z_f}}{P_n}\right)^{1-\sigma} \sum_{i\in I_f}\left(\frac{\frac{p_{if}}{z_{ifn}}}{P_{fn}}\right)^{1-\eta} = z_g L_n w_n \left(\frac{\frac{P_{fn}}{z_f}}{P_n}\right)^{1-\sigma}.$$

Let $Lw^N L_n w_n$ denote total income of all consumers, and $\alpha_n = L_n w_n / \sum_{n=1}^N L_n w_n$ denote the share of income for consumers in segment n. Firm f total revenues R_f are given by:

$$R_f = \sum_{n=1}^N \sum_{i \in I_f} L_n p_{ifn} q_{ifn} = z_g L w \sum_{n=1}^N \alpha_n z_f^{\sigma-1} \left(\frac{P_{fn}}{P_n} \right)^{1-\sigma}. \tag{22}$$

We also keep track of the total variable costs of a firm, which are:

$$TC_f = \sum_{n=1}^N \sum_{i \in I_f} L_n c_{if} q_{ifn} = z_g z_f^{\sigma-1} L w \sum_{n=1}^N \alpha_n P_n^{\sigma-1} P_{fn}^{\eta-\sigma} \sum_{i \in I_f} z_{ifn}^{\eta-1} c_{if} p_{if}^{-\eta}.$$

A.1.3 First Order Condition

To simplify the derivations of the first-order condition, let us consider the derivative of revenues with respect to prices p_{if}. In Bertrand competition, firms take into account the effects of changing prices both on the firm and segment-specific price index P_{fn}, and on the product group and segment-specific price index P_n:

$$\frac{\partial R_f}{\partial p_{if}} = z_g L w \sum_{n=1}^N \alpha_n z_f^{\sigma-1} \left(\frac{P_{fn}}{P_n} \right)^{1-\sigma} (1-\sigma) \left[\frac{\partial P_{fn}}{\partial p_{if}} \frac{1}{P_{fn}} - \frac{\partial P_n}{\partial P_{fn}} \frac{P_{fn}}{P_n} \frac{\partial P_{fn}}{\partial p_{if}} \frac{1}{P_{fn}} \right].$$

Using the definitions of market shares in (7) and (8), we obtain:

$$\frac{\partial R_f}{\partial p_{if}} = z_g L w \sum_{n=1}^N \alpha_n s_{fn} (1-\sigma) \left[\frac{s_{ifn}}{p_{if}} - \frac{s_{ifn}}{p_{if}} s_{fn} \right] = (1-\sigma) \frac{z_g L w}{p_{if}} \sum_{n=1}^N \alpha_n s_{fn} s_{ifn} (1 - s_{fn}).$$

Taking the derivative of total costs with respect to p_{if} yields:

$$\begin{aligned}
\frac{\partial TC_f}{\partial p_{if}} &= z_g z_f^{\sigma-1} L w \sum_{n=1}^N \alpha_n P_n^{\sigma-1} P_{fn}^{\eta-\sigma} \left[(\sigma-1) \frac{s_{ifn}}{p_{if}} s_{fn} + (\eta-\sigma) \frac{s_{ifn}}{p_{if}} \right] \sum_{i \in I_f} z_{ifn}^{\eta-1} c_{if} p_{if}^{-\eta} \\
&+ z_g z_f^{\sigma-1} L w \sum_{n=1}^N \alpha_n P_n^{\sigma-1} P_{fn}^{\eta-\sigma} z_{ifn}^{\eta-1} (-\eta) c_{if} p_{if}^{-\eta-1} \\
&= z_g L w \sum_{n=1}^N \alpha_n \left(\frac{\frac{P_{fn}}{z_f}}{P_n} \right)^{1-\sigma} P_{fn}^{\eta-1} \left[(\sigma-1) \frac{s_{ifn}}{p_{if}} s_{fn} + (\eta-\sigma) \frac{s_{ifn}}{p_{if}} \right] \sum_{i \in I_f} z_{ifn}^{\eta-1} c_{if} p_{if}^{-\eta} \\
&+ z_g L w \sum_{n=1}^N \alpha_n \left(\frac{\frac{P_{fn}}{z_f}}{P_n} \right)^{1-\sigma} P_{fn}^{\eta-1} z_{ifn}^{\eta-1} (-\eta) c_{if} p_{if}^{-\eta-1}.
\end{aligned}$$

Using the definitions of market shares, we obtain:

$$\frac{\partial TC_f}{\partial p_{if}} = \frac{z_g L w}{p_{if}} \sum_{n=1}^{N} \alpha_n s_{fn} \left[(\sigma - 1) s_{ifn} s_{fn} + (\eta - \sigma) s_{ifn} \right] \frac{\sum_{i \in I_f} z_{ifn}^{\eta-1} c_{if} p_{if}^{-\eta}}{\sum_{i \in I_f} z_{ifn}^{\eta-1} p_{if}^{1-\eta}} +$$
$$- \eta \frac{z_g L w}{p_{if}} \sum_{n=1}^{N} \alpha_n s_{fn} \frac{s_{ifn}}{p_{if}} c_{if}.$$

Using the definition of market shares of UPC within a firm, we can rewrite the term above as:

$$\frac{\sum_{i \in I_f} z_{ifn}^{\eta-1} c_{if} p_{if}^{-\eta}}{\sum_{i \in I_f} z_{ifn}^{\eta-1} p_{if}^{1-\eta}} = \frac{\sum_{i \in I_f} s_{ifn} \frac{c_{if}}{p_{if}}}{\sum_{i \in I_f} s_{ifn}} = \sum_{i \in I_f} s_{ifn} \frac{c_{if}}{p_{if}}.$$

It follows that the derivative of the firm total costs is equal to:

$$\frac{\partial TC_f}{\partial p_{if}} = \frac{z_g L w}{p_{if}} \left[\sum_{n=1}^{N} \alpha_n s_{fn} s_{ifn} \left[(\sigma - 1) s_{fn} + (\eta - \sigma) \right] \sum_{i \in I_f} s_{ifn} \frac{c_{if}}{p_{if}} - \eta \sum_{n=1}^{N} \alpha_n s_{fn} s_{ifn} \frac{c_{if}}{p_{if}} \right].$$

We, therefore, obtain the following first-order condition:

$$\frac{\partial R_f}{\partial p_{if}} = \frac{\partial TC_f}{\partial p_{if}} \implies$$
$$(\sigma - 1) \sum_{n=1}^{N} \alpha_n s_{fn} s_{ifn} (1 - s_{fn}) = \eta \sum_{n=1}^{N} \alpha_n s_{fn} s_{ifn} \frac{c_{if}}{p_{if}}$$
$$- \sum_{n=1}^{N} \alpha_n s_{fn} s_{ifn} \left[(\sigma - 1) s_{fn} + (\eta - \sigma) \right] \sum_{i \in I_f} s_{ifn} \frac{c_{if}}{p_{if}}.$$

It is convenient to rewrite the equation as a function of markups $\mu_{if} = p_{if}/c_{if}$:

$$(\sigma - 1) \sum_{n=1}^{N} \alpha_n s_{fn} s_{ifn} (1 - s_{fn}) = \frac{\eta}{\mu_{if}} \sum_{n=1}^{N} \alpha_n s_{fn} s_{ifn} - \sum_{n=1}^{N} \alpha_n s_{fn} s_{ifn} \left[(\sigma - 1) s_{fn} + (\eta - \sigma) \right] \sum_{i \in I_f} \frac{s_{ifn}}{\mu_{if}}$$

Notice that $\sum_{n=1}^{N} \alpha_n s_{fn} s_{ifn} = s_{if}$ equals the market share of UPC i of firm f within the product group. We divide both sides of the equation by $\sum_{n=1}^{N} \alpha_n s_{fn} s_{ifn}$, to obtain the expression (10) shown in the main text.

A.1.4 Proof of Proposition 1

Proposition 1 is stated for all product groups $g = 1, ..., G$, but in this proof we drop the subscript g. Jensen's inequality states that $\sum_{n=1}^{N} \alpha_n s_{fn}^2 \geq \left(\sum_{n=1}^{N} \alpha_n s_{fn} \right)^2$, and this inequality holds as equality if and only if $s_{fn} = s_{fm} \;\forall\; m, n$. So the inequality is strict if $s_{fn} \neq s_{fm}$ for some firm f and segments $m \neq n$. Let us consider the ratio s_{fn}/s_{fm} from (5) and (7), using the fact that the firm shifter z_f (appearing as z_{fg} in (5)) cancels because s_{fn} and s_{fm} are for the same firm f:

$$\frac{s_{fn}}{s_{fm}} = \left(\frac{P_{fn}/P_n}{P_{fm}/P_m}\right)^{1-\sigma} = \left(\frac{z_{ifm}}{z_{ifn}}\frac{P_m}{P_n}\right)^{1-\sigma}, \tag{23}$$

where the second equality follows from (6) and using the assumption of a single product firm in the *HetC* model, so that $I_{fn} = \{i\}$. We need to show that $s_{fn} \neq s_{fm}$ for two consumer segments m and n and some firm f to ensure that Jensen's inequality holds strictly.

Suppose to the contrary that $s_{fn} = s_{fm}$ in (23), which implies that $P_m/P_n = z_{ifn}/z_{ifm}$. Then we make use $z_{jf'm}/z_{jf'n} \neq z_{ifm}/z_{ifn}$ from Definition 1, to conclude that

$$1 \neq \left(\frac{z_{jf'm}}{z_{jf'n}}\frac{z_{ifn}}{z_{ifm}}\right)^{1-\sigma} = \left(\frac{z_{jf'm}}{z_{jf'n}}\frac{P_m}{P_n}\right)^{1-\sigma} = \frac{s_{f'n}}{s_{f'm}},$$

where the second equality follows from our assumption that $s_{fn} = s_{fm}$, which implies that $P_m/P_n = z_{ifn}/z_{ifm}$, and the third equality follows by computing $s_{f'n}/s_{f'm}$ just like in (23). So we find that $s_{f'n} \neq s_{f'm}$, which guarantees that Jensen's inequality holds strictly.

A.1.5 Demand Elasticities and Markups

Figure A.1 plots the percentage point difference in markups between the *HetC* model and the *RepC* model. We consider a numerical exercise with four segments with the following expenditure shares $\alpha = [0.5, 0.2, 0.2, 0.1]$ and a firm with the following firm-segment market share across segments: $s_{fn} = [0.4, 0.05, 0.1, 0.2]$. We fixed these values to minimize the number of simulations needed to in the exercise. The firm sells five products, whose market shares in each segment are drawn from a uniform distribution. In particular, for each product we draw a value from the uniform distribution $\tilde{u}_{ifn}$ and normalize it so that the draws yield the market share of the items within a firm-segment ($s_{ifn} = \tilde{u}_{ifn}/\sum_i \tilde{u}_{ifn}$). For each value of σ and η shown in the figure, we draw the market share of the firm's products 1000 times and compute the percentage point difference between μ_{HetC} and μ_{RepC}. Across realizations, we compute the average difference and the 5th and 95th percentiles, which we plot in the figures. Higher demand elasticities are associated with lower markups and the relationship is steeper for the elasticity across-firm (σ) than within-firm (η).

Figure A.1: Markups and Demand Elasticities

(a) Within-Firm Elasticity η

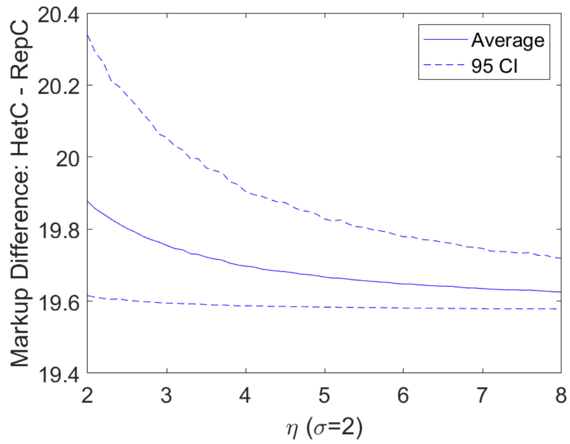

(b) Within-Firm Elasticity η

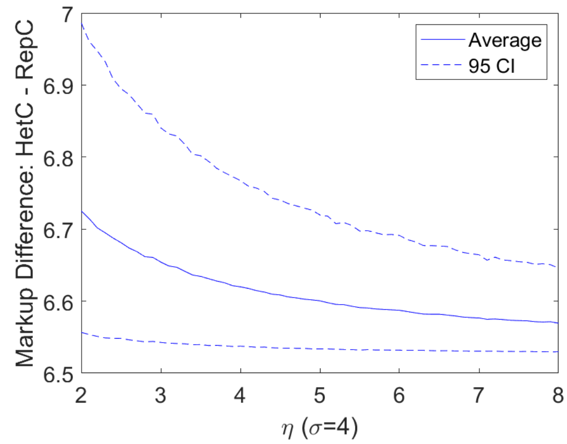

(c) Across-Firm Elasticity σ

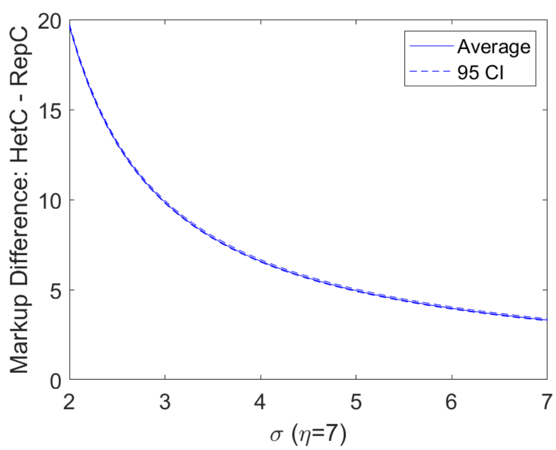

(d) Across-Firm Elasticity σ (Zoom)

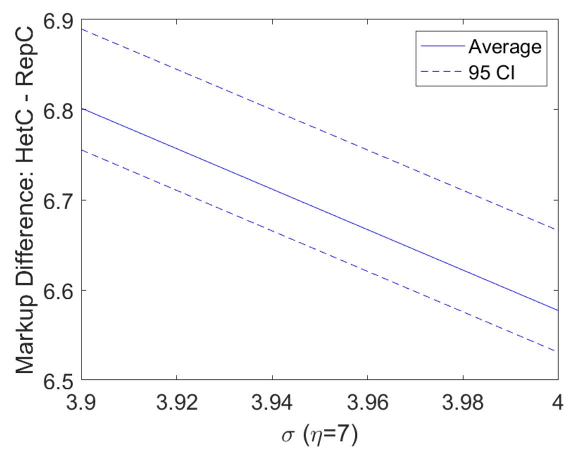

(e) Across-Firm Elasticity σ

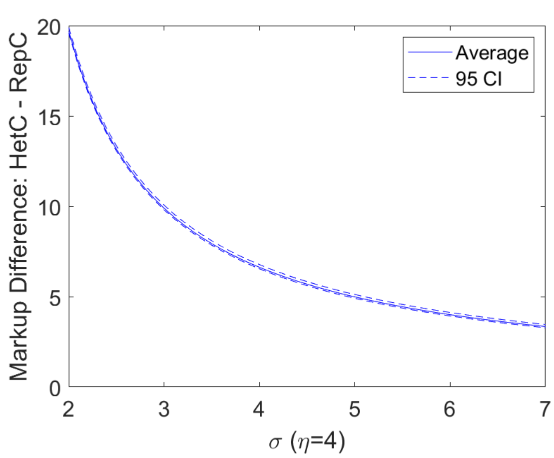

(f) Across-Firm Elasticity σ (Zoom)

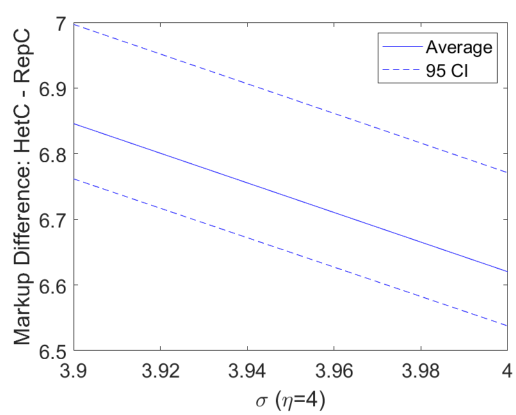

A.1.6 Cannibalization Rate

In a slight abuse of notation (because we omit the subscript f), let $\bar{I} \equiv \left|I_f\right|$ denote the number of items in the set I_f, and suppose that we rank these items by decreasing demand. Following Hottman et al. (2016), we can compute the cannibalization rate defined as the elasticity of demand for the last UPC, $q_{\bar{I}fn}$, with respect to the number of UPCs $\bar{I}$. Such an elasticity can be derived analytically, assuming that the number of products is large enough to be approximated by a continuum, so that:

$$\frac{\partial P_{fn}}{\partial \bar{I}} \approx \frac{P_{fn} s_{\bar{I}fn}}{1-\eta}.$$

As a result, we can compute the cannibalization rate as follows:

$$\begin{aligned}
\frac{\partial q_{\bar{I}fn}}{\partial \bar{I}} &= (\sigma-1)\frac{q_{\bar{I}fn}}{P_n}\frac{\partial P_n}{\partial P_{fn}}\frac{\partial P_{fn}}{\partial \bar{I}} + (\eta-\sigma)\frac{q_{\bar{I}fn}}{P_n}\frac{\partial P_{fn}}{\partial \bar{I}} \\
&= (\sigma-1)\frac{q_{\bar{I}fn}}{P_f n}\frac{\partial P_n}{\partial P_{fn}}\frac{P_{fn} s_{\bar{I}fn}}{1-\eta} + (\eta-\sigma)\frac{q_{\bar{I}fn}}{P_n}\frac{P_{fn} s_{\bar{I}fn}}{1-\eta} \\
&= (\sigma-1)\frac{q_{\bar{I}fn} s_{fn} s_{\bar{I}fn}}{1-\eta} + (\eta-\sigma)\frac{q_{\bar{I}fn} s_{\bar{I}fn}}{1-\eta},
\end{aligned}$$

with which we can obtain the cannibalization rate by segment:

$$\frac{\partial q_{\bar{I}fn}}{\partial \bar{I}}\frac{\bar{I}}{q_{\bar{I}fn}} = -\left(\frac{\sigma-1}{\eta-1} s_{fn} + \frac{\sigma-\eta}{\eta-1}\right) s_{\bar{I}fn}\bar{I}. \tag{24}$$

Thus, expanding the scope reduces the demand for an individual UPC from a segment and such a reduction is proportional to the market share of the firm in the niche (s_{fn}), the market share of the UPC in the segment-firm ($s_{\bar{I}fn}$), and the number of UPCs ($\bar{I}$). The difference in demand elasticities within and across firms generates part of the cannibalization effects: when a firm introduces a new variety, the new variety cannibalizes the sales of the firm's existing varieties *and* the sales of other firms. The magnitude of the own-cannibalization effects depends on the firm market size and on the substitutability within-firm relative to across-firm. Following Hottman et al. (2016), let us define the cannibalization rate as (24) evaluated at a particular product that has a market share in the segment equal to the average market share, so that $s_{\bar{I}fn}\bar{I} = 1$ and $s_{\bar{I}fn} = 1/\bar{I}$.

To account for the fact that introducing a new UPC cannibalizes demand across other segments, we compute an average cannibalization rate across segments where we use segment size (α_n) as weights:

$$\sum_{n=1}^{N} \alpha_n \frac{\partial q_{\bar{I}fn}}{\partial \bar{I}}\frac{\bar{I}}{q_{\bar{I}fn}} = -\left(\frac{\sigma-1}{\eta-1}\sum_{n=1}^{N}\alpha_n s_{fn} + \frac{\sigma-\eta}{\eta-1}\right), \tag{25}$$

which has the convenient property of being identical to the formula in Hottman et al. (2016), since $\sum_{n=1}^{N}\alpha_n s_{fn}$ is the market share of the firm. Although the average cannibalization rate is independent of consumer heterogeneity, the rate is heterogeneous across segments, which could be a crucial factor for firms deciding to introduce new products.

We compute the cannibalization rate by segment and the average cannibalization rate for firms defined in equations (24) and (25). Then, for each firm, we compute the difference between the maximum cannibalization rate and its minimum, across segments. In Figure A.2, we report the average difference by product group and the difference for the top one and five firms. For instance, in the oral hygiene case, the difference between the maximum and minimum cannibalization rate is 0.02, which suggests minimal differences across segments. However, when we consider the top five firms and the top one firm, the rate jumps to 0.2 and 0.24. As shown in the picture, considering the largest firms, the heterogeneity across segments for cannibalization rate is sizable

Figure A.2: Max-Min Difference of Cannibalization Rates across Segments

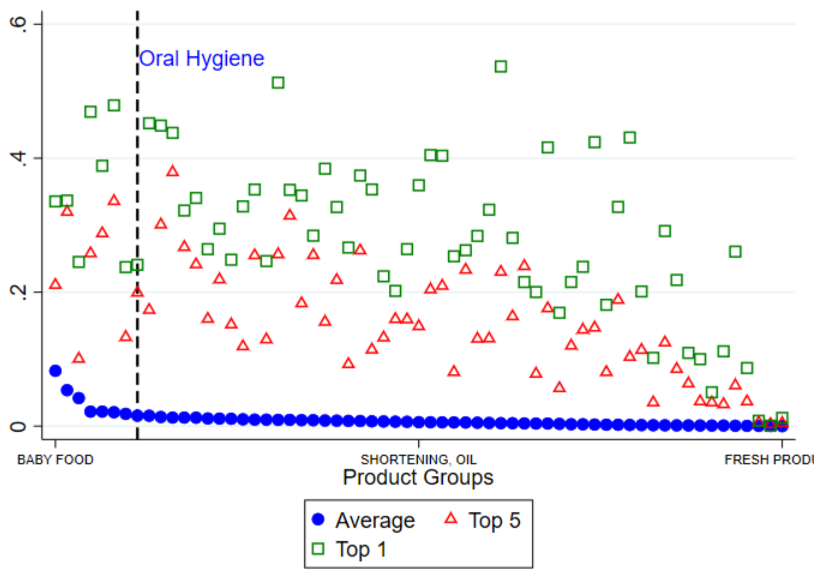

A.2 Consumer Segments

Table A.1 reports the number of consumer types by product group using our clustering methodology of section 4.1 for 2012. Table A.2 reports the percentage point difference in markups between the *HetC* and the *RepC* models across product groups.

Table A.1: Number of Customer Types by Product Group (2012)

	Number of Consumer Types
Baby Food	4
Baking Supplies	4
Batteries And Flashlights	3
Beer	3
Bread And Baked Goods	6
Breakfast Food	2
Butter And Margarine	4
Candy	3
Cereal	4
Cheese	2
Coffee	6
Condiments, Gravies, And Sauces	2
Cookies	3
Cosmetics	7
Cot Cheese, Sour Cream, Toppings	4
Cough And Cold Remedies	7
Crackers	5
Deodorant	7
Detergents	2
Disposable Diapers	3
Dressings/Salads/Prep Foods-Deli	4
Eggs	6
First Aid	7
Fresh Meat	6
Fresh Produce	2
Hair Care	3
Household Cleaners	5
Household Supplies	5
Ice Cream, Novelties	7
Jams, Jellies, Spreads	3
Juice, Drinks - Canned, Bottled	7
Kitchen Gadgets	3
Laundry Supplies	7
Liquor	4
Medications/Remedies/Health Aids	2
Milk	6
Nuts	7
Oral Hygiene	7
Packaged Meats-Deli	4
Packaged Milk And Modifiers	2
Paper Products	6
Pet Care	6
Pet Food	5
Pizza/Snacks/Hors Doeurves-Frzn	7
Prepared Food-Dry Mixes	7
Prepared Food-Ready-To-Serve	4
Prepared Foods-Frozen	4
Salad Dressings, Mayo, Toppings	2
Shaving Needs	2
Shortening, Oil	4
Snacks	2
Soft Drinks-Non-Carbonated	7
Soup	3
Spices, Seasoning, Extracts	2
Tea	7
Tobacco & Accessories	7
Unprep Meat/Poultry/Seafood-Frzn	2
Vegetables - Canned	5
Vegetables-Frozen	7
Vitamins	2
Wine	2
Wrapping Materials And Bags	6
Yogurt	3

Table A.2: Percentage Point Difference between *HetC* and *RepC* markup formula

	Average	Top 5 Firms	Top 1 Firm
Eggs	128.5	223.7	230.7
Butter And Margarine	83.1	97.8	165.6
Fresh Meat	65.8	114.5	259.2
Tobacco & Accessories	64.2	84.8	119.3
Vegetables-Frozen	56.4	75.9	106.7
Shaving Needs	53.6	56.6	89.1
Cereal	52.9	55.1	74.8
Disposable Diapers	49.8	49.9	55.0
Deodorant	47.3	50.5	53.5
Soup	42.4	48.4	62.5
Milk	42.2	78.4	131.0
Beer	40.6	52.4	63.7
Crackers	39.4	46.3	45.5
Coffee	37.3	50.8	76.2
Pizza/Snacks/Hors Doeurves-Frzn	35.0	43.4	46.9
Yogurt	34.1	38.0	57.0
Jams, Jellies, Spreads	33.4	42.1	59.0
Pet Food	31.9	38.1	62.4
Packaged Milk And Modifiers	27.3	33.0	49.0
Wrapping Materials And Bags	25.4	28.4	26.1
Paper Products	25.2	28.6	35.6
Oral Hygiene	25.2	34.0	37.5
Detergents	23.5	26.1	44.9
Cot Cheese, Sour Cream, Toppings	22.0	38.9	76.0
First Aid	19.7	33.3	46.6
Juice, Drinks - Canned, Bottled	19.4	33.1	21.5
Nuts	18.3	25.9	34.4
Shortening, Oil	18.3	31.5	59.1
Laundry Supplies	18.0	23.4	32.9
Breakfast Food	17.3	22.5	38.0
Prepared Food-Dry Mixes	16.2	22.8	27.3
Tea	14.7	24.0	41.7
Soft Drinks-Non-Carbonated	14.1	24.8	45.7
Batteries And Flashlights	14.0	15.5	21.4
Household Cleaners	13.2	20.7	32.9
Baby Food	12.7	13.6	24.0
Vegetables - Canned	11.4	17.2	31.5
Cosmetics	11.3	15.1	17.4
Household Supplies	10.9	20.4	31.4
Pet Care	10.2	19.8	38.6
Prepared Food-Ready-To-Serve	9.9	17.1	24.1
Ice Cream, Novelties	8.9	14.4	16.9
Cough And Cold Remedies	8.7	12.5	26.4
Prepared Foods-Frozen	7.6	12.9	21.1
Bread And Baked Goods	7.4	11.3	15.3
Hair Care	7.1	10.0	18.8
Snacks	6.6	10.8	15.8
Spices, Seasoning, Extracts	6.6	10.1	13.6
Baking Supplies	4.6	8.3	16.3
Candy	3.7	5.6	9.3
Liquor	3.5	6.5	18.8
Kitchen Gadgets	3.1	7.1	18.2
Dressings/Salads/Prep Foods-Deli	2.6	5.6	8.5
Salad Dressings, Mayo, Toppings	1.5	1.9	4.3
Cheese	1.2	2.3	5.1
Packaged Meats-Deli	1.1	2.1	4.3
Cookies	0.9	1.7	2.8
Condiments, Gravies, And Sauces	0.6	1.2	2.6
Wine	0.6	1.6	2.7
Vitamins	0.3	0.6	0.8
Medications/Remedies/Health Aids	0.1	0.0	0.0
Unprep Meat/Poultry/Seafood-Frzn	0.0	0.0	0.0
Fresh Produce	0.0	0.0	0.0

In Table A.3, we provide summary statistics for the demographic heterogeneity across segments. In particular, for each demographic group described in the rows, we compute its share in each segment and product group. Then, we compute the coefficient of variation (defined as the ratio between standard deviation and average), average, standard deviation, difference between maximum and minimum value, maximum, and minimum of such a share across segments within a product group. We then average across years and across product groups. Take the row "Young". The coefficient of variation for young households is 0.13. Young households constitute on average 31% of a segment and the standard deviation across segments is, on average, 0.04, and the coefficient of variation is 10.8. The segment that is typically most represented by young households has a share of 36%, while the segment that is typically underrepresented by young households has a share of 27%. On average, the difference between the most "Young" segment and the least "Young" segment is 9 percentage points.

Figures A.3 and A.4 plot the distribution across product groups of a few quantitatively relevant demographic characteristics. Figures A.5 and A.6 show the changes over time of these demographic characteristics.

Table A.3: Demographic Characteristics: Summary Statistics

	Coeff. Var.	Mean	Std. Dev.	Max-Min Diff.	Max	Min
Young	0.13	0.31	0.04	0.09	0.36	0.27
Mid-Age	0.05	0.21	0.01	0.02	0.22	0.20
High-Income	0.11	0.38	0.04	0.09	0.43	0.34
Low-Income	0.12	0.28	0.03	0.08	0.32	0.24
With Kid	0.01	0.93	0.01	0.02	0.94	0.92
Single	0.05	0.65	0.03	0.07	0.68	0.61
White	0.02	0.90	0.02	0.04	0.92	0.88
Black	0.18	0.07	0.01	0.03	0.08	0.05
High School	0.13	0.19	0.03	0.06	0.22	0.16
College	0.02	0.64	0.01	0.03	0.65	0.63

The table reports the coefficient of variation, average, standard deviation, max-min difference, maximum, and minimum across years and product groups of the share of households that belong to the demographic groups reported in the first column. All values are multiplied by 100. Notice that for ease of exposition, we have simplified the list of demographic characteristics. In particular, we only show the results for a subset of characteristics. For age, households are divided into Young, Mid-Age, and Old. For income, households are divided into high-income, middle-income, and low-income.

Figure A.3: Segment Demographic Characteristics: Heterogeneity Across Product Groups (Max-Min Difference)

(a) Age

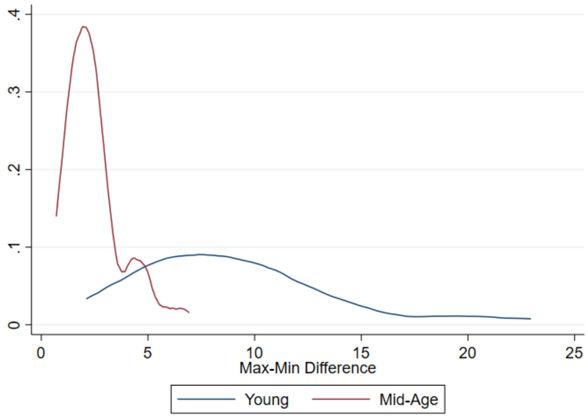

(b) Income

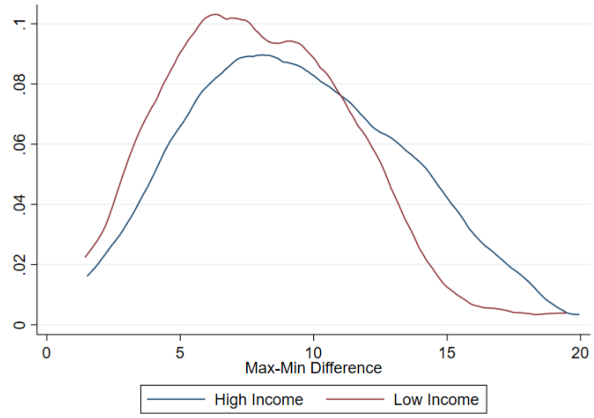

(c) Race

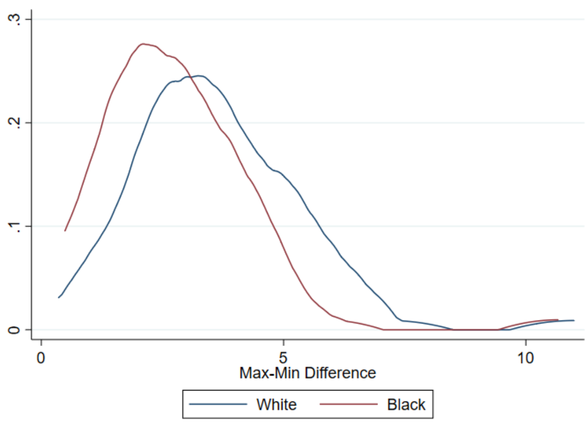

(d) Education

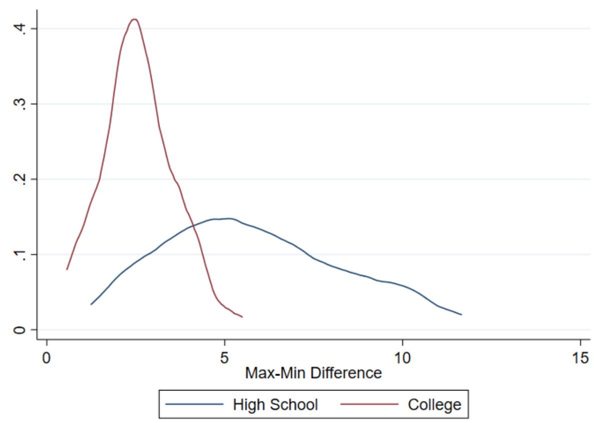

Figure A.4: Segment Demographic Characteristics: Heterogeneity Across Product Groups (Coefficient of Variation)

(a) Age

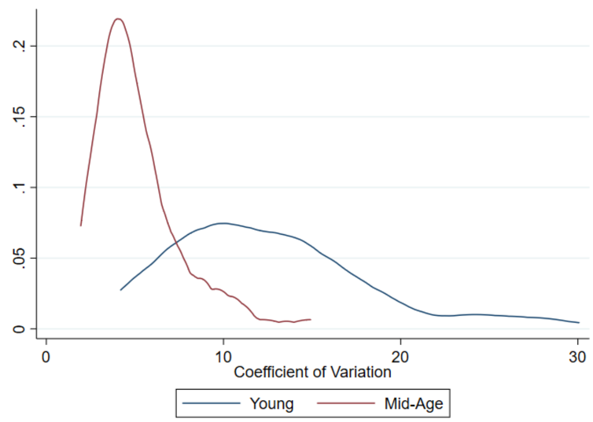

(b) Income

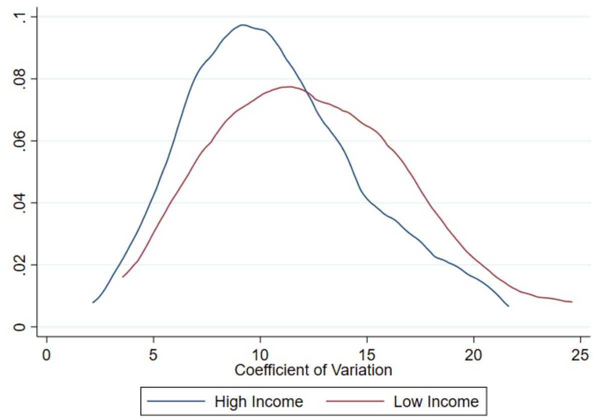

(c) Race

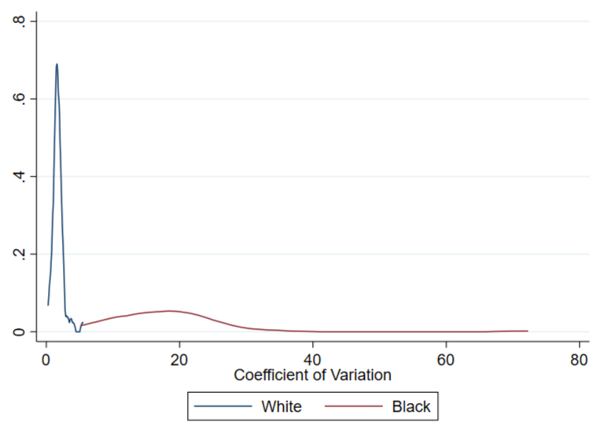

(d) Education

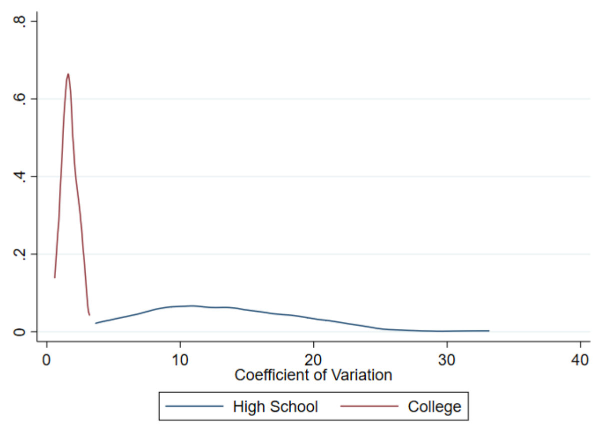

Figure A.5: Segment Demographic Characteristics: Changes over Time (Max-Min Difference)

(a) Age

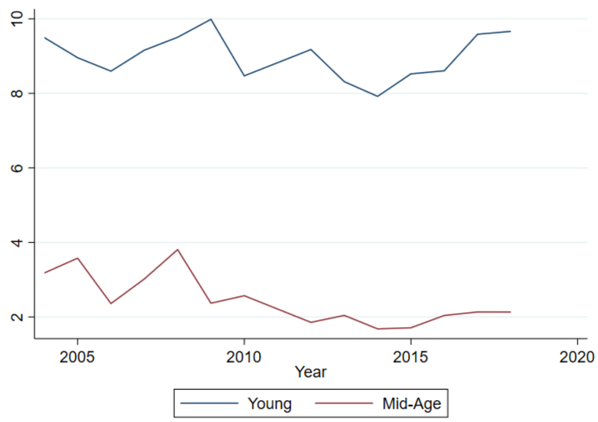

(b) Income

(c) Race

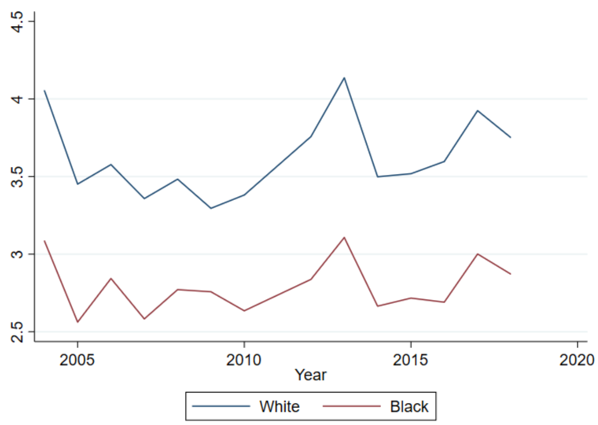

(d) Education

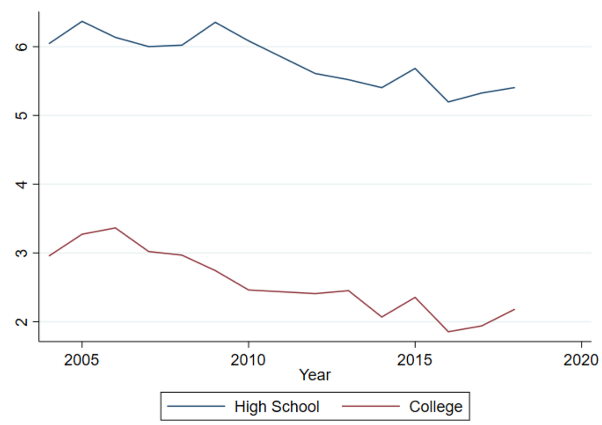

Figure A.6: Segment Demographic Characteristics: Changes over Time (Coefficient of Variation)

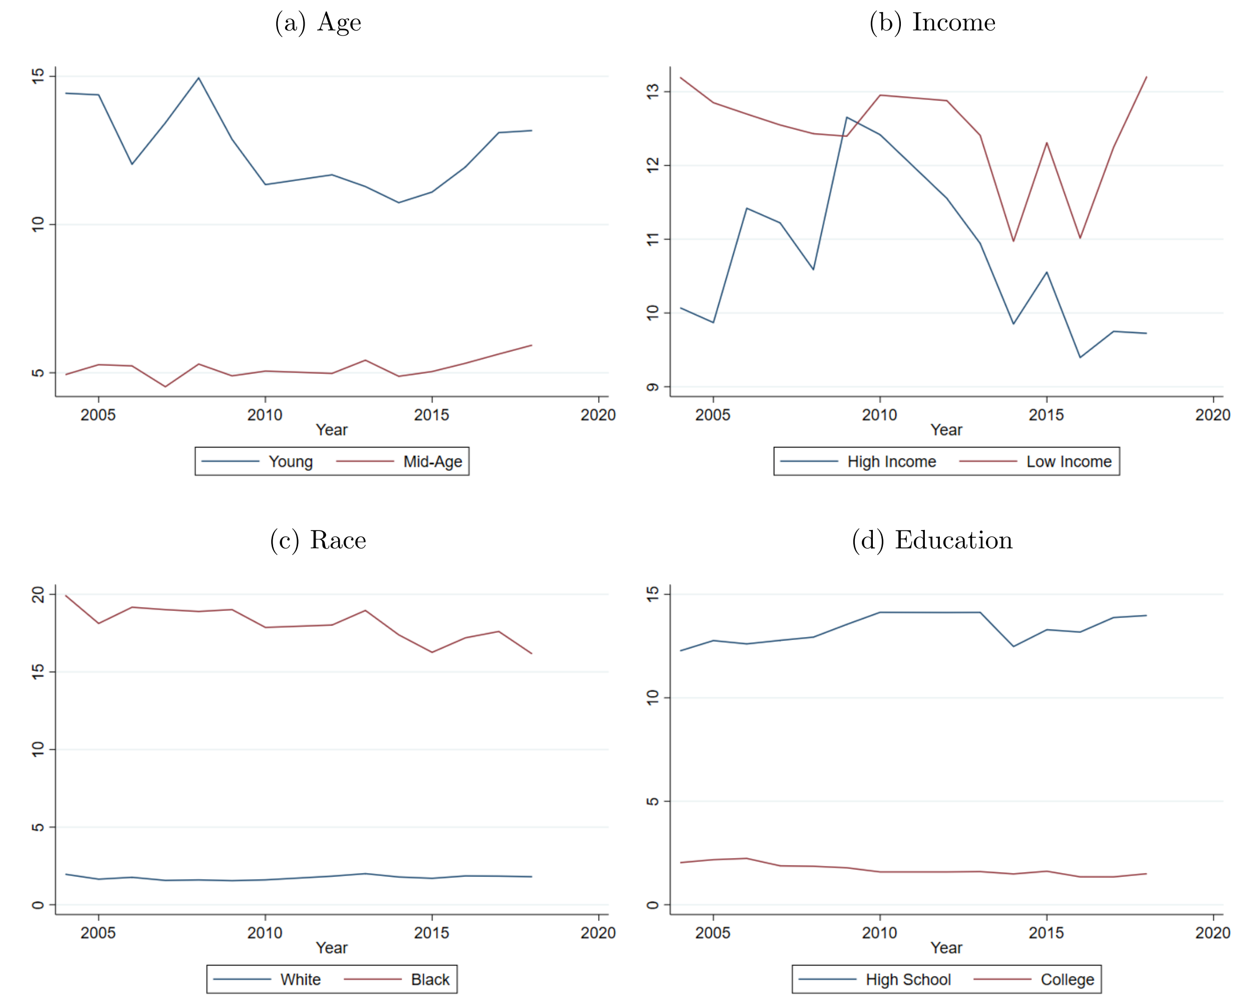

A.3 Geographic Distribution of Consumer Segments

Let $Popsh_{n|c}$ and $Popsh_g$ be the population share of consumers of type n in location c and at the national level, respectively. Under the null hypothesis that consumers of type n is independently uniformly distributed across regions $c = 1, ..., C$, the Bayesian rule suggests that:

$$H_0 : \quad Popsh_{n|c} = \frac{Popsh_n \times Popsh_c}{Popsh_c} = Popsh_n \tag{26}$$

where $Popsh_n$ is the population share of group n at the national level, and $Popsh_c$ is the population of consumers in region c as a share of national total. Rejection of (26) may raise the concern that consumers are concentrated in particular regions and firms may practice price discrimination across segments by setting different prices in different locations, which is inconsistent with the model and empirical specifications. In each year and product group, we test equation (26) by consumer segment, where we consider consumer segments are based upon their taste. Since the number of certain types of consumers is somewhat sparse at the county level, we consider a region as a state.

Panel A of Table A.4 reports the results for testing the random distribution of consumer segments across space. The last column shows low rejection rates (i.e., less than 14%) through all years, suggesting that the probability that certain types of consumers are clustered in some states is low. Table A.5 lists product groups where our null hypothesis has been rejected in the test. To rule out the possibility that our results are not driven by some outlier states, especially those with the least represented states in the consumer survey, we repeat the exercise and drop the lowest represented states with the lowest number of consumers. This set of results is provided in panel B of Table A.4, and product groups where our null hypothesis has been rejected in the test are provided in Table A.6.

Table A.4: Test for Geographic Distribution of Consumer Segments

Year	Average # of Segments / Group	Number of Tests	Number of Rejections	Rejection %
		Panel A: All States		
2004	4.31	276	51	18.48%
2005	4.36	279	38	13.62%
2006	4.22	270	38	14.07%
2007	4.30	275	33	12.00%
2008	4.56	292	36	12.33%
2009	4.53	290	35	12.07%
2010	4.45	285	33	11.58%
2011	4.13	264	41	15.53%
2012	4.41	282	32	11.35%
2013	4.30	275	42	15.27%
2014	4.06	260	45	17.31%
2015	4.27	273	33	12.09%
2016	3.92	251	36	14.34%
2017	4.05	259	39	15.06%
2018	4.25	272	39	14.34%
	Panel B: Drop the Bottom Five States with Lowest Number of Consumers			
2004	3.08	197	27	13.71%
2005	3.30	211	24	11.37%
2006	3.17	203	24	11.82%
2007	3.25	208	17	8.17%
2008	3.25	208	15	7.21%
2009	3.50	224	12	5.36%
2010	3.02	193	12	6.22%
2011	3.48	223	11	4.93%
2012	3.34	214	20	9.35%
2013	3.31	212	19	8.96%
2014	4.06	260	24	9.23%
2015	4.27	273	21	7.69%
2016	3.92	251	21	8.37%
2017	4.05	259	25	9.65%
2018	4.25	272	24	8.82%

Note: The table reports statistics tests for the null hypothesis displayed in (26), which is conducted for each of the 64 product groups by year independently, and the third column reports the total number of tests. The fourth and fifth columns report the number of rejections and the percentage rate.

Table A.5: Product Group with Rejections (2004-2018)

Group Code	Group Description	Number of Rejections
6004	DEODORANT	28
2503	COT CHEESE, SOUR CREAM, TOPPINGS	27
5003	WINE	27
510	PREPARED FOOD-READY-TO-SERVE	22
2005	ICE CREAM, NOVELTIES	21
6002	COSMETICS	20
4001	FRESH PRODUCE	18
6003	COUGH AND COLD REMEDIES	18
2510	YOGURT	16
1506	CRACKERS	15
1508	SOFT DRINKS-NON-CARBONATED	14
1005	CEREAL	13
1002	BAKING SUPPLIES	13
3001	DRESSINGS/SALADS/PREP FOODS-DELI	13
2007	PIZZA/SNACKS/HORS DOEURVES-FRZN	13
5502	BATTERIES AND FLASHLIGHTS	12
6016	SHAVING NEEDS	12
6014	ORAL HYGIENE	12
2008	PREPARED FOODS-FROZEN	11
514	VEGETABLES - CANNED	11
4502	DISPOSABLE DIAPERS	11
1016	SHORTENING, OIL	11
2506	MILK	10
6011	HAIR CARE	10
1007	CONDIMENTS, GRAVIES, AND SAUCES	9
1507	SNACKS	9
507	JUICE, DRINKS - CANNED, BOTTLED	9
3501	FRESH MEAT	9
1012	PACKAGED MILK AND MODIFIERS	9
6012	MEDICATIONS/REMEDIES/HEALTH AIDS	9
1011	NUTS	8
4507	PAPER PRODUCTS	8
4510	TOBACCO & ACCESSORIES	8
2505	EGGS	8
1501	BREAD AND BAKED GOODS	8
5001	BEER	7
4504	HOUSEHOLD CLEANERS	7
1004	BREAKFAST FOOD	7
501	BABY FOOD	7
5002	LIQUOR	6
5507	ELECTRONICS, RECORDS, TAPES	6
1505	COOKIES	6
2501	BUTTER AND MARGARINE	6
4511	WRAPPING MATERIALS AND BAGS	5
511	PREPARED FOOD-DRY MIXES	5
2009	UNPREP MEAT/POULTRY/SEAFOOD-FRZN	5
2010	VEGETABLES-FROZEN	5
6008	FIRST AID	4
1017	SPICES, SEASONING, EXTRACTS	3
503	CANDY	3
5515	KITCHEN GADGETS	3
4505	HOUSEHOLD SUPPLIES	2
513	SOUP	2
4506	LAUNDRY SUPPLIES	2
508	PET FOOD	2
3002	PACKAGED MEATS-DELI	2
1020	TEA	1
4501	DETERGENTS	1
506	JAMS, JELLIES, SPREADS	1
1015	SALAD DRESSINGS, MAYO, TOPPINGS	1

Table A.6: Robustness: Drop the Bottom Five States with Lowest Number of Consumers (2004-2018)

Group Code	Group Description	Number of Rejections
5003	WINE	32
3501	FRESH MEAT	20
6012	MEDICATIONS/REMEDIES/HEALTH AIDS	17
6004	DEODORANT	15
5502	BATTERIES AND FLASHLIGHTS	14
4001	FRESH PRODUCE	12
1012	PACKAGED MILK AND MODIFIERS	10
6011	HAIR CARE	10
6003	COUGH AND COLD REMEDIES	9
2505	EGGS	8
4502	DISPOSABLE DIAPERS	8
6018	VITAMINS	7
1508	SOFT DRINKS-NON-CARBONATED	7
6008	FIRST AID	7
508	PET FOOD	6
2005	ICE CREAM, NOVELTIES	6
2503	COT CHEESE, SOUR CREAM, TOPPINGS	6
2007	PIZZA/SNACKS/HORS DOEURVES-FRZN	6
4501	DETERGENTS	6
1016	SHORTENING, OIL	5
506	JAMS, JELLIES, SPREADS	5
5002	LIQUOR	5
510	PREPARED FOOD-READY-TO-SERVE	4
2009	UNPREP MEAT/POULTRY/SEAFOOD-FRZN	4
1004	BREAKFAST FOOD	4
6002	COSMETICS	4
3001	DRESSINGS/SALADS/PREP FOODS-DELI	4
1506	CRACKERS	4
4510	TOBACCO & ACCESSORIES	4
5001	BEER	4
4507	PAPER PRODUCTS	4
5515	KITCHEN GADGETS	4
514	VEGETABLES - CANNED	3
2510	YOGURT	3
1002	BAKING SUPPLIES	3
1005	CEREAL	3
5507	ELECTRONICS, RECORDS, TAPES	3
2010	VEGETABLES-FROZEN	2
503	CANDY	2
4506	LAUNDRY SUPPLIES	2
1007	CONDIMENTS, GRAVIES, AND SAUCES	2
507	JUICE, DRINKS - CANNED, BOTTLED	2
1507	SNACKS	1
501	BABY FOOD	1
1011	NUTS	1
1006	COFFEE	1
1505	COOKIES	1
2008	PREPARED FOODS-FROZEN	1
2501	BUTTER AND MARGARINE	1
1017	SPICES, SEASONING, EXTRACTS	1
1020	TEA	1
1501	BREAD AND BAKED GOODS	1

A.4 Robustness Checks of Clustering Algorithm

We provide an example of how we compute the probability that a consumer is classified into a segment where the majority of consumers had also appeared in a given segment in the baseline case. Suppose we would like to categorize 12 consumers into several groups based on different criteria. Table A.7 provides the baseline and alternative classification results, with the headline indicating the segment name. For instance, according to the baseline classification, segment A comprises consumers indexed as 1, 2, 3, and 4. In the alternative method, segment α consists of consumers indexed as 2, 3, 4, 5, 7, and 12.

Table A.7: Procedure Example

Baseline Classification				
Segment:	A	B	C	D
	1	5	8	10
	2	6	9	11
	3	7		12
	4			
Alternative Classification				
Segment:	α	β	γ	
	2	1	9	
	3	6	10	
	4	8	11	
	5			
	7			
	12			

As segment A and α are hardly comparable (due to the nature of the classification algorithm), it is difficult to tell whether consumers (who used to be) in segment A are classified in the same segment under the new criteria. To resolve this ambiguity, we consider that a consumer is classified into the "same segment" as before if the new segment (where the consumer is in the new criteria) comprises the majority of segment members belonging to her/his old segment. In the above example, consumer 1 is in segment β, while the majority of her/his group members belong to segment α (i.e., 2, 3, 4). So, we think consumer 1 is not classified into the "same segment" in baseline criteria (i.e., segment A). In contrast, consumers 2, 3, and 4 are considered to remain in the same group, as the majority of her/his segment members stay with her/him.

In this way, seven consumers are classified into the "same segment" as before, and they are consumers 2, 3, 4, 5, 7, 10, 11. The rest are classified into different segments: consumers 1, 6, 8, 9, 12. Suppose members of a specific segment in the baseline criteria are evenly distributed in new segments. In that case, we consider all these consumers are not in the same segment as before, corresponding to consumers 8 and 9 in the above example. The probability that a consumer is classified into a segment consisting of most consumers who had been put into the same segment based on the baseline criteria is thus $7/12 = 0.58$. We compute such a probability for each year and product category, and their distribution is reported in Figure A.7.

Figure A.7: Probability of being Classified into the Same Group

(a) Share of UPC Number by Brand

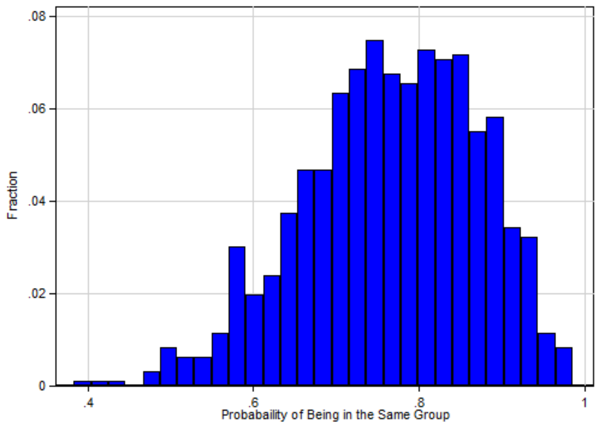

(b) Brand Indicator Variable

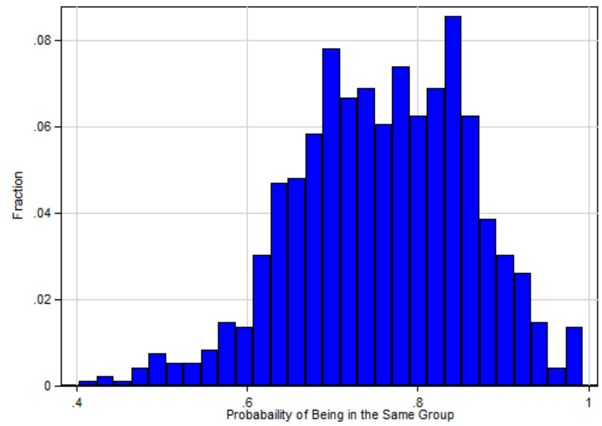

Figure A.8: Average Number of Segments by Product Group Comparision (2004 - 2018)

(a) Baseline & Share of UPC Number by Brand

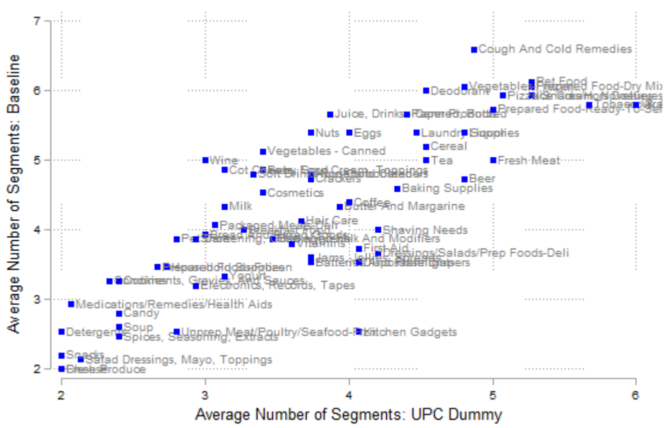

(b) Baseline & Brand Indicator

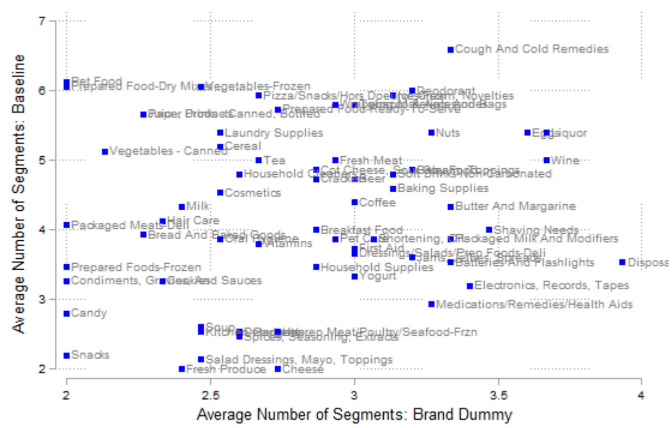

A.5 Profit Share

Figure A.9 reports the market share of the top 1, 5, and 10 firms in each product group in 2012. Figure A.10 reports the profit share computed in the three models and the difference across models holding constant the expenditure shares for each product group at its 2004 level. As a result, the profit share predicted by the model with small firms is constant. In Figure A.11, we replicate dropping a few outlier product groups that experience a drastic drop in the profit share after 2010. The product groups are baby food; condiments, gravies, and sauces; unprepared meat/poultry/seafood frozen; dressings/salads/prepared foods - deli;

fresh produce; vitamins. Figure A.12 shows the result using alternative clustering criteria as discussed in section 4.2.3.

Figure A.9: Market Share of the Largest Firms by Product Group

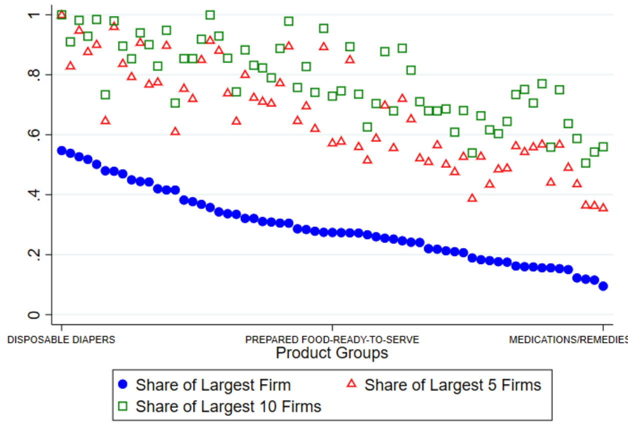

The figure plots the market share of the largest firm, largest 5 firms, and largest 10 firms across product groups.

Figure A.10: Profit Share Over Time - Constant Expenditure Shares

(a) Aggregate Profit Share Across Models

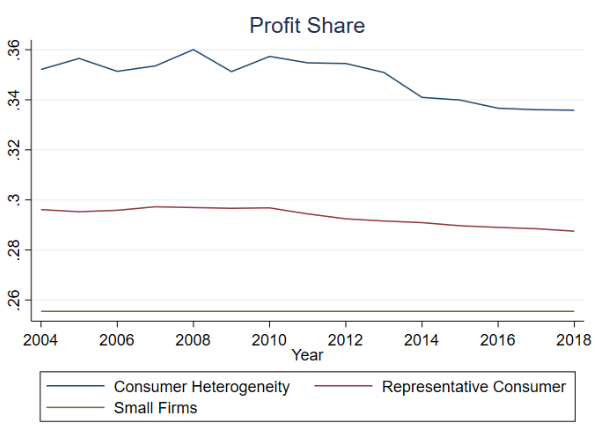

(b) Percentage Point Difference Across Models

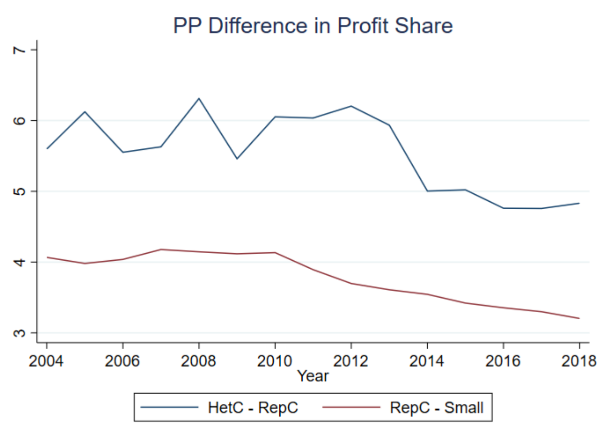

Figure A.11: Profit Share Over Time - Drop Outliers

(a) Aggregate Profit Share Across Models

(b) Percentage Point Difference Across Models

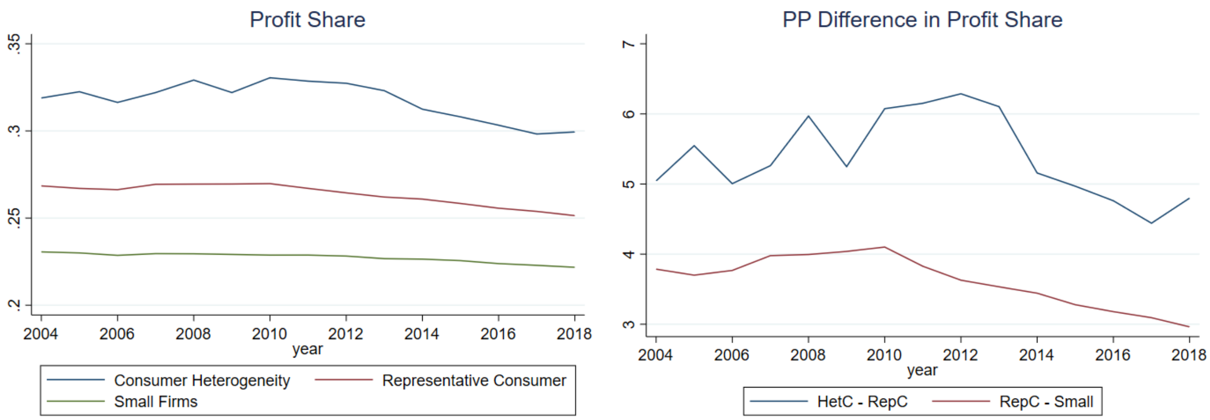

Figure A.12: Profit Share Over Time - Alternative Clustering Criteria

(a) Share of UPC Number by Brand

(b) Brand Indicator

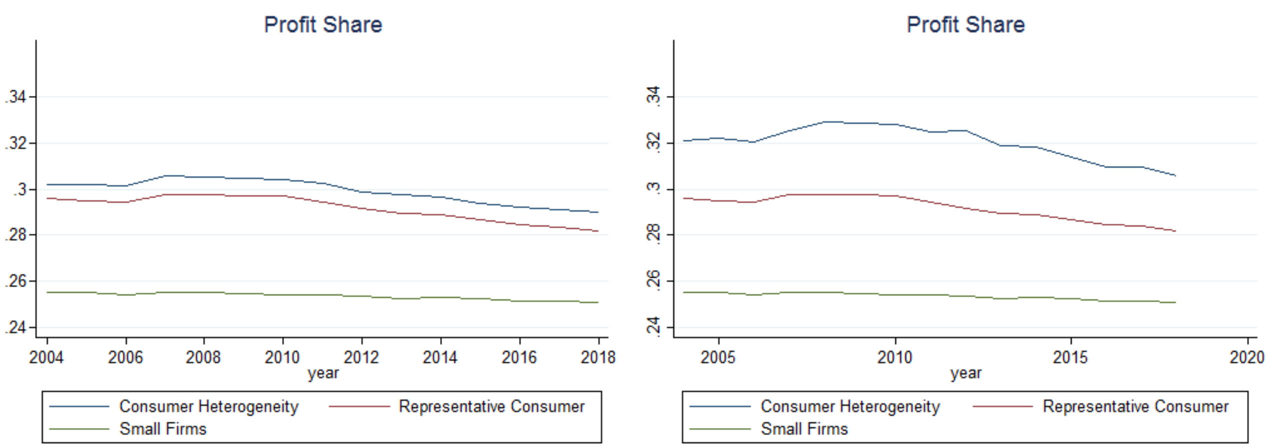

Figure A.13 plots the average number of firms per product group and the average number of products per firm (per product group) from 2004 to 2018. These changes in the number of firms and in their scope can contribute to the decline in profits after 2010 for the *RepC* model, in particular, where consumer heterogeneity does not play a role.

Figure A.13: Number of Firms and Number of Products Per Firm

(a) Number of Firms

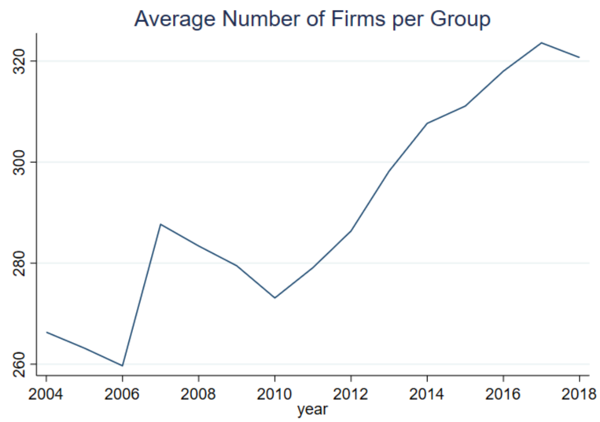

(b) Number of Products Per Firm

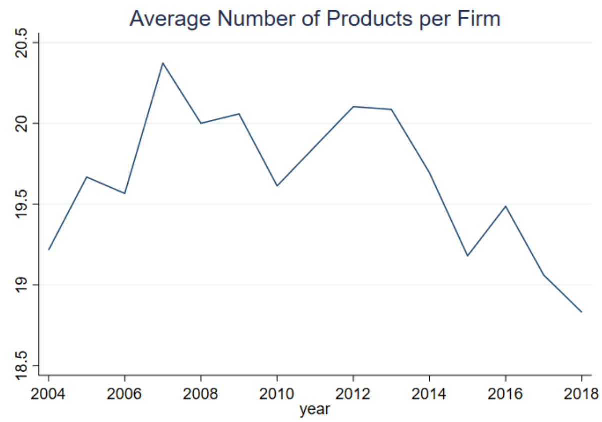

A.6 Profit Share with Confidence Interval

Figure A.14 reports the 95% confidence intervals for the profit share computed in the three models and the difference across models holding constant the expenditure shares for each product group at its 2004 level. Robust standard errors are calculated using bootstrapping. Firms are sampled 30 times randomly within strata based on size and product categories, where a bootstrapped sample has the same sample size as the raw data.[45] We re-compute their market shares, then markups and profit shares. Panel (b) shows the difference in the profit shares between the *HetC* and *RepC* models and between the *RepC* and small firm models (with 0.1 on the vertical axis representing a 10 percentage point difference, for example). The profit shares of the *HetC* and *RepC* models are significantly different in all years, and likewise for the *RepC* and small firm models except in 2011.

Figure A.14: Profit Share with 95% Confidence Interval

(a) Aggregate Profit Share Across Models

(b) Percentage Point Difference Across Models

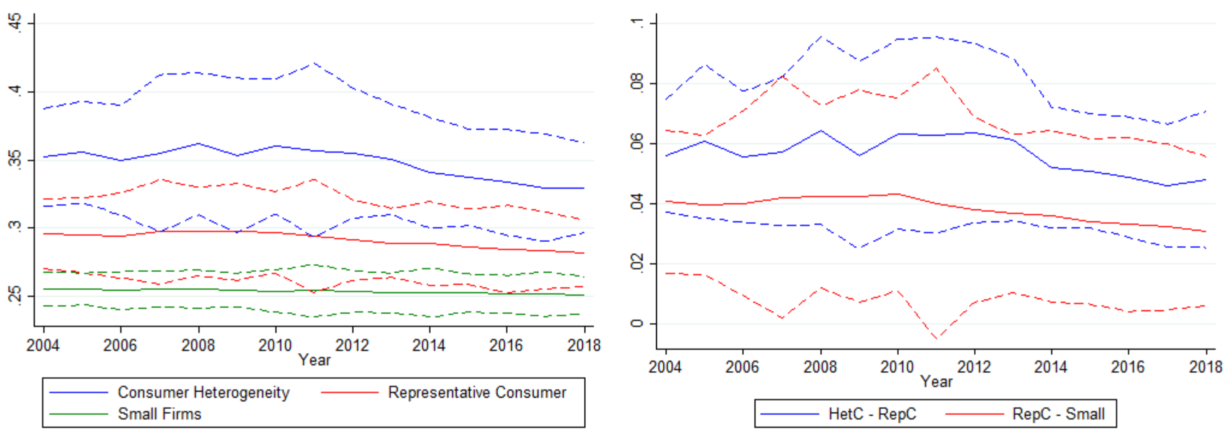

A.7 Robustness

For robustness, we also divide consumers into segments based on geography and distribution channel. In the next paragraphs, we describe how we divide consumers into segments according to the two dimensions.

Geography. We consider the division by state, using the home scanner data to divide consumers according to the states in which they reside.

Distribution Channel. We consider the division by retail chain. We assign consumers to a retail chain according to the store belonging to the chain in which they conduct the majority of their purchases. That is, both for the geographical and the distribution channel

[45] In each product category, we divide firms into 20 size categories, with firm size measured as total sales each year.

criteria, the demand from each segment n consists of the average demand of consumers in each geographical area or retail chain when we use the home scanner data.

In Figures A.15, A.16, A.17, and A.18 we use the two criteria described in this section along with consumer type to divide consumers in segments. Hence, one segment is a geography-retail chain-type triplet. Figure A.15 confirms the results of Figure A.20, since we find a positive correlation between sales and the common component of the demand shifter (panel (a)) and a negative correlation between the common component of the demand shifter and the variance of the segment-specific component (panel (b)). In Figure A.16, we plot the coefficient of a regression of IP_{fg} on firm sales (panel (a)) and scope (panel (b)). The results are in line with those shown in Figure A.21, in which consumers are divided in segments only by their type. In Figure A.17, we confirm the results of Figure A.22 as we find a positive correlation between EP_{fg} and firm sales (panel (a)) and scope (panel (b)). The results indicate that larger firms use their products to target different segments, defined at the state-retail chain-type level. Finally, Figure A.18 replicates the results shown in section A.8.1. The figure shows the distribution of the coefficients of a univariate regression of B_{fgn}, which is defined in (27), on the average household expenditures and expenditure share of each segment.

Figure A.15: Demand Shifters and Segments

(a) $\bar{\phi}_{ifg}$ and UPC sales

(b) $\bar{\phi}_{ifg}$ and variance of ϕ_{ifgn}

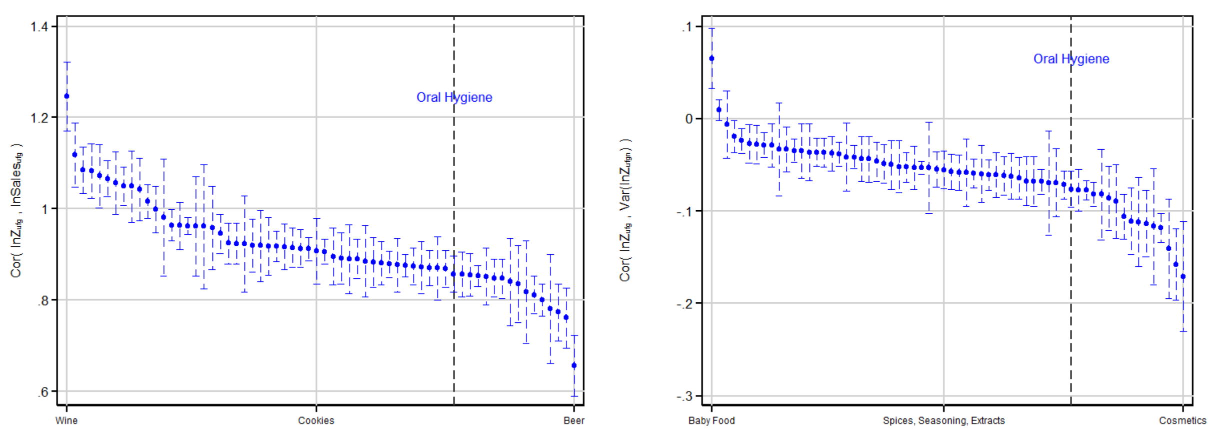

Figure A.16: Intensiveness of Personalization and Firm Characteristics

(a) Firm Sales

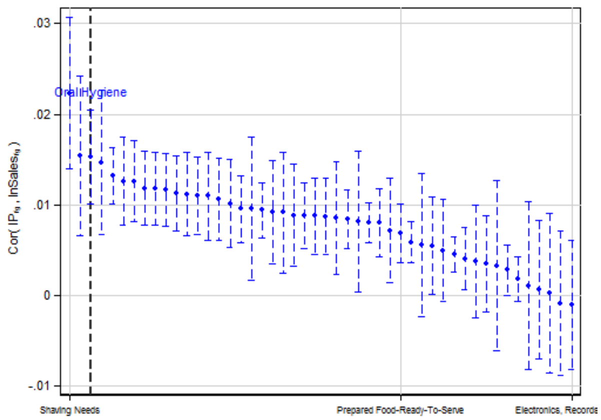

(b) Firm Scope

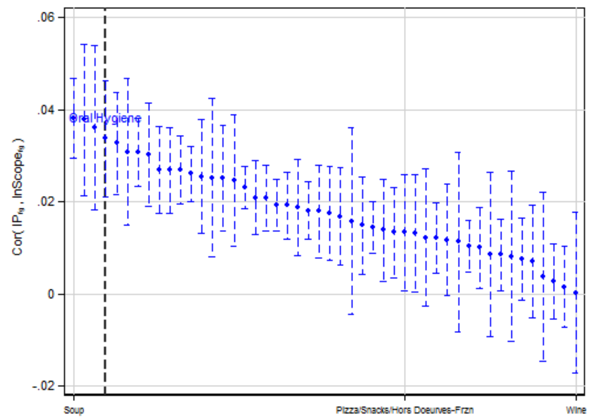

Figure A.17: Extensiveness of Personalization and Firm Characteristics

(a) Firm Sales

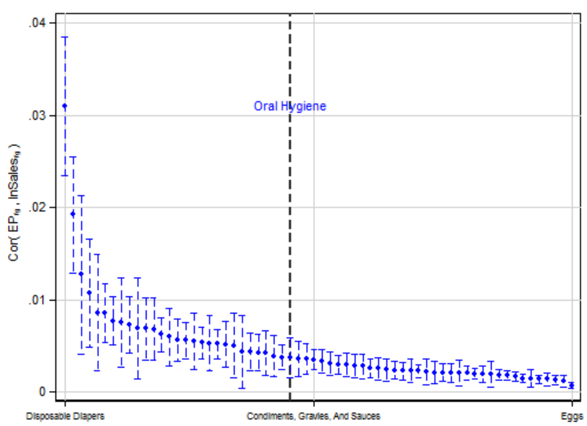

(b) Firm Scope

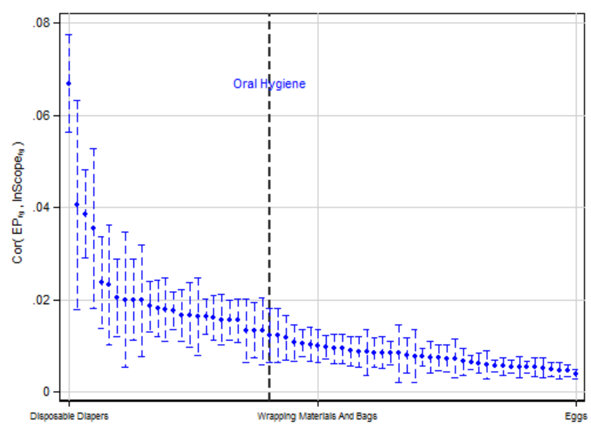

Figure A.18: Consumer Characteristics and Product Personalization

(a) Average Household Expenditure

(b) Expenditure Share

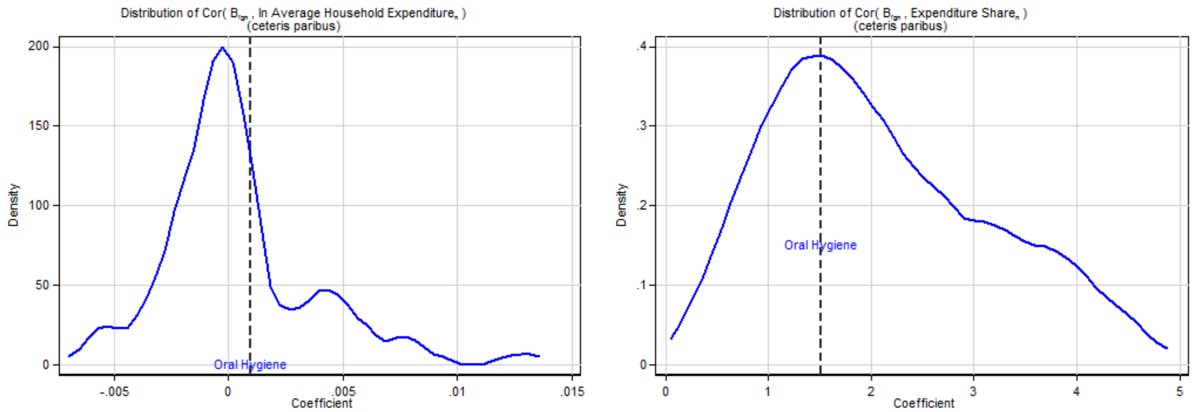

A.8 Firm Strategy to Target Segments

While in the main text we document the strategies firms follow to deal with consumer heterogeneity over time, in this section, we consider the year 2012 and examine the heterogeneity across product groups and across firms. In Figure A.19, we plot the results from the variance decomposition discussed in section 6.1.1 across product groups. In particular, for each product group, we report how much of the variance of the demand shifter is explained by the segment-specific component.

Panel (a) of Figure A.20 reports the coefficients obtained by regressing UPC log sales on the estimated UPC taste parameter in logs common across segments. Panel (b) of Figure A.20 displays the slope of a regression of $\ln \bar{\phi}_{ifg}$ on the variance of $\ln \phi_{ifgn}$ across segments for all product groups. Both figures report the point estimate of the coefficient and the 95% confidence interval, where we control for firm fixed effects in each regression and cluster standard errors at the firm level. The results indicate that products that have larger common component of the demand shifters have larger sales and have lower variance of the segment-specific component.

Figure A.21 reports the coefficient of a regression of IP_{fg} on firm sales (panel (a)) and firm scope (panel (b)). The coefficients are generally close to zero and often statistically insignificant. Figure A.17 reports the coefficient of a regression of EP_{fg} on firm sales (panel (a)) and firm scope (panel (b)). Contrary to the results on the intensiveness of personalization, here we find a positive and statistically significant coefficient for all product groups. The robust pattern indicates that larger firms target more segments than smaller firms.

Figure A.19: Contribution of Segments $\ln \phi_{ifgn}$ to Explain Product Appeal

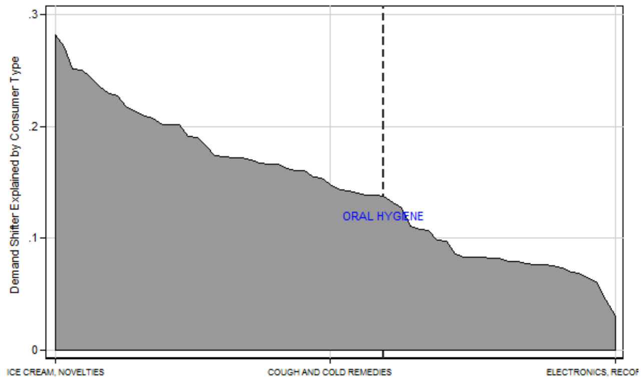

Figure A.20: Demand Shifters and Segments (All UPC) - Type (Consumer Panel)

(a) $\bar{\phi}_{ifg}$ and UPC sales

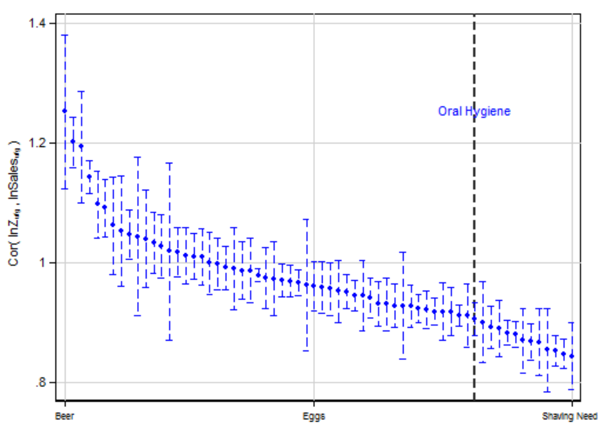

(b) $\bar{\phi}_{ifg}$ and variance of ϕ_{ifgn}

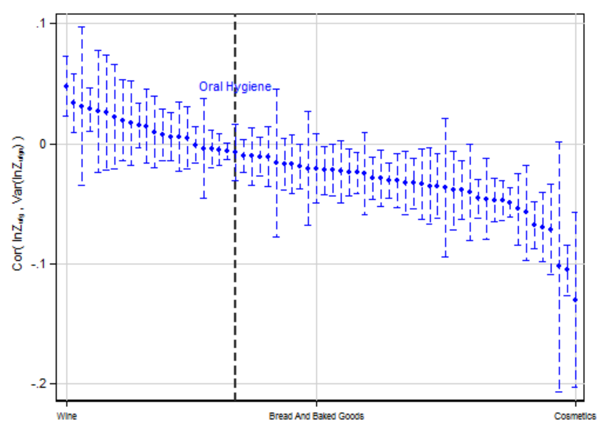

Figure A.21: Intensiveness of Personalization and Firm Characteristics (All UPC) - Type (Consumer Panel Data)

(a) Firm Sales

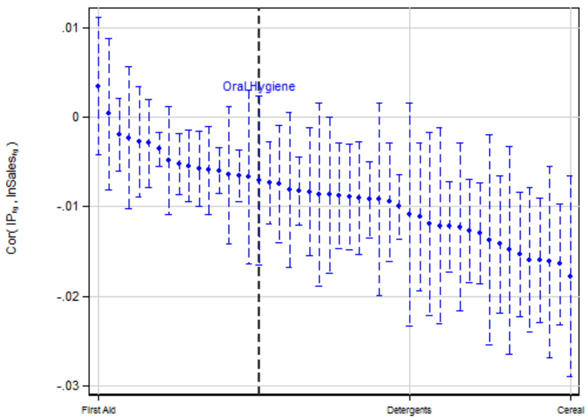

(b) Firm Scope

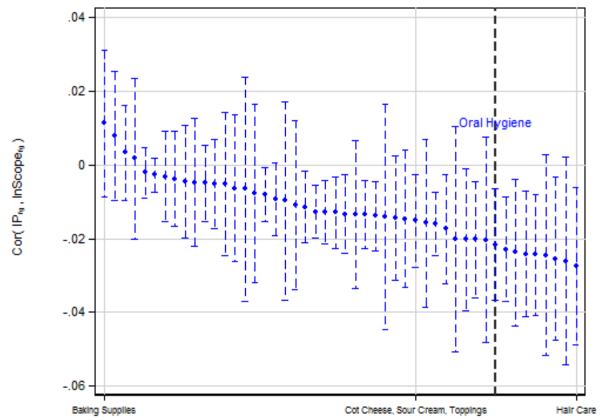

Figure A.22: Extensiveness of Personalization and Firm Characteristics (All UPC) - Type (Consumer Panel Data)

(a) Firm Sales

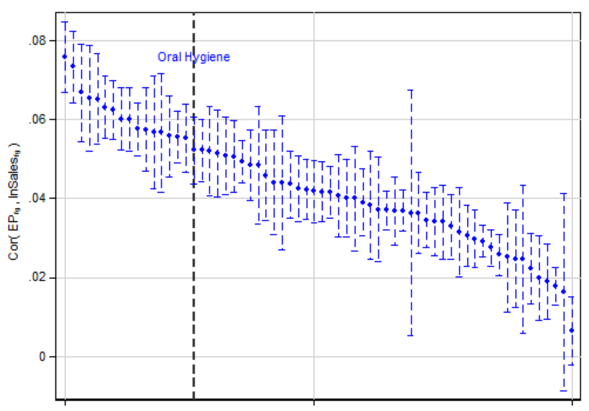

(b) Firm Scope

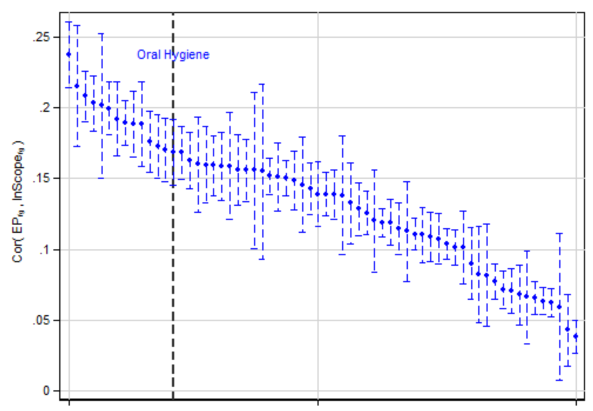

A.8.1 Firm Personalization Strategies

In this section, we determine the segment characteristics are mostly correlated with firm personalization strategies. Consider the share of best-ranked products of a firm f in segment n:

$$B_{fgn} = \frac{\text{\# best-ranked products of firm f in segment n}}{\text{Total \# of products of firm f}} \tag{27}$$

By definition, $B_{fgn} \in [0, 1]$. If a firm has no products targeting a given segment then $B_{fgn} = 0$. In contrast, if all products of a firm target a segment n then $B_{fgn} = 1$. To examine the main drivers of firms sorting into different segments, we run a univariate regression of

B_{fgn} on segment size, defined as the share of total expenditures in the product group, and average expenditure per household.

Figure A.23 shows the distribution of the coefficients of univariate regression of B_{fgn} on the average household expenditures and expenditure share of each segment. On average, the effect of household expenditure is zero. In contrast, the size of a segment measured by its expenditure share tends to have a positive effect on B_{fgn}. This means that firms tend to mainly target the segment with the largest expenditure share.

Figure A.23: Consumer Characteristics and Product Personalization (All UPC) - Type (Consumer Panel Data)

(a) Average Household Expenditure

(b) Expenditure Share

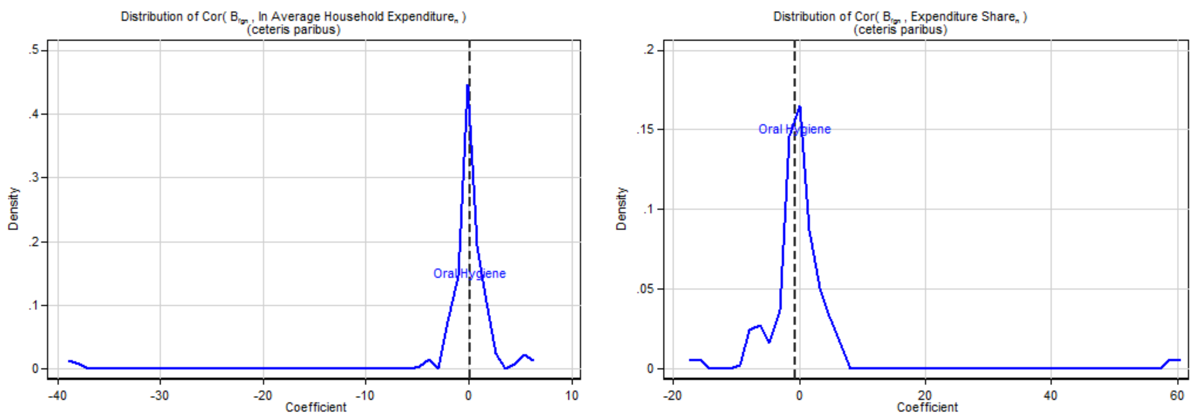

This fact, combined with the positive relationship between the extensiveness of personalization and firm size, suggests the presence of a segment ladder: the smaller firms tend to target a smaller number of segments, and these segments tend to be the larger ones. Larger firms are able to expand the number of segments they target, by further targeting the smaller segments. Although examining such a pattern for all product groups can be hard to illustrate, in the Appendix A.9.2 on the oral hygiene product group, we confirm and show such a segment ladder.

A.9 The Oral Hygiene Product Case

A.9.1 The Oral Hygiene Product Case

The tables in this section show the differences in consumption patterns and demographic characteristics of the five segments described in Table 1 for the product group of oral hygiene products. Table A.8 shows the share of expenditures within segment on different firms. For instance, consumers of Segment 4 disproportionally purchase Procter & Gamble products, which account for 64% of their expenditure share on oral hygiene products. The second most purchased firm is Colgate-Palmolive, which accounts for 8% of expenditures. Table A.9 report the distribution across segments of different consumer demographic characteristics. The first two columns show the population shares of different race/ethnicity for each segment.

The segmentation across segments does not follow these demographic variables, as the share of each race/ethnicity is very similar across segments. For example, the share of white households ranges between 0.88 and 0.92. The remaining columns show the share of the population by income, age, and education. We find that Segment 4 is overly represented by high-income households, since they account for 41% of the segment population, while in the other segments the share of high-income households is 30-37%.

Table A.8: Segments Top Firms

	Top Brand		Second Top Brand	
Segment #	Name	Share	Name	Share
1	Procter & Gamble	0.09	Colgate-Palmolive	0.08
2	Warner-Lambert	0.38	Philips	0.15
3	Glaxosmithkline	0.40	Johnson & Johnson	0.10
4	Procter & Gamble	0.64	Colgate-Palmolive	0.08
5	Block Drug	0.52	Procter & Gamble	0.10
6	Colgate-Palmolive	0.64	Procter & Gamble	0.10
7	Church & Dwight	0.53	Colgate-Palmolive	0.09

Table A.9: Segments characteristics

Segment #	White	Black	Low-Income	High-Income	Young	Middle Age	High School	College
1	0.91	0.06	0.26	0.37	0.32	0.17	0.17	0.65
2	0.88	0.08	0.26	0.37	0.39	0.16	0.16	0.66
3	0.88	0.09	0.24	0.41	0.30	0.16	0.15	0.65
4	0.92	0.05	0.34	0.30	0.20	0.14	0.22	0.63
5	0.91	0.06	0.28	0.34	0.35	0.16	0.17	0.66
6	0.90	0.07	0.32	0.31	0.30	0.17	0.18	0.65
7	0.88	0.07	0.33	0.33	0.31	0.16	0.16	0.66

A.9.2 Segment Ladder

Firms tend to mainly target the segments with the largest average expenditures. There is, in fact, a positive relationship between the average expenditure of a segment and the share of firm products targeting such segment. In addition, larger firms tend to target multiple segments, but they predominantly focus their products on catering to the tastes of the same segments targeted by the smaller firms.

We investigate a possible segment ladder for consumer type segments. We divide firms into three bins indexed by b according to their product scope. In particular, the first bin contains firms with a narrow scope (bottom 33 percentile with respect to the number of UPC products), the second firms with a middle scope (33 to 66 percentile), and the third firms with a wide scope (top 33 percentile). In each bin, there are F_b firms. For each bin and segment, we compute the average B_{fgn} previously defined in (27). Table A.10 summarizes the average B_{fn} for firms with narrow, middle, and wide scope.[46]

[46]The average numbers of UPC products per firm are 2.6, 5.8, and 72 for the narrow-, middle-, and

Table A.10: segment Ladder and Firm Scope

Segment #	Narrow-scope	Middle-scope	Wide-scope
1	0.227	0.170	0.127
2	0.042	0.071	0.099
3	0.019	0.069	0.109
4	0.139	0.155	0.185
5	0.065	0.073	0.120
6	0.060	0.089	0.128
7	0.449	0.372	0.232

Narrow-scope firms tend to disproportionally offer large shares of their product varieties to segments 1 and 7. Segment 7 is also mainly targeted by middle- and wide-scope firms, and segment 1 by middle-scope firms. However, middle- and especially wide-scope firms tend to target other segments as well. The comparison of targeted segments across firms highlights a ladder in the scope expansion. Firms tend to first target segments 7 and 1. As firms add more products, these are targeted to the remaining segments.

A.10 Inequality and Profit Share

Figure A.24 reports the positive correlation between the level of income inequality and the gap between the profit share of *HetC* and *RepC* models.

Figure A.24: Difference between Profit Shares of *HetC* and *RepC* models and Inequality (2004-2018)

(Δ lnProfit Share: Gini Index)

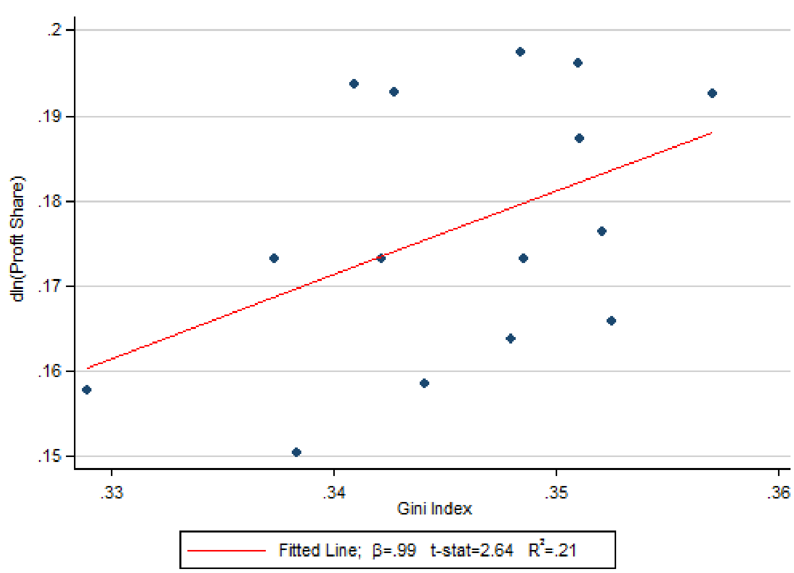

wide-scope firms.

A.11 Fixed Costs of Targeting Consumers

In this section, we extend our earlier model to allow for fixed costs for each UPC item. We think of these fixed costs as reflecting marketing costs for targeting an item to a particular consumer segment. As in the previous section, we let $n(i)$ denote the consumer segment with the highest demand parameter for item i, meaning that $z_{ifn(i)} > z_{ifm}$ for all $m = 1, ..., N$ with $m \neq n(i)$. Each consumer segment. We assume that targeting a particular segment n is associated with product design and marketing costs FC_n, which can differ across n because some segments might be easier to reach, or appeal to, than others. For each UPC i, we then denote the fixed cost of targeting that associated consumer segment by $FC_{n(i)}$. In addition, there can be any other fixed costs associated with firm f that we denote by FC_f.

Including the fixed marketing costs, the profits from UPC i are written as:

$$\tilde{\pi}_{if} = \sum_{n=1}^{N} L_n y_n s_{fn} s_{ifn} \left(1 - \frac{1}{\mu_{if}}\right) - FC_{n(i)}. \tag{28}$$

Firm f's total profits over all its UPCs are then:

$$\tilde{\pi}_f(I_f) = \sum_{i \in I_f} \tilde{\pi}_{if}(I_f) - FC_f, \tag{29}$$

where FC_f are any additional fixed costs of the firm.

Note that the profits $\tilde{\pi}_{if}(I_f)$ appearing on the right of (29) refer only to the profits earned on UPC item i, but they depend on the entire set of products I_f sold by firm f, as well as on the set of products sold by other firms. This dependence means that there are multiple equilibrium in the optimal set of products offered by each firm, i.e., what one firm chooses to offer depends on what other have already offered. Regardless of this multiplicity, the set of products sold by each firm in equilibrium must satisfy three conditions. First, profits must be non-negative, so $\tilde{\pi}_f(I_f) \geq 0$ in (29). Second, profits cannot increase by dropping a product from the set being produced and sold (with no change in the products sold by other firms):

$$\tilde{\pi}_f(I_f) \geq \tilde{\pi}_f(I_f/\{i\}) \qquad \forall i \in I_f, \tag{30}$$

where $I_f/\{i\}$ denotes the set of products I_f but without item i. Third, profits cannot increase by adding a new product to the set being produced and sold (with no change in the products sold by other firms):

$$\tilde{\pi}_f(I_f) \geq \tilde{\pi}_f(I_f \cup \{j\}) \qquad \forall j \notin I_f.$$

It is convenient to express the second of these conditions using variable profits π_{if}, so that $\tilde{\pi}_{if}(I_f) = \pi_{if}(I_f) - FC_{n(i)}$. Then it follows from (29) that (30) is rewritten as:

$$\pi_{if}(I_f) - \sum_{j \in I_f/\{i\}} \left[\pi_{jf}(I_f/\{i\}) - \pi_{jf}(I_f)\right] \geq FC_{n(i)} \qquad \forall i \in I_f. \tag{31}$$

To interpret this expression, $\pi_{if}(I_f)$ is the variable profits earned from item i. The next term on the left reflects the *drop* in profits obtained on the other $I_f/\{i\}$ items already sold by the firm from adding i and cannibalizing their sales, so that $\left[\pi_{jf}(I_f/\{i\}) - \pi_{jf}(I_f)\right] \geq 0$.[47] Using this inequality in (31), we obtain:

$$FC_{n(i)} \leq \pi_{if}(I_f). \tag{32}$$

Expression (32) states that an upper bound on the marketing and design costs associated with segment n are the variable profits earned on each UPC item i targeting that segment. Let us denote the total set of UPCs that target segment n by I_n. We have allowed the costs of targeting consumer segment $n(i)$ to depend on the characteristics of that set of households, but not on the item i itself. It follows that we can tighten the inequality in (32) as

$$FC_n \leq min_{i \in I_n} \left\{ \pi_{if}(I_f) \right\}. \tag{33}$$

In practice, we estimate these fixed marketing costs by treating the inequality above as equality. We stress that these marketing costs are only one component of the fixed costs of each firm, which also appear as FC_f on the right of (29), and we do not attempt to estimate these broader firm-specific fixed costs.[48] In Figure A.25, we plot the bounds for the fixed costs against the expenditure share of segments. We find a negative relationship: the marketing costs associated with targeting larger segments are lower than those for targeting niche segments. The result suggests that firms predominantly target those segments that have the lowest marketing costs. As a result, these segments are also the ones with the largest total revenues. Furthermore, the results rationalize the cross-sectional findings on the extensiveness of personalization. Smaller firms can only cover low marketing costs for the largest segments. Larger firms, in contrast, are able to afford the payment of larger marketing costs and thus can target smaller, niche segments, which tend to have larger marketing costs.

Using the estimated product-firm profits (see section 5), our procedure to estimate fixed marketing costs yields a vector of fixed costs (one per segment) for each product group in each year. For each product group in each year, we compute the average fixed marketing cost and the minimum and maximum across segments within the product group. We normalize these fixed costs by dividing them by the average fixed cost in 2004 for a given product group. Then, we take the average of these three measures across all product groups for each year and plot the results in Figure A.26. In 2004, the minimum fixed cost is 50% of the average and the maximum fixed cost is 60% larger than the average.

With the exception of a spike in 2005, the estimated average fixed costs remain fairly constant until 2011. After 2011, the average fixed cost increases, attaining a 2018 value

[47] In the Appendix, we also consider the cannibalization effects that firms face. Following (Hottman et al., 2016), we focus on cannibalization rates, namely the elasticity of demand of a segment for the last UPC with respect to the number of UPCs of a firm. Intuitively, cannibalization rates vary across segments. For each firm, we compute the difference between the maximum cannibalization rate and its minimum, across segments, finding that the heterogeneity across segments for the largest firms is sizable, while for the smallest firms it is negligible.

[48] While the resulting estimate of $F_{n(i)}$ from treating (33) as equality constitutes an upper bound for the fixed marketing cost of targeting the segment, in practice we find that this upper bound is still quite small, simply because the minimum of the variable profits on the right of (33) is small. If we deduct these fixed costs $F_{n(i)}$ from the variable profits of the firm, we find that the profit shares computed as in section 5 differ only in the third significant digit.

Figure A.25: Fixed Marketing Costs and Segment Characteristics

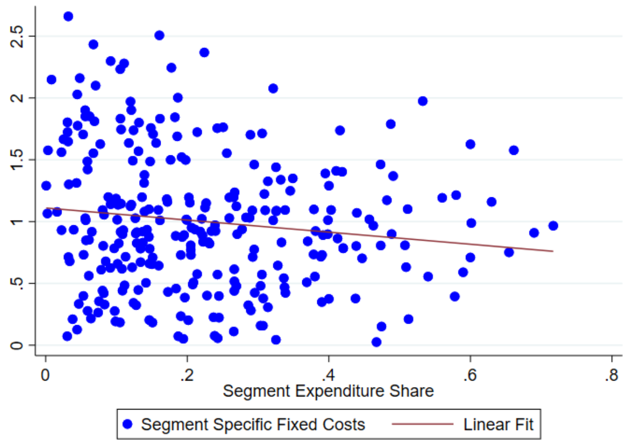

which is 54% larger than its value in 2004. This increase in the fixed cost corresponds to the decline in the profit share. This increase in the fixed cost can partially account for the decline in the role of consumer heterogeneity: as the fixed costs of marketing increase, it is more difficult for firms to place their products in a given segment. A similar pattern arises when we examine both the minimum and maximum fixed costs (averaged over product groups). We also consider the cross-section variation of fixed costs across segments for the year 2012.

Figure A.26: Fixed Marketing Costs Over Time