

NBER WORKING PAPER SERIES

FIRM HETEROGENEITY IN CONSUMPTION BASKETS:
EVIDENCE FROM HOME AND STORE SCANNER DATA

Benjamin Faber
Thibault Fally

Working Paper 23101
<http://www.nber.org/papers/w23101>

NATIONAL BUREAU OF ECONOMIC RESEARCH
1050 Massachusetts Avenue
Cambridge, MA 02138
January 2017, Revised January 2020

We would like to thank Mary Amiti, Costas Arkolakis, David Atkin, Kirill Borusyak, Pablo Fajgelbaum, Cecile Gaubert, Jessie Handbury, James Harrigan, Xavier Jaravel, Thierry Mayer, Volker Nocke, Steve Redding, Andres Rodriguez-Clare, Nicolas Schultz, Eric Verhoogen, Jonathan Vogel and participants at multiple seminars and conferences for helpful comments and discussions. Dmitri Koustas and May-Lyn Cheah provided outstanding research assistance. The analysis is based on data from The Nielsen Company (US), LLC and marketing databases provided by the Kilts Center for Marketing Data Center at The University of Chicago Booth School of Business. The conclusions drawn from the Nielsen data are those of the researchers and do not reflect the views of Nielsen, nor of the National Bureau of Economic Research. Nielsen is not responsible for, had no role in, and was not involved in analyzing and preparing the results reported herein.

NBER working papers are circulated for discussion and comment purposes. They have not been peer-reviewed or been subject to the review by the NBER Board of Directors that accompanies official NBER publications.

© 2017 by Benjamin Faber and Thibault Fally. All rights reserved. Short sections of text, not to exceed two paragraphs, may be quoted without explicit permission provided that full credit, including © notice, is given to the source.

Firm Heterogeneity in Consumption Baskets: Evidence from Home and Store Scanner Data
Benjamin Faber and Thibault Fally
NBER Working Paper No. 23101
January 2017, Revised January 2020
JEL No. E31,F61,O51

ABSTRACT

A growing literature has documented the role of firm heterogeneity within sectors for nominal income inequality. This paper explores the implications for household price indices across the income distribution. Using detailed matched US home and store scanner microdata, we present evidence that rich and poor households source their consumption from different parts of the firm size distribution within disaggregated product groups. We use the data to examine alternative explanations, propose a tractable quantitative model with two-sided heterogeneity that rationalizes the observed moments, and calibrate it to explore general-equilibrium counterfactuals. We find that larger, more productive firms endogenously sort into catering to the taste of richer households, and that this gives rise to asymmetric effects on household price indices. We quantify these effects in the context of policy counterfactuals that affect the distribution of disposable incomes on the demand side or profits across firms on the supply side.

Benjamin Faber
Department of Economics
University of California, Berkeley
697A Evans Hall
Berkeley, CA 94720
and NBER
benfaber@econ.berkeley.edu

Thibault Fally
Department of Agricultural and Resource Economics
University of California, Berkeley
University Hall, 7th Floor
Berkeley, CA 94720-3310
and NBER
fally@berkeley.edu

1 Introduction

Income inequality has been on the rise in the US and many other countries, attracting the sustained attention of policy makers and the general public. In this context, a growing literature has documented the role of Melitz-type firm heterogeneity within sectors in accounting for nominal income inequality.¹ In this paper, we explore the implications of firm heterogeneity for household price indices across the income distribution. We aim to answer three central questions: i) to what extent do rich and poor households source their consumption baskets from different parts of the firm size distribution?; ii) what forces explain these differences?; and iii) what are the implications of the answers to i) and ii) for the impact of policies or economic shocks on real income inequality?

In answering these questions, the paper makes three main contributions. First, using detailed matched home and store scanner consumption microdata, we document significant differences in the weighted average firm sizes that rich and poor US households source their consumption from, and explore alternative explanations.² Second, to rationalize these moments we develop a tractable quantitative model of product quality choice with two-sided heterogeneity across firms and households, and use the microdata to calibrate it and quantify the forces underlying the stylized facts. Third, we explore general-equilibrium (GE) policy counterfactuals to illustrate how, in a setting where rich and poor households source their consumption from heterogeneous firms, policies and economic shocks give rise to asymmetric price index effects across the income distribution.

At the center of the analysis lies a detailed collection of microdata that allow us to trace the firm size distribution into the consumption baskets of households across the income distribution. We combine a dataset of 345 million consumer transactions when aggregated to the household-by-retailer-by-barcode-by-half-year level from the AC Nielsen US Home Scanner data over the period 2006-2014, with a dataset of 12.2 billion store transactions when aggregated to the store-by-barcode-by-half-year level from the AC Nielsen US Retail Scanner data covering the same period. The combination of home and store-level scanner data allows us to trace the size distribution of producers of brands—measured in terms of national sales that we aggregate across on average 27,000 retail establishments each half year in the store scanner data—into the consumption baskets of on average 59,000 individual households per half year in the home scanner data within more than 1000 disaggregated retail product modules, such as carbonated drinks, shampoos, pain killers, desktop printers or microwaves.

The analysis proceeds in four steps. In step 1, we use the data to document a new set of stylized facts. We find that the richest 20 (resp. 10) percent of US households source their consumption from on average 20 (resp. 27) percent larger producers of brands within disaggregated product groups compared to the poorest 20 (resp. 10) percent of US households. We also document that these differences in weighted-average firm sizes arise in a setting where the rank order of household budget shares on different producers within a product group is preserved across income groups—i.e. the largest firms command the highest budget shares for all income groups. We interpret these stylized facts as equilibrium outcomes in a setting where both consumers and firms optimally choose their product attributes. We also use the microdata to explore a number of alternative explanations, and find that differential coverage of retail consumption or differences

¹e.g. Card et al. (2013); Helpman et al. (2017); Song et al. (2018). See related literature at the end of this section.

²We refer to firm size in terms of relative firm sales within product groups.

in product supply and pricing across the income distribution are unlikely driving these results. We also document that the relationship between incomes and firm sizes holds within each of the product departments in the data, and explore whether consumption categories not fully represented in the Nielsen data, such as consumer durables, health expenditures or digital media, show similar patterns. Focusing on subsets of these categories that are covered in the data—household appliances, pharmaceuticals and video, audio and software purchases respectively—we find similar or more pronounced differences in firm sizes across incomes.

In step 2 we write down a tractable model that rationalizes the stylized facts in the data. On the consumption side, we specify non-homothetic preferences allowing households across the income distribution to differ both in terms of their price elasticities as well as in their evaluations of product quality attributes. On the production side, we introduce quality choice into a model of heterogeneous firms within sectors. Both marginal and fixed costs can be functions of output quality, allowing for economies of scale in production. Markups can vary across firms due to both oligopolistic competition and selling to heterogeneous consumers. Firms now operate in a setting where their pricing and quality choices affect the composition of market demand that they face. Modeling product choices with two-sided heterogeneity implies that shocks that affect producers differently, such as trade integration, corporate taxes and regulations, can feed into the consumption baskets of rich and poor households asymmetrically. In turn, changes in the distribution of disposable incomes, due to e.g. income tax reform, affect firms differently across the size distribution with GE knock-on effects on household price indices.

In step 3, we use the microdata to estimate the preference and technology parameters. On the demand side, we find that rich and poor households differ both in terms of price elasticities and their valuation of product quality attributes. We find that poorer households have higher price elasticities relative to richer households, but that these differences, while statistically significant, are relatively minor. We also find that, while households on average agree on the ranking of quality evaluations across producers, richer households value higher quality significantly more. On the production side, we estimate that producing higher quality increases both the marginal and the fixed costs of production, giving rise to economies of scale in quality production. To estimate these technology parameters, we follow two different estimation strategies. The first follows existing work, and is based on cross-sectional variation in firm scale and output quality. The second exploits within-firm changes in brand quality and scale over time. To identify the effect of firm scale on output quality in the panel estimation, we use state-level measures of changes in brand quality on the left-hand side, and construct an instrument (IV) for national brand scale on the right-hand side. The IV exploits pre-existing differences in the geography of brand sales across other US states interacted with state-level variation in average sales growth observed in other product groups.

The parameter estimates from step 3 reveal two opposing forces that in equilibrium determine both firm sizes across consumption baskets and the sorting of firms across product quality attributes. On the one hand, larger firms offer lower quality-adjusted prices, which increases the share of their sales coming from more price-sensitive lower-income consumers. Since these consumers value quality relatively less, this channel, *ceteris paribus*, leads poorer households to source their consumption from on average larger firms, which in turn pushes these firms to produce at lower output quality. On the other hand, the estimated economies of scale in quality production

give larger firms incentives to sort into higher product quality, catering to the taste of wealthier households. Empirically, we find that this second channel dominates the first, giving rise to the endogenous sorting of larger, more productive firms into products that are valued relatively more by richer households. Armed with these estimates, we find that the observed stylized facts from step 1 translate into significant differences in the weighted-average product quality and quality-adjusted prices embodied in consumption baskets across the income distribution. The richest 20 percent of US households source their consumption from on average 22 percent higher-quality producers compared to the poorest quintile of households. At the same time, we find that the richest income quintile source their consumption at on average 10 percent lower quality-adjusted prices. Overall, we find that the calibrated model based on the estimates from step 3 fits the observed differences in firm sizes across consumption baskets from step 1 both qualitatively and quantitatively.

In the final step 4, we use the calibrated model to quantify a new set of price index implications in the context of two policy counterfactuals. The first counterfactual affects the distribution of disposable incomes on the demand side, and is motivated by the recent debate about progressive income taxation in the US and elsewhere. We evaluate the price index effects of returning from current US tax rates to a more progressive post-WWII US tax schedule, increasing the effective rate on the richest household group in our calibration from currently around 30 to 50 percent. This policy also closely corresponds to the counterfactual of moving the US to the current average effective rate on this group among Northern European countries, and it is in line with the proposed tax reforms of two leading presidential candidates for the 2020 US elections (Sanders and Warren). We find that the resulting compression of disposable incomes gives rise to GE knock-on effects—through changes in firm scale, output quality, variable markups and exit/entry across the firm size distribution—that affect rich and poor households differently, amplifying the progressivity of the reform. As a result, the poorest 20 percent of US households experience a 3 percentage point lower cost of living inflation for retail consumption compared the richest 20 percent.³

The second policy counterfactual affects the profits of firms across the size distribution on the supply side. The counterfactual is motivated by the ongoing debate about closing loopholes in corporate taxation. A growing literature in public finance has documented larger possibilities for tax avoidance or evasion among large US corporations (e.g. [Bao & Romeo \(2013\)](#), [Güvener et al. \(2017\)](#), [Wright & Zucman \(2018\)](#)). We use these findings to evaluate the implications of eliminating the kink that has been documented at the 95th percentile of the firm size distribution in the otherwise smoothly increasing relationship between firm size and effective corporate tax rates. This policy would lead to an increase of on average 5 percent in corporate taxes paid by the largest 5 percent of producers, ranging between on average 1 percent at the 95th percentile to 11 percent at the 99th. We find that even such a relatively modest adjustment in corporate taxation leads to a meaningful GE effect on differences in cost of living inflation between rich and poor households. This is in the order of a 2 percentage point lower inflation for retail consumption among the bottom 20 percent of US households compared to the top 20 percent. We document that the direct incidence of this policy—holding initial household and firm decisions fixed—accounts for about 40 percent of the inflation difference, with the remainder due to endogenous firm adjustments

³Given the data allow us to calibrate the model at the level of five broad income groups, while both tax changes and taste for quality increase with incomes within these groups, these results provide conservative estimates.

that affect consumption baskets differently.

We also explore the implications for the distribution of the gains from trade, extending a textbook Melitz model with two symmetric countries to quality choice with two-sided heterogeneity within countries, and calibrating it to the US microdata. We find that allowing for firm heterogeneity across consumption baskets makes the gains from trade more unequal: increasing bilateral import penetration by 10 percentage points leads to a 1.5-2.5 percentage point lower cost of living inflation in retail consumption for the richest 20 percent of households compared to the poorest 20 percent. Taken together, these findings illustrate that firm heterogeneity affects inequality in more complex ways than through the nominal earnings of workers, that have been the focus of the existing literature. These insights arise after introducing a basic set of features that we observe in the microdata—allowing for product attribute choice by heterogeneous firms and households—into an otherwise standard economic environment.

This paper is related to the growing literature on the extent, causes and consequences of firm heterogeneity within sectors that has spanned different fields in economics, including international trade (Bernard et al., 2007; Melitz, 2003), industrial organization (Bartelsman et al., 2013), macroeconomics (Hsieh & Klenow, 2009), development (Peters, 2013), labor economics (Card et al., 2013) and management (Bloom & Van Reenen, 2007). Within this literature, this paper is most closely related to existing work on the implications of firm heterogeneity for nominal income inequality (Burstein & Vogel, 2017; Card et al., 2013; Davis & Harrigan, 2011; Frias et al., 2009; Helpman et al., 2017; Sampson, 2014; Song et al., 2018).

Our theoretical framework builds on existing work, both on quality choice across the domestic income distribution with non-homothetic preferences on the demand side (Fajgelbaum et al., 2011; Handbury, 2019) and on quality choice across heterogeneous firms on the supply side (e.g. Bastos et al. (2018); Feenstra & Romalis (2014); Fielser et al. (2018); Johnson (2012); Kugler & Verhoogen (2012)). While these building blocks have been used separately—i.e. heterogeneous firms facing a representative agent in each country or non-homothetic preferences facing homogeneous firms—this paper combines them in a quantitative GE model. This two-sided heterogeneity within countries allows us to rationalize the observed differences in firm sizes across consumption baskets, and gives rise to new price index implications across the income distribution in the light of economic shocks or policy changes.⁴

The paper is also related to a growing empirical literature using the Nielsen data (Broda & Weinstein, 2010; Handbury, 2019; Handbury & Weinstein, 2014; Hottman et al., 2016). Most of this literature has been based on the home scanner data. More recently, Argente & Lee (2016) and Jaravel (2019) use the home and store scanner data to document that lower-income households have experienced higher cost of living inflation over the past decade and beyond. Argente & Lee (2016) relate this finding to a higher possibility for quality-downgrading among higher-income households

⁴A notable exception is Feenstra & Romalis (2014) who combine quality choice by heterogeneous firms under monopolistic competition with non-homothetic preferences on the demand side. Since their objective is to infer quality from observed unit values in international trade flows, their model features a representative agent within countries on the demand side. On the supply side, they follow e.g. Baldwin & Harrigan (2011); Kugler & Verhoogen (2012) modeling quality choice as a deterministic function of a firm’s productivity draw. To be able to rationalize our empirical results on the effect of firm scale on quality upgrading in firm-level panel data, we instead allow fixed costs to be a function of output quality, following earlier work by Sutton (1998), Hallak & Sivadasan (2013) and the second model variant in Kugler & Verhoogen (2012).

during the Great Recession, and Jaravel (2019) to more innovation and competition in product groups consumed more by richer households. In this paper, we present new empirical evidence suggesting that the widely documented presence of firm heterogeneity within sectors translates asymmetrically into the consumption baskets of rich and poor households, quantify the underlying forces and explore the GE implications in the context of policy counterfactuals.

The remainder proceeds as follows. Section 2 describes the data. Section 3 documents stylized facts and explores alternative explanations. Section 4 presents the model. Section 5 presents parameter estimation and model calibration. Section 6 presents the counterfactual analysis. Section 7 concludes.

2 Data

Retail Scanner Data The Nielsen retail scanner data consist of weekly price and quantity information generated by point-of-sale systems for more than 100 participating retail chains across all US markets between January 2006 and December 2014. When a retail chain agrees to share their data, all of their stores enter the database. As a result, the database includes more than 50,000 retail establishments. The stores in the database vary widely in terms of formats and product types: e.g. grocery, drug, mass merchandising, appliances, liquor, or convenience stores.

Data entries can be linked to a store identifier and a chain identifier so a given store can be tracked over time and can be linked to a specific chain. While each chain has a unique identifier, no information is provided that directly links the chain identifier to the name of the chain. This also holds for the home scanner dataset described below. The implication of this is that the product descriptions and barcodes for generic store brands within product modules have been anonymized. However, both numeric barcode and brand identifiers are still uniquely identified, which allows us to observe sales for individual barcodes of generic store brands within each product module in the same way we observe sales for non-generic products.

In appendix Table A.1 we aggregate the raw microdata to the store-by-barcode-by-half-year level. On average each half year covers \$113 billion worth of retail sales across 27,000 individual stores in more than 1000 disaggregated product modules, 2500 US counties and across more than 730,000 barcodes belonging to 175,000 producers of brands.⁵ As described in more detail in the following section, we use these data in combination with the home scanner data described below in order to trace the distribution of firm sizes (measured in terms of national sales measured across on average 27,000 stores per half year) into the consumption baskets of individual households.

Home Scanner Data The Nielsen home scanner data is collected through hand-held scanner devices that households use at home after their shopping in order to scan each individual transaction they have made. Importantly, the home and store level scanner datasets can be linked: they use the same codes to identify retailers, product modules, product brands as well as barcodes. As described in more detail in the following section, we use this feature of the database to estimate weighted average differences in firm sizes across consumption baskets.

In Table A.1 we aggregate the raw microdata to the household-by-barcode-by-half-year level.

⁵We do not make use of Nielsen’s “Magnet” database that covers non-barcoded products, such as fresh produce.

On average each six-month period covers \$109 million worth of retail sales across 59,000 individual households in more than 1000 product modules, 2600 US counties and close to 600,000 barcodes belonging to 185,000 producers of brands. One shortcoming of the home scanner dataset is that nominal household incomes are measured imprecisely. First, incomes are reported only across discrete income ranges. And those income bins are measured with a two-year lag relative to the observed shopping transactions in the dataset. To address this issue, we divide households in any given half year into percentiles of total retail expenditure per capita.⁶ To address potential concerns about a non-monotonic (or decreasing) relationship between total retail outlays and incomes, we also confirm in appendix Figure A.1 that our measure of total retail expenditure per capita is monotonically increasing in reported nominal incomes two years prior (confirming existing evidence that retail expenditure has a positive income elasticity).

Table A.1 also clarifies the relative strengths and weaknesses of the two Nielsen datasets. The strength of the home scanner database is the detailed level of budget share information that it provides alongside household characteristics. Its relative weakness in comparison to the store-level retail scanner data is that the home scanner sample of households only covers a small fraction of the US retail market in any given period. Relative to the home scanner data, the store-level retail scanner data cover more than 1000 times the retail sales in each half year. This paper takes advantage of both datasets for the empirical analysis, by combining national sales by product from the store scanner data with the detailed information on individual household consumption shares in the home scanner data.

3 Stylized Facts

This section draws on the combination of the home scanner and retail scanner data to document a set of stylized facts about firm heterogeneity embodied in the consumption baskets of households across the income distribution.

3.1 Definition of Firms

We begin in appendix Figure A.2 to show, using both datasets, what has been shown many times in manufacturing establishment microdata (Bartelsman et al., 2013; Bernard et al., 2007): firm sizes (measured in terms of sales) differ substantially within disaggregated product groups. In this and the subsequent figures and tables, we define a firm as a producer of a brand within product modules in the Nielsen data. This leads to an average of about 150 active firms within a given product module. Two possible alternatives would be to define a firm as a barcode product (leading to on average 700 firms per module), or as a holding company (leading to on average less than 40 firms per module). To fix ideas, an example for the product module Shampoo would be the barcode product Ultimate Hydration Shampoo (22 oz bottle) belonging to brand TRESemmé that

⁶Per capita expenditure can be misleading due to non-linearities in per capita outlays with respect to household size (e.g. Subramanian & Deaton (1996)). To address this concern, we non-parametrically adjust for household size by first regressing log total expenditure on dummies for each household size with a household size of 1 being the reference category and a full set of household socio-economic controls. We then deflate observed household total expenditure to per capita equivalent expenditure by subtracting the point estimate of the household size dummy (which is non-zero and positive for all households with more than one member).

is owned by the holding company Unilever.

We choose the definition of firms as brands within product modules for two main reasons, and then check the robustness of our findings to alternative definitions. First, our objective is to define a producer within a given module as closely as possible to an establishment in commonly used manufacturing microdata, following the literature on firm heterogeneity (Melitz, 2003). The definition of firms as holding companies (e.g. Procter&Gamble) would be problematic as these conglomerates operate across hundreds of brands produced in different establishments. The definition of firms at the barcode level would be problematic for the opposite reason, because the same establishment produces different pack sizes of the same product that are marked by different barcodes. In this light, defining producers of brands within disaggregated product modules as firms is likely the closest equivalent to observing several different establishments operating in the same disaggregated product group. Second, our theoretical framework features endogenous investments in product quality across firms, and it is at the level of brands within product groups that these decisions appear to be most plausible.⁷

Figure A.2 also points to an interesting difference between the home and store scanner datasets: the distribution of national market shares measured using the home scanner data (on average 59,000 households per half year) appears to be compressed relative to that measured using the store scanner data (on average 27,000 thousand stores per half year). This compression is stronger before applying the Nielsen household sampling weights, but still clearly visible after applying the weights. It also holds when restricting attention to producers of brands observed in both datasets. In the following, we report the main new stylized fact using the firm size distributions computed from both datasets. Given that the retail scanner data capture more than 1000 times the amount of transactions compared to the home scanner data, we then choose the store scanner data as our preferred measure of brand-level national market shares.

3.2 Firm Heterogeneity Across Consumption Baskets

Figure 1 depicts the main stylized fact of the paper. Pooling repeated cross-sections across 18 half-year periods, we depict percentiles of household per capita expenditure on the x-axis and weighted average deviations of log firm sales from the product module-by-half-year means on the y-axis.⁸ The weights correspond to each household’s retail consumption shares across all brands in all product modules consumed during the six-month period. When collapsed to five per capita expenditure quintiles on the right panel of Figure 1, we find that the richest 20 percent of US households source their consumption from on average 20 percent larger producers of brands within disaggregated product modules compared to the poorest 20 percent. These figures are our preferred measure of the national firm size distribution using the store scanner data. But, as the figure shows, a very similar relationship holds when using the firm size distribution from

⁷To corroborate this, we confirm in the data that 95 percent of variation in the average unit value paid for barcodes within product modules is accounted for by brand-by-pack-size fixed effects. Similarly, on average 80 percent of the sum of absolute differences in expenditure shares between the richest and poorest household quintiles across all barcodes within product modules are accounted for by differences in budget shares across brands (leaving 20 percent to be explained by differences within brands).

⁸To avoid measurement error from exiting or entering households in the consumer panel, we restrict attention to households for each half-year period that we observe to make purchases in both the first and the final month of the half year.

the home scanner data instead. This relationship is monotonic across the income distribution, and the firm size difference increases to 27 percent when comparing the richest and poorest 10 percent of households. As discussed above, Figure 1 is also robust to alternative definitions of firms in the Nielsen data: appendix Figure A.3 shows close to identical results when defining firms instead as holding companies within product groups. In Figure A.5, we also show an alternative representation plotting sales shares from different quintiles of the income distribution on the y-axis across firm sizes on the x-axis.⁹

What types of shopping decisions are driving these differences in weighted average firm sizes across the income distribution? In appendix Table A.2, we present the brands with the most positive and most negative differences in consumption shares between rich and poor household quintiles across three popular product modules for each of the eight product departments in our consumption microdata. Alongside the two brand names, we also list the difference in their log average unit values (price per physical unit) as well as the difference in their national market shares within that product module. Two features stand out. First, the brand that is most disproportionately consumed by the rich has a higher unit value and a larger market share relative to the brand that is most disproportionately consumed by the poor.¹⁰ Second, looking at the brand names it appears to be the case that richer households have a tendency to consume from the leading premium brands whereas the poorest quintile of households have a higher tendency to pick either generic store brands, or cheaper second and third-tier brands in the product group (e.g. Tropicana vs generic OJ, Pepsi vs generic Cola, Duracell vs Rayovac, Tide vs Purex, Dove vs Dial, Heinz vs Hunt’s).

Finally, we investigate whether these observed differences in product choices are driven by a fundamental disagreement about relative product appeal across rich and poor households. Do we see rich households consuming a large share of their expenditure from the largest producers while poor households spend close to none of their budget on those same producers? Or do households from different income groups agree on their relative evaluations of value-for-money across producers, such that the rank order of their budget shares is preserved across the income distribution? Appendix Figure A.6 documents that the latter appears to be the case in the data. The figure depicts—separately for each income group—non-parametric estimates of the relationship between income group-specific deviations in log expenditures across brands within product modules on the y-axis and deviations in log total market sales of those same producers in the store scanner data on the x-axis. The fact that expenditure shares within each income group increase monotonically with total firm size suggests that households on average agree on their evaluation of product quality attributes given prices, and that the rank order of budget shares across producers is preserved to a striking extent across all income groups. To express this in a single statistic, we find that the rank order correlation between the richest income quintile and the poorest for rankings of brand market shares within product modules is .89 when pooled across all product modules in the data. However, it is also apparent from the difference in slopes depicted in Figure A.6 that while all households spend most of their budget on the largest firms within product modules, richer households spend relatively more of their total budget on these largest producers relative to poorer households.

⁹Figure A.4 also shows the relationship in Figure 1 after computing firm sizes in terms of quantities (units sold).

¹⁰The scanner data allow the comparison of prices per unit (unit values) for identically measured units across products within product module (e.g. liters of milk, units of microwaves, grams of cereal, etc).

Related to this, we also investigate the role of the extensive margin in product choice underlying Figure 1. As depicted in Figure A.7, we find that while the average number of both UPCs and brands consumed within a module increase with incomes, the fraction of total retail expenditure accounted for by products consumed across all 5 income groups is very large (> 95 percent) and close to constant across the income distribution. Consistent with the rank order correlations above, this suggests that the extensive margin, while visible in the data, is unlikely to play a significant role in accounting for differences in average firm sizes across groups of rich and poor households.¹¹

3.3 Interpretation and Alternative Explanations

One natural interpretation of these stylized facts is that they arise as equilibrium outcomes in a setting where both heterogeneous firms and households choose the product attributes they produce or consume. However, there are a number of alternative and somewhat more mechanical explanations that we explore using the microdata before moving on to the model. Another question is how representative these findings are for consumption categories that are not well represented in the Nielsen data.

Representativeness and Data-Driven Explanations One concern is that the relationship documented in Figure 1 could in part be driven by shortcomings of the data. First, we explore the representativeness of the stylized fact in Figure 1. To this end, we estimate the relationship separately for each of the 18 half-year periods and for each of eight broad product categories in the Nielsen data: beverages, dairy products, dry grocery, frozen foods, general merchandise, non-food grocery, health and beauty, and packaged meat.¹² As depicted in appendix Figure A.9, we find that the pattern of firm size differences across consumption baskets holds across these different product segments and is not driven by one particular type of consumer products. We also find that the stylized fact in Figure 1 holds in each of the 18 half-year periods in our data.

As reported by e.g. Broda & Weinstein (2010) and more recently in Jaravel (2019), the product groups covered in the Nielsen data can be matched to product groups in the US CEX expenditure surveys that account for about 40 percent of goods consumption, which in turn accounts for about one third of total household outlays. To further explore to what extent the stylized fact in Figure 1 holds among consumption categories not fully represented in the Nielsen data, such as durables consumption, health expenditures or digital media, we use parts of these categories that are covered in the data. In particular, appendix Figure A.10 estimates Figure 1 separately for household appliances (e.g. air conditioners, refrigerators, kitchen appliances, desktop printers), pharmaceuticals purchases and audio, video and software purchases. Reassuringly, we find that the relationship holds similarly or is more pronounced.¹³

¹¹In support of this, Figure A.8 confirms that we find close to identical results after limiting the product space to only products consumed by all income quintiles in a given period (or across all periods).

¹²We combine observations for alcoholic and non-alcoholic beverages as one department in these graphs. Our reported findings above hold separately for both of these departments. We pool them here to be consistent with Section 5, where having one combined group for Beverages addresses sparsity in the parameter estimation.

¹³The Nielsen data unfortunately do not cover services consumption. For future research, newly available credit card microdata could potentially be used to investigate the relationship between card holder incomes and the firm size of service providers they source their consumption from.

Second, it could be the case that generic store brands are produced by the same (large) producers and sold under different labels across retail chains. If poorer households source more of their consumption from generics, then we could under-estimate their weighted average producer size due to this labeling issue. To address this concern, appendix Figure A.11 re-estimates Figure 1 after restricting consumption to sum to 100% for all non-generic product consumption for each household. We find close to identical results compared to the baseline in Figure 1.

Third, it could be the case that we are missing systematically different shares of retail consumption across rich and poor households due to the exclusion of products sold by retail chains that are not participating in the store-level retail scanner data (that we use to compute national market sales across producers). For example, it could be the case that richer households purchase a larger share of their retail consumption from independent boutique retailers selling small specialized brands, in which case we would be over-stating the weighted average firm sizes among richer households. Conversely, it could be true that richer households spend on average more on retail chains (compared to e.g. corner stores). To address such concerns, we make use of the fact that the home scanner data do not restrict reporting to participating retail chains in the store scanner data.¹⁴ In appendix Figure A.12, we then re-estimate Figure 1 after only including households for which we observe more than 90, 95 or 97.5 percent of their total reported retail expenditure on brands that are matched across both data sets, and find that the relationship holds close to identically in all specifications. Figure A.13 further corroborates this finding by showing that the fraction of total expenditure on brands in both datasets differs by less than 3 percent across the income groups.¹⁵

Another data-related concern is that the Nielsen data do not allow us to observe firm sales outside the US market. For both US-based exporters and imported brands, we are thus mismeasuring total firm sales relative to domestic-only US producers. Given existing evidence on the selection of firms into trade, it is likely that the resulting measurement error in firm sizes is positively related to the observed US market shares in the Nielsen data (understating true differences in firm sizes). To corroborate this, we report differences in weighted-average firm sizes across incomes separately for product groups both below and above-median import penetration or export shares.¹⁶ As depicted in appendix Figures A.14 and A.15, we find that the differences in firm sizes across rich and poor households are indeed slightly more pronounced in the below-median sectors for both import or export shares.

Differences in Supply and Pricing Across Incomes Another explanation could be that rich and poor households live in geographically segmented markets and/or shop across segmented store formats, so that differential access to producers, rather than heterogeneous household preferences, could be driving the results. In appendix Figure A.16 we explore to what extent differences in household geographical location as well as differences in retail formats within locations play in accounting for Figure 1. We first re-estimate the same relationship after conditioning on county-by-

¹⁴Since retail chain participation in the store scanner data is not made public by Nielsen, home scanner participants are not made aware of this either.

¹⁵Further reassurance against the “missing retailers” concern is also apparent in Figure 1 that depicts very similar patterns when using 100 percent of household retail consumption as reported in the home scanner data.

¹⁶To this end, we match the Nielsen product groups to 4-digit SIC codes in 2005 US trade data. See appendix Table A.3. Below median is equivalent to less than 10 percent for both import penetration and export shares.

half-year fixed effects when plotting the firm size deviations on the y-axis. Second, we additionally condition on individual household consumption shares across 79 different retail store formats (e.g. supermarkets, price clubs, convenience stores, pharmacies, liquor stores).¹⁷ We find a very similar relationship compared to Figure 1 in both cases, suggesting that differential access to producers is unlikely to be the driver.

Related to differential access to producers, differential pricing across income groups could be another potential explanation. For this to matter in our context, it would have to be the case that larger and smaller producers differ in their extent of price discrimination across income groups, such that larger firms become relatively more attractive for richer households. To explore this in the data, appendix Figure A.17 plots the relationship between unit values and firm sizes separately for prices paid by the poorest and the richest income quintiles. If it were true that larger firms offer differentially lower or higher prices compared to smaller firms across high and low-income groups, we would expect the slope of these relationships to differ between the income groups. Using unit values and firm sizes defined either at the UPC or at the brand-level, we find that this relationship is close to identical for both rich and poor consumers, providing some reassurance that richer households do not pick larger firms because of differential pricing.

Fixed Product Attributes Finally, we explore the notion that large firms are large because they sell to richer households. If firms were born with fixed product attributes and/or brand perceptions, and some firms got lucky to appeal to the rich, while other producers cannot respond over time by altering their own product attributes or brand perceptions, this would mechanically lead to richer households sourcing from larger firms (as the rich account for a larger share of total sales).¹⁸

Here, we document that in the medium or long run this notion seems hard to reconcile with either the data or the existing literature on endogenous quality choice by firms. First, a body of empirical work has documented that firms endogenously choose their product attributes as a function of market demand in a variety of different empirical settings (e.g. Bastos et al. (2018), Dingel (2016)).¹⁹ Second, the scanner data suggest that producers of brands frequently alter the physical characteristics and/or presentation of their products over time. Appendix Table A.4 documents that each half year close to 10 percent of producers of brands replace their products with changed product characteristics (e.g. packaging or product improvements) that have the identical pack sizes to the previous replaced varieties on offer by the same brand—suggesting that producers are indeed capable of choosing their product attributes as a function of market conditions.²⁰ In

¹⁷We condition on 79 store formats within the same county to capture potential differences in access across inner-city vs. suburbs or due to car ownership. Conditioning on individual stores would give rise to the concern that households choose to shop at different retailers precisely due to the product mix on offer, rather than capturing differences in access.

¹⁸This also relates to the original note in Melitz (2003) that the heterogeneity parameter can either be thought of as a marginal cost draw in a setting with horizontal differentiation, or as a quality draw in a setting with vertical differentiation.

¹⁹Another literature in support of this is the marketing literature on firm strategies using advertising to affect brand perceptions over time (e.g. Keller et al. (2011)).

²⁰It could still be the case that our 18 repeated cross-sections (half years) depicted in Figure 1 are partly capturing the result of short-term taste shocks across products that differ between rich and poor households while hitting a fixed number of producers with fixed product attributes. To further investigate this possibility, we re-estimate the relationship in Figure 1 after replacing contemporary differences in firm sales by either sales three years before or

support of these descriptive moments, we also provide more direct empirical evidence in Section 5 as part of our technology parameter estimation, documenting that an exogenous increase in the scale of production leads to brand-level quality upgrading over time (see Tables 2 and 3).

4 Theoretical Framework

This section develops a tractable quantitative GE model that rationalizes the observed moments in the microdata. We introduce two basic features into an otherwise standard Melitz model of heterogeneous firms. On the demand side, we allow for non-homothetic preferences so that consumers across the income distribution can differ in both their price elasticities and in their product quality evaluations. On the producer side, firms with different productivities face the observed distribution of consumer preferences and optimally choose their product attributes and markups. Reflecting the asymmetry of market shares within sectors we document above, we also depart from the assumption of monopolistic competition, and allow for variable markups under oligopolistic competition. The exposition below abstracts from time subscripts. Appendices 2-5 provide additional details on the model and its extensions.

4.1 Model Setup

Consumption The economy consists of two broad sectors: retail consumption (goods available in stores and supermarkets) and an outside sector. As in Handbury (2019), we consider a two-tier utility where the upper-tier depends on utility from retail shopping U_G and the consumption of an outside good z :

$$U = U(U_G(z), z) \quad (1)$$

For the sake of exposition, we do not explicitly specify the allocation of expenditures in retail vs. non-retail items, but assume that the outside good is normal.²¹ We denote by $H(z)$ the cumulative distribution of z across households in equilibrium, and normalize to one the population of consumers. Utility from retail consumption is then defined by:

$$U_G(z) = \prod_n \left[\sum_{i \in G_n} (q_{ni} \varphi_{ni}(z))^{\frac{\sigma_n(z)-1}{\sigma_n(z)}} \right]^{\alpha_n(z) \cdot \frac{\sigma_n(z)}{\sigma_n(z)-1}} \quad (2)$$

where n refers to a product module in the Nielsen data and i refers to a brand producer within the product module.²² $\varphi_{ni}(z)$ is the perceived quality (product appeal) of brand i at income level z , and $\sigma_n(z)$ is the elasticity of substitution between brands in module n at income level z . Allowing demand parameters for retail consumption to be functions of outside good consumption

three years in the future of the current period. If the distribution of firm sizes was subject to significant temporary swings over time, then we would expect the two counterfactual relationships to slope quite differently from our baseline estimate in Figure 1. Instead, appendix Figure A.18 suggests that the estimated differences in producer sizes are practically identical.

²¹Appendix 3 provides an envelope theorem result that holding z fixed yields a first-order approximation of the compensating variation from retail price shocks under any arbitrary (unspecified) upper-tier utility function. In the counterfactual analysis in Section 6, we report results both holding z fixed and allowing for endogenous changes in z after specifying (1).

²²We show in Appendix 3.D that these preferences can be derived from the aggregation of discrete-choice preferences across many individual agents within group z .

z introduces non-homotheticities in a very flexible manner.²³ As we focus most of our attention on within-product module allocations, we model the choice over product modules with a Cobb-Douglas upper-tier, where $\alpha_n(z)$ refers to the fraction of expenditures spent on product module n at income level z (assuming $\sum_n \alpha_n(z) = 1$ for all z).²⁴ Comparing two goods i and j within the same module n , relative expenditures by consumers of income level z are then given by:

$$\log \frac{x_{ni}(z)}{x_{nj}(z)} = (\sigma_n(z) - 1) \left[\log \frac{\varphi_{ni}(z)}{\varphi_{nj}(z)} - \log \frac{p_{ni}}{p_{nj}} \right] \quad (3)$$

Motivated by the evidence discussed in Figure A.6 above, we let household quality evaluations $\log \varphi_{ni}(z)$ depend on an intrinsic quality term $\log \phi_{ni}$ associated with brand i and a multiplicative term $\gamma_n(z)$ depending on income level z :

$$\text{Intrinsic Quality Assumption:} \quad \log \varphi_{ni}(z) = \gamma_n(z) \log \phi_{ni} \quad (4)$$

With the normalization $\int_{\Omega_z} \gamma_n(z) dz = 1$ (where Ω_z is a set of z household types),²⁵ this intrinsic quality term also corresponds to the democratic average quality evaluation across households:

$$\log \phi_{ni} = \int_{\Omega_z} \log \varphi_{ni}(z) dz \quad (5)$$

In the empirical estimation below, we estimate perceived quality $\varphi_{ni}(z)$ separately for each income group to verify whether relative quality evaluations are indeed preserved across income levels before imposing the above restriction. Finally, the retail price index is income-specific and given by $P_G(z) = \prod_n P_n(z)^{\alpha_n(z)}$, where the price index $P_n(z)$ for each product module n is defined as:

$$P_n(z) = \left[\sum_{i \in G_n} p_{ni}^{1-\sigma_n(z)} \varphi_{ni}(z)^{\sigma_n(z)-1} \right]^{\frac{1}{1-\sigma_n(z)}}. \quad (6)$$

Production For each product group n , entrepreneurs draw their productivity a from a cumulative distribution $G_n(a)$ upon paying a sunk entry cost F_{nE} , as in Melitz (2003). For the remainder of this section, we index firms by a instead of i , since all relevant firm-level decisions are uniquely determined by firm productivity a . The timing of events is as follows. First, entrepreneurs pay the entry cost F_{nE} and discover their productivity a . Second, each entrepreneur decides at which level of quality to produce or exit. Third, production occurs and markets clear. Firms compete in

²³For instance, demand systems with a choke price can generate price elasticities that depend on income (Arkolakis et al., 2018), but offer significantly less flexibility in that relationship. Related work by Feenstra & Romalis (2014) features a single demand elasticity across incomes. Handbury (2019) derives conditions under which such preferences are rational and well-defined, and Appendix 3.C provides an equivalent demand structure, based on results in Fally (2018), in which non-homotheticities do not rely on outside good consumption z .

²⁴We abstract from within-brand product substitution by summing up sales across potentially multiple barcodes. Appendix 4.A presents an extension of our model to multi-product firms following Hottman et al. (2016). In our setting, we show that as long as the ratio of between- to within-brand elasticities of substitution does not significantly differ across income groups z , firm variation in within-brand product variety does not play a role for differences in firm sizes across z (Figure 1). In Section 5.1 we find minor differences in the cross-brand elasticities across income groups, and in Table A.5 we find no evidence of significant differences in the between-to-within ratios.

²⁵This normalization sets the simple mean of preference parameters $\gamma_n(z)$ equal to unity across a fixed set of household types z . As shown in Appendix 3.E, this normalization to unity across household types z is without loss of generality.

prices, allowing for oligopolistic interactions as in [Hottman et al. \(2016\)](#).²⁶

We normalize the cost of labor (wage w) to unity. There are two cost components: a variable and a fixed cost (in terms of labor). We allow for the possibility that both the marginal and the fixed cost of production are a function of output quality. The latter captures potential overhead costs such as design, R&D and marketing which do not directly depend on the quantities being produced but affect the quality of the product. In turn, variable costs depend on the level of quality of the production as well as the entrepreneur's productivity, as in [Melitz \(2003\)](#). Hence, the total cost associated with the production of a quantity q with quality ϕ and productivity a is:

$$c_n(\phi)q/a + f_n(\phi) + f_{0n} \quad (7)$$

where $f_n(\phi)$ is the part of fixed costs that directly depend on output quality. For tractability, we adopt a simple log-linear parameterization for fixed costs as in [Hallak & Sivadasan \(2013\)](#) and the second model variant of [Kugler & Verhoogen \(2012\)](#):

$$f_n(\phi) = b_n \beta_n \phi^{\frac{1}{\beta_n}} \quad (8)$$

For instance, fixed costs increase with quality ($\beta_n > 0$) if higher quality entails higher expenditures on research and development or marketing, for which the cost is not directly dependent of the scale of production. On the contrary, it could be the case that fixed costs are decreasing in quality if such fixed investments are most suitable for mass-producing low-quality output at lower marginal costs.²⁷ Similarly, we assume that variable costs depend log-linearly on quality, with parameter ξ_n to capture the elasticity of the cost increase to the level of output quality:²⁸

$$c_n(\phi) = \phi^{\xi_n} \quad (9)$$

As long as ξ_n is smaller than the minimum quality evaluation $\gamma_n(z)$, firms choose positive levels of quality in equilibrium, as we further discuss below.

4.2 Equilibrium and Counterfactuals

In equilibrium, consumers maximize their utility, expected profits upon entry equal the sunk entry cost, and firms choose their price, quality and quantity to maximize profits. There are two sources for variable markups across firms in our setting. They are determined by both their market power in oligopoly as well as the composition of the consumers that each firm sells to (i.e. the sales-weighted average price elasticity they face across consumer groups z). To see this, prices are given by:

$$p_n(a) = \frac{\phi(a)^{\xi_n} \mu_n(a)}{a} \quad (10)$$

²⁶We also consider a generalized version of Cournot quantity competition following [Atkeson and Burstein \(2008\)](#), whereby firms compete in quantity. This yields very similar results, both for the cross-section of implied markups as well as for the counterfactual analysis.

²⁷In Appendix 4.B, we also present a model extension allowing for heterogeneous fixed costs that may be correlated with productivity, as in [Hallak & Sivadasan \(2013\)](#).

²⁸There is no need for a constant term as it would be isomorphic to a common productivity shifter after redefining $G_n(a)$.

where μ_n is the markup over marginal cost:

$$\mu_n(a) \equiv \frac{p_n(a)}{c_n(a)} = 1 + \frac{\int_z x_n(z, a) dH(z)}{\int_z (\sigma_n(z) - 1) (1 - s_n(z, a)) x_n(z, a) dH(z)} \quad (11)$$

$x_n(z, a)$ denotes sales of firm with productivity a to consumers of income level z . $s_n(z, a)$ are market shares in product module n . Due to oligopolistic competition, this markup is larger than $\frac{\tilde{\sigma}_n(a)}{\tilde{\sigma}_n(a)-1}$, the markup under monopolistic competition, where $\tilde{\sigma}_n(a)$ is the weighted average elasticity of substitution across consumers that the firm sells to:

$$\tilde{\sigma}_n(a) = \frac{\int_z \sigma_n(z) x_n(z, a) dH(z)}{\int_z x_n(z, a) dH(z)}$$

The first-order condition in ϕ characterizes optimal quality $\phi_n(a)$ for firms with productivity a . Existence of a non-degenerate equilibrium choice for quality requires that we fit in either of two cases.²⁹ In the first case, higher quality entails higher fixed costs ($\beta_n > 0$ and $b_n > 0$) and the increase in marginal costs does not exceed consumer valuation for quality ($\xi_n < \gamma_n(z)$ for all z). In this case, large firms sort into producing higher quality. In the second case, marginal costs increase more strongly with higher quality output ($\xi_n > \gamma_n(z)$ for all z), and fixed costs are decreasing ($\beta_n < 0$ and $b_n < 0$). In this case, relatively small “boutique” producers sort into producing higher quality products while mass producers produce low-cost, low-quality products.

In both cases, optimal quality satisfies:

$$\phi_n(a) = \left(\frac{\tilde{\gamma}_n(a) - \xi_n}{b_n \mu_n(a)} \cdot X_n(a) \right)^{\beta_n} \quad (12)$$

where $X_n(a) = \int_z x(a, z) dH(z)$ denotes total sales of firm a in product module n and $\tilde{\gamma}_n(a)$ is the weighted average quality valuation $\gamma_n(z)$ for firm with productivity a , weighted by sales and price elasticities across its customer base:

$$\tilde{\gamma}_n(a) = \frac{\int_z \gamma_n(z) (\sigma_n(z) - 1) (1 - s_n(z, a)) x_n(z, a) dH(z)}{\int_z (\sigma_n(z) - 1) (1 - s_n(z, a)) x_n(z, a) dH(z)} \quad (13)$$

Optimal quality is determined by several forces that are apparent in equation (12). First, when $\beta_n > 0$, larger sales induce higher optimal quality, as reflected in the term $X_n(a)^{\beta_n}$. This is the scale effect due to the higher fixed costs of producing at higher quality. If we compare two firms with the same customer base, the larger one would more profitably invest in upgrading quality if $\beta_n > 0$. Second, optimal quality depends on how much the firm-specific customer base value quality, captured by $\tilde{\gamma}_n(a)$. Firms that tend to sell to consumers with high $\gamma_n(z)$ also tend to have higher returns to quality upgrading. Third, optimal quality depends on technology and the cost structure. A higher elasticity of marginal costs to quality, ξ_n , induces lower optimal quality. If instead $\beta_n < 0$, higher quality is associated with smaller firm scale $X_n(a)$, a higher average quality valuation $\tilde{\gamma}_n(a)$ and a lower elasticity of marginal costs to quality ξ_n .

²⁹If $\beta_n > 0$, $b_n > 0$ and $\xi_n > \gamma_n(z)$ optimal quality is zero. If $\beta_n < 0$, $b_n < 0$ and $\xi_n < \gamma_n(z)$ optimal quality is infinite.

Finally, when firms chose prices and quality to maximize profits, those profits are given by:

$$\pi_n(a) = \left(1 - \frac{1}{\mu_n(a)}\right) \left[\int_z (1 - \beta_n (\gamma_n(z) - \xi_n)(\sigma_n(z) - 1)(1 - s_n(z, a))) x_n(a, z) dH(z) \right] - f_{0n} \quad (14)$$

where $\beta_n (\gamma_n(z) - \xi_n)(\sigma_n(z) - 1)(1 - s_n(z, a))$ is the share of operating profits that are invested in the fixed costs of quality upgrading.³⁰ In equilibrium, average profits (assuming equality between averages and expectations) across all entrants equal sunk entry costs, and surviving firms are those with positive profits.

An equilibrium is defined as a set of sales, quality choices and markups that satisfy equations (3), (10), (12) and (11) for each firm, such that price indices are given by (6) and such that profits are given by (14). Additional details are provided in Appendix 2.³¹

Firm Heterogeneity across Consumption Baskets To rationalize the observed stylized facts through the lens of the model, we examine the weighted average of log firm size $X_n(a)$ for each income group z , which corresponds to what we plot on the y-axis of Figure 1:

$$\log \widetilde{X}_n(z) = \frac{\int_a x_n(z, a) \log X_n(a) dG_n(a)}{\int_a x_n(z, a) dG_n(a)}$$

How $\widetilde{X}_n(z)$ varies with income (i.e. the slope of the estimated relationship in Figure 1) reflects how $x_{ni}(z, a)$ varies across firms i and consumer income z . For the sake of exposition, let us assume for now that quality valuation $\gamma_n(z)$ and price elasticities $\sigma_n(z)$ are continuous and differentiable w.r.t income z . We can then express the derivative $\frac{\partial \log \widetilde{X}_n(z)}{\partial z}$ as a function of two covariance terms (where Cov_z denotes a covariance weighted by sales to consumers z):

$$\begin{aligned} \frac{\partial \log \widetilde{X}_n(z)}{\partial z} &= \frac{\partial \gamma_n(z)}{\partial z} (\sigma_n(z) - 1) Cov_z(\log X_n(a), \log \phi_n(a)) \\ &\quad - \frac{\partial \sigma_n(z)}{\partial z} Cov_z(\log X_n(a), \log(p_n(a)/\phi_n(a))^{\gamma_n(z)}) \end{aligned} \quad (15)$$

From this expression, we see that the difference in weighted-average firm size in consumption baskets across the income distribution is driven by how preference parameters depend on income ($\frac{\partial \gamma_n}{\partial z}$ and $\frac{\partial \sigma_n}{\partial z}$), and by how firm size correlates with quality and quality-adjusted prices. The first line in equation (15) reflects a quality channel. It is positive if firm size increases with quality and if richer households care relatively more about intrinsic product quality ($\frac{\partial \gamma_n}{\partial z} > 0$). The second term captures a price effect, which would work in the same direction as the quality channel if, and only if, richer households were more price elastic compared to poorer households, as the final covariance term between firm size and quality-adjusted prices is negative (lower quality-adjusted prices lead to larger sales when $\sigma_n(z) > 1$). If, instead, higher income consumers were less price elastic, but attached greater value to product quality, the two channels in (15) would be opposing one another

³⁰This term must be smaller than unity. Second-order conditions require $\beta_n (\gamma_n(z) - \xi_n)(\sigma_n(z) - 1) < 1$ to ensure a well-defined equilibrium for all firms.

³¹Appendix 2 also examines second-order conditions, discusses conditions for uniqueness, and describes a special case with closed-form solutions.

in generating the observed heterogeneity in firm sizes across income groups in Figure 1.

The decomposition in equation (15) relies primarily on our demand-side structure and does not yet impose assumptions on the production side. In turn, the supply-side structure can shed light on the potential sources of the covariance terms. Prices are given by equation (10) while equilibrium product quality satisfies equation (12). In particular, the correlation between firm size and quality appearing in the first term can be expressed as:

$$Cov_z(\log X_n(a), \log \phi_n(a)) = \beta_n Var_z(\log X_n(a)) + \beta_n Cov_z(\log X_n(a), \log((\tilde{\gamma}_n(a) - \xi_n)/\mu_n(a))) \quad (16)$$

As part of the estimation that follows, we can quantify each of these terms and decompose the observed firm heterogeneity across the consumption baskets in Figure 1 into the underlying channels.

Counterfactual Analysis Our framework naturally lends itself to quantifying GE counterfactuals. In Appendix 5.A, we write the equilibrium equations in terms of counterfactual changes that govern adjustments in firm sales, quality, variable markups, entry, exit and price indices. Solving for counterfactual equilibria requires data on initial sales $x_{n0}(z, a)$ for each firm across different consumer groups, in addition to estimates of demand and supply parameters: $\sigma_n(z)$, $\gamma_n(z)$, β_n and ξ_n . With these moments in hand, we can solve for changes in quality $\frac{\phi_{n1}(a)}{\phi_{n0}(a)}$, sales $\frac{x_{n1}(z, a)}{x_{n0}(z, a)}$, markups $\frac{\mu_{n1}(a)}{\mu_{n0}(a)}$, the mass of firms $\frac{N_{n1}}{N_{n0}}$, firm survival $\delta_{nD}(a)$ and consumer price indices $\frac{P_{n1}(z)}{P_{n0}(z)}$. In particular, equilibrium changes in quality can be derived by taking ratios of equation (12), changes in sales are derived from equations (3) and (10), changes in markups from equation (11), changes in profits from equation (14), and changes in cost of living from equation (6). As described in Appendix 5.A, we do not require estimates of firm productivity a or initial firm quality $\phi(a)$ to conduct our counterfactual exercise.³² Appendix 5.B also derives a six-fold decomposition of the counterfactual price index effects that we use and further discuss as part of the counterfactual analysis in Section 6.

In our baseline counterfactual analysis, we hold initial household types z fixed, without specifying the upper-tier utility function in (1).³³ As shown in Appendix 3, this provides a first-order approximation of the compensating variation due to retail price changes for any arbitrary upper-tier utility. To investigate potential second-order effects on price indices (through changes in demand parameters $\alpha_n(z)$, $\sigma_n(z)$ and $\gamma_n(z)$), we then also examine counterfactuals allowing for endogenous changes in z after specifying the upper-tier in (1).

5 Estimation

This section presents parameter estimation and model calibration. To bring the model to the data, we define five consumer income groups z in terms of quintiles of the US income distribution, and introduce time subscripts to reflect 18 half-year cross-sections of data. We begin by estimating the preference parameters, σ_{nz} and γ_{nz} , that combined with detailed data on sales and unit values,

³²This approach follows Dekle et al. (2007) among others.

³³This is similar to e.g. Handbury (2019), Atkin et al. (2018) and Redding & Weinstein (2020) and follows earlier work by McFadden & Train (2000).

allow us to quantify the distribution of product quality, quality-adjusted prices and markups across producers of brands and household consumption baskets. With these estimates in hand, we then proceed to estimate the technology parameters, β_n and ξ_n . As well as being of interest in their own right, these estimates allow us to quantify the forces underlying Figure 1 at the end of this section, and to explore policy counterfactuals in the final section.

5.1 Price Elasticities

We begin by estimating the demand elasticity ($1 - \sigma_{nz}$) that we allow to vary across household income groups and product groups. From equation (3) we get the following estimation equation:

$$\Delta \log (s_{nzict}) = (1 - \sigma_{nz}) \Delta \log (p_{nict}) + \eta_{nzct} + \epsilon_{nzict} \quad (17)$$

where as before z , n and i denote household groups, product modules and brands. c and t indicate US counties and 18 half years (17 changes), and s_{nzict} are budget shares within product module n . η_{nzct} are household group-by-product module-by-county-by-half-year fixed effects that capture the local CES price index. Consistent with our specification of preferences at the level of household groups, we estimate (17) after aggregating consumption shares in the home scanner data for the period 2006-2014 to the level of household quintile-by-county-by-module-by-half-year bins.³⁴ To address concerns about autocorrelation in the error term ϵ_{nzict} , we cluster standard errors at the county level.³⁵

To address the standard simultaneity concern that taste shocks in the error term are correlated with observed price changes, we follow the empirical literature in industrial organization (e.g. Hausman (1999), Nevo (2000) and Hausman & Leibtag (2007)) and make the identifying assumption that consumer taste shocks are idiosyncratic across counties whereas supply-side cost shocks are correlated across space. For the supply-side variation needed to identify σ_{nz} , we exploit the fact that store chains frequently price nationally or regionally without taking into consideration changes in local demand conditions (DellaVigna & Gentzkow, 2019). In particular, we instrument for local consumer price changes across brands $\Delta \log (p_{nict})$ with either national or state-level leave-out mean price changes: $\frac{1}{N-1} \sum_{j \neq c} \Delta \log (p_{nijt})$. These two instruments identify potentially different local average treatment effects. The national leave-out means IV estimates σ_{nz} off retail chains that price their products nationally, whereas the state-level leave-out means may extend the complier group to regional and local retailers.

A potentially remaining concern that this IV strategy would not be able to address are demand shocks at the national or state-level that are correlated with observed product price changes. Advertisement campaigns would be a natural candidate for this concern. For this to lead to a downward bias in the σ_{nz} estimates, it would have to be the case that the advertisement campaign

³⁴We aggregate household purchases to income groups using sampling weights (“projection factor” in Nielsen) to compute $\Delta \log (s_{nzict})$, and limit the sample to income group-by-county-by-half-year cells with at least 25 households. To compute brand-level log price changes we first compute projection-factor-weighted price means for each barcode-by-county-by-half-year cell, and then compute $\Delta \log (p_{nict})$ as a brand-level Tornqvist price index across all barcodes belonging to the same brand. As reported in Appendix Table A.6, point estimates are not sensitive to either the decision to take mean prices (rather than medians) or the decision to take a Tornqvist price index (rather than Laspeyres or a simple average).

³⁵Clustering at this level yields slightly more conservative standard errors than sensible alternatives (clustering at the level of brands, product modules, county-by-income groups, county-by-half-years or county-by-product modules).

first affects demand, but then also leads to higher prices. We would argue that this is not likely to be the case for most national or state-level advertisement campaigns. For example, an “informative” advertisement campaign containing price information would not lead to a bias in our estimation of σ_{nz} , as the variation is driven by consumers reacting to a change in prices (promotions etc.). A second type of “persuasive” campaign could be aimed at improving the brand’s perception instead, which would be more problematic for the exogeneity of the IV. For identification, we require that it is not the case that firms on average launch persuasive advertisement campaigns and simultaneously increase their prices. Given the longer-term objective of most image-oriented advertisement campaigns (e.g. Keller et al. (2011)), and the fact that we use half-yearly variation in prices and consumption decisions in our estimations, we believe this to be a plausible baseline assumption.

To document sensitivity, we also report counterfactual results in Section 6 across alternative parameter assumptions. Finally, the key empirical moment in our welfare quantification does not rely on the levels of σ_{nz} , but on the observed heterogeneity across different income groups. And while it is possible that some of the discussed endogeneity concerns may affect rich and poor households differently, such concerns would require somewhat more elaborate stories compared to the traditional simultaneity bias in demand estimation.³⁶

Estimation Results Panel A of Table 1 shows the pooled estimation results across all household and product groups. In support of the IV strategy, we find that the point estimates change from slightly positive in the OLS specification to negative and statistically significant in both IV estimations as well as the joint IV column. The estimates from the two different instruments are very similar and suggest a pooled elasticity of substitution of about 2.2. Panel B presents IV estimates separately across product departments that range between 1.5 and 3.5. These estimates fall at the center of a large existing literature in IO and Marketing using brand-level consumption data to estimate the sales-to-price elasticity of demand.³⁷ They are, however, somewhat lower than empirical work based on moment conditions of the double-differenced residuals in demand and supply using the GMM estimation approach pioneered by Feenstra (1994) (e.g. Broda & Weinstein (2010), Hottman et al. (2016)).³⁸ As a robustness exercise, we report counterfactuals in the final section for both our baseline estimates and assuming larger values of the price elasticity.

In the final column of Panel A, we take the pooled sample, but interact the log price changes with household income group identifiers to estimate to what extent there are statistically significant differences between household quintiles. The most convincing way to estimate such household differences in σ_{nz} is to additionally include brand-by-period-by-county fixed effects, so that we identify differences in the elasticity of substitution by comparing how different households react to the identical price change. We choose the richest income group as the reference category absorbed by the additional fixed effects. We find that poorer households have significantly higher price

³⁶Also note that in the absence of additional linked household data on times spent shopping, it would be impossible to disentangle the potential for differences in the opportunity cost of search/shopping time across rich and poor households (e.g. (Aguiar & Hurst, 2007)).

³⁷See e.g. a meta-analysis by Bijmolt et al. (2005). Reviewing close to 2000 estimates, they find a median elasticity of -2.2, with one half of the estimates falling in the range between -1 and -3.

³⁸See also Soderbery (2015) and more recent work by Ray (2019) addressing potential upward bias due to many weak instruments in this setting.

elasticities compared to richer households. In terms of magnitude, however, these differences are relatively minor. We estimate that the price elasticity for the poorest two quintiles is about 0.4 larger than that for the richest quintile.

Panel C reports results for each of the product departments across two income groups: the bottom two quintiles and the top 3 quintiles. These sixteen σ_{nz} estimates reported in Panel C are the point estimates that we use as our baseline parameter values in the analysis that follows. This is motivated by the income group heterogeneity reported in the final column of Panel A, and due to the fact that statistical power starts to become an issue when estimating these parameters separately across individual product departments. The trade-off that we face here is one between relatively precisely estimated point estimates relative to allowing for richer patterns of heterogeneity. For completeness, appendix Table A.7 reports the results when estimating forty σ_{nz} parameters (5 across each of the 8 product departments). As becomes clear from that table, a larger number of parameters start having large standard errors and lack statistical significance compared to our preferred set of estimates in Panel C of Table 1. An alternative approach is to estimate the heterogeneity parametrically by interacting changes in log prices with either the average of log household total expenditure per capita in a given *zct* bin or the average percentile of total per capita expenditure in that bin. We report those results in appendix Table A.8, and use these alternative estimates as part of the sensitivity analysis in the counterfactual analysis below.

5.2 Brand Quality, Quality-Adjusted Prices and Markups

Armed with estimates of σ_{nz} , we can use the scanner data following equations (3) and (5) to estimate product quality evaluations $\log \varphi_{nzi}$, $\log \phi_{ni} = \frac{1}{N_z} \sum_z \log \varphi_{nzi}$ and quality-adjusted prices, $\log \left(\frac{p_{ni}}{\phi_{ni}} \right)$, across producers of brands and household consumption baskets. To do this, we use an additional empirical moment from the data, product unit values, in combination with observed product sales and the estimated σ_{nz} .

Appendix Figure A.19 depicts the distribution of mean deviations in log product unit values within product module-by-half-year cells across the income distribution (aggregated as expenditure-weighted averages for each household.)³⁹ As shown in the figure, the richest quintile of US households source their consumption from firms that have on average 12 percent higher unit values within product modules compared to the poorest quintile. In turn, Figure A.20 depicts the relationship between firm sizes and unit values.

The left panel of Figure 2 proceeds to present the distribution of the estimated weighted average product quality deviations across household consumption baskets. We find that the documented differences in terms of firm sizes translate into significant differences in the weighted average product quality as well as quality-adjusted prices embodied in consumption baskets across the income distribution. The richest 20 percent of US households source their consumption from on average 22 percent higher quality producers compared to the poorest 20 percent of households. Using the estimates of income-group-specific product quality shifters, appendix Figure A.21 confirms the motivating evidence from Figure A.6 above: rich and poor households agree that product quality increases across the firm size distribution, but this relationship is steeper among richer households.

³⁹We compute brand-level unit values as sales-weighted means across barcode transactions at the level of brands-by-half-year cells.

Moving from differences in product quality to quality-adjusted prices, the right panel of Figure 2 documents that the richest quintile source their consumption at on average 10 percent lower quality-adjusted prices.

The parameter estimates for σ_{nz} in combination with data on firm sales by income group also allow us to compute the distribution of the effective (weighted average) elasticities of substitution faced by individual producers, $\left(\tilde{\sigma}_{ni} = \frac{\sum_z \sigma_{nz} x_{nzi}}{\sum_z x_{nzi}}\right)$. Following equation (11), this together with firm market shares under oligopolistic competition determines the distribution of markups across firms. The left panel of Figure 3 presents the estimation results for variation in μ_{ni} as a function of firm size deviations within product modules. Larger firms charge significantly higher markups due to both their market power within product modules and the fact that they face lower price elasticities due to selling a higher share of their output to higher-income households (who, in turn, have lower parameter values for σ_{nz}).

With estimates of φ_{nzi} in hand, we proceed to estimate the final set of preference parameters, γ_{nz} , that govern the valuation of product quality characteristics across the household income distribution. Following expression (4), the estimation equation is:

$$\log(\varphi_{nzit}) = \gamma_{nz} \log(\phi_{nit}) + \eta_{nzt} + \epsilon_{nzit} \quad (18)$$

where η_{nzt} are income group-by-product module-by-half-year fixed effects. To address the concern of correlated measurement errors when moving from model to data, that appear both on the left hand side (the income group specific product quality evaluations) and the right hand side (the democratic average product quality evaluation), we instrument for $\log(\phi_{nit})$ with two half-year lagged values of product quality. To address autocorrelation in the error term ϵ_{nzit} , we cluster standard errors at the level of product modules.

Appendix Table A.9 presents the estimation results across bins of household groups and product departments. In accordance with the documented stylized facts, richer household groups value increases in product quality relatively more across each of the product departments. However, there we also find differences in the extent of this heterogeneity across different product departments. For example, beverages, dairy products and packaged meat are among the departments with the highest difference in the taste for quality, whereas general merchandise and health and beauty care have the lowest differences in household taste for quality across income groups.

As we do above for the firm-level parameter $\tilde{\sigma}_{ni}$, we can use these estimates in combination with the sales data to compute the weighted average product quality evaluations faced by each brand producer, following expression (13). The right panel of Figure 3 reports these results across the firm size distribution within product groups. We find that larger producers of brands face a composition of market demand with significantly higher marginal valuations for product quality.

5.3 Technology Parameters

In this subsection, we propose two approaches to estimate the technology parameters.

Estimation in the Cross-Section Armed with estimates of μ_{ni} and $\tilde{\gamma}_{ni}$, we proceed to estimate the technology parameters β_n and ξ_n : the first determines the presence and size of economies

of scale in the production of product quality. The second determines the extent to which marginal costs increase with higher product quality. An intuitive way to estimate β_n is through the relationship between unit values and market shares within product modules. If we imposed the assumptions of monopolistic competition and homogeneous consumer preferences (representative agent), we would get the following estimation equation from (3) and (12) above:

$$\log(p_{nit}) = \left(\beta_n - \frac{1}{\sigma_n - 1} \right) \log(X_{nit}) + \eta_{nt} + \epsilon_{nit} \quad (19)$$

where η_{nt} are product module-by-half-year fixed effects. Intuitively, if brands were of the same quality then the relationship between unit values (that would be identical to prices in this case) and market shares would be governed by the slope of the demand curve $-\frac{1}{\sigma_n - 1}$. The extent to which firms of larger scale sort into producing higher product quality would then be captured by the production function parameter β_n . To see this more clearly, we can re-write (19) with product quality on the left hand side: $\log(\phi_{nit}) = \beta_n \log(X_{nit}) + \eta_{nt} + \epsilon_{nit}$, where following (3) and (5) $\log(\phi_{nit}) = \log(p_{nit}) + \frac{1}{\sigma_n - 1} \log(X_{nit})$. This same logic and estimation equation have been used in the existing literature on quality choice across heterogeneous firms under the representative agent assumption (e.g. Kugler & Verhoogen (2012)).

When allowing for variable markups in oligopolistic competition and heterogeneity in tastes for quality and price elasticities across consumers—giving rise to firm-specific demand compositions for $\tilde{\gamma}_{ni}$ and $\tilde{\sigma}_{ni}$ —this estimation equation requires two additional correction terms. From (5) and (12) we get:

$$\begin{aligned} \log(p_{nit}) = & \left(\beta_n - \frac{1}{\bar{\sigma}_n - 1} \right) \log(X_{nit}) - \frac{1}{N_z} \sum_z \frac{1}{\sigma_{nz} - 1} \log\left(\frac{X_{nzit}}{X_{nit}}\right) \\ & + \beta_n \log((\tilde{\gamma}_{nit} - \xi_n) / \mu_{nit}) + \eta_{nt} + \epsilon_{nit} \end{aligned} \quad (20)$$

where N_z is the number of consumer groups (5 in our application), $\frac{1}{\bar{\sigma}_n - 1} = \frac{1}{N_z} \sum_z \frac{1}{\sigma_{nz} - 1}$. The first additional term on the right generalizes the downward-sloping demand relationship $\left(-\frac{1}{\sigma_n - 1} \log(X_{nit})\right)$ in equation (19), to allow for the fact that different producers face different market demand elasticities due to differences in the composition of their customers. The second additional term captures the fact that different producers may sort into higher or lower product quality due to variable markups and differences in the composition of their customer base (valuing quality more or less given prices).

As above, following (3) and (5), we can re-write equation (20) for estimation as: $\log(\phi_{nit}) = \beta_n \log(X_{nit} (\tilde{\gamma}_{nit} - \xi_n) / \mu_{nit}) + \eta_{nt} + \epsilon_{nit}$. Conditional on η_{nt} fixed effects, this allows us to jointly estimate the technology parameters β_n and ξ_n for each product department by estimating β_n using IV regressions across iterations of ξ_n , and selecting the best-fitting parameter combination. We use iterations of ξ_n in steps of 0.01 in the range between 0 and twice the maximum estimated γ_{nz} , and do not impose an ex ante assumption about the existence of economies of scale in quality production ($\beta_n > 0$).

Two identification concerns in (20) are correlated measurement errors on the left and right hand sides, and temporary consumer taste shocks: deviations around ϕ_{ni} over time that would

mechanically lead to a biased estimate $\beta_n = \frac{1}{\bar{\sigma}_n - 1}$ if unit values and firm quality (but not sales) remain unchanged in response to the temporary taste shock. To address both of these concerns, we instrument for markup and composition-adjusted firm scale $\log(X_{nit}(\tilde{\gamma}_{nit} - \xi_n)/\mu_{nit})$ with two half-year lags. To address concerns about autocorrelation in the error term, we cluster the standard errors at the level of product modules as before. A deeper concern with the cross-sectional estimation is that omitted factors, such as firm-specific quality upgrading costs discussed in Section 4 and Appendix 4.B, that affect both firm scale and product quality choices are difficult to rule out.

Finally, cross-sectional variation would be insufficient to distinguish between alternative micro-foundations to rationalize the observed sorting of larger firms into higher output quality. In our model we follow earlier work by e.g. Sutton (1998) allowing fixed costs to increase with output quality. Alternatively, e.g. Baldwin & Harrigan (2011) and Feenstra & Romalis (2014) model quality choices as a direct function of a firm’s productivity draw. While both model variants give rise to a cross-sectional relationship between firm scale and product quality, only the former would be consistent with changes in output quality as a function of exogenous changes in firm scale in firm-level panel data.

Panel Estimation Estimation equation (20) extends the existing literature on quality choice across firms to a setting that also allows for heterogeneity on the consumption side. But it also follows the existing literature in that it is based on cross-sectional variation across firms. An alternative approach is to use within-brand variation over time. The natural panel data approach would be to write (20) in log changes instead of log levels on both the left and right hand sides. However, the estimation of β_n would still likely be biased, even when assuming that changes in firm scale on the right-hand side were perfectly exogenous. To see this, imagine we helicopter-dropped a random sales shock onto a firm that does not adjust either product quality or prices: even though product attributes stay unchanged, we would mechanically conclude that there are economies of scale in quality production ($\widehat{\beta}_n = \frac{1}{\bar{\sigma}_n - 1} > 0$). The reason is that demand shocks that one would usually want to exploit as IV for firm sales to estimate economies of scale in production, would in our setting, holding firm prices and quality constant, be mechanically interpreted as an increase in product quality.

To address this concern, we propose the following panel estimation strategy. Re-writing expression (3) for state-level demand instead of national-level, and again substituting for product quality from the optimal quality choice equation (12), we get:

$$\begin{aligned} \Delta \log(p_{nist}) = & \beta_n \Delta \log(X_{nit}) - \frac{1}{N_z} \sum_z \frac{1}{\sigma_{nz} - 1} \Delta \log(X_{nizst}) \\ & + \beta_n \Delta \log((\tilde{\gamma}_{nit} - \xi_n)/\mu_{nit}) + \eta_{nst} + \epsilon_{nit} \end{aligned} \quad (21)$$

where subscript s indexes US states, η_{nst} are state-by-product module-by-half-year fixed effects, and Δ indicates a two-year change (4 changes in our database starting from the first half year in 2006 until the end of 2014). As before, the second term on the right captures the demand-side relationship between sales and product unit values conditional on product quality, but this time

at the state level. For instance, changes in firm productivity (and thus unit values on the left) conditional on product quality are captured by this term. The first and third terms capture the relationship between unit values and sales that is driven by changes in product quality. Following (12), firm changes in product quality are a function of national firm scale and the firm’s composition of consumer taste parameters (where $\tilde{\sigma}_{nit}$ is part of μ_{nit}).

The advantage of writing the estimation equation in terms of state-level unit values on the left is that a helicopter drop of sales on a brand producer in another region of the US does not lead to a mechanical bias in β_n , unlike in the example above. The reason is that unless the firm changes its product quality in response, shocks to firm scale in other states have no effect on local unit values. The estimation also does not confound conventional economies of scale in production with economies of scale in product quality: if marginal costs fell with larger scale—holding quality constant—, this would be accounted for by the conventional demand relationship (second term on the right) discussed above. As above, we can re-write (21) for estimation as: $\Delta \log(\phi_{nist}) = \beta_n \Delta \log(X_{nit}(\tilde{\gamma}_{nit} - \xi_n) / \mu_{nit}) + \eta_{nst} + \epsilon_{nit}$.

The first remaining identification concern in (21) is correlated measurement errors between the left and right hand sides. A second concern is that firm changes in national sales are partly driven by taste shocks that could be correlated across states, which—holding constant product quality and unit values but not sales—would bias the estimate of β_n . To exploit plausibly exogenous variation in shocks to firm-level scale (21), we use leave-out mean changes in log firm sales across other states ($s' \neq s$), and using other product modules ($n' \neq n$). We then construct a weighted average of these leave-out mean changes in log firm sales using each firm’s pre-existing share of total sales across different states.

This shift-share instrument for markup and composition-adjusted firm scale ($\Delta \log(X_{nit}(\tilde{\gamma}_{nit} - \xi_n) / \mu_{nit})$) is thus based on average changes in firm scale over time that exclude both the product group of the firm and the state in which $\Delta \log(\phi_{nist})$ on the left-hand side is observed. The identifying assumption of this strategy is that plausibly exogenous shocks to firm scale from other regions of the US do not affect changes in state-level brand quality through other channels but firm scale.

Estimation Results We start in Table 2 by presenting reduced-form estimation results of the relationship between unit values or product quality on the left hand side and national firm sales on the right hand side. The raw empirical moment that is most directly informative of the degree of quality sorting across firm sizes is the fact that product unit values increase with national brand sales. This holds for both the cross-section of firms and for within-firm changes over time. It also holds in both OLS and IV estimations. In the cross-section, the IV addresses concerns about correlated measurement errors between unit values and firm scale and temporary taste shocks that could drive both left and right hand sides. In the panel data estimation, we have two-year changes in state-level log unit values on the left hand side, and we instrument the right hand side using plausibly exogenous changes in national firm sales (computed using the shift-share instrument described above). The IV point estimate of this panel IV specification in column 6 of Table 2 suggests that a 10 percent increase in a firm’s national sales leads to a 0.7 percent increase in its unit value.

The same pattern of results holds when we replace unit values with our model-based measure of product quality on the left hand side. In both the cross-section and the within-brand estimation product quality increases with national firm scale, and again this holds before and after addressing identification concerns using our instruments. In the panel IV specification in column 8 we find that a 10 percent increase in national firm sales leads to a 5.7 percent change in brand quality.

Table 3 proceeds to the estimation of β_n and ξ_n . To estimate these, the main difference to the previous reduced-form table lies in the additional inclusion of brand-level consumer compositions and firm markups on the right-hand side, as shown above in equations (20) and (21). The first panel reports the results when pooling all product groups, and the IV point estimates of the best-fitting parameter combination of β_n and ξ_n are not far from the reduced form results reported in Table 2. The second panel reports the technology parameter estimates separately for grocery and non-grocery product groups, and Appendix Table A.10 reports the estimation results separately for each product department. An interesting pattern emerges from the parameter estimates: in both the cross-sectional specification and the panel data approach, the IV point estimates of the economies of scale parameter in quality production are significantly larger for non-grocery product groups (e.g. health & beauty, merchandise) compared to grocery product groups. As indicated by the first stage F-statistics in the appendix table, the panel data estimation does not have sufficient power to precisely estimate β_n and ξ_n separately for each product department. For this reason, we use the precisely estimated parameters for grocery and non-grocery product groups reported in Table 3 for the counterfactual quantification in the following section. Following the discussion of identification concerns in the cross-section above, we use the IV panel estimation as our preferred parameter estimates, and report counterfactuals using the cross-section IV estimates as part of the robustness checks.

5.4 Quantification of Forces

Armed with the preference and technology parameter estimates, we can check whether the calibrated model quantitatively replicates the main stylized fact documented in Figure 1. We can also use the calibrated model to quantify the forces underlying this observed relationship. Following expressions (15) and (16), we can decompose the observed differences in weighted average firm sizes across consumption baskets into different forces on the demand and supply sides. In Figure 4, we depict different calibrated distributions of weighted average firm sizes across the aggregate consumption baskets of the five income groups alongside the observed moments in the data. We do this 18 times for each of the half-year periods in our dataset, and plot the mean outcomes for both the actual and calibrated moments.

In the first calibration, we only make use of the first part of expression (15) to predict the consumption choices of rich and poor income groups in a model world where the only source of heterogeneity between them is that they are subject to different estimated demand elasticities. That is, we predict the consumption shares of rich and poor income groups within product groups taking the quality and quality-adjusted prices of brands as given on the supply side in the data, assigning all households the same average taste-for-quality parameters $\bar{\gamma}_n$, but making use of the observed differences in their $\sigma_n(z)$ estimates. As depicted in Figure 4, household heterogeneity in price elasticities would, *ceteris paribus*, push poor households to consume from significantly

larger firms compared to rich households—the opposite direction to what we observe in the data in Figure 1 across individual households and in Figure 4 across the aggregate demand of different income groups in the data.

In the next calibration, we predict household consumption shares across the 5 income groups after also taking into account the second source of heterogeneity on the consumption side in expression (15): the fact that rich and poor households are estimated to value product quality differently. Again, we take as given the product quality and quality-adjusted prices on the supply side across brands in the data, and predict income group-specific consumption shares that are now taking into account both heterogeneity in σ_{nz} and in γ_{nz} . As shown in Figure 4, the fact that higher-income households are estimated to have significantly stronger tastes for product quality pushes in the opposite direction to the heterogeneity due to price elasticities, and dominates that first effect. The sum of the two effects in expression (15) closely replicates the differences in firm sizes across income quintiles documented in Figure 1.

In the final calibration, we fully endogenize both product choices on the consumer side and product choices on the firm side. That is, rather than predicting the consumption shares of income groups within product groups conditional on the available mix of product quality and quality-adjusted prices on offer across producers, we first predict the product quality choice across the firm size distribution using the equilibrium expression (16), and then let consumers optimally allocate budget shares on the demand side based on these predicted firm product choices. The only raw moments we use in these calibrations from the data is the observed distribution of firm sales across income groups for each of the 18 half years that we combine with the structure of the model and the estimated parameters to make predictions about the equilibrium differences in firm sizes across consumption baskets.

In addition to quantifying the (opposing) forces underlying the observed stylized fact in Figure 1, this exercise is useful to validate to what extent the calibrated model can capture the observed moments in the data, before proceeding to the counterfactual quantifications in the following section. Reassuringly, as depicted in Figure 4, the calibrated model is able to closely replicate the observed differences in weighted-average firm sizes across the income distribution.

6 Counterfactuals

In this section, we use the calibrated model to quantify a new set of GE effects on household price indices in two policy counterfactuals. The first policy directly affects the distribution of disposable incomes on the demand side. The second policy directly affects profits across the firm size distribution on the supply side. In our framework, these policies can affect household price indices differently across the income distribution—both through direct effects and endogenous GE adjustments in firm scale, product quality choices, variable markups and exit/entry across the firm size distribution that affect rich and poor households differently. In the final section, we investigate the sensitivity of these counterfactual results across a number of alternative model and parameter assumptions. We also discuss the implications for the distribution of the gains from trade relative to a conventional Melitz (2003) model of heterogeneous firms.

6.1 Counterfactual 1: Progressive Income Tax Reform

One implication of the stylized fact in Figure 1 and the model we use to rationalize it in Section 4 is that changes to the distribution of disposable nominal incomes lead to GE effects on consumer price indices that tend to amplify the observed change in nominal inequality. To make this theoretical result concrete and put numbers on the mechanisms, we evaluate the implications of increasing the effective tax rate on the richest household group in our calibration (incomes above the 80th percentile) by 20 percentage points (from currently around 30 to 50 percent).

Through the lens of our model, this counterfactual allows us to relate to current policy debates in three main respects. First, this policy captures the historical change in US effective rates on incomes of the top 20 percent moving back to previous levels in the 1950s and 60s (Saez & Zucman, 2019).⁴⁰ Second, this policy closely corresponds to the counterfactual of moving the US from the current effective rate on the top 20 percent to the average effective rate on this group among Northern European countries.⁴¹ Third, this policy change is also in line with the proposed tax reforms by two leading presidential candidates for the 2020 US elections (Sanders and Warren).⁴²

The direct effect of this policy is a compression of the distribution of disposable nominal incomes. In line with the motivation behind this policy, we redistribute the tax revenues to the poorest 20 percent of US households in our baseline counterfactual. Alternatively, we report results without redistributing the tax revenues to household disposable incomes. We solve for counterfactual changes 18 times, based on the observed brand sales to the five income groups for each half-year cross-section in the scanner data and our estimates for the parameters $\sigma_n(z)$, $\gamma_n(z)$, β_n and ξ_n .⁴³

We solve for the counterfactual equilibrium as described in Section 4.2, and decompose the mechanisms as derived in Appendix 5. To compute confidence intervals that account for both sampling variation in the sales data across the 18 cross-sections and in the parameter estimates, we bootstrap the quantification 200 times for each half year of data. In each bootstrap, we draw the parameters $\sigma_n(z)$, $\gamma_n(z)$, β_n and ξ_n from a normal distribution with a mean equal to the point estimate and a standard deviation equal to the standard error of the estimate.⁴⁴

Figure 5 presents the counterfactual results. The left panel depicts the difference in the retail price index effect of the policy reform across the five US income groups. As a result of the GE forces, we find that progressive income taxes gives rise to a meaningful amplification of the policy’s direct

⁴⁰While effective rates vary at a more granular level than the income quintiles we calibrate the model to, the counterfactual captures the historical change in effective rates on the sum of incomes above the 80th percentile (the income-weighted average change). Given both tax changes and taste for quality increase with incomes within this group, the results provide conservative estimates in this respect.

⁴¹The average effective rate on incomes above the 80th percentile reported in the Luxembourg Income Study for Austria, Denmark, Finland, Germany, the Netherlands, Norway and Sweden is 52 percent.

⁴²See the tax sections of their campaign webpages (<https://berniesanders.com>, <https://elizabethwarren.com>) and effective tax calculator of their proposals at TaxJusticeNow.org. The Sanders proposal is close to the 20 percentage point increase above. The Warren proposal is closer to a 15 percentage point increase. We do not count health reform as part of taxation here, which would reinforce the progressivity.

⁴³For each of the 18 half-year periods, we verify that the moments in the data and estimated parameter values satisfy the uniqueness conditions discussed in Section 4.2, Appendix 2 and Appendix 5. Our fitted model cannot perfectly match sales and quality valuations for each brand and each quintile. We can, however, perfectly fit the data by adding a multiplicative adjustment term ϵ_{niz} to quality valuations specific to each income quintile in expression (4). We then obtain the exact same counterfactual equations for *changes* in quality, sales and price indices as long as we hold these adjustment terms constant.

⁴⁴This is a parametric bootstrap (Horowitz, 2001). See e.g. Atkin et al. (2018) for a recent application.

effect, increasing its progressivity through price index effects. In particular, the bottom income quintile experience a slightly more than 3 percentage point lower cost of living inflation for retail consumption compared to the richest quintile (point estimate 3.14 with bootstrapped standard error of 0.63).

Table A.11 presents the decomposition of the inflation difference between the richest and the poorest quintiles. The main drivers are related to both average and differential changes in firm scale and product quality across the distribution of initial firm sizes. As the policy leads to a reduction in the share of total sales to richer households in the economy, who value product quality the most, firms on average have incentives to downgrade product quality so that initial sales-weighted average product quality decreases. Since rich households value quality more, this average effect increases retail cost of living inflation among the rich relative to lower incomes, accounting for about 10 percent of the total difference in inflation.

The main effect is that initial sales to rich households are particularly concentrated among initially large firms producing at higher product quality. The tax reform thus also leads to a compression of the firm size distribution. With economies of scale in quality production ($\beta_n > 0$), this translates into asymmetric effects on quality and quality-adjusted prices across the consumption baskets of rich and poor households. The right panel of 5 depicts this asymmetric effect on quality downgrading across the initial firm size distribution within product groups (with positive changes indicating reductions in product quality on the y-axis). On average, firms above the median size within product groups downgrade their product quality and vice-versa for firms below the median size. As shown in Table A.11, this effect accounts for about 80 percent of the overall inflation differential.

There is also a smaller effect operating through firm exit and love of variety. Since richer households have slightly higher estimated elasticities of substitution, firm exit increases their retail price index relative to poorer consumers (accounting for 6 percent of the overall difference). Figure A.23 also breaks up the price index effects across the different departments in the scanner data. In appendix Figure A.22 we quantify the implications of the same policy change, but without redistribution of the tax proceeds to household disposable incomes. Reassuringly, we find that the effect on household inflation differences is quite similar, with about a 2.5 percentage point inflation differential between the richest and poorest quintiles. In Section A.22, we tend to find similar or slightly stronger results when revisiting the counterfactual across a number of alternative model or parameter assumptions.

6.2 Counterfactual 2: Closing Loopholes in Corporate Taxation

A second implication of our framework is that regulations that affect the profits of large and small firms differently give rise to differential price index effects (even in absence of GE adjustments), as these firms enter consumption baskets across the income distribution asymmetrically. To make this theoretical result concrete and put numbers on the mechanisms, we relate to the ongoing debate about closing loopholes in corporate taxation. A growing literature in public finance has documented larger possibilities for tax evasion among large US corporations (e.g. Guvenen et al. (2017), Wright & Zucman (2018)). A study by Bao & Romeo (2013) documents that the relationship between firm sales and effective corporate tax rates in the US—which is monotonically

and smoothly increasing up to the 95th percentile of firm sales—shows a sharp kink with a switch in the sign of the slope for the largest 5 percent of producers.

We use these findings to evaluate the impact of eliminating the kink in the otherwise smooth relationship between effective tax rates and firm size.⁴⁵ This policy change would lead to an increase of on average 5 percent in corporate taxes paid by the largest 5 percent of producers. This increase ranges from on average 1 percent at the 95th percentile to 11 percent at the 99th percentile of firm sizes. In our baseline counterfactual we follow the definition of firms as brands from Section 3 to define tax entities. Alternatively, we report results after using holding companies as the unit of taxation.

Figure 6 presents the counterfactual results. The left panel depicts the difference in the retail price index effect of the policy across the five US income groups. We find that even this relatively modest adjustment in corporate taxation leads to a meaningful GE effect on differences in price indices between rich and poor households. This is in the order of a 2 percentage point lower inflation for retail consumption among the bottom 20 percent of US households compared to the top 20 percent (point estimate of 1.9, bootstrapped standard error of 0.23).

The right panel of Table A.11 presents the decomposition of this differential impact. The first channel is the direct incidence of the policy, holding fixed initial product choices by firms and consumers. Since the largest firms sort into producing higher product quality, they represent a larger fraction of retail consumption sourced by the richest quintile compared to the poorest. This direct effect accounts for 40 percent of the overall difference in consumer inflation.

The second channel is again related to changes in firm scale and product quality across the firm size distribution. The largest firms are faced with a decrease in effective revenues and experience a reduction in their sales relative to firms not subject to the tax increase. Figure A.24 depicts this compression of the firm size distribution, plotting counterfactual changes in log firm sales as a function of initial percentiles of the firm size distribution. Given economies of scale in quality production, and the large share of sales that the right tail of the size distribution represents, this shows up in the reduction of both the sales-weighted average product quality as well as the covariance of quality downgrading and initial firm size. The resulting changes in quality and quality-adjusted prices for consumption increases cost of living inflation relatively more among richer households, accounting for another 40 percent of the overall difference in price indices. As in the previous counterfactual, slight differences in love of variety imply that firm exit increases the retail price index more among richer households (accounting for 8 percent of the overall difference). Figure A.26 also breaks up the price index effect across the different departments in the scanner data. Finally, in the context of corporate taxation, tax entities may be holding companies rather than establishments producing brands. Reassuringly, in Figure A.25 we find that the counterfactual results are very similar (slightly stronger at 2.2 percentage points) when using holding companies instead of brands in the data to assign counterfactual tax changes.

6.3 Robustness and Additional Results

In the final section, we explore the sensitivity of these findings to alternative model and parameter assumptions. First, as discussed in Section 5, we use alternative values for the demand elasticities

⁴⁵To do so, we use estimates documented in Table 2 of Bao & Romeo (2013).

and technology parameters compared to our baseline estimates for $\sigma_n(z)$, β_n and ξ_n . Second, we re-estimate counterfactual outcomes after allowing household types (z) to endogenously change as a result of the counterfactual shocks. Third, we report counterfactuals under the more common assumption on market structure of monopolistic competition. Finally, we investigate the implications for the distribution of the gains from trade relative to a conventional Melitz (2003) model with two symmetric countries.

Alternative Parameter Values In line with the empirical literature in Industrial Organization and Quantitative Marketing, we find somewhat lower sales-to-price elasticities compared to the recent trade literature. To explore the sensitivity of our counterfactuals, we thus re-estimate each of them after multiplying our preferred estimates in Table 1 by a factor of 1.5 or 2 (keeping relative elasticities unchanged between income groups), or after assuming $\sigma_n(z)+1$ and $\sigma_n(z)+2$. Reassuringly, we tend to find somewhat stronger effects compared to our baseline counterfactuals. Instead of using separate $\sigma_n(z)$ estimates by income group, we can also use the parametric specification in Section 5.1 as a function of income percentiles, shown in Panel C of Table A.8. This yields more continuous variation in $\sigma_n(z)$ across income groups compared to the baseline estimation for below and above median income groups. We find very similar counterfactual results as in our baseline specification. On the production side, we also report results after using the cross-sectional technology parameter estimates, instead of our preferred panel IV estimates in Section 5.3. We find similar though somewhat larger effects for both counterfactuals (3.5 vs 3.1 and 2.3 vs 1.9 respectively).

Endogenous Changes in Outside Good Consumption (z) As discussed in Section 4, we hold household income types z fixed in our baseline counterfactuals, ignoring that changes in real incomes may push some households across z group boundaries and thereby affecting demand parameters $\alpha_n(z)$, $\sigma_n(z)$ and $\gamma_n(z)$ (see envelope theorem result in Appendix 3). While appealing for not having to take a stance on the upper-tier utility function (1), this shuts down potential second-order effects on price indices changes in these parameters. In our theoretical framework, the change in consumer type is captured by changes in the consumption of the outside good (z). In Appendix 3, we specify the upper-tier between retail consumption and the outside good, and quantify the change in outside good consumption induced by price index changes in the two counterfactuals above. Reassuringly, we find that the second-order adjustment channel has a negligible effect on outside good consumption (< 0.3 percent), and thus a negligible effect on the counterfactual results. This may not be surprising, as the price index differentials we quantify above are in the order of a 3 percentage point difference for retail consumption between rich and poor households. The fact that such a shock does not give rise to meaningful second-order effects through changes in household taste-for-quality or price elasticities—through pushing some households across quintiles—appears reasonable.

Monopolistic Competition Our model allows for large firms to take into account their market power through oligopolistic competition. A more common case with heterogeneous firms following Melitz (2003) features a continuum of firms interacting under monopolistic competition. Even in this case, our framework with two-sided heterogeneity allows for variable markups across firms

as a function of differences in their composition of consumer types. For completeness, we compute the two counterfactual equilibria in a model extension with monopolistic competition instead. Appendix 5.A derives the system of counterfactual equations and appendix Figure A.27 presents the counterfactual results. The point estimates are similar compared to the baseline results (allowing for oligopoly), and slightly stronger in the first counterfactual (3.5 vs 3.1 and 1.9 vs 1.9 respectively).

Implications for the Distribution of the Gains from Trade Finally, we investigate the implications for the distribution of the gains from trade. To do so, we introduce quality choice under two-sided heterogeneity into an otherwise standard Melitz model with monopolistic competition and two symmetric countries, and calibrate the model to the US data we use above (see Appendix 5.C).⁴⁶ As in Melitz (2003), a decrease in trade costs induces a reallocation in which the largest firms expand through trade while less productive firms either shrink or exit. In our framework, better access to imported varieties and exit of domestic producers affect the price indices of rich and poor households asymmetrically. In addition, lower trade costs lead to heterogeneous changes in product quality and markups across firms. As a result, we find that a 10 percentage point bilateral increase in import penetration leads to a 1.7-2.6 percentage point lower retail price inflation for the richest 20 percent of households compared to the poorest 20 percent. Relative to the conventional case—where heterogeneous firms are evenly represented across consumption baskets—we find that the stylized fact in Figure 1 and model we use to rationalize it imply a more unequal distribution of the gains from trade.

7 Conclusion

This paper presents empirical evidence that the widely documented presence of Melitz-type firm heterogeneity within sectors translates asymmetrically into the consumption baskets of households across the income distribution. To do so, we bring to bear detailed home and store scanner microdata that allow us to trace the national firm size distribution into the consumption baskets of individual households. We use the data to explore the underlying forces, and propose a quantitative GE model of quality choice under two-sided heterogeneity to rationalize the observed moments and explore policy counterfactuals.

The analysis provides several insights. We document large and statistically significant differences in the weighted average firm sizes that rich and poor households source their consumption from. After quantifying a set of opposing forces, we find that this pattern is mainly explained by two features of household preferences and firm technology. On the demand side, rich and poor households on average agree on their ranking of product quality evaluations within product groups. However, richer households value higher quality attributes significantly more compared to poorer households. On the production side, we estimate that producing higher output quality increases both the marginal and the fixed costs of production. Combined, these forces give rise to the en-

⁴⁶We calibrate fixed trade costs such that half of of output is produced by exporting firms. We calibrate variable trade costs such that export sales of exporters equal 20 percent of their output. Combining these two statistics, about 10 percent of aggregate output is traded. The counterfactual is to reduce variable trade costs from an equilibrium with no trade to this new trade equilibrium.

dogenous sorting of larger, more productive firms into products that are valued relatively more by wealthier households.

These results have implications for the effect of policies and economic shocks on real income inequality. We find that the direct effect of progressive income taxation on inequality is amplified through asymmetric GE effects on household price indices, and that business regulations or trade liberalization that affect large and small firms differently give rise to new distributional implications. Underlying these findings is a rich interplay of firm adjustments in scale, product quality, variable markups, exit and entry that vary across the firm size distribution and thus affect rich and poor households differently.

Overall, our findings suggest that firm heterogeneity affects real income inequality in more complex ways than through the nominal earnings of workers, which have been the focus of the existing literature. These findings arise after introducing a basic set of features that we observe in the microdata—allowing for choice in product attributes by heterogeneous firms and households with non-homothetic preferences—into an otherwise standard economic environment. Empirically, these findings emphasize the importance of capturing changes in price indices at a granular level of product aggregation for both the measurement of changes in real income inequality over time and for studying the effects of policies and economic shocks.

References

- Aguiar, M., & Hurst, E. (2007). Measuring trends in leisure: the allocation of time over five decades. *The Quarterly Journal of Economics*, 122(3), 969–1006.
- Argente, D., & Lee, M. (2016). Cost of living inequality during the great recession.
- Arkolakis, C., Costinot, A., Donaldson, D., & Rodríguez-Clare, A. (2018). The elusive pro-competitive effects of trade. *The Review of Economic Studies*, 86(1), 46–80.
- Atkin, D., Faber, B., & Gonzalez-Navarro, M. (2018). Retail globalization and household welfare: Evidence from Mexico. *Journal of Political Economy*, 126(1), 1–73.
- Baldwin, R., & Harrigan, J. (2011). Zeros, quality, and space: Trade theory and trade evidence. *American Economic Journal: Microeconomics*, 3(2), 60–88.
- Bao, D.-H., & Romeo, G. C. (2013). Tax avoidance and corporations in the United State – the effective tax rate abnormality for the top five percent by corporate size. *Journal of Applied Business and Economics*, 14(4), 88–100.
- Bartelsman, E., Haltiwanger, J., & Scarpetta, S. (2013). Cross-country differences in productivity: The role of allocation and selection. *The American Economic Review*, 103(1), 305–334.
- Bastos, P., Silva, J., & Verhoogen, E. (2018). Export destinations and input prices. *American Economic Review*, 108(2), 353–92.
- Bernard, A. B., Jensen, J. B., Redding, S. J., & Schott, P. K. (2007). Firms in international trade. *The Journal of Economic Perspectives*, 105–130.
- Bijmolt, T. H., Van Heerde, H. J., & Pieters, R. G. (2005). New empirical generalizations on the determinants of price elasticity. *Journal of Marketing Research*, 42(2), 141–156.


- Bloom, N., & Van Reenen, J. (2007). Measuring and explaining management practices across firms and countries. *The Quarterly Journal of Economics*, 1351–1408.
- Broda, C., & Weinstein, D. (2010). Product creation and destruction: Evidence and price implications. *American Economic Review*, 100, 691–723.
- Burstein, A., & Vogel, J. (2017). International trade, technology, and the skill premium. *Journal of Political Economy*, 125(5), 1356–1412.
- Card, D., Heining, J., & Kline, P. (2013). Workplace heterogeneity and the rise of West German wage inequality. *The Quarterly Journal of Economics*, 128(3), 967–1015.
- Davis, D. R., & Harrigan, J. (2011). Good jobs, bad jobs, and trade liberalization. *Journal of International Economics*, 84(1), 26–36.
- Dekle, R., Eaton, J., & Kortum, S. (2007). Unbalanced trade. *The American Economic Review*.
- DellaVigna, S., & Gentzkow, M. (2019). Uniform pricing in us retail chains. *The Quarterly Journal of Economics*, 134(4), 2011–2084.
- Dingel, J. I. (2016). The determinants of quality specialization. *The Review of Economic Studies*, 84(4), 1551–1582.
- Fajgelbaum, P., Grossman, G., & Helpman, E. (2011). Income distribution, product quality, and international trade. *Journal of Political Economy*, 119(4), 721–765.
- Fally, T. (2018). Integrability and generalized separability. *NBER Working Paper 25025*.
- Feenstra, R. C. (1994). New product varieties and the measurement of international prices. *The American Economic Review*, 157–177.
- Feenstra, R. C., & Romalis, J. (2014). International prices and endogenous quality. *The Quarterly Journal of Economics*, 129(2), 477–527.
- Fieler, A. C., Eslava, M., & Xu, D. Y. (2018). Trade, quality upgrading, and input linkages: Theory and evidence from colombia. *American Economic Review*, 108(1), 109–46.
- Frias, J. A., Kaplan, D. S., & Verhoogen, E. A. (2009). Exports and wage premia: Evidence from mexican employer-employee data. *Working Paper, Columbia University*.
- Guvenen, F., Mataloni Jr, R. J., Rassier, D. G., & Ruhl, K. J. (2017). Offshore profit shifting and domestic productivity measurement. *NBER Working Paper 23324*.
- Hallak, J. C., & Sivadasan, J. (2013). Product and process productivity: Implications for quality choice and conditional exporter premia. *Journal of International Economics*, 91(1), 53–67.
- Handbury, J. (2019). Are poor cities cheap for everyone? Non-homotheticity and the cost of living across US cities. *NBER Working Paper 26574*.
- Handbury, J., & Weinstein, D. E. (2014). Goods prices and availability in cities. *The Review of Economic Studies*, 82(1), 258–296.
- Hausman, J. (1999). Cellular telephone, new products, and the CPI. *Journal of business & economic statistics*, 17(2), 188–194.
- Hausman, J., & Leibtag, E. (2007). Consumer benefits from increased competition in shopping outlets: Measuring the effect of wal-mart. *Journal of Applied Econometrics*, 22(7), 1157–1177.

- Helpman, E., Itskhoki, O., Muendler, M.-A., & Redding, S. J. (2017). Trade and inequality: From theory to estimation. *The Review of Economic Studies*, 84(1), 357–405.
- Horowitz, J. L. (2001). The bootstrap. In J. J. Heckman & E. Leamer (Eds.), *Handbook of econometrics* (Vol. 5, p. 3159 - 3228).
- Hottman, C. J., Redding, S. J., & Weinstein, D. E. (2016). Quantifying the sources of firm heterogeneity. *The Quarterly Journal of Economics*, 131(3), 1291–1364.
- Hsieh, C.-T., & Klenow, P. J. (2009). Misallocation and manufacturing TFP in China and India. *Quarterly Journal of Economics*, 124(4).
- Jaravel, X. (2019). The unequal gains from product innovations: Evidence from the us retail sector. *The Quarterly Journal of Economics*, 134(2), 715–783.
- Johnson, R. (2012). Trade and prices with heterogeneous firms. *Journal of International Economics*.
- Keller, K. L., Parameswaran, M., & Jacob, I. (2011). *Strategic brand management: Building, measuring, and managing brand equity*. Pearson Education India.
- Kugler, M., & Verhoogen, E. (2012). Prices, plant size, and product quality. *The Review of Economic Studies*, 79(1), 307–339.
- McFadden, D., & Train, K. (2000). Mixed MNL models for discrete response. *Journal of applied Econometrics*, 447–470.
- Melitz, M. (2003). The impact of trade on intra-industry reallocations and aggregate industry productivity. *Econometrica*, 71(6), 1695–1725.
- Nevo, A. (2000). Mergers with differentiated products: The case of the ready-to-eat cereal industry. *The RAND Journal of Economics*, 395–421.
- Peters, M. (2013). Heterogeneous mark-ups, growth and endogenous misallocation.
- Ray, W. (2019). Cyclicity of markups: Evidence from heteroskedasticity-identified estimation. *Working Paper, London School of Economics*.
- Redding, S. J., & Weinstein, D. E. (2020). A unified approach to estimating demand and welfare changes. *The Quarterly Journal of Economics*, 135(1), 503–560.
- Saez, E., & Zucman, G. (2019). *The triumph of injustice: How the rich dodge taxes and how to make them pay*. New York: WW Norton & Co.
- Sampson, T. (2014). Selection into trade and wage inequality. *American Economic Journal: Microeconomics*, 6(3), 157–202.
- Soderbery, A. (2015). Estimating import supply and demand elasticities: Analysis and implications. *Journal of International Economics*, 96(1), 1–17.
- Song, J., Price, D. J., Guvenen, F., Bloom, N., & Von Wachter, T. (2018). Firming up inequality. *The Quarterly journal of economics*, 134(1), 1–50.
- Subramanian, S., & Deaton, A. (1996). The demand for food and calories. *Journal of political economy*, 133–162.
- Sutton, J. (1998). *Technology and market structure: Theory and history*. The MIT Press.
- Wright, T., & Zucman, G. (2018). The exorbitant tax privilege. *NBER Working Paper 24983*.

Figures and Tables


Figures

Figure 1: Richer Households Source Their Consumption from Larger Firms




Notes: The figure depicts deviations in weighted average log firm sales embodied in the consumption baskets of on average 59,000 US households during 18 half-year periods between 2006-14. The y-axis in both graphs displays weighted average deviations in log producer sales within more than 1000 product modules where the weights are household expenditure shares across producers of brands. In the first step, we compute brand-level deviations from mean log national sales within product module-by-half-year cells from either the home or the store-level scanner data. In the second step, these deviations are then matched to brand-level half-yearly household expenditure weights in the home scanner data. The final step is to collapse these data to weighted average log firm size deviations embodied in household consumption baskets. The x-axis displays national percentiles of per capita total household retail expenditure per half year period (see Section 2). The fitted relationships in the left graph correspond to local polynomial regressions. Standard errors in both graphs are clustered at the county level, and the displayed confidence intervals are at the 95% level. Table A.1 provides descriptive statistics. See Section 3 for discussion.

Figure 2: Distribution of Weighted Average Product Quality and Quality-Adjusted Prices across Consumption Baskets




Notes: The figure depicts deviations in weighted average log brand quality embodied in the consumption baskets of on average 59,000 US households during 18 half-year periods between 2006-14. The y-axis on the left (right) displays weighted average deviations in log brand quality (quality-adjusted prices) within more than 1000 product modules where the weights are household expenditure shares across producers of brands. The x-axis in both graphs displays national percentiles of per capita total household retail expenditure per half-year period (see Section 2). The fitted relationships correspond to local polynomial regressions. Standard errors in both graphs are clustered at the county level, and the displayed confidence intervals are at the 95% level. See Section 5 for discussion.

Figure 3: Producers Charge Variable Markups and Face Different Tastes for Quality




Notes: The figure depicts deviations in the markup (μ_{ni}) and taste-for-quality parameters ($\tilde{\gamma}_{ni}$) across the firm size distribution for 18 half-yearly cross-sections between 2006-2014. The y-axis displays deviations in μ_{ni} and $\tilde{\gamma}_{ni}$ relative to product module-by-half-year means. The x-axis displays deviations of log firm sales at the same level. The fitted relationships correspond to local polynomial regressions. Standard errors in both graphs are clustered at the level of product modules, and the displayed confidence intervals are at the 95% level. See Section 5 for discussion.

Figure 4: Decomposition of the Underlying Forces



Notes: The figure depicts predicted (model-based) and observed deviations in firm sizes across consumption baskets. See Section 5.4 for discussion.


Figure 5: Counterfactual 1: Inflation Differences and Quality Downgrading Due to More Progressive Income Taxes



39

Notes: Both graphs display counterfactual changes averaged across 18 half-year cross-sections of data. 95% confidence intervals are based on robust standard errors. The left panel displays counterfactual differences in cost of living inflation for retail consumption across income groups. The right panel displays deviations of output quality downgrading across percentiles of initial firm sales within product-module-by-period cells. See Section 6 for discussion.

Figure 6: Counterfactual 2: Inflation Differences and Quality Downgrading Due to Closing Loopholes in Corporate Taxation



Notes: Both graphs display counterfactual changes averaged across 18 half-year cross-sections of data. 95% confidence intervals are based on robust standard errors. The left panel displays counterfactual differences in cost of living inflation for retail consumption across income groups. The right panel displays deviations of output quality downgrading across percentiles of initial firm sales within product-module-by-period cells. See Section 6 for discussion.

Tables

Table 1: Price Elasticities

<i>Panel A: Pooled Estimates</i>					
Dependent Variable: Change in Log Budget Shares	OLS	National IV	State IV	Both IVs	Both IVs
(1- σ) All Households	0.257*** (0.0288)	-1.184*** (0.0356)	-1.090*** (0.0415)	-1.181*** (0.0316)	
(1- σ) Poorest Quintile (Relative to Richest)					-0.375*** (0.131)
(1- σ) 2nd Poorest Quintile (Relative to Richest)					-0.391*** (0.116)
(1- σ) Median Quintile (Relative to Richest)					-0.163** (0.0674)
(1- σ) 2nd Richest Quintile (Relative to Richest)					-0.271** (0.104)
Quintile-by-Module-by-County-by-Period FX	✓	✓	✓	✓	✓
Brand-by-County-by-Period FX	✗	✗	✗	✗	✓
Observations	9,989,508	9,989,508	9,283,699	9,283,699	9,283,699
First Stage F-Stat		718.6	314.7	420.0	84.33

<i>Panel B: By Product Department</i>								
	Beverages	Dairy	Dry Grocery	Frozen Foods	General Merchandise	Health and Beauty	Non-Food Grocery	Packaged Meat
Dependent Variable: Change in Log Budget Shares	Both IVs	Both IVs	Both IVs	Both IVs	Both IVs	Both IVs	Both IVs	Both IVs
(1- σ) All Households	-1.091*** (0.149)	-0.716*** (0.0559)	-1.324*** (0.0405)	-1.336*** (0.0672)	-2.353*** (0.222)	-0.504*** (0.0878)	-1.100*** (0.0911)	-1.318*** (0.151)
Quintile-by-Module-by-County-by-Period FX	✓	✓	✓	✓	✓	✓	✓	✓
Observations	755,648	775,238	4,570,372	945,956	205,830	778,667	982,261	269,726
First Stage F-Stat	542.2	253.0	407.8	126.7	169.2	217.0	731.7	56.63

<i>Panel C: By Department and Household Group</i>								
	Beverages	Dairy	Dry Grocery	Frozen Foods	General Merchandise	Health and Beauty	Non-Food Grocery	Packaged Meat
Dependent Variable: Change in Log Budget Shares	Both IVs	Both IVs	Both IVs	Both IVs	Both IVs	Both IVs	Both IVs	Both IVs
(1- σ) Below Median Quintiles	-1.272*** (0.252)	-0.809*** (0.142)	-1.481*** (0.105)	-1.341*** (0.148)	-2.436*** (0.368)	-0.506* (0.272)	-1.383*** (0.239)	-1.329*** (0.261)
(1- σ) Median and Above Quintiles	-1.041*** (0.147)	-0.689*** (0.0569)	-1.288*** (0.0462)	-1.336*** (0.0721)	-2.339*** (0.249)	-0.501*** (0.107)	-1.048*** (0.0757)	-1.316*** (0.155)
Quintile-by-Module-by-County-by-Period FX	✓	✓	✓	✓	✓	✓	✓	✓
Observations	755,648	775,238	4,570,372	945,956	205,830	778,667	982,261	269,726
First Stage F-Stat	139.0	347.5	254.1	50.17	131.4	109.4	298.0	37.68

Notes: See Section 5 for discussion. Standard errors are in parentheses below point estimates and clustered at the level of counties. ***, **, * indicate 1, 5 and 10 percent confidence levels.

Table 2: Product Quality and Firm Scale: Reduced-Form Evidence

Dependent Variables:	ALL PRODUCT GROUPS							
	Cross-Section				Panel Data			
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
	Log Unit Value		Log Quality		Δ Log Unit Value		Δ Log Quality	
OLS	IV	OLS	IV	OLS	IV	OLS	IV	
Log National Firm Sales	0.0280*** (0.00339)	0.0253*** (0.00390)	1.128*** (0.0312)	1.142*** (0.0309)				
Δ Log National Firm Sales					0.0365*** (0.00320)	0.0705*** (0.0138)	1.131*** (0.0415)	0.569*** (0.0589)
Product Module-by-Period FX	✓	✓	✓	✓	✗	✗	✗	✗
State-by-Product Module-by- Period FX	✗	✗	✗	✗	✓	✓	✓	✓
Observations	1,330,976	1,330,976	1,330,976	1,330,976	1,789,078	1,789,078	1,789,078	1,789,078
Number of Product Module Clusters	1031	1031	1031	1031	1004	1004	1004	1004
First Stage F-Stat		322552		322552		251.1		251.1

Notes: See Section 5 for discussion. Standard errors are in parentheses below point estimates and clustered at the level of product modules. ***, **, * indicate 1, 5 and 10 percent confidence levels.

Table 3: Technology Parameter Estimates

Dependent Variable:	ALL PRODUCT GROUPS							
Log Product Quality or Changes in Log Quality	Cross-Section		Panel Data		Cross-Section		Panel Data	
	OLS	IV	OLS	IV	OLS	IV	OLS	IV
Log Firm Scale or Changes in Log Firm Scale (β)	1.1132*** (0.0309)	1.1352*** (0.0307)	1.1638*** (0.0466)	0.4887*** (0.0653)				
ξ Parameter	0.82	0.82	0.32	0.32				
Observations	1,330,976	1,330,976	1,422,244	1,422,244				
Number of Clusters	1,031	1,031	994	994				
First Stage F-Stat		311103.21		247.6				

Dependent Variable:	GROCERY				NON-GROCERY			
Log Product Quality or Changes in Log Quality	Cross-Section		Panel Data		Cross-Section		Panel Data	
	OLS	IV	OLS	IV	OLS	IV	OLS	IV
Log Firm Scale or Changes in Log Firm Scale (β)	0.9462*** (0.0124)	0.958*** (0.0129)	0.9387*** (0.0113)	0.2898*** (0.0745)	1.5357*** (0.0851)	1.5785*** (0.084)	1.525*** (0.1006)	0.9485*** (0.1345)
ξ Parameter	0.45	0.45	0.45	0.45	0.79	0.79	0.85	0.85
Observations	1,002,542	1,002,542	1,031,295	1,031,295	328,434	328,434	390,949	390,949
Number of Clusters	719	719	696	696	312	312	298	298
First Stage F-Stat		279437.12		181.31		66343.15		103.65

Notes: See Section 5 for discussion. Standard errors are in parentheses below point estimates and clustered at the level of product modules. ***, **, * indicate 1, 5 and 10 percent confidence levels.


Appendix

Appendix 1: Additional Figures and Tables

Table A.1: Descriptive Statistics


Home Scanner Data		Retail Scanner Data	
Number of Half Year Periods 2006-14	18	Number of Half Year Periods 2006-14	18
Number of Observations (Summed up to Household-Period-Barcode-Retailer)	344,533,688	Number of Observations (Summed up to Store-Period-Barcode)	12,206,598,912
Number of Households per Period	58,769	Number of Stores per Period	27,290
Number of Product Modules per Period	1,090	Number of Product Modules per Period	1,092
Number of Brands per Period	185,286	Number of Brands per Period	175,095
Number of Barcodes per Period	594,504	Number of Barcodes per Period	727,932
Number of Retailers per Period	774	Number of Retailers per Period	102
Number of Counties per Period	2,671	Number of Counties per Period	2,500
Total Sales per Period (Using Projection Weights)	108,580,633 (211,447,813,471)	Total Sales per Period	113,315,047,442

Figure A.1: Observed Expenditure Per Capita and Reported Income Brackets



Notes: The figure depicts the relationship between our measure of log expenditure per capita and reported nominal income brackets two years before across 18 half-yearly cross-sections between 2006-2014. The y-axis displays within-half-year deviations in log reported incomes after assigning households the mid-point of their reported income bracket. The x-axis displays percentiles of per-capita expenditure within a given half year (see Section 2). Standard errors in both graphs are clustered at the county level, and the displayed confidence intervals are at the 95% level. See Section 2 for discussion.

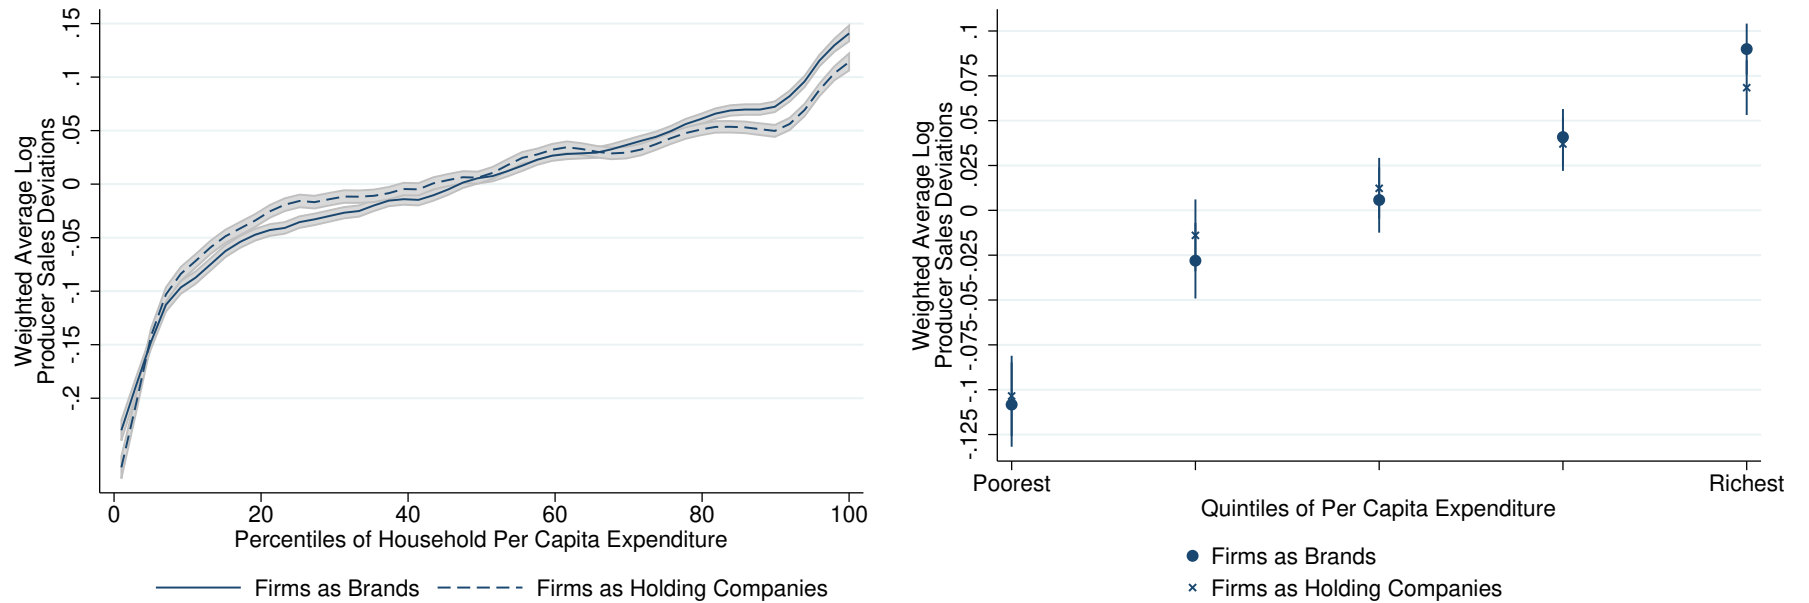
Figure A.2: Firm Heterogeneity in the Home and Retail Scanner Data



3


Notes: The figure on the left depicts the firm size distribution for all brands present in either the home or store scanner data. The figure on the right restricts attention to producers of brands that are present in both datasets. Table A.1 provides descriptive statistics. See Section 3 for discussion.

Figure A.3: Firm Heterogeneity Across Consumption Baskets - Robustness to Alternative Firm Definition



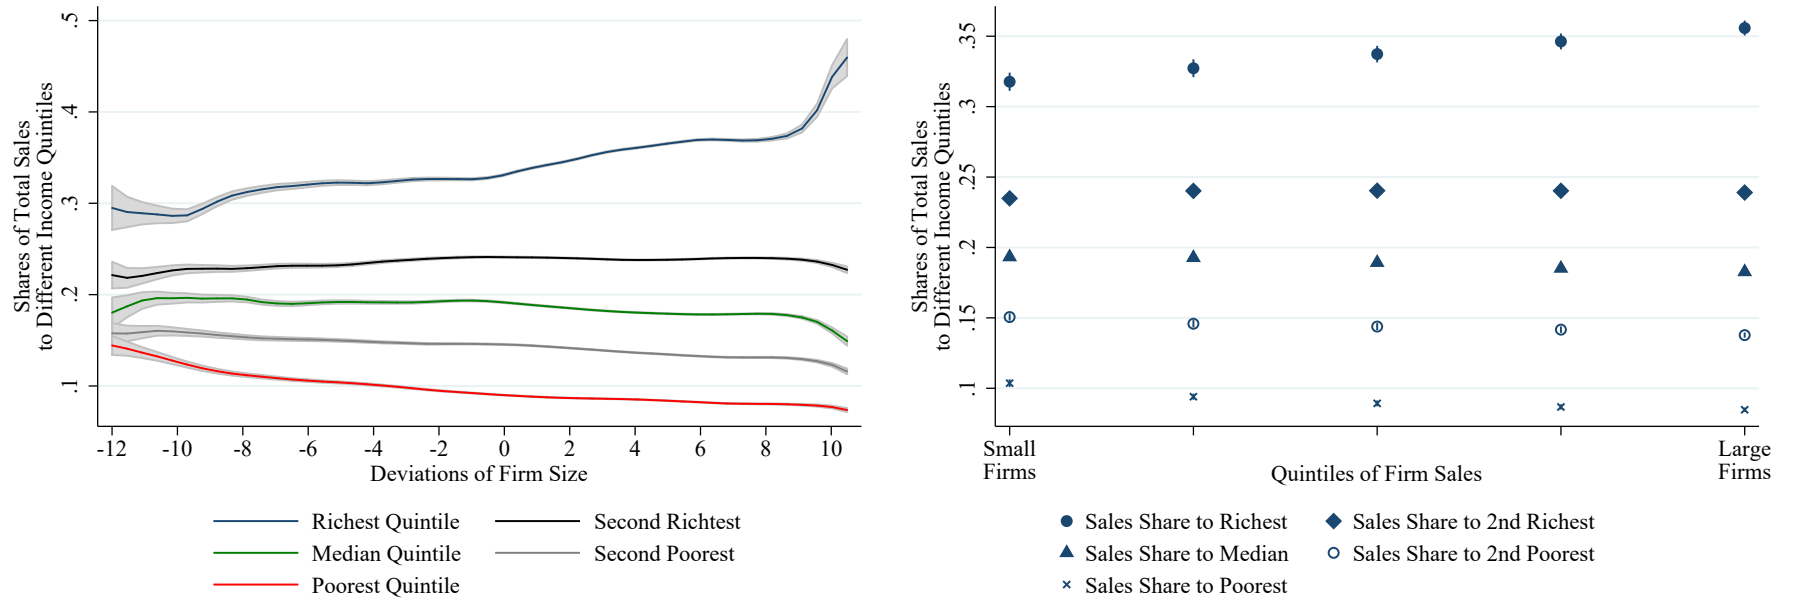
Notes: The figure depicts deviations in weighted average log firm sales embodied in the consumption baskets of on average 59,000 US households during 18 half-year periods between 2006-14. The y-axis in both graphs displays weighted average deviations in log producer sales within product modules (“firms as brands”) or product groups (“firms as holding companies”) where the weights are household expenditure shares across firms. Firms are defined either as brands (232 k in the dataset) or holding companies (145 k in the dataset). To define holding companies in the data, we follow Broda and Weinstein (2010) and take the first 6 digits of the EAN barcode. Following Hottman et al. (2016), this correctly identifies holding companies in about 80 percent of the cases. For the remainder, this method will tend to over-aggregate the brands into holding companies, so that this robustness check should be seen as conservative. National firm size deviations are based on the store scanner data. These firm size deviations are depicted across consumption baskets conditional on half-year fixed effects. The x-axis displays national percentiles of per capita total household retail expenditure per half year (see Section 2). The fitted relationship in the left graph corresponds to a local polynomial regression. Standard errors in both graphs are clustered at the county level, and the displayed confidence intervals are at the 95% level. Table A.1 provides descriptive statistics. See Section 3 for discussion.

Figure A.4: Firm Heterogeneity Across Consumption Baskets - Firm Size in Terms of Quantities



Notes: The graph replicates the right panel of Figure 1 in the text after replacing firm size deviations in terms of log revenues by log quantities (units sold). Units of output are measured identically across products within a product module (e.g. liters of milk, units of microwaves, grams of cereal, etc). Standard errors are clustered at the county level, and the displayed confidence intervals are at the 95% level. See Section 3 for discussion.

Figure A.5: Firm Heterogeneity Across Consumption Baskets - Flipped Axes




Notes: The figure depicts the share of firm sales to each of the 5 income quintiles across the firm size distribution within product modules averaged across 18 half-year periods and product modules. The x-axis in the left panel are deviations in log firm total sales relative to product module-by-half-year bins. The x-axis in the right panel are quintiles of the firm size distribution within product modules. The fitted relationship in the left graph corresponds to a local polynomial regression. Standard errors in both graphs are clustered at the product module level, and the displayed confidence intervals are at the 95% level. See Section 3 for discussion.

Table A.2: Examples for Popular Product Modules across Different Departments

Product Department	Product Module	Brand with Highest Budget Share Difference (Rich Minus Poor)	Brand with Lowest Budget Share Difference (Rich Minus Poor)	Brands' Difference in Market Shares (Highest Minus Lowest)	Brands' Difference in Log Unit Values (Highest Minus Lowest)
ALCOHOLIC BEVERAGES	BEER	BUDWEISER	MILLER HIGH LIFE	0.129	0.302
ALCOHOLIC BEVERAGES	BOURBON-STRAIGHT/BONDED	MAKER'S MARK	TEN HIGH	0.055	0.246
ALCOHOLIC BEVERAGES	SCOTCH	DEWAR'S WHITE LABEL	GLENFIDDICH	0.111	2.832
DAIRY	CHEESE-PROCESSED SLICES-AMERICAN	KRAFT DELI DELUXE	BORDEN	0.042	0.452
DAIRY	DAIRY-FLAVORED MILK-REFRIGERATED	NESTLE NESQUIK	GENERIC STORE BRAND	0.078	1.117
DAIRY	YOGURT-REFRIGERATED	DANNON	GENERIC STORE BRAND	0.225	0.469
DRY GROCERY	CATSUP	HEINZ	HUNT'S	0.513	0.307
DRY GROCERY	FRUIT JUICE - ORANGE - OTHER CONTAINER	TROPICANA	GENERIC STORE BRAND	0.314	0.590
DRY GROCERY	SOFT DRINKS - CARBONATED	PEPSI R	GENERIC STORE BRAND	0.069	0.362
FROZEN FOODS	FROZEN NOVELTIES	WEIGHT WATCHERS	GENERIC STORE BRAND	0.025	0.986
FROZEN FOODS	FROZEN WAFFLES & PANCAKES & FRENCH TOAST	KELLOGG'S EGGO	AUNT JEMIMA	0.491	0.129
FROZEN FOODS	PIZZA-FROZEN	DIGIORNO	TOTINO'S	0.147	0.607
GENERAL MERCHANDISE	BATTERIES	DURACELL	RAYOVAC	0.321	0.350
GENERAL MERCHANDISE	PRINTERS	HEWLETT PACKARD OFFICEJET	CANON PIXMA	0.062	0.338
GENERAL MERCHANDISE	VACUUM AND CARPET CLEANER APPLIANCE	DYSON	BISSELL POWER FORCE	0.065	2.084
HEALTH & BEAUTY CARE	PAIN REMEDIES - HEADACHE	ADVIL	GENERIC STORE BRAND	0.078	0.086
HEALTH & BEAUTY CARE	SANITARY NAPKINS	ALWAYS MX PD/WG ULTR THN OVRNT	GENERIC STORE BRAND	0.030	1.591
HEALTH & BEAUTY CARE	SHAMPOO-AEROSOL/ LIQUID/ LOTION/ POWDER	PANTENE PRO-V	ALBERTO VO5	0.109	1.444
NON-FOOD GROCERY	CIGARS	HAV-A-TAMPA	POM POM OPERAS	0.023	0.375
NON-FOOD GROCERY	DETERGENTS - HEAVY DUTY - LIQUID	TIDE - H-D LIQ	PUREX - H-D LIQ	0.283	0.779
NON-FOOD GROCERY	SOAP - BAR	DOVE	DIAL	0.221	0.772
PACKAGED MEAT	BACON-REFRIGERATED	OSCAR MAYER	BAR S	0.214	0.961
PACKAGED MEAT	BRATWURST & KNOCKWURST	JOHNSONVILLE	KLEMENT'S	0.678	0.141
PACKAGED MEAT	FRANKS-COCKTAIL-REFRIGERATED	HILLSHIRE FARM	CAROLINA PRIDE	0.388	0.243

Notes: See Section 3 for discussion.


Figure A.6: Households on Average Agree on Relative Product Quality Evaluations



8

Notes: The figure depicts the relationship between income group-specific deviations in log expenditures spent across producers within more than 1000 product modules (y-axis), and deviations in log total market sales of those same producers in the store scanner data (x-axis) for on average 59,000 US households during 18 half-year periods between 2006-14. The left panel shows the full sample, and the right panel restricts attention to firm size deviations on the x-axis between -2 to 2 log points. The fitted relationships in both graphs correspond to local polynomial regressions. See Section 3 for discussion.


Figure A.7: Firm Heterogeneity Across Consumption Baskets - Role of Extensive Margin (1)



6

Notes: The left panel depicts average expenditure shares on brands consumed by all income quintiles in a given half-year period or ever across the per-capita expenditure distribution. The right panel depicts the average number of brands or UPCs consumed per household per product module in a given half-year period across the income distribution. The fitted relationships in both graphs correspond to local polynomial regressions. Standard errors in both graphs are clustered at the county level, and the displayed confidence intervals are at the 95% level. See Section 3 for discussion.


Figure A.8: Firm Heterogeneity Across Consumption Baskets - Role of Extensive Margin (2)



10

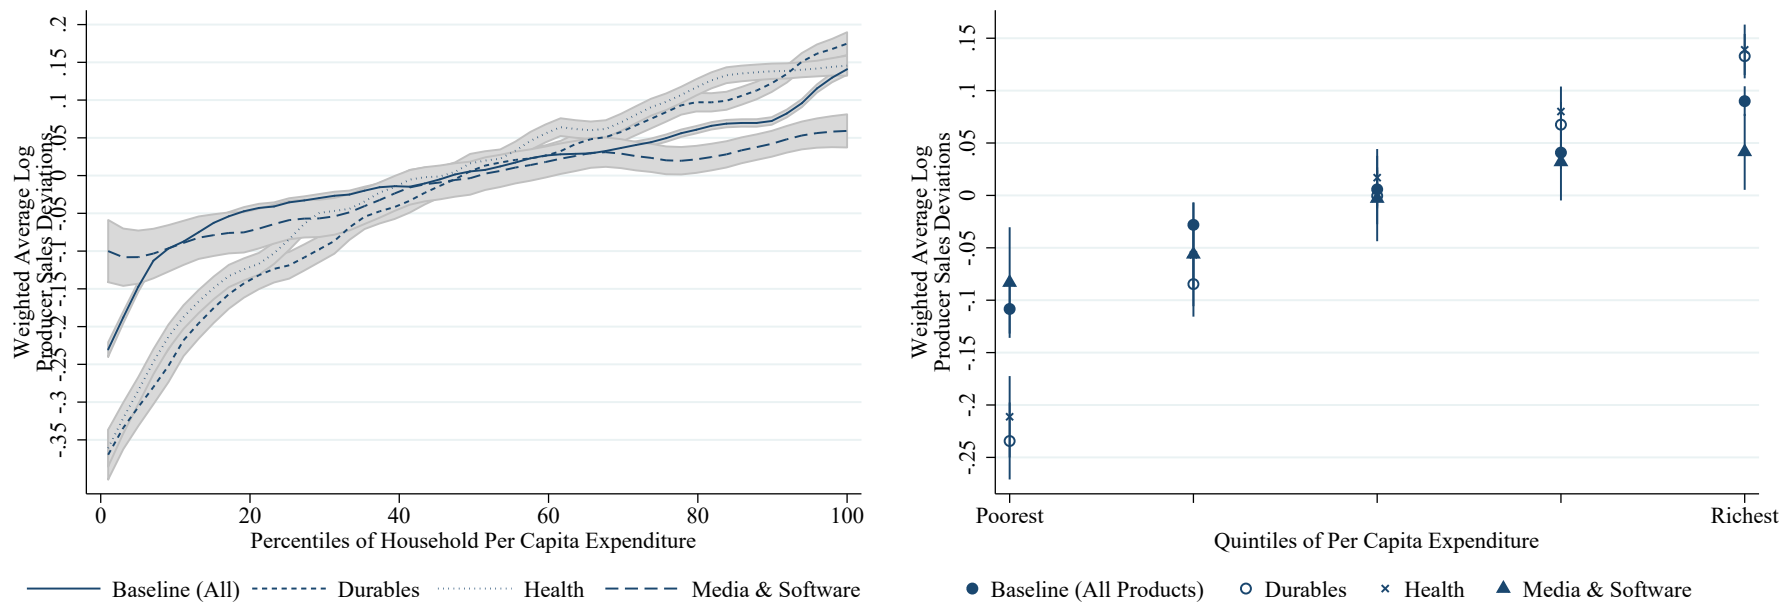
Notes: The figures replicate the relationships estimated in Figure 1 after restricting retail consumption to only those brands that are consumed by all income quintiles in a given period or ever. The fitted relationship in the left panel corresponds to a local polynomial regression. Standard errors in both graphs are clustered at the county level, and the displayed confidence intervals are at the 95% level. See Section 3 for discussion.

Figure A.9: Heterogeneity across Product Departments



Notes: The fitted relationships correspond to local polynomial regressions. See Section 3 for discussion.

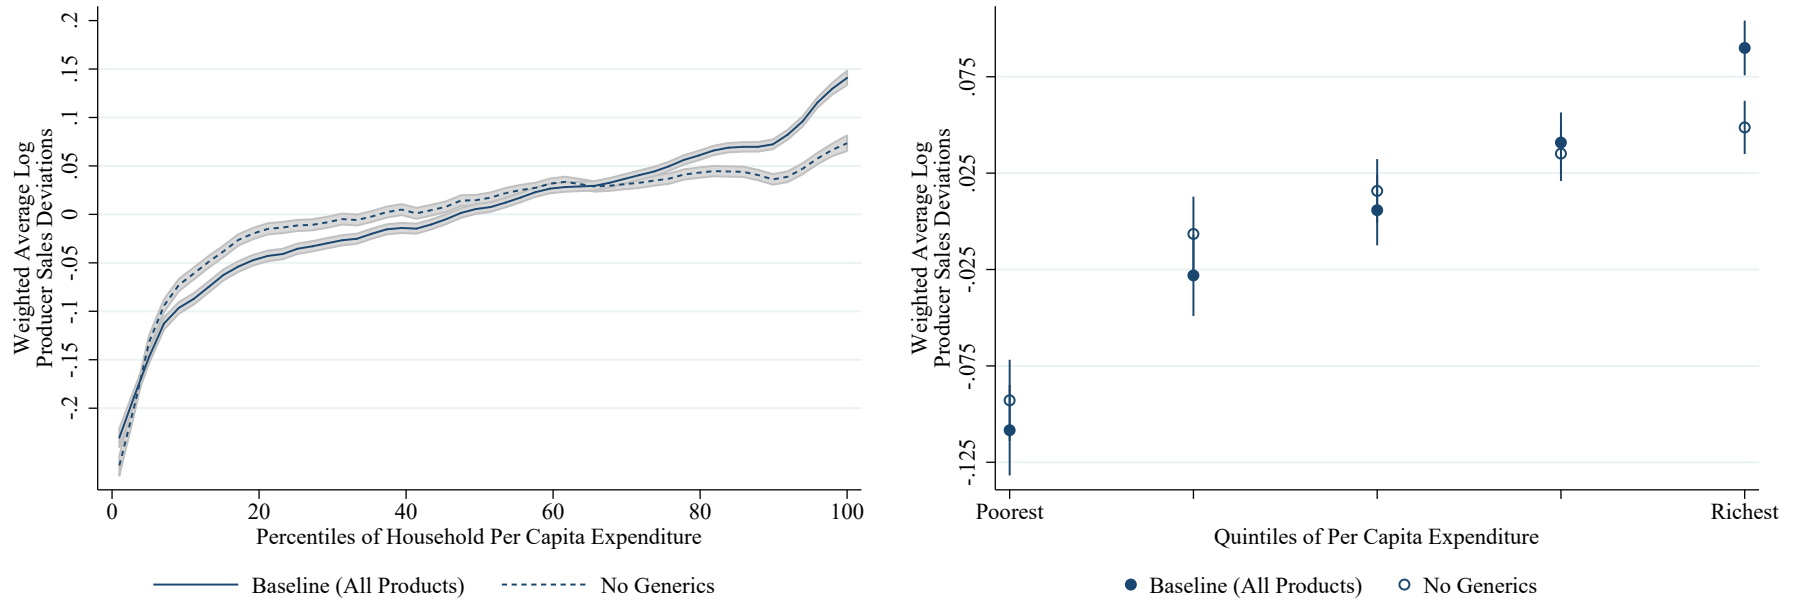
Figure A.10: Replicating Figure 1 for Product Categories Not Fully Represented in Nielsen



12

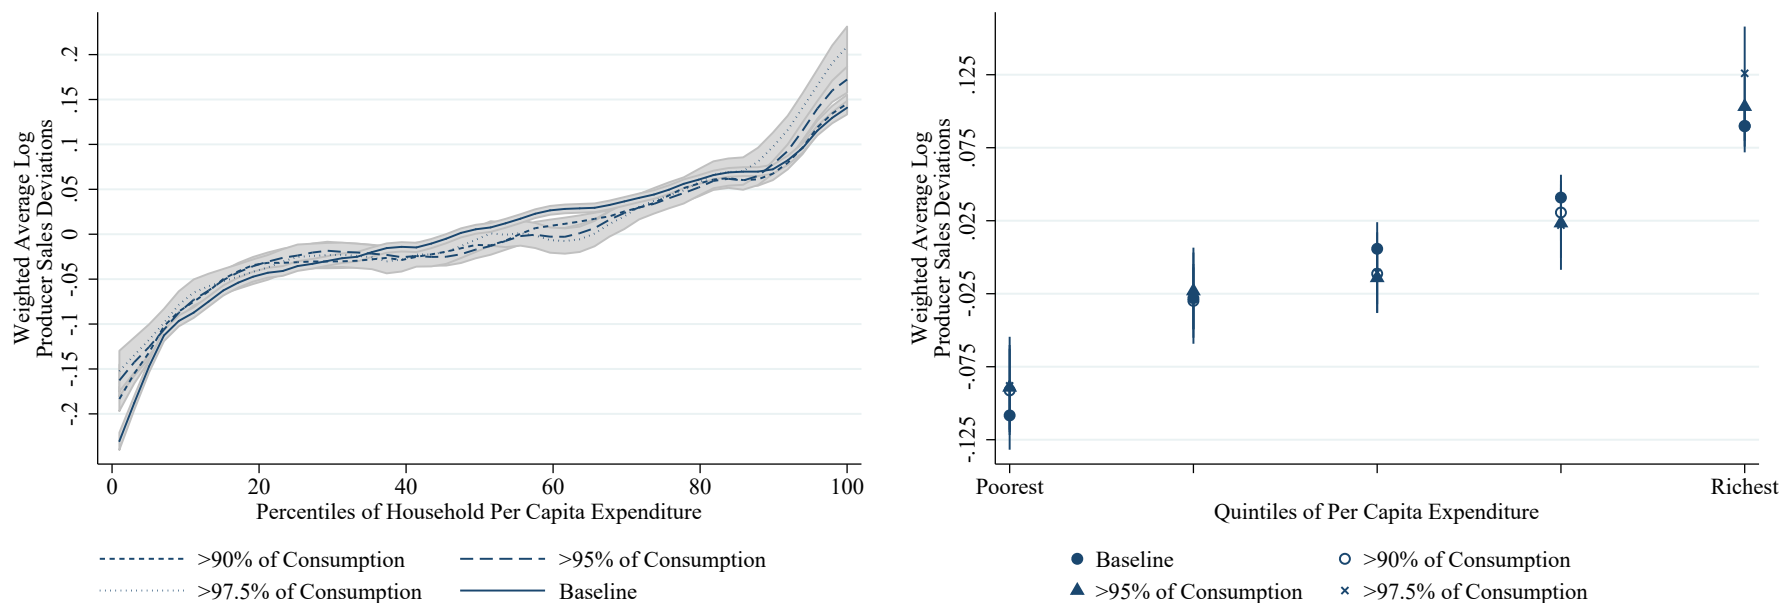
Notes: The figure computes the same relationship as in Figure 1, but separately for household consumption of appliances, pharmaceuticals and audio, video and software purchases respectively. The fitted relationship in the left graph corresponds to a local polynomial regression. Standard errors in both graphs are clustered at the county level, and the displayed confidence intervals are at the 95% level. See Section 3 for discussion.

Figure A.11: The Role of Generic Retailer Brands



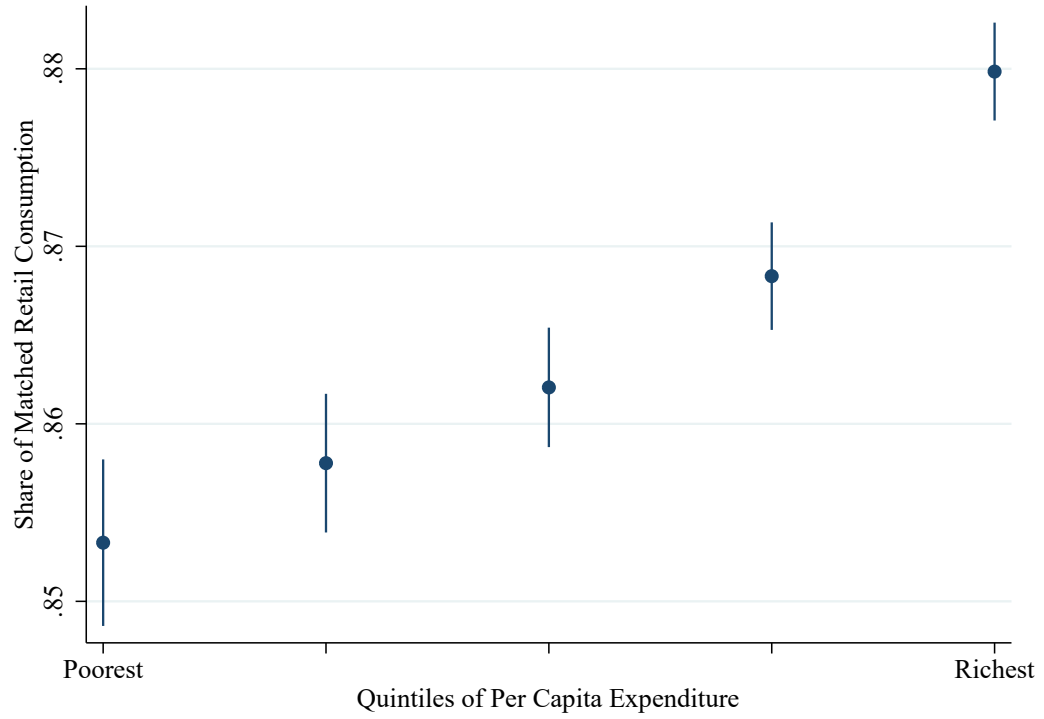
Notes: The figure depicts deviations in weighted average log firm sales embodied in the consumption baskets of on average 59,000 US households during 18 half-year periods between 2006-14. The y-axis in both graphs displays weighted average deviations in log producer sales within more than 1000 product modules where the weights are household expenditure shares across producers of brands. National firm size deviations are based on the store scanner data. These firm size deviations are depicted across consumption baskets conditional on half-year fixed effects for i) the full sample of households and products, ii) only for households with matched firm size deviations for more than 90% of total consumption, and iii) only for consumption spent on brands that are not generic store brands. The x-axis displays national percentiles of per capita total household retail expenditure per half year period (see Section 2). The fitted relationship in the left graph corresponds to a local polynomial regression. Standard errors in both graphs are clustered at the county level, and the displayed confidence intervals are at the 95% level. See Section 3 for discussion.

Figure A.12: The Role of Retail Coverage Across Incomes (1)



Notes: The figure depicts deviations in weighted average log firm sales embodied in the consumption baskets of on average 59,000 US households during 18 half-year periods between 2006-14. The y-axis in both graphs displays weighted average deviations in log producer sales within more than 1000 product modules where the weights are household expenditure shares across producers of brands. National firm size deviations are based on the store scanner data. These firm size deviations are depicted across consumption baskets conditional on half-year fixed effects for i) the full sample of households, ii) only households with more than 90% of their total reported retail consumption matched to national firm size deviations, iii) more than 95% and iv) more than 97.5%. The x-axis displays national percentiles of per capita total household retail expenditure per half year period (see Section 2). The fitted relationship in the left graph corresponds to a local polynomial regression. Standard errors in both graphs are clustered at the county level, and the displayed confidence intervals are at the 95% level. See Section 3 for discussion.

Figure A.13: The Role of Retail Coverage Across Incomes (2)



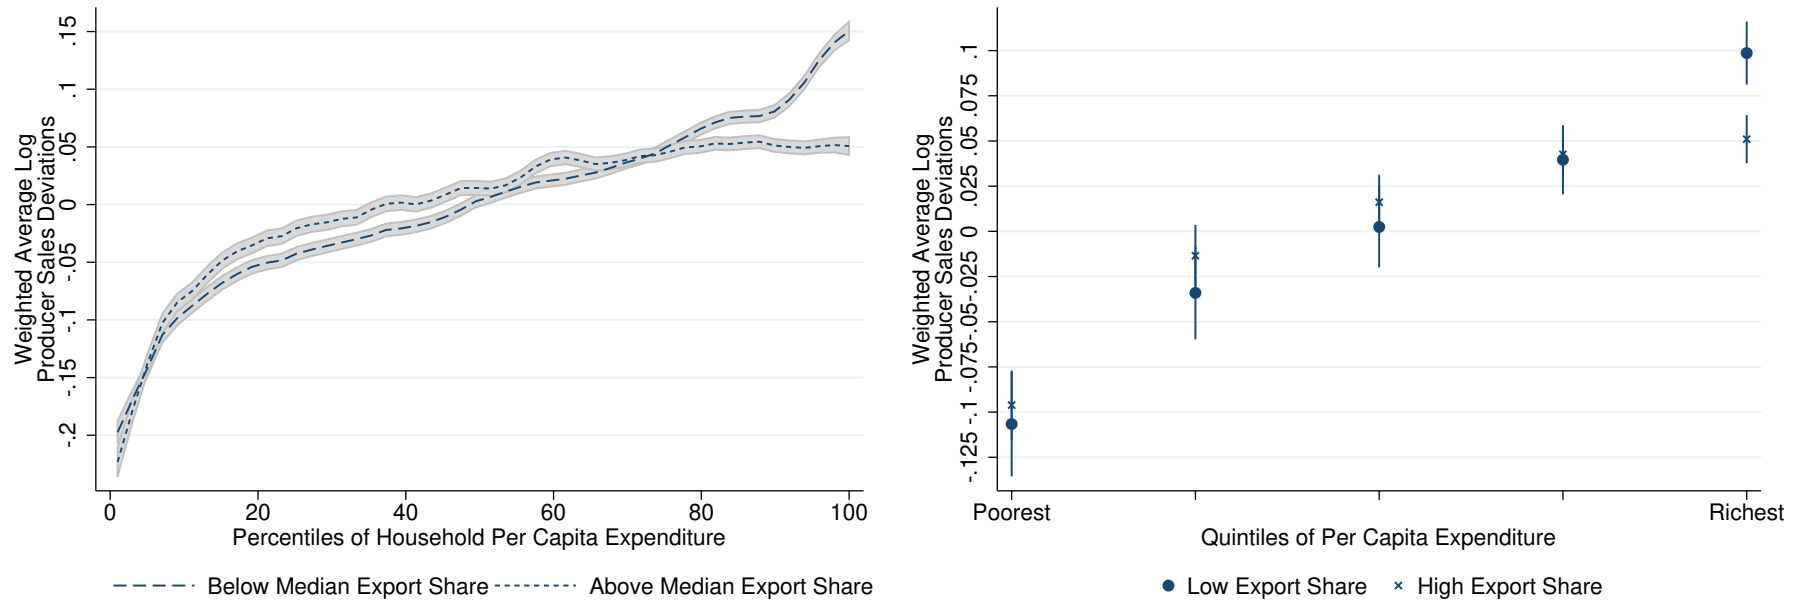
Notes: The graph depicts the average share of total reported retail consumption that is matched to the national firm size distribution (i.e. brands in the store scanner data) for each income quintile. See Section 3 for discussion.

Table A.3: Import Penetration and Export Shares across Nielsen Product Groups

Product Group	Import Penetration	Export Share	Product Group	Import Penetration	Export Share	Product Group	Import Penetration	Export Share
AUTOMOTIVE	0.3617847	0.2486256	GLASSWARE, TABLEWARE	0.2245018	0.1524106	SOFT DRINKS-NON-CARBONATED	0.0336043	0.0086848
BABY FOOD	0.1118066	0.0966579	GROOMING AIDS	0.1167166	0.1128873	SOFT GOODS	0.7952904	0.2435451
BABY NEEDS	0.1784844	0.1161126	GUM	0.1388452	0.0668655	SOUP	0.1121759	0.1031848
BAKED GOODS-FROZEN	0.0670658	0.0171679	HAIR CARE	0.1167166	0.1128873	SPICES, SEASONING, EXTRACTS	0.1891926	0.0300319
BAKING MIXES	0.0468524	0.0408634	HARDWARE, TOOLS	0.3464647	0.1938586	STATIONERY, SCHOOL SUPPLIES	0.2941211	0.0844581
BAKING SUPPLIES	0.0468524	0.0408634	HOUSEHOLD CLEANERS	0.0433067	0.0643612	SUGAR, SWEETENERS	0.273909	0.0688637
BATTERIES AND FLASHLIGHTS	0.2123691	0.2624945	HOUSEHOLD SUPPLIES	0.4019963	0.1488679	TABLE SYRUPS, MOLASSES	0.1118066	0.0966579
BEER	0.1441532	0.0169751	HOUSEWARES, APPLIANCES	0.2362423	0.1303144	TEA	0.1118066	0.0966579
BREAD AND BAKED GOODS	0.0670658	0.0171679	ICE	0.0351416	0.0361945	TOBACCO & ACCESSORIES	0.013934	0.0312302
BREAKFAST FOOD	0.0257685	0.0564193	ICE CREAM, NOVELTIES	0.0052825	0.0069596	TOYS & SPORTING GOODS	0.6588651	0.2392575
BREAKFAST FOODS-FROZEN	0.0861109	0.0842245	INSECTICIDS/PESTICIDS/RODENTICIDS	0.104745	0.1432804	UNPREP MEAT/POULTRY/SEAFOOD-FRZN	0.0913397	0.0872797
BUTTER AND MARGARINE	0.0680775	0.0123938	JAMS, JELLIES, SPREADS	0.1309384	0.0582835	VEGETABLES - CANNED	0.0960721	0.0443963
CANDY	0.0710144	0.0077201	JUICE, DRINKS - CANNED, BOTTLED	0.0336043	0.0086848	VEGETABLES AND GRAINS - DRIED	0.133023	0.2293571
CANNING, FREEZING SUPPLIES	0.0132684	0.0114468	JUICES, DRINKS-FROZEN	0.2088982	0.1062177	VEGETABLES-FROZEN	0.2088982	0.1062177
CARBONATED BEVERAGES	0.0336043	0.0086848	KITCHEN GADGETS	0.3589703	0.1280043	VITAMINS	0.2337178	0.0619368
CEREAL	0.0257685	0.0564193	LAUNDRY SUPPLIES	0.0562014	0.0770468	WINE	0.3268787	0.0572849
CHARCOAL, LOGS, ACCESSORIES	0.1214207	0.1521547	LIGHT BULBS, ELECTRIC GOODS	0.4504395	0.2750829	WRAPPING MATERIALS AND BAGS	0.1692727	0.0573695
CHEESE	0.0380319	0.0075332	LIQUOR	0.4521224	0.1505386	YEAST	0.0468524	0.0408634
COFFEE	0.0857798	0.069193	MEDICATIONS/REMEDIES/HEALTH AIDS	0.2935742	0.0980508	YOGURT	0.1191513	0.1491615
CONDIMENTS, GRAVIES, AND SAUCES	0.0915703	0.0452809	MEN'S TOILETRIES	0.1167166	0.1128873			
COOKIES	0.0710144	0.0077201	MILK	0.001964	0.0016345			
COOKWARE	0.6298456	0.1369722	NUTS	0.1585044	0.4884069			
COSMETICS	0.1167166	0.1128873	ORAL HYGIENE	0.1167166	0.1128873			
COT CHEESE, SOUR CREAM, TOPPING	0.0380319	0.0075332	PACKAGED MEATS-DELI	0.0641855	0.0614173			
COUGH AND COLD REMEDIES	0.2337178	0.0619368	PACKAGED MILK AND MODIFIERS	0.1191513	0.1491615			
CRACKERS	0.0128479	0.0210505	PAPER PRODUCTS	0.0799848	0.0817885			
DEODORANT	0.1167166	0.1128873	PASTA	0.1622465	0.0441698			
DESSERTS, GELATINS, SYRUP	0.120178	0.0163943	PERSONAL SOAP AND BATH ADDITIVES	0.0562014	0.0770468			
DESSERTS/FRUITS/TOPPINGS-FROZEN	0.1184343	0.0590694	PET CARE	0.4042157	0.4886733			
DETERGENTS	0.0562014	0.0770468	PET FOOD	0.0159851	0.0612656			
DIET AIDS	0.2337178	0.0619368	PHOTOGRAPHIC SUPPLIES	0.3283408	0.3013209			
DISPOSABLE DIAPERS	0.0991881	0.076225	PICKLES, OLIVES, AND RELISH	0.0915703	0.0452809			
DOUGH PRODUCTS	0.0468524	0.0408634	PIZZA/SNACKS/HORS D'OEUVRES-FRZN	0.0128479	0.0210505			
DRESSINGS/SALADS/PREP FOODS-DELI	0.1118066	0.0966579	PREPARED FOOD-DRY MIXES	0.1142222	0.1378016			
EGGS	0.1191513	0.1491615	PREPARED FOOD-READY-TO-SERVE	0.1118066	0.0966579			
ELECTRONICS, RECORDS, TAPES	0.7363864	0.5216842	PREPARED FOODS-FROZEN	0.1151989	0.0849323			
ELECTRONICS, RECORDS, TAPES	0.4504395	0.2750829	PUDDING, DESSERTS-DAIRY	0.1191513	0.1491615			
ETHNIC HABA	0.1167166	0.1128873	SALAD DRESSINGS, MAYO, TOPPING	0.0915703	0.0452809			
FEMININE HYGIENE	0.1167166	0.1128873	SANITARY PROTECTION	0.1167166	0.1128873			
FIRST AID	0.2337178	0.0619368	SEAFOOD - CANNED	0.6691664	0.3021904			
FLORAL, GARDENING	0.4937129	0.3243936	SEASONAL	0.1641153	0.2400124			
FLOUR	0.0224125	0.0223004	SEWING NOTIONS	0.3320014	0.3418083			
FRAGRANCES - WOMEN	0.1167166	0.1128873	SHAVING NEEDS	0.1167166	0.1128873			
FRESH MEAT	0.0641855	0.0614173	SHOE CARE	0.9280522	0.2199842			
FRESHENERS AND DEODORIZERS	0.1167166	0.1128873	SHORTENING, OIL	0.0996696	0.1159477			
FRUIT - CANNED	0.0960721	0.0443963	SKIN CARE PREPARATIONS	0.1167166	0.1128873			
FRUIT - DRIED	0.1142222	0.1378016	SNACKS	0.0128479	0.0210505			
			SNACKS, SPREADS, DIPS-DAIRY	0.0548211	0.0724421			


Notes: See Section 3 for discussion. Based on US trade data for 2005 at the 4-digit SIC product level. Export shares and import penetration are the shares of exports or imports in total production.

Figure A.14: Firm Heterogeneity Across Consumption Baskets - Robustness to Export Shares



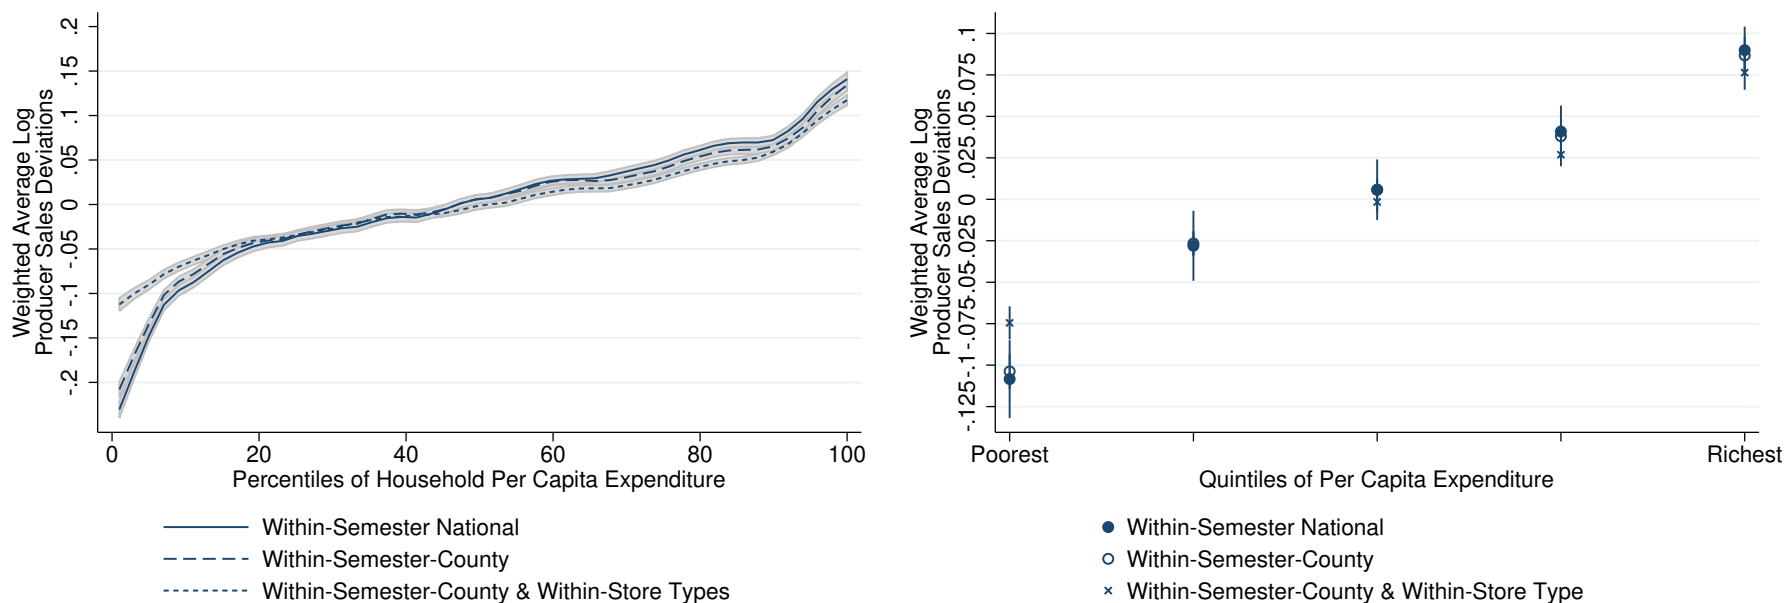
Notes: The fitted relationship in the left graph corresponds to a local polynomial regression. Standard errors in both graphs are clustered at the county level, and the displayed confidence intervals are at the 95% level. See Section 3 for discussion.

Figure A.15: Firm Heterogeneity Across Consumption Baskets - Robustness to Import Penetration



Notes: The fitted relationship in the left graph corresponds to a local polynomial regression. Standard errors in both graphs are clustered at the county level, and the displayed confidence intervals are at the 95% level. See Section 3 for discussion.

Figure A.16: The Role of Differential Access to Producers



61

Notes: The figure depicts deviations in weighted average log firm sales embodied in the consumption baskets of on average 59,000 US households during 18 half-year periods between 2006-14. The y-axis in both graphs displays weighted average deviations in log producer sales within more than 1000 product modules where the weights are household expenditure shares across producers of brands. National firm size deviations are based on the store scanner data. These firm size deviations are depicted across consumption baskets i) conditional on half-year fixed effects, ii) conditional on half-year-by-county fixed effects, and iii) conditional on half-year-by-county fixed effects and household consumption shares across 79 different store formats. The x-axis displays national percentiles of per capita total household retail expenditure per half year period (see Section 2). The fitted relationship in the left graph corresponds to a local polynomial regression. Standard errors in both graphs are clustered at the county level, and the displayed confidence intervals are at the 95% level. See Section 3 for discussion.

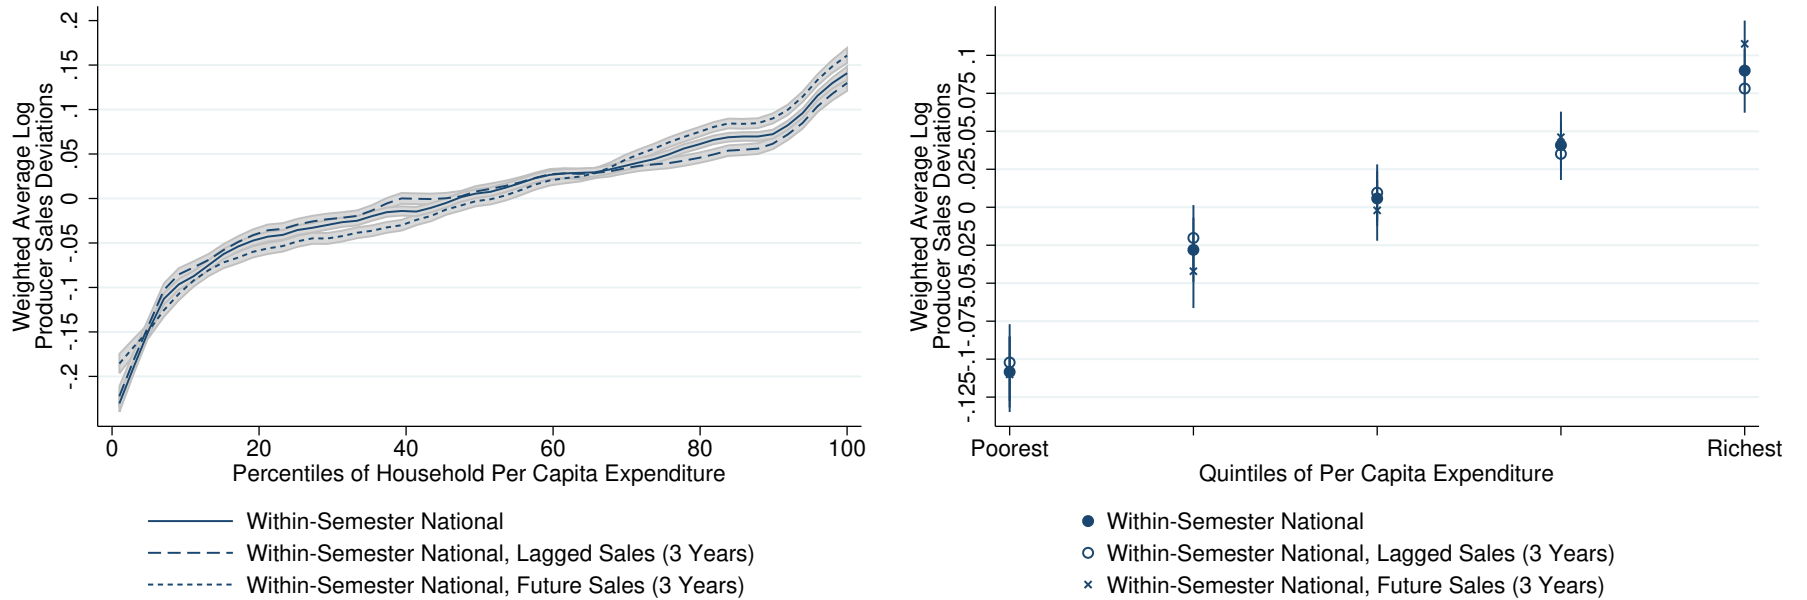
Figure A.17: The Role of Differential Pricing



20

Notes: The figure depicts deviations in log prices per unit paid across deviations in log brand sales (left) or UPC sales (right) relative to product module-by-half-year means. The figures plot this relationship separately for the average prices per unit paid for brands or UPCs as reported by the poorest and richest income quintiles respectively. The fitted relationships in both graph correspond to local polynomial regression. Standard errors are clustered at the product module level, and the displayed confidence intervals are at the 95% level. See Section 3 for discussion.

Figure A.18: The Role of Temporary Taste Shocks that Differ across Rich and Poor Households



Notes: The figure depicts deviations in weighted average log firm sales embodied in the consumption baskets of on average 59,000 US households during 18 half-year periods between 2006-14. The y-axis in both graphs displays weighted average deviations in log producer sales within more than 1000 product modules where the weights are household expenditure shares across producers of brands. National firm size deviations are based on the store scanner data. These firm size deviations are depicted across consumption baskets conditional on half-year fixed effects for i) same period firm size differences, ii) three-year lagged firm size differences, and iii) three-year future firm size differences. The x-axis displays national percentiles of per capita total household retail expenditure per half year period (see Section 2). The fitted relationship in the left graph corresponds to a local polynomial regression. Standard errors in both graphs are clustered at the county level, and the displayed confidence intervals are at the 95% level. Table A.1 provides descriptive statistics. See Section 3 for discussion.

Table A.4: Firms Alter Their Product Attributes

Fraction of Barcodes Replaced with New Barcodes with Identical Pack Sizes of Same Brand	
1st Half 2006	-
2nd Half 2006	0.108
1st Half 2007	0.077
2nd Half 2007	0.076
1st Half 2008	0.068
2nd Half 2008	0.064
1st Half 2009	0.052
2nd Half 2009	0.057
1st Half 2010	0.049
2nd Half 2010	0.067
1st Half 2011	0.053
2nd Half 2011	0.070
1st Half 2012	0.074
2nd Half 2012	-

Notes: See Section 3 for discussion.

Table A.5: Income Group Ratios of Within and Cross-Brand Price Elasticities

Dependent Variable: Change in Log Budget Shares	Cross-Brand Both IVs	Within-Brand Both IVs
(1- σ) All Households		
(1- σ) Below Median Quintiles	-1.288*** (0.0624)	-0.945*** (0.0804)
(1- σ) Median and Above Quintiles	-1.151*** (0.0325)	-1.019*** (0.0631)
Quintile-by-Module-by-County-by-Period FX	✓	✗
Quintile-by-Module-by-Brand-by-County-by-Period FX	✗	✓
Observations	9,285,679	16,582,717
First Stage F-Stat	312.5	410.8
Estimate of Ratio of σ 's (Poor/Rich)	1.063 (0.0303)	0.963 (0.0417)
95% Confidence Interval of Ratio	[1.0044, 1.123]	[0.882, 1.045]

Notes: See Section 4 for discussion. Standard errors are in parentheses below point estimates and clustered at the county level. ***, **, * indicate 1, 5 and 10 percent confidence levels.

Table A.6: Alternative Specifications for Estimating Price Elasticities

<i>Panel A: Pooled Estimates - Tornqvist Price Index</i>	Based on Mean Price (Baseline Estimate)				Based on Median Price			
	OLS	National IV	State IV	Both IVs	OLS	National IV	State IV	Both IVs
Dependent Variable: Change in Log Budget Shares								
(1- σ) All Households	0.257*** (0.0288)	-1.184*** (0.0356)	-1.090*** (0.0415)	-1.181*** (0.0316)	0.165*** (0.0241)	-1.153*** (0.0359)	-1.045*** (0.0379)	-1.145*** (0.0312)
Quintile-by-Module-by-County-by-Period FX	✓	✓	✓	✓	✓	✓	✓	✓
Observations	9,989,508	9,989,508	9,283,699	9,283,699	9,989,508	9,989,508	9,283,699	9,283,699
First Stage F-Stat		718.6	314.7	420.0		761.9	328.5	451.8
<i>Panel B: Pooled Estimates - Laspeyres Price Index</i>	Based on Mean Price				Based on Median Price			
Dependent Variable: Change in Log Budget Shares	OLS	National IV	State IV	Both IVs	OLS	National IV	State IV	Both IVs
(1- σ) All Households	0.259*** (0.0276)	-1.093*** (0.0366)	-1.016*** (0.0443)	-1.096*** (0.0334)	0.163*** (0.0232)	-1.079*** (0.0366)	-0.987*** (0.0404)	-1.075*** (0.0325)
Quintile-by-Module-by-County-by-Period FX	✓	✓	✓	✓	✓	✓	✓	✓
Observations	9,989,508	9,989,508	9,283,699	9,283,699	9,989,508	9,989,508	9,283,699	9,283,699
First Stage F-Stat		659.4	307.7	396.0		735.5	319.5	451.1
<i>Panel C: Pooled Estimates - Simple Mean Price Index</i>	Based on Mean Price				Based on Median Price			
Dependent Variable: Change in Log Budget Shares	OLS	National IV	State IV	Both IVs	OLS	National IV	State IV	Both IVs
(1- σ) All Households	0.259*** (0.0275)	-1.218*** (0.0397)	-1.118*** (0.0471)	-1.215*** (0.0365)	0.172*** (0.0234)	-1.190*** (0.0394)	-1.080*** (0.0432)	-1.183*** (0.0355)
Quintile-by-Module-by-County-by-Period FX	✓	✓	✓	✓	✓	✓	✓	✓
Observations	9,989,508	9,989,508	9,283,699	9,283,699	9,989,508	9,989,508	9,283,699	9,283,699
First Stage F-Stat		601.3	288.6	343.2		642.1	301.9	374.0

Notes: See Section 5 for discussion. Standard errors are in parentheses below point estimates and clustered at the county level. ***, **, * indicate 1, 5 and 10 percent confidence levels.

Table A.7: Full Cross of Price Elasticity Estimates by Household and Product Groups

<i>By Department and Household Group</i>	Beverages	Dairy	Dry Grocery	Frozen Foods	General Merchandise	Health and Beauty	Non-Food Grocery	Packaged Meat
Dependent Variable: Change in Log Budget Shares	Both IVs	Both IVs	Both IVs	Both IVs	Both IVs	Both IVs	Both IVs	Both IVs
(1- σ) Poorest Quintile	-1.137*** (0.320)	-0.753*** (0.155)	-1.520*** (0.165)	-1.426*** (0.302)	-1.126 (0.832)	-0.817** (0.326)	-1.168*** (0.356)	-1.486*** (0.489)
(1- σ) 2nd Poorest Quintile	-1.348*** (0.296)	-0.845*** (0.236)	-1.463*** (0.104)	-1.308*** (0.147)	-2.818*** (0.413)	-0.173 (0.324)	-1.504*** (0.250)	-1.233*** (0.296)
(1- σ) Median Quintile	-0.821** (0.336)	-0.667*** (0.103)	-1.322*** (0.0888)	-1.171*** (0.162)	-2.011*** (0.445)	-0.341 (0.207)	-0.819*** (0.168)	-1.293*** (0.404)
(1- σ) 2nd Richest Quintile	-1.112*** (0.209)	-0.901*** (0.0912)	-1.377*** (0.0759)	-1.306*** (0.190)	-2.943*** (0.512)	-0.274 (0.208)	-1.170*** (0.163)	-1.401*** (0.244)
(1- σ) Richest Quintile	-1.101*** (0.145)	-0.544*** (0.0924)	-1.211*** (0.0641)	-1.424*** (0.136)	-2.126*** (0.227)	-0.650*** (0.166)	-1.064*** (0.116)	-1.274*** (0.209)
Quintile-by-Module-by-County-by-Period FX	✓	✓	✓	✓	✓	✓	✓	✓
Observations	755,648	775,238	4,570,372	945,956	205,830	778,667	982,261	269,726
First Stage F-Stat	139.0	347.5	254.1	50.17	131.4	109.4	298.0	37.68

Notes: See Section 5 for discussion. Standard errors are in parentheses below point estimates and clustered at the county level. ***, **, * indicate 1, 5 and 10 percent confidence levels.

Table A.8: Parametric Specification of Heterogeneity in Price Elasticities


<i>Panel A: Pooled Estimates</i>					
Dependent Variable: Change in Log Budget Shares	Both IVs	Both IVs	Both IVs	Both IVs	Both IVs
Change in Log Price	-1.181*** (0.0316)	-1.940*** (0.439)	-1.314*** (0.0796)		
Change in Log Price X Log Per Capita Expenditure		0.105* (0.0600)		0.278*** (0.0595)	
Change in Log Price X Percentile of Per Capita Expenditure			0.00203** (0.00102)		0.00490*** (0.00103)
Quintile-by-Module-by-County-by-Period FX	✓	✓	✓	✓	✓
Brand-by-County-by-Period FX	✗	✗	✗	✓	✓
Observations	9,283,699	9,283,699	9,283,699	9,283,699	9,283,699
First Stage F-Stat	420	311.4	317.9	207.2	211.8

<i>Panel B: By Department Across Log Expenditure</i>								
Dependent Variable: Change in Log Budget Shares	Beverages	Dairy	Dry Grocery	Frozen Foods	General Merchandise	Health and Beauty	Non-Food Grocery	Packaged Meat
	Both IVs	Both IVs	Both IVs	Both IVs	Both IVs	Both IVs	Both IVs	Both IVs
Change in Log Price	-1.457 (1.552)	-1.774** (0.697)	-2.732*** (0.841)	-0.220 (0.882)	-3.451 (2.851)	0.420 (1.797)	-1.605 (1.440)	-1.911 (1.838)
Change in Log Price X Log Per Capita Expenditure	0.0512 (0.206)	0.148 (0.0973)	0.195* (0.116)	-0.154 (0.123)	0.151 (0.386)	-0.127 (0.248)	0.0698 (0.193)	0.0821 (0.248)
Quintile-by-Module-by-County-by-Period FX	✓	✓	✓	✓	✓	✓	✓	✓
Observations	755,648	775,238	4,570,372	945,956	205,830	778,667	982,261	269,726
First Stage F-Stat	331.8	136.8	235.4	64.55	79.21	142	404.4	61.54

<i>Panel C: By Department Across Percentiles</i>								
Dependent Variable: Change in Log Budget Shares	Beverages	Dairy	Dry Grocery	Frozen Foods	General Merchandise	Health and Beauty	Non-Food Grocery	Packaged Meat
	Both IVs	Both IVs	Both IVs	Both IVs	Both IVs	Both IVs	Both IVs	Both IVs
Change in Log Price	-1.158*** (0.341)	-0.888*** (0.140)	-1.565*** (0.145)	-1.181*** (0.154)	-2.468*** (0.550)	-0.369 (0.316)	-1.246*** (0.283)	-1.372*** (0.377)
Change in Log Price X Percentile of Per Capita Expenditure	0.00105 (0.00387)	0.00275 (0.00196)	0.00366* (0.00205)	-0.00228 (0.00221)	0.00165 (0.00676)	-0.00196 (0.00460)	0.00214 (0.00340)	0.000795 (0.00466)
Quintile-by-Module-by-County-by-Period FX	✓	✓	✓	✓	✓	✓	✓	✓
Observations	755,648	775,238	4,570,372	945,956	205,830	778,667	982,261	269,726
First Stage F-Stat	319.9	135.2	235.9	74.18	75.59	137.9	416.5	69.86


Notes: See Section 5 for discussion. Standard errors are in parentheses below point estimates and clustered at the level of counties. ***, **, * indicate 1, 5 and 10 percent confidence levels.

Figure A.19: Distribution of Weighted Average Unit Values across Consumption Baskets




Notes: The figure depicts deviations in weighted average log firm unit values embodied in the consumption baskets of on average 59,000 US households during 18 half-year periods between 2006-14. The y-axis in both graphs displays weighted average deviations in log producer unit values within more than 1000 product modules where the weights are household expenditure shares across producers of brands. In the first step, we calculate brand-level deviations from mean log national unit values within product module-by-half-year cells from the store-level scanner data, where brand-level unit values are expenditure weighted means across multiple barcodes within the brand. In the second step, these are then matched to brand-level half yearly household expenditure weights in the home scanner data. The final step is to collapse these data to weighted average log unit value deviations embodied in household consumption baskets. The x-axis displays national percentiles of per capita total household retail expenditure per half year period (see Section 2). The fitted relationship in the left graph corresponds to a local polynomial regression. Standard errors in both graphs are clustered at the county level, and the displayed confidence intervals are at the 95% level. See Section 5 for discussion.

Figure A.20: Unit Values and Firm Size



Notes: The figure depicts deviations in log firm unit values across the firm size distribution for 18 half-yearly cross-sections between 2006-2014. The y-axis displays deviations of log unit values relative to product module-by-half-year means. The x-axis displays deviations of log firm sales at the same level. The fitted relationships correspond to local polynomial regressions. Standard errors in both graphs are clustered at the level of product modules, and the displayed confidence intervals are at the 95% level. See Section 5 for discussion.

Figure A.21: Households Agree on Product Quality Evaluations (But Rich Households Value Quality Relatively More)



Notes: The figure depicts the relationship between deviations in log brand quality or quality-adjusted prices and deviations in log firm total sales for on average more than 150,000 producers of brands during 18 half-year periods between 2006-14. We estimate brand-level quality and quality-adjusted prices as evaluated by each quintile of total household per capita expenditure as discussed in Sections 4 and 5.

Table A.9: Heterogeneous Quality Evaluations

Dependent Variable:	Log Brand Sales by Household Group	Poorest Quintile		2nd Poorest Quintile		Median Quintile		2nd Richest Quintile		Richest Quintile	
		OLS	IV	OLS	IV	OLS	IV	OLS	IV	OLS	IV
ALL PRODUCT MODULES	Log Average Brand Sales	0.923*** (0.00359)	0.923*** (0.00391)	0.957*** (0.00412)	0.959*** (0.00433)	1.037*** (0.00261)	1.036*** (0.00230)	1.044*** (0.00281)	1.042*** (0.00277)	1.037*** (0.00410)	1.040*** (0.00469)
	Product Module-by-Period FX	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓
	Observations	1,854,522	1,330,947	1,854,522	1,330,947	1,854,522	1,330,947	1,854,522	1,330,947	1,854,522	1,330,947
	Number of Product Module Clusters	1046	1030	1046	1030	1046	1030	1046	1030	1046	1030
BEVERAGES	Log Average Brand Sales	0.840*** (0.00423)	0.836*** (0.00495)	0.892*** (0.00402)	0.902*** (0.00383)	1.082*** (0.00478)	1.078*** (0.00500)	1.102*** (0.00548)	1.097*** (0.00470)	1.084*** (0.00380)	1.087*** (0.00438)
	Product Module-by-Period FX	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓
	Observations	182,279	123,506	182,279	123,506	182,279	123,506	182,279	123,506	182,279	123,506
	Number of Product Module Clusters	69	68	69	68	69	68	69	68	69	68
DAIRY	Log Average Brand Sales	0.886*** (0.00420)	0.888*** (0.00490)	0.905*** (0.00247)	0.907*** (0.00300)	1.067*** (0.00218)	1.062*** (0.00262)	1.071*** (0.00247)	1.067*** (0.00245)	1.072*** (0.00430)	1.076*** (0.00494)
	Product Module-by-Period FX	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓
	Observations	116,853	90,097	116,853	90,097	116,853	90,097	116,853	90,097	116,853	90,097
	Number of Product Module Clusters	46	45	46	45	46	45	46	45	46	45
DRY GROCERY	Log Average Brand Sales	0.883*** (0.00225)	0.883*** (0.00235)	0.913*** (0.00126)	0.917*** (0.00131)	1.061*** (0.00106)	1.056*** (0.00107)	1.069*** (0.00121)	1.067*** (0.00119)	1.074*** (0.00187)	1.077*** (0.00229)
	Product Module-by-Period FX	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓
	Observations	718,629	530,010	718,629	530,010	718,629	530,010	718,629	530,010	718,629	530,010
	Number of Product Module Clusters	398	392	398	392	398	392	398	392	398	392
FROZEN FOODS	Log Average Brand Sales	0.969*** (0.00411)	0.966*** (0.00584)	1.000*** (0.00259)	0.999*** (0.00284)	1.006*** (0.00225)	1.003*** (0.00275)	1.011*** (0.00230)	1.012*** (0.00280)	1.014*** (0.00362)	1.020*** (0.00471)
	Product Module-by-Period FX	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓
	Observations	126,928	93,633	126,928	93,633	126,928	93,633	126,928	93,633	126,928	93,633
	Number of Product Module Clusters	78	76	78	76	78	76	78	76	78	76
GENERAL MERCHANDISE	Log Average Brand Sales	0.942*** (0.00243)	0.941*** (0.00312)	0.976*** (0.00180)	0.978*** (0.00215)	1.021*** (0.00139)	1.018*** (0.00136)	1.027*** (0.00167)	1.027*** (0.00201)	1.034*** (0.00248)	1.036*** (0.00274)
	Product Module-by-Period FX	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓
	Observations	197,828	132,375	197,828	132,375	197,828	132,375	197,828	132,375	197,828	132,375
	Number of Product Module Clusters	143	140	143	140	143	140	143	140	143	140
HEALTH & BEAUTY CARE	Log Average Brand Sales	0.959*** (0.00248)	0.960*** (0.00282)	0.996*** (0.00260)	0.998*** (0.00326)	1.015*** (0.00400)	1.014*** (0.00335)	1.020*** (0.00211)	1.018*** (0.00214)	1.009*** (0.00453)	1.010*** (0.00549)
	Product Module-by-Period FX	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓
	Observations	284,425	196,046	284,425	196,046	284,425	196,046	284,425	196,046	284,425	196,046
	Number of Product Module Clusters	173	172	173	172	173	172	173	172	173	172
NON-FOOD GROCERY	Log Average Brand Sales	0.816*** (0.00186)	0.820*** (0.00228)	0.840*** (0.00244)	0.846*** (0.00297)	1.109*** (0.00275)	1.105*** (0.00276)	1.121*** (0.00162)	1.117*** (0.00209)	1.113*** (0.00372)	1.113*** (0.00426)
	Product Module-by-Period FX	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓
	Observations	190,086	136,350	190,086	136,350	190,086	136,350	190,086	136,350	190,086	136,350
	Number of Product Module Clusters	128	126	128	126	128	126	128	126	128	126
PACKAGED MEAT	Log Average Brand Sales	0.978*** (0.00510)	0.978*** (0.00499)	0.990*** (0.00392)	0.991*** (0.00315)	1.003*** (0.00526)	1.002*** (0.00573)	1.013*** (0.00316)	1.010*** (0.00400)	1.016*** (0.00578)	1.020*** (0.00736)
	Product Module-by-Period FX	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓
	Observations	37,494	28,930	37,494	28,930	37,494	28,930	37,494	28,930	37,494	28,930
	Number of Product Module Clusters	11	11	11	11	11	11	11	11	11	11

Notes: See Section 5 for discussion. Standard errors are in parentheses below point estimates and clustered at the level of product modules. ***, **, * indicate 1, 5 and 10 percent confidence levels.

Table A.10: Technology Parameter Estimates

ALL PRODUCT GROUPS				
Dependent Variable:	Cross-Section		Panel Data	
Log Product Quality or Changes in Log Quality	OLS	IV	OLS	IV
Log Firm Scale or Changes in Log Firm Scale (β)	1.1132*** (0.0309)	1.1352*** (0.0307)	1.1588*** (0.0461)	0.4805*** (0.0643)
ξ Parameter	0.82	0.82	0.31	0.31
Observations	1,330,976	1,330,976	1,422,244	1,422,244
Number of Clusters	1,031	1,031	994	994
First Stage F-Stat	311103.21		253.09	

Dependent Variable:	BEVERAGES				DAIRY				DRY GROCERY				FROZEN FOODS			
	Cross-Section		Panel Data		Cross-Section		Panel Data		Cross-Section		Panel Data		Cross-Section		Panel Data	
	OLS	IV	OLS	IV	OLS	IV	OLS	IV	OLS	IV	OLS	IV	OLS	IV	OLS	IV
Log Firm Scale or Changes in Log Firm Scale (β)	1.0374*** (0.0212)	1.0612*** (0.0261)	1.0609*** (0.0288)	-0.6287 (1.0353)	1.487*** (0.0054)	1.5221*** (0.008)	1.3815*** (0.0547)	0.3034 (0.3499)	0.8012*** (0.0026)	0.8248*** (0.0031)	0.863*** (0.0072)	0.4053*** (0.0655)	0.8412*** (0.0034)	0.8517*** (0.0044)	0.8183*** (0.0121)	0.1563 (0.1676)
ξ Parameter	0.77	0.77	0.01	0.01	0.88	0.88	0.88	0.88	0.88	0.88	0.01	0.01	0.19	0.19	0.96	0.96
Observations	123,509	123,509	102,141	102,141	90,097	90,097	67,235	67,235	530,020	530,020	573,503	573,503	93,635	93,635	92,038	92,038
Number of Clusters	68	68	66	66	45	45	44	44	393	393	381	381	76	76	74	74
First Stage F-Stat	48927.91		4.39		29611.79		20.18		134277.8		140.98		21249.68		26.77	

Dependent Variable:	GENERAL MERCHANDISE				HEALTH & BEAUTY CARE				NON-FOOD GROCERY				PACKAGED MEAT			
	Cross-Section		Panel Data		Cross-Section		Panel Data		Cross-Section		Panel Data		Cross-Section		Panel Data	
	OLS	IV	OLS	IV	OLS	IV	OLS	IV	OLS	IV	OLS	IV	OLS	IV	OLS	IV
Log Firm Scale or Changes in Log Firm Scale (β)	0.5275*** (0.005)	0.5367*** (0.0057)	0.5204*** (0.0096)	0.2098*** (0.0798)	2.2156*** (0.0069)	2.2666*** (0.0078)	2.1702*** (0.0153)	1.6298*** (0.1428)	0.976*** (0.0069)	0.9783*** (0.0071)	1.0201*** (0.0123)	0.2347 (0.3613)	0.8584*** (0.0062)	0.8736*** (0.0059)	0.7816*** (0.0294)	-0.0258 (0.367)
ξ Parameter	0.94	0.94	0.94	0.94	0.96	0.96	0.96	0.96	0.42	0.42	0.05	0.05	0.97	0.97	0.97	0.97
Observations	132,383	132,383	129,686	129,686	196,051	196,051	261,263	261,263	136,351	136,351	171,067	171,067	28,930	28,930	25,311	25,311
Number of Clusters	140	140	130	130	172	172	168	168	126	126	121	121	11	11	10	10
First Stage F-Stat	29222.36		46.75		36443.05		56.72		39746		9.74		40138.38		12.38	

Notes: See Section 5 for discussion. Standard errors are in parentheses below point estimates and clustered at the level of product modules. ***, **, * indicate 1, 5 and 10 percent confidence levels.

Table A.11: Decomposition of Counterfactuals

Difference in Retail Inflation (Richest - Poorest Quintile)	Counterfactual 1: Progressive Income Taxes		Counterfactual 2: Closing Corporate Tax Loopholes	
(0) Direct Incidence of Corporate Taxes	0.000 (0) [0]	(0%)	0.764 (0.107) [0]	(40%)
(1) Change in Weighted Average Product Quality	0.311 (0.021) [0.086]	(10%)	0.566 (0.024) [0.119]	(30%)
(2) Asymmetric Scale Effect	2.474 (0.158) [0.223]	(79%)	0.209 (0.022) [0.047]	(11%)
(3) Asymmetric Changes in Markups	-0.027 (0.01) [0.101]	(-1%)	-0.099 (0.01) [0.013]	(-5%)
(4) Love of Variety	0.183 (0.007) [0.622]	(6%)	0.144 (0.006) [0.068]	(8%)
(5) Asymmetric Effect of Exit	0.000 (0) [0]	(0%)	0.000 (0) [0]	(0%)
(6) Between-Group Effect	0.195 (0.036) [0.046]	(6%)	0.320 (0.031) [0.035]	(17%)
Total Effect	3.136 (0.174) [0.632]	(100%)	1.903 (0.11) [0.23]	(100%)


Notes: See Section 6 for discussion and Appendix 5 for details on decomposition. Robust standard errors across 18 six-month periods are in parentheses below point estimates. Bootstrapped standard errors are in square brackets.

Table A.12: Robustness to Alternative Parameters

Difference in Inflation (Richest - Poorest Quintiles)	Baseline	Tech Parameter Estimates from Cross-Section	Price Elasticities x 1.5	Price Elasticities x 2	Price Elasticities - 1	Price Elasticities - 2	Parametric Specification of Heterogeneity in $\sigma_n(z)$
Counterfactual 1: Progressive Income Taxes	3.136 (0.174)	3.485 (0.120)	3.536 (0.188)	3.792 (0.195)	4.056 (0.212)	6.682 (0.268)	3.245 (0.164)
Counterfactual 2: Closing Corporate Tax	1.903 (0.110)	2.336 (0.188)	2.097 (0.118)	2.150 (0.133)	2.127 (0.123)	1.788 (0.156)	1.776 (0.112)


Notes: See Section 6 for discussion. Robust standard errors across 18 six-month periods are in parentheses below point estimates.

Figure A.22: Counterfactual 1: Alternative Use of Tax Proceeds



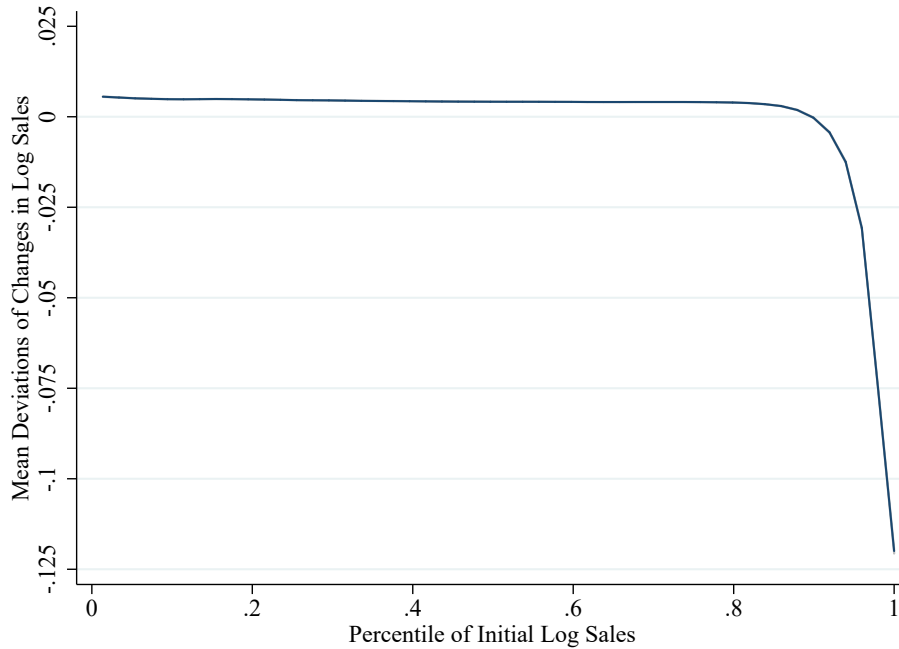
Notes: Graph displays counterfactual differences in retail inflation averaged across 18 half-year cross-sections of data. 95% confidence intervals are based on robust standard errors. See Section 6 for discussion.

Figure A.23: Counterfactual 1: Inflation Differences across Product Departments




Notes: Graph displays counterfactual differences in retail inflation averaged across 18 half-year cross-sections of data. 95% confidence intervals are based on robust standard errors. See Section 6 for discussion.

Figure A.24: Counterfactual 2: Compression of Firm Size Distribution




Notes: The graph plots mean deviations in sales growth as a function of initial firm sales within product modules, averaged across 18 half year periods. 95% confidence intervals based on standard errors that are clustered at the level of product modules. Section 6 for discussion.

Figure A.25: Counterfactual 2: Alternative Definition of Tax Entities




Notes: Graph displays counterfactual differences in retail inflation averaged across 18 half-year cross-sections of data. 95% confidence intervals are based on robust standard errors. See Section 6 for discussion.

Figure A.26: Counterfactual 2: Inflation Differences across Product Departments



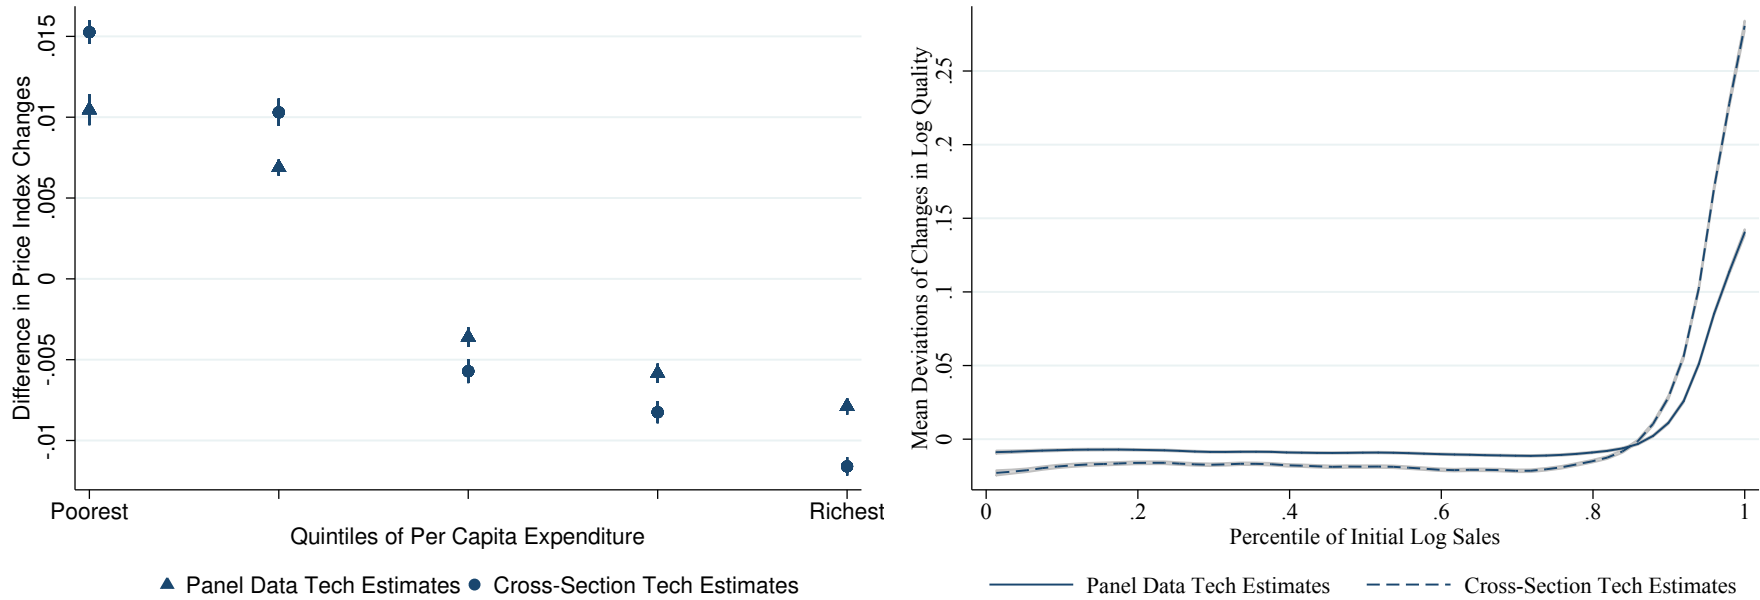
Notes: Graph displays counterfactual differences in retail inflation averaged across 18 half-year cross-sections of data. 95% confidence intervals are based on robust standard errors. See Section 6 for discussion.

Figure A.27: Counterfactuals Under Monopolistic Competition



Notes: The left panel displays differences in retail inflation in the first counterfactual (progressive taxes). The right panel displays differences in retail inflation in the second counterfactual (corporate taxes). 95% confidence intervals are based on robust standard errors across 18 half-year periods. See Section 6 for discussion.

Figure A.28: Inflation Differences and Quality Upgrading Due to Bilateral Trade Liberalization



Notes: 95% confidence intervals in the left panel are based on robust standard errors across 18 half-year periods. The right panel is based on counterfactual changes across producers of brands within product modules and within 18 six-month periods, and the confidence intervals are based on standard errors that are clustered at the level of product modules. See Section 6 and Appendix 5.C for discussion.

Appendix 2: Mathematical Appendix

2.A) First-Order Conditions in Markups and Quality

For a given firm with productivity a , we can write profits as a function of markups μ (ratio of price p to marginal cost c) and quality ϕ as follows:

$$\pi_n(a, \mu, \phi) = \left(1 - \frac{1}{\mu}\right) \int_z x_n(a, z, \mu, \phi) dH(z) - f_n(\phi) - f_{n0}$$

where fixed costs depend on quality such that:

$$f_n(\phi) = \beta_n b_n \phi^{\frac{1}{\beta_n}}$$

and sales to income group z satisfy:

$$x_n(a, z, \mu, \phi) = \alpha_n(z) E(z) P_n(z)^{\sigma_n(z)-1} a^{\sigma_n(z)-1} \mu^{1-\sigma_n(z)} \phi^{(\sigma_n(z)-1)(\gamma_n(z)-\xi_n)}$$

with $E(z)$ referring to total retail expenditure by consumer of income group z . Note that $\frac{\partial \log P_n}{\partial \log \mu} = 1 - s_n$, i.e. decreases with market share $s_n(a, z, \mu, \phi) = \frac{x_n}{\alpha_n E}$ of the firm within income group z .

Markups: For markups μ , the first-order condition yields:

$$0 = \frac{\partial \pi_n}{\partial \log \mu} = - \left(1 - \frac{1}{\mu}\right) \int_z (\sigma_n(z) - 1)(1 - s_n) x_n dH(z) + \frac{1}{\mu} \int_z x_n dH(z) \quad (22)$$

Hence optimal markups satisfy:

$$\mu_n(a) = 1 + \frac{\int_z x_n(a, z) dH(z)}{\int_z (\sigma_n(z) - 1)(1 - s_n(a, z)) x_n(a, z) dH(z)}$$

where $x_n(a, z)$ and $s_n(a, z)$ refer to sales and market share of firm a among consumers z .

Quality: For quality ϕ , we obtain the following first-order condition:

$$0 = \frac{\partial \pi_n}{\partial \log \phi} \quad (23)$$

$$= \left(1 - \frac{1}{\mu}\right) \int_z (\sigma_n(z) - 1)(\gamma_n(z) - \xi_n)(1 - s_n(z, a)) x_n(z, a) dH(z) - b_n \phi^{\frac{1}{\beta_n} - 1} \quad (24)$$

With μ as above and with $\tilde{\gamma}_n(a)$ defined as:

$$\tilde{\gamma}_n(a) = \frac{\int_z \gamma_n(z) (\sigma_n(z) - 1)(1 - s_n) x_n dH(z)}{\int_z (\sigma_n(z) - 1)(1 - s_n) x_n dH(z)}$$

we obtain the expression in the text for optimal quality:

$$\phi_n(a) = \left(\frac{\tilde{\gamma}_n(a) - \xi_n}{b_n \mu_n(a)} \cdot X_n(a) \right)^{\beta_n}$$

where $X_n(a)$ denotes total sales of firm with productivity a .

2.B) Second-Order Conditions in Markups and Quality

To ensure the uniqueness of equilibrium in prices and quality, we need to verify that the Hessian is definite negative in markups and quality. The Hessian is definite negative if these two conditions are satisfied:

$$\frac{\partial^2 \pi_n}{\partial \log \mu^2} < 0$$

and

$$\frac{\partial^2 \pi_n}{\partial \log \mu^2} \frac{\partial^2 \pi_n}{\partial \log \phi^2} > \left(\frac{\partial^2 \pi_n}{\partial \log \phi \partial \log \mu} \right)^2$$

We first examine the first inequality, which ensures that the first-order condition for markup μ leads to a unique solution for a given level of quality ϕ .

Second-Order Condition in Markups: For markups μ , the derivative of 22 yields:

$$\begin{aligned} \frac{\partial^2 \pi}{\partial \log \mu^2} &= \left(1 - \frac{1}{\mu}\right) \int_z (\sigma_n(z) - 1)^2 (1 - s_n)(1 - 2s_n) x_n dH(z) \\ &\quad - \frac{1}{\mu} \int_z [1 + 2(\sigma_n(z) - 1)(1 - s_n)] x_n dH(z) \end{aligned}$$

The second order condition in markups is satisfied if:

$$\left(1 - \frac{1}{\mu}\right) \int_z (\sigma_n(z) - 1)^2 (1 - s_n)(1 - 2s_n) x_n dH(z) < \frac{1}{\mu} \int_z [1 + 2(\sigma_n(z) - 1)(1 - s_n)] x_n dH(z)$$

where μ satisfies the first-order condition:

$$\left(1 - \frac{1}{\mu}\right) \int_z (\sigma_n(z) - 1)(1 - s_n) x_n dH(z) = \frac{1}{\mu} \int_z x_n dH(z)$$

Hence, the second-order condition in μ holds if:

$$\frac{\int_z (\sigma_n(z) - 1)^2 (1 - s_n)(1 - 2s_n) x_n dH(z)}{\int_z (\sigma_n(z) - 1)(1 - s_n) x_n dH(z)} < 1 + 2 \frac{\int_z (\sigma_n(z) - 1)(1 - s_n) x_n dH(z)}{\int_z x_n dH(z)}$$

A sufficient condition is:

$$\frac{\int_z (\sigma_n(z) - 1)^2 (1 - s_n)^2 x_n dH(z)}{\int_z (\sigma_n(z) - 1)(1 - s_n) x_n dH(z)} < 1 + 2 \frac{\int_z (\sigma_n(z) - 1)(1 - s_n) x_n dH(z)}{\int_z x_n dH(z)}$$

In this last inequality, the left-hand side is smaller than $\max_z (\sigma_n(z) - 1)(1 - s_n)$ while the right-hand side is not smaller than $1 + 2 \min_z (\sigma_n(z) - 1)(1 - s_n)$. With $\max_z (\sigma_n(z) - 1)(1 - s_n) < 1 + 2 \min_z (\sigma_n(z) - 1)(1 - s_n)$, the inequality above is satisfied. We verify empirically that this inequality holds given our estimates. More generally, this holds when $\sigma_n(z)$ is not too heterogeneous across consumers. Note also that this inequality is always satisfied under monopolistic competition when σ_n is identical across income groups.

Second-Order Condition in Quality: Using equation (24) and $\frac{\partial \log x_n}{\partial \log \phi} = \frac{\partial \log s_n}{\partial \log \phi} = (\sigma_n(z) - 1)(\gamma_n(z) - \xi_n)(1 - s_n)$, the second derivative in quality ϕ is:

$$\frac{\partial^2 \pi_n}{\partial \log \phi^2} = \left(1 - \frac{1}{\mu}\right) \int_z (\sigma_n(z) - 1)^2 (\gamma_n(z) - \xi_n)^2 [(1 - s_n)^2 - (1 - s_n)s_n] x_n dH(z) - \frac{b_n}{\beta_n} \phi^{\frac{1}{\beta_n}}$$

$$= \left(1 - \frac{1}{\mu}\right) \int_z (\sigma_n(z) - 1)^2 (\gamma_n(z) - \xi_n)^2 (1 - s_n)(1 - 2s_n)x_n dH(z) - \frac{b_n}{\beta_n} \phi^{\frac{1}{\beta_n}}$$

This second derivative is negative when β_n is small enough. More specifically, when quality satisfies the first order condition (24), this second derivative is negative if and only if:

$$\beta_n \cdot \frac{\int_z (\sigma_n(z) - 1)^2 (\gamma_n(z) - \xi_n)^2 (1 - s_n)(1 - 2s_n)x_n(z, a) dH(z)}{\int_z (\sigma_n(z) - 1)(\gamma_n(z) - \xi_n)(1 - s_n)x_n(z, a) dH(z)} < 1$$

A sufficient condition is that $\beta_n(\sigma_n(z) - 1)(\gamma_n(z) - \xi_n) < 1$ for all income groups z . Condition $\beta_n(\sigma_n(z) - 1)(\gamma_n(z) - \xi_n) < 1$ (for all z) is also a necessary condition to ensure that the second derivative in quality is negative irrespective of the patterns of sales and market shares across income groups.

Joint Second-Order Condition in Quality and Markups: The cross derivative in quality and markups is:

$$\begin{aligned} \frac{\partial^2 \pi_n}{\partial \log \phi \partial \log \mu} &= - \left(1 - \frac{1}{\mu}\right) \int_z (\sigma_n(z) - 1)^2 (\gamma_n(z) - \xi_n)(1 - s_n)(1 - 2s_n)x_n dH(z) \\ &\quad + \frac{1}{\mu} \int_z (\sigma_n(z) - 1)(\gamma_n(z) - \xi_n)(1 - s_n)x_n dH(z) \end{aligned}$$

When μ satisfied the first order condition, this yields:

$$\begin{aligned} \frac{\partial^2 \pi_n}{\partial \log \phi \partial \log \mu} &= \left(\frac{\int_z x_n dH(z)}{\mu} \right) \left(\frac{(\int_z (\sigma_n(z) - 1)(\gamma_n(z) - \xi_n)(1 - s_n)x_n dH(z))}{\int_z x_n dH(z)} \right. \\ &\quad \left. - \frac{\int_z (\sigma_n(z) - 1)^2 (\gamma_n(z) - \xi_n)(1 - s_n)(1 - 2s_n)x_n dH(z)}{\int_z (\sigma_n(z) - 1)(1 - s_n)x_n dH(z)} \right) \end{aligned}$$

In addition to the second-order conditions in markups and quality, $\frac{\partial^2 \pi_n}{\partial \log \mu^2} < 0$ and $\frac{\partial^2 \pi_n}{\partial \log \phi^2} < 0$, the Hessian is definite negative only if its determinant is positive, i.e. if $\frac{\partial^2 \pi_n}{\partial \log \mu^2} \frac{\partial^2 \pi_n}{\partial \log \phi^2} > \left(\frac{\partial^2 \pi_n}{\partial \log \phi \partial \log \mu}\right)^2$. Using the expressions for second and cross derivative, using $\frac{1}{\mu-1} = \frac{\int_z (\sigma_n(z) - 1)(1 - s_n)x_n dH(z)}{\int_z x_n dH(z)}$ and multiplying each side by $\beta_n^2(\mu - 1)$, this inequality is equivalent to:

$$\begin{aligned} &\left(\frac{(\int_z \theta_n(z)(1 - s_n)x_n dH(z)) (\int_z (\sigma_n(z) - 1)(1 - s_n)x_n dH(z))}{\int_z x_n dH(z)} - \frac{\int_z \theta_n(z)(\sigma_n(z) - 1)(1 - s_n)(1 - 2s_n)x_n dH(z)}{\int_z x_n dH(z)} \right)^2 \\ &< \left[\frac{\int_z \theta_n(z)(1 - s_n)x_n dH(z)}{\int_z x_n dH(z)} - \frac{\int_z \theta_n(z)^2(1 - s_n)(1 - 2s_n)x_n dH(z)}{\int_z x_n dH(z)} \right] \left[\frac{\mu + 1}{(\mu - 1)^2} - \frac{\int_z (\sigma_n(z) - 1)^2(1 - s_n)(1 - 2s_n)x_n dH(z)}{\int_z x_n dH(z)} \right] \end{aligned}$$

where we denote $\theta_n(z) = \beta_n(\sigma_n(z) - 1)(\gamma_n(z) - \xi_n)$.

We check numerically that this condition holds for all brands given our parameter estimates. Again, this inequality holds when β_n is not too large (i.e. $\theta_n(z)$ is not too large) or when there is a small covariance between $\sigma_n(z)$ and $\theta_n(z) = \beta_n(\sigma_n(z) - 1)(\gamma_n(z) - \xi_n)$ (the left-hand side of this inequality is similar to a covariance term). Under monopolistic competition, this inequality can be more explicitly rewritten as an upper bound on the covariance between $\sigma_n(z) - 1$ and $\theta_n(z)$:

$$\begin{aligned} &\left(\frac{(\int_z \theta_n(z)x_n dH(z)) \int_z (\sigma_n(z) - 1)x_n dH(z)}{\int_z x_n dH(z)} - \frac{\int_z \theta_n(z)(\sigma_n(z) - 1)x_n dH(z)}{\int_z x_n dH(z)} \right)^2 \\ &< \left(\frac{\int_z \theta_n(z)(1 - \theta_n(z))x_n dH(z)}{\int_z x_n dH(z)} \right) \left[\frac{\mu + 1}{(\mu - 1)^2} - \frac{\int_z (\sigma_n(z) - 1)^2 x_n dH(z)}{\int_z x_n dH(z)} \right] \end{aligned}$$

2.C) A Simple Case with Closed-Form Solutions

A simple and tractable case with closed-form solution assumes two groups z_0 and z_1 with homogeneous price elasticity $\sigma(z_1) = \sigma(z_0)$ under monopolistic competition, and a minimal taste for quality among income group z_0 (e.g. poorest income group): $\gamma_n(z_0) = \xi_n$. Another case with closed-form solution is one with homogeneous consumers (single z_1) is a special subcase of this two-group case.

Denote by $H_n(z)$ the fraction of total expenditures corresponding to consumers z within product module n , and by E_n total expenditures (normalizing the mass of consumers to one in each category). In that case, we have:

$$\tilde{\gamma}_n(a) = \xi_n + \frac{(\gamma_n(z_1) - \xi_n) x_n(z_1, a)}{X_n(a)}$$

where $X_n(a) = x_n(z_1, a) + x_n(z_0, a)$ denotes total sales of firm a . Let $\mu_n = \frac{\sigma_n}{\sigma_n - 1}$ denote markups. We get quality:

$$\begin{aligned} \phi_n(a) &= \left(\frac{\gamma_n(z_1) - \xi_n}{b_n \mu_n} x_n(z_1, a) \right)^{\beta_n} \\ &= \left(\frac{\gamma_n(z_1) - \xi_n}{b_n \mu_n} a^{\sigma_n - 1} \phi_{n1}(a)^{(\sigma_n - 1)(\gamma_n(z_1) - \xi_n)} H_n(z_1) E_n P_n(z_1)^{\sigma_n - 1} \right)^{\beta_n} \\ &= \left(\frac{\gamma_n(z_1) - \xi_n}{b_n \mu_n} a^{\sigma_n - 1} H_n(z_1) E_n P_n(z_1)^{\sigma_n - 1} \right)^{\frac{\beta_n}{1 - \beta_n(\sigma_n - 1)(\gamma_n(z_1) - \xi_n)}} \end{aligned}$$

Sales to income group z_1 are then:

$$x_n(z_1, a) = \left(\frac{a P_n(z_1)}{\mu_n} \right)^{\frac{\sigma_n - 1}{1 - \beta_n(\sigma_n - 1)(\gamma_n(z_1) - \xi_n)}} \left(\frac{\gamma_n(z_1) - \xi_n}{b_n \mu_n} H_n(z_1) E_n \right)^{\frac{\beta_n(\sigma_n - 1)(\gamma_n(z_1) - \xi_n)}{1 - \beta_n(\sigma_n - 1)(\gamma_n(z_1) - \xi_n)}} H_n(z_1) E_n$$

while sales to income group z_0 are:

$$x_n(z_0, a) = (a P_n(z_0) / \mu_n)^{\sigma_n - 1} H_n(z_0) E_n$$

We can see for instance that sales to groups z_1 (consumers who care about quality) are more elastic to productivity than for group z_0 :

$$\frac{d \log x_n(a, z_1)}{d \log a} = \frac{\sigma_n - 1}{1 - \beta_n(\sigma_n - 1)(\gamma_n(z_1) - \xi_n)} > \sigma_n - 1 = \frac{d \log x_n(a, z_0)}{d \log a}$$

hence larger firms (more productive firms) tend to sell relatively more to group z_1 : $\frac{d \log(x_n(a, z_1) / x_n(a, z_0))}{d \log a} > 0$, though more productive firms sell more to each group z than other firms.

For group z_1 , the price index is then:

$$\begin{aligned} P_n(z_1)^{1 - \sigma_n} &= N_n \mu_n^{1 - \sigma_n} \int_a a^{\sigma_n - 1} \phi_{n1}(a)^{(\sigma_n - 1)(\gamma_n(z_1) - \xi_n)} dG_n(a) \\ &= N_n \mu_n^{1 - \sigma_n} \int_a a^{\frac{\sigma_n - 1}{1 - \beta_n(\sigma_n - 1)(\gamma_n(z_1) - \xi_n)}} dG_n(a) \left(\frac{\gamma_n(z_1) - \xi_n}{b_n \mu_n} H_n(z_1) E_n P_n(z_1)^{\sigma_n - 1} \right)^{\frac{\beta_n(\sigma_n - 1)(\gamma_n(z_1) - \xi_n)}{1 - \beta_n(\sigma_n - 1)(\gamma_n(z_1) - \xi_n)}} \\ &= \left(N_n \mu_n^{1 - \sigma_n} \int_a a^{\frac{\sigma_n - 1}{1 - \beta_n(\sigma_n - 1)(\gamma_n(z_1) - \xi_n)}} dG_n(a) \right)^{1 - \beta_n(\sigma_n - 1)(\gamma_n(z_1) - \xi_n)} \left(\frac{\gamma_n(z_1) - \xi_n}{b_n \mu_n} H_n(z_1) E_n \right)^{\beta_n(\sigma_n - 1)(\gamma_n(z_1) - \xi_n)} \end{aligned}$$

This expression highlights a key feedback effect that is quantified in our counterfactuals: a higher

share $H_n(z_1)$ of expenditures by richer households z_1 (with higher taste for quality) leads to a lower price index for this income group. For the other income group z_0 with $\gamma_n(z_0) = \xi_n$, there is no feedback effect: $P_n(z_0)^{1-\sigma_n} = N_n \mu_n^{1-\sigma_n} \int_a a^{\sigma_n-1} dG_n(a)$.

2.D) Other Expressions for Sales and Profits in the General Case

Profits (Equation 14):

As shown above:

$$\phi_n(a) = \left(\frac{\tilde{\gamma}_n(a) - \xi_n}{b_n \mu_n(a)} \cdot X_n(a) \right)^{\beta_n}$$

where $\tilde{\gamma}_n(a)$ is a weighted average quality valuation $\gamma_n(z)$ for firm with productivity a

$$\tilde{\gamma}_n(a) = \frac{\int_z \gamma_n(z) (\sigma_n(z) - 1)(1 - s_n(z, a)) x_n(z, a) dH(z)}{\int_z (\sigma_n(z) - 1)(1 - s_n(z, a)) x_n(z, a) dH(z)}$$

and where X_n denotes total sales

We obtain that fixed costs spent on quality upgrading equal:

$$f_n(\phi_n(a)) = \beta_n b_n \phi_n(a)^{\frac{1}{\beta_n}} = \beta_n (\tilde{\gamma}_n(a) - \xi_n) X_n(a) / \mu_n(a)$$

Given that variable costs correspond to a share $1 - 1/\mu_n(a)$ of total sales, we obtain that profits equal:

$$\pi_n(a) = \left[1 - 1/\mu_n(a) - \beta_n (\tilde{\gamma}_n(a) - \xi_n) / \mu_n(a) \right] X_n(a) - f_{0n}$$

where f_{0n} corresponds to fixed costs are independent of quality. Equivalently, using the expressions for $\mu_n(a)$ and $\tilde{\gamma}_n(a)$, we can obtain profits more directly as a function of consumer taste for quality $\gamma_n(z)$:

$$\pi_n(a) = \left(1 - \frac{1}{\mu_n(a)} \right) \left[\int_z (1 - \beta_n (\gamma_n(z) - \xi_n) (\sigma_n(z) - 1) (1 - s_n(a, z))) x_n(a, z) dH(z) \right] - f_{0n}$$

Decomposition of Average Firm Size Differences Across Baskets (Equation 15): The weighted average of firm size for each income group z is defined as:

$$\log \widetilde{X}_n(z) = \frac{\int_a x_n(z, a) \log X_n(a) dG_n(a)}{\int_a x_n(z, a) dG_n(a)}$$

Hence the slope in Figure 1 corresponds to:

$$\frac{\partial \log \widetilde{X}_n(z)}{\partial z} = \frac{\int_a x_n(z, a) (\log X_n(a)) \frac{\partial \log x_n}{\partial z} dG_n(a)}{\int_a x_n(z, a) dG_n(a)} - \left(\frac{\int_a x_n(z, a) \log X_n(a) dG_n(a)}{\int_a x_n(z, a) dG_n(a)} \right) \left(\frac{\int_a x_n(z, a) \frac{\partial \log x_n}{\partial z} dG_n(a)}{\int_a x_n(z, a) dG_n(a)} \right)$$

In turn, the derivatives of sales to each income group w.r.t z equal:

$$\frac{\partial \log x_n(z, a)}{\partial z} = \frac{\partial \gamma_n(z)}{\partial z} (\sigma_n(z) - 1) \log \phi_n(a) - \frac{\partial \sigma_n(z)}{\partial z} \log \left(\frac{p_n(a)}{\phi_n(a)^{\gamma_n(z)}} \right) + cst(z)$$

where $cst(z)$ denotes a term that is common across all firms (only depends on price elasticities and price indices) and cancels out in the next expression.

If we plug this into the expression above for $\frac{\partial \log \widetilde{X}_n}{\partial z}$, we obtain:

$$\frac{\partial \log \widetilde{X}_n(z)}{\partial z} = \frac{\partial \gamma_n}{\partial z} (\sigma_n(z) - 1) \left[\frac{\int_a x_n(z, a) (\log X_n(a)) (\log \phi_n(a)) dG_n(a)}{\int_a x_n(z, a) dG_n(a)} \right]$$

$$\begin{aligned}
& - \left(\frac{\int_a x_n(z, a) \log X_n(a) dG_n(a)}{\int_a x_n(z, a) dG_n(a)} \right) \left(\frac{\int_a x_n(z, a) \log \phi_n(a) dG_n(a)}{\int_a x_n(z, a) dG_n(a)} \right) \Big] \\
& - \frac{\partial \sigma_z}{\partial z} \cdot \left[\frac{\int_a x_n(z, a) (\log X_n(a)) (\log(p_n(a)/\phi_n(a)^{\gamma_n(z)}))}{\int_a x_n(z, a) dG_n(a)} \right. \\
& \quad \left. - \left(\frac{\int_a x_n(z, a) \log X_n(a) dG_n(a)}{\int_a x_n(z, a) dG_n(a)} \right) \left(\frac{\int_a x_n(z, a) \log(p_n(a)/\phi_n(a)^{\gamma_n(z)}) dG_n(a)}{\int_a x_n(z, a) dG_n(a)} \right) \right]
\end{aligned}$$

which can be rewritten as two covariance terms as described in the main text.

Estimation Equation for β_n and ξ_n (Equation 20): Here we work with discrete consumer groups, indexing z as a subscript, and we index firms (brands or holding companies) by subscript i .

Starting from the following equality that we use to estimate φ_{bz} :

$$\log X_{niz} = (1 - \sigma_{nz}) \log p_{ni} + (\sigma_{nz} - 1) \log \varphi_{niz}$$

and using the definition of democratic quality $\log \phi_{ni} = \frac{1}{5} \sum_z \log \varphi_{niz}$ (again, by construction), we get:

$$\log p_{ni} = -\frac{1}{\bar{\sigma}_n - 1} \log X_{ni} + \log \phi_{ni} - \frac{1}{5} \sum_z \frac{1}{\sigma_{nz} - 1} \log \left(\frac{X_{niz}}{X_{ni}} \right)$$

where we define $\frac{1}{\bar{\sigma}_n - 1}$ as an arithmetic average:

$$\frac{1}{\bar{\sigma}_n - 1} = \frac{1}{5} \sum_z \frac{1}{\sigma_{nz} - 1}$$

Next, we can use our expression for optimal quality which gives, up to some error ε_{ni} :

$$\log \phi_{ni} = \beta_n \log X_{ni} + \beta_n \log ((\tilde{\gamma}_{ni} - \xi_n) / \mu_{ni}) - \beta_n \log b_n + \varepsilon_{ni}$$

which can be incorporated into the above expression in order to obtain the estimation equation:

$$\log p_{ni} = \left(\beta_n - \frac{1}{\bar{\sigma}_n - 1} \right) \log X_{ni} + \beta_n \log ((\tilde{\gamma}_{ni} - \xi_n) / \mu_{ni}) - \frac{1}{5} \sum_z \frac{1}{\sigma_{nz} - 1} \log \left(\frac{X_{niz}}{X_{ni}} \right) + \eta_n + \varepsilon_{ni}$$

where ε_{ni} is the error in predicting quality and η_n is an industry constant.

Appendix 3: On Consumer Preferences and Outside Consumption (z)

3.A) Compensating Variations with Unspecified Upper-Tier Utility for Outside Consumption

Lemma *For an individual of initial outside consumption z , the compensating variation (as a first-order approximation) is given by:*

$$\frac{CV}{E_G} = - \left. \frac{\Delta_p P_G(z, p)}{P_G(z, p)} \right|_z$$

where $dP_G(z, p)$ reflects the change in prices index P_G , holding z constant.

Proof: This can be seen as a consequence of Shephard's Lemma. Consider the expenditure function: $e(P_G, U)$. The compensating variation (CV) is defined implicitly such that:

$$e(p', \varphi', U) - CV = w = e(p, \varphi, U)$$

Using Shephard's Lemma, we obtain a first-order approximation:

$$d \log e = \sum_i \frac{p_{ni} q_{ni}}{w} d \log p_{ni} - \sum_i \frac{p_{ni} q_{ni}}{w} d \log \varphi_{niz}$$

since quality φ_{niz} for each brand i (valued by income group z) is defined as a price-equivalent demand shifter. Hence, with $d \log e = CV/w$, we have:

$$\frac{CV}{w} = \sum_i \frac{p_{ni} q_{ni}}{w} (d \log p_{ni} - d \log \varphi_{niz})$$

One can also verify that, holding z constant, the change in the retail price corresponds to:

$$d \log P_G(z) = \sum_i \frac{p_{ni} q_{ni}}{E_G} (d \log p_{ni} - d \log \varphi_{niz})$$

where E_G denotes expenditures in retail shopping. We obtain the result in the lemma above by combining this equality with the first-order approximation of CV.

3.B) Changes in Outside Consumption with Multiplicative Upper-Tier Utility

In this appendix section, we examine how the change in price indices across households in our counterfactuals may have induced a meaningful change in the consumption of the outside good z , holding nominal income constant. Such a change could give rise to second-order effects on household price indices in the counterfactuals—since demand parameters α , γ and σ are allowed to vary with consumer types z . In our main counterfactuals, retail prices tend to increase faster for rich relative to poor households, which may affect the consumption of outside goods z differentially. Here we quantify the magnitude of such potential endogenous changes in z , and find that they are negligible relative to the counterfactual changes holding z constant.

In this exercise, we assume that upper-tier utility in (1) takes the form:

$$U = A(z) U_G(z)^{\alpha_G} \tag{25}$$

for some constant α_G and a function $A(z)$ of the outside good z , and with $U_G(z)$ defined as in the main text. This includes the special case where $A(z) = z^{1-\alpha_G}$. This utility function is flexible enough to yield various patterns of income elasticities for grocery and outside good consumption, as we show below.

With such utility, consumers choose z to maximize:

$$\log U = \max_z [\alpha_G \log(w - z) - \alpha_G \log P_G(z) + \log A(z)]$$

given their income w and price indices $P_G(z)$. With this specification, one can see that the overall level of prices $P_G(z)$ does not affect the share of income spent on z : multiplying all prices by the same constant does not affect z . In fact, any preferences that have this property can be written as in equation 25. Yet, a change in the patterns of prices across z 's still influences the consumption of the outside good z . Moreover, equation 25 does not impose any constraint on the income elasticity of the outside good z .

In what follows, we first examine the income elasticity of z and then express the changes in z induced by the change in prices as a function of our counterfactual results, the share of expenditures in retail shopping and the income elasticity of the outside good.

Income Elasticity of the Outside Good The first-order condition in z can be written:

$$\varepsilon_A(z) = \alpha_G \varepsilon_P(z) + \frac{z \alpha_G}{w - z} \quad (26)$$

where we define $\varepsilon_A(z) = \frac{zA'(z)}{A(z)}$ and $\varepsilon_P(z) = \frac{zP'_G(z)}{P_G(z)}$. Moreover, the second-order condition in z imposes:

$$\alpha_G \frac{d\varepsilon_P}{d \log z} + \frac{zw\alpha_G}{(w - z)^2} - \frac{d\varepsilon_A}{d \log z} > 0$$

We will assume that the second-order condition is satisfied, which holds as long as the retail price index either increases or does not decrease too fast with z . Note that $\frac{d\varepsilon_A}{d \log z} = 0$ in the case where $A(z) = z^{1-\alpha_G}$.

Suppose that \log income ($\log w$) increases by small amount $d \log w$. Differentiating the condition 26, we obtain:

$$\frac{d\varepsilon_A}{d \log z} d \log z = \alpha_G \frac{d\varepsilon_P}{d \log z} d \log z + \frac{zw\alpha_G}{(w - z)^2} d \log z - \frac{zw\alpha_G}{(w - z)^2} d \log w$$

Hence the income elasticity of z is:

$$\eta_w(z) \equiv \frac{d \log z}{d \log w} = \frac{\frac{zw\alpha_G}{(w - z)^2}}{\frac{zw\alpha_G}{(w - z)^2} + \alpha_G \frac{d\varepsilon_P}{d \log z} - \frac{d\varepsilon_A}{d \log z}}$$

As described in the main text, we assume throughout the paper that the income elasticity of z is positive. Under the preferences above, this condition is satisfied as long as the second-order condition above is satisfied.

Effect of Counterfactual Change in Prices on z While an uniform change $d \log P_G$ does not affect the consumption of the outside good z , a differential change in price indices $d\varepsilon_P(z) = d\left(\frac{zP'_G(z)}{P_G(z)}\right)$, like the one resulting from our counterfactual exercise, will affect z .

Suppose that $\varepsilon_P(z) = \frac{zP'_G(z)}{P_G(z)}$ increases by $d\varepsilon_P^{CF}$ uniformly. As a first-order approximation, taking the derivative of 26, we obtain:

$$\frac{d\varepsilon_A}{d \log z} d \log z = \alpha_G d\varepsilon_P^{CF} + \alpha_G \frac{d\varepsilon_P}{d \log z} d \log z + \frac{zw\alpha_G}{(w - z)^2} d \log z$$

Hence, solving for the change in z , we obtain:

$$\frac{d \log z}{d \varepsilon_P^{CF}} = - \frac{\alpha_G}{\frac{zw\alpha_G}{(w-z)^2} + \alpha_G \frac{d \varepsilon_P}{d \log z} - \frac{d \varepsilon_A}{d \log z}}$$

Note that this term is negative, which implies that, in this framework, retail consumption and outside consumption are complements: if retail shopping becomes relatively more expensive for higher z 's (as we obtain in our main counterfactual exercises), the optimal z is lower. We can further re-write this term using the income elasticity of z to obtain a much more simple expression:

$$\begin{aligned} \frac{d \log z}{d \varepsilon_P} &= - \frac{\left(\frac{w-z}{w}\right)^2}{\frac{z}{w}} \frac{\frac{zw\alpha_G}{(w-z)^2}}{\frac{zw\alpha_G}{(w-z)^2} + \alpha_G \frac{d \varepsilon_P}{d \log z} - \frac{d \varepsilon_A}{d \log z}} \\ &= - \frac{s_G^2}{1 - s_G} \eta_w \end{aligned}$$

where $s_G = \frac{w-z}{w}$ denotes the share of retail in consumer expenditures. As both s_G and η_w can be readily observed or estimated, we can use this expression to quantify how a change in the slope of the retail price index with respect to z will induce a change in z .

Quantifying the Change in Price Schedule In our counterfactual, we estimate a double-difference:

$$\Delta \Delta \log P = [\log P_{t1}(z_1) - \log P_{t1}(z_0)] - [\log P_{t0}(z_1) - \log P_{t0}(z_0)]$$

comparing different quintiles of consumers, z_1 and z_0 , across different sets of prices, $P_{t1}(\cdot)$ and $P_{t0}(\cdot)$ (new and initial prices). For the purpose of estimation, we have not estimated the price changes over a continuum of z , but our counterfactual results are not too far from a log-linear relationship between price changes and income across quintiles. Hence, a good approximation of the change in price schedule is to compute the straight slope of the change in log prices between the two quintiles (slopes in Figures 6a, 7a and 8a):

$$\varepsilon_P(z_1) - \varepsilon_P(z_0) = \frac{z P'_{t1}(z)}{P_{t1}(z)} - \frac{z P'_{t0}(z)}{P_{t0}(z)} \approx \frac{\Delta_t \Delta_z \log P}{\log z_0 - \log z_1}$$

for any two quintiles of consumers, z_1 and z_0 .

Plugging it into the induced change in z , we obtain how a change in the price schedule affects outside good consumption z :

$$d \log z \approx - \frac{\eta_w s_G^2}{1 - s_G} \frac{\Delta_t \Delta_z \log P}{\log z_0 - \log z_1}$$

Numerical Application: Using the following estimates:

- $\eta_w \approx 2$ (upper bound to be conservative)
- $\Delta_t \Delta_z \log P \approx -0.03$ (larger differential effect of two counterfactuals)
- $\log z_0 - \log z_1 \approx \log 75,000 - \log 15,000 \approx 1.6$ (small to be conservative)
- Share of retail shopping $s_G \approx 0.25$ (large to be conservative)

we obtain the following change in z that would be caused by the changes in prices in the counterfactuals:

$$d \log z \approx 2 \times 0.08 \times 0.03 / 1.6 \approx 0.003$$

Hence, quantitatively, the changes in z induced by the price changes in our counterfactuals are very small and negligible relative to existing differences in outside good consumption between rich and poor income groups. Such changes in z would also have negligible GE impacts on sales, quality and prices on the supply side relative to the changes obtained in our counterfactual shocks.

3.C) Non-Homotheticity without Outside Consumption

Here we lay out a very similar way to flexibly model non-homotheticity in consumer preferences for retail shopping, but without relying on an outside good z , drawing on recent work by Fally (2018). Suppose that utility is defined implicitly by:

$$\sum_i \left(\frac{q_i}{f(U)\phi_i^{\gamma(U)}} \right)^{\frac{\sigma(U)-1}{\sigma(U)}} = 1 \quad (27)$$

With such utility, the share of good i in expenditures is given by:

$$\frac{p_i q_i}{\sum_j p_j q_j} = \frac{(p_i/\phi_i^{\gamma(U)})^{1-\sigma(U)}}{\sum_j (p_j/\phi_j^{\gamma(U)})^{1-\sigma(U)}}$$

which is identical to the expression for market share in the baseline equilibrium with price elasticity $\sigma(U)$ and quality valuation elasticity $\gamma(U)$ varying with utility U instead of outside consumption z . Suppose that $\sigma(U)$ decreases with utility (empirically-relevant case) and that quality taste parameter $\gamma(U)$ increases with utility. As shown in Fally (2018), such implicitly-defined utility is well-defined if:

$$\frac{f'(U)}{f(U)} > \log \left(\sum_i \phi_i^{\gamma'(U)/\eta(U)} \right)^{\eta(U)}$$

with $\eta(U) \equiv \frac{\partial}{\partial U} \left(\frac{1}{\sigma(U)-1} \right) = -\frac{\sigma'(U)}{(\sigma(U)-1)^2} > 0$ (see Proposition 4 and condition 13 in Fally, 2018). Intuitively, this condition is satisfied if there is not too much heterogeneity in quality valuations and price elasticities, i.e. if $\gamma'(u)$ and $\eta(U)$ are not too large relative to $f'(U)$. Conditional on ϕ , $\gamma(U)$ and $\sigma(U)$, we can always choose $f(U)$ to satisfy this condition for all levels of utility. Note that a price index is given by:

$$P(p, U) = \left(\sum_j (p_j/\phi_j^{\gamma(U)})^{1-\sigma(U)} \right)^{\frac{1}{1-\sigma(U)}}$$

This can also be extended to multiple product groups n , defining the expenditure function as:

$$P(p, U) = \prod_n \left(\sum_{j \in G_n} (p_j/\phi_j^{\gamma_n(U)})^{1-\sigma_n(U)} \right)^{\frac{\alpha_n(U)}{1-\sigma_n(U)}}$$

which is identical to the expression in the text with U instead of outside consumption z . The change in the price index as a function of utility U would then be observationally identical to the change in the price index in the baseline model expressed as a function of z . As a first-order approximation, the welfare implications of changes in prices and quality are also identical to those in the baseline model, conditional on simulated changes in quality and price:

$$\frac{CV}{E} \approx -\Delta \log P(U, p) .$$

3.D) Equivalent Discrete-Choice Model for Grocery Shopping

In this appendix section, we describe a discrete choice model as in Anderson et al. (1987) to describe how aggregation of heterogeneous consumers buying only one good by product module can be equivalent to utility in Equation 1 in the main text:

$$U_G(z) = \max_q \prod_n \left[\sum_{i \in G_n} (q_{ni} \varphi_{ni}(z))^{\frac{\sigma_n(z)-1}{\sigma_n(z)}} \right]^{\alpha_n(z) \cdot \frac{\sigma_n(z)}{\sigma_n(z)-1}} \quad (28)$$

The proof follows Handbury (2019), except that the Cobb-Douglas upper tier simplifies the equivalence result.

So, Instead of preferences 28, suppose that individual j from income group z has utility:

$$U_j(z) = \sum_n \alpha_n(z) \max_{i \in G_n, q_{jni}} [\log q_{jni} + \log \varphi_{ni}(z) + \mu_n(z) \epsilon_{jni}] \quad (29)$$

maximizing over the vector $\{y_{jn}\}$ of income allocated to each module n and goods i in module n , the chosen good i and its quantity q_{jni} for each product module n , under the budget constraints:

$$\begin{aligned} \sum_n y_{jn} &\leq E(z) \\ \sum_{i \in G_n} q_{jni} p_{ni} &\leq y_{jzn} \end{aligned}$$

where $E(z)$ refers to total income allocated to grocery shopping for consumers of income group z . In expression 29 above, $\log \varphi_{ni}(z)$ is a quality shifter associated with product z in module n that is specific to income group z . In turn, the last term $\mu_n(z) \epsilon_{jni}$ is a specific taste shock for each individual j and good i .

With these preferences, each consumer j consumes a unique good i^* in product module n . Given the vector $\{y_{jn}\}_n$ of expenditures in each module n , the good i^* being chosen maximizes:

$$i^* = \operatorname{argmax}_{i \in G_n} [\log y_{jn} - \log p_{ni} + \log \varphi_{ni}(z) + \mu_n(z) \epsilon_{jni}]$$

Hence we can see that the choice of the good i by consumer j in income group z does not depend on income y_{jn} that is allocated to a specific product module n . The good that is consumed simply maximizes:

$$i^* = \operatorname{argmax}_{i \in G_n} [-\log p_{ni} + \log \varphi_{ni}(z) + \mu_n(z) \epsilon_{jni}] \quad (30)$$

If, within income group z , the choice of good i^* does not depend on the allocation of income y_{jn} , a key implication is that the allocation of income across product modules n does not depend on the specific draws ϵ_{jni} :

$$\begin{aligned} U_j(z) &= \max_{\{y_{jn}\}} \left\{ \sum_n \alpha_n(z) \max_{i \in G_n} [\log y_{jn} - \log p_{ni} + \log \varphi_{ni}(z) + \mu_n(z) \epsilon_{jni}] \right\} \\ &= \max_{\{y_{jn}\}} \left\{ \sum_n \alpha_n(z) \log y_{jn} \right\} + \sum_n \alpha_n(z) \max_{i \in G_n} [-\log p_{ni} + \log \varphi_{ni}(z) + \mu_n(z) \epsilon_{jni}] \end{aligned}$$

which leads to y_{jn} being equal to a fraction $\alpha_n(z)$ of income $E(z)$ spent on grocery shopping (for consumers j belonging to income group z):

$$y_{jn} = \alpha_n(z) E(z)$$

Using this property and additional assumptions on the distribution of shocks ϵ_{jni} , we can now examine aggregate consumption patterns, aggregating across individuals j within each income group z . Suppose that we have a large number of consumers and that ϵ_{jni} is i.i.d. and drawn from a Gumbel distribution (type-II extreme value distribution) as in Anderson et al. (1987). Equation 30 implies that a share:

$$s_{ni}(z) = \frac{\left(\frac{\varphi_{ni}(z)}{p_{ni}}\right)^{\frac{1}{\mu_n(z)}}}{\sum_{i' \in G_n} \left(\frac{\varphi_{ni'}(z)}{p_{ni'}}\right)^{\frac{1}{\mu_n(z)}}}$$

of consumers will choose good i among all goods in G_n . Given that all consumers within income group z spend an amount $y_{jn} = \alpha_n(z) E(z)$ on module n , we obtain the following expenditures for income group z on good i :

$$x_{zni} = \frac{\left(\frac{\varphi_{ni}(z)}{p_{ni}}\right)^{\sigma_n(z)-1}}{\sum_{i' \in G_n} \left(\frac{\varphi_{ni'}(z)}{p_{ni'}}\right)^{\sigma_n(z)-1}} \alpha_n(z) E(z)$$

where $\sigma_n(z) = 1 + \frac{1}{\mu_n(z)}$ denotes the elasticity of substitution between goods i on aggregate for consumers of income group z . This shows that utility described in equation 29 is exactly equivalent to the consumption patterns obtained with the preferences described in equation 28 above and equation 2 in the main text.

3.E) Normalization of $\int_{\Omega_z} \gamma_n(z) dz = 1$

Equation (4) in the main text specifies:

$$\log \varphi_{ni}(z) = \gamma_n(z) \log \phi_{ni}$$

We adopt the normalization $\int_{\Omega_z} \gamma_n(z) dz = 1$ where Ω_z is a given set of z 's. In our empirical application, we work with 5 household groups indexed by z , which we define as quintiles in total retail per capita expenditure. In our current specification, we define these quintiles in each half-year period, because this brings greatest clarity when talking about quintiles of US households in any given period. In the following, we describe how to interpret our parameter values, equilibrium conditions and estimation equations with an alternative normalization (other than unity). Suppose instead that $\int_{\Omega_z} \gamma_n(z) dz$ is normalized to λ . With this normalization, ϕ_{ni} satisfies:

$$\log \phi_{ni} = \frac{1}{\lambda} \int_{\Omega_z} \log \varphi_{ni}(z) dz \quad (31)$$

Other equations characterizing sales, price indices and optimal quality remain the same:

$$\frac{x_{ni}(z)}{x_{nj}(z)} = \left(\frac{\phi_{ni}}{\phi_{nj}}\right)^{\gamma_n(z)(\sigma_n(z)-1)} \left(\frac{p_{ni}}{p_{nj}}\right)^{1-\sigma_n(z)} \quad (32)$$

$$P_n(z) = \left[\sum_{i \in G_n} p_{ni}^{1-\sigma_n(z)} \phi_{ni}^{\gamma_n(z)(\sigma_n(z)-1)} \right]^{\frac{1}{1-\sigma_n(z)}} \quad (33)$$

$$\phi_n(a) = \left(\frac{\tilde{\gamma}_n(a) - \xi_n}{b_n \mu_n(a)} \cdot X_n(a) \right)^{\beta_n} \quad (34)$$

where $\tilde{\gamma}_n(a)$ equals the sales-weighted average of $\gamma_n(z)$ across consumers (see text).

Does that model lead to different estimation equations and counterfactual simulations? The answer is no: while the interpretation of some parameters depends on the normalization, the quantitative implications for sales, price indices and welfare do not depend on this normalization. To be more precise, the model is isomorphic to an alternative normalization where the new model parameters are:

$$\gamma'_n(z) = \gamma_n(z)/\lambda \quad (35)$$

$$\beta'_n = \beta_n\lambda \quad (36)$$

$$b'_n = b_n/\lambda \quad (37)$$

$$\xi'_n = \xi_n/\lambda \quad (38)$$

$$\phi'_{ni} = (\phi_{ni})^\lambda \quad (39)$$

while other parameters and variables remain identical. With these, we can check that:

- Average firm quality can now be defined as a simple average: $\log \phi'_{ni} = \int_{\Omega_z} \log \varphi_{ni}(z) dz$
- The above equilibrium equations 32-34 hold. For instance, quality satisfies:

$$\phi'_n(a) = \phi_n(a)^\lambda = \left(\frac{\tilde{\gamma}_n(a) - \xi_n}{b_n \mu_n(a)} \cdot X_n(a) \right)^{\lambda \beta_n} = \left(\frac{\tilde{\gamma}'_n(a) - \xi'_n}{b'_n \mu_n(a)} \cdot X_n(a) \right)^{\beta'_n} \quad (40)$$

- Profits do not depend on the normalization. Equation (14) for profits yields:

$$\begin{aligned} \pi_n(a) &= \left(1 - \frac{1}{\mu_n(a)} \right) \left[\int_z (1 - \beta_n(\gamma_n(z) - \xi_n)(\sigma_n(z) - 1)(1 - s_n(a, z))) x_n(a, z) dH(z) \right] - f_{0n} \\ &= \left(1 - \frac{1}{\mu_n(a)} \right) \left[\int_z (1 - \beta'_n(\gamma'_n(z) - \xi'_n)(\sigma_n(z) - 1)(1 - s_n(a, z))) x_n(a, z) dH(z) \right] - f_{0n} \end{aligned}$$

- The estimation of supply side parameters also remains the same since they are directly obtained from equation 40.
- Counterfactual equilibrium conditions (see below) also remain the same. For instance, for quality upgrading and sales, we have:

$$\begin{aligned} \frac{\phi'_{n1}(a)}{\phi'_{n0}(a)} &= \left(\frac{\phi_{n1}(a)}{\phi_{n0}(a)} \right)^\lambda = \left[\frac{\int_z (\gamma_n(z) - \xi_n)(\sigma_n(z) - 1)(1 - s_{n1}(a, z)) x_{n1}(z, a) \mu_{n0}(a) dH(z)}{\int_z (\gamma_n(z) - \xi_n)(\sigma_n(z) - 1)(1 - s_{n0}(a, z)) x_{n0}(z, a) \mu_{n1}(a) dH(z)} \right]^{\lambda \beta_n} \\ &= \left[\frac{\int_z (\gamma'_n(z) - \xi'_n)(\sigma_n(z) - 1)(1 - s_{n1}(a, z)) x_{n1}(z, a) \mu_{n0}(a) dH(z)}{\int_z (\gamma'_n(z) - \xi'_n)(\sigma_n(z) - 1)(1 - s_{n0}(a, z)) x_{n0}(z, a) \mu_{n1}(a) dH(z)} \right]^{\beta'_n} \end{aligned}$$

Finally, we could have also defined the 5 quintile groups across all years, instead of defining them in terms of quintiles within each period (after converting nominal expenditures across semesters into real expenditure). The same normalization to unity across these five groups could have been applied without loss of generality, as we show below in the remainder of this appendix. The slight difference between these two approaches (quintiles by semester vs quintiles across semesters) is that a small number of households would have been assigned to different z groups across semesters. To ensure the robustness of our preferred approach to defining z groups globally instead (across the full distribution of observed per-semester expenditures), we have re-run the parameter estimation and all four counterfactuals under this alternative classification. Reassuringly, the point estimates are virtually identical in all cases. This is due to the fact that i) the US did not experience strong real income growth across the semesters over this period, and ii) relative household expenditure positions are relatively stable across semesters.

Appendix 4: Supply-Side Model Extensions

4.A) Model Extension with Multi-Product Firms

Sales Across Brands and Products

Let us index each product by subscript i and each brand by subscript b . We denote by $\varphi_{nb}^{Tot}(z)$ the average quality of a brand, while we denote by $\varphi_{nbi}^{MP}(z)$ additional idiosyncratic quality shocks at the product level, so that product quality of each product i of brand b corresponds to the product $\varphi_{nbi}^{MP}(z)\varphi_{nb}^{Tot}(z)$. As in Hottman et al. (2016), we normalize the average idiosyncratic quality shock to zero: $\sum_i \log \varphi_{nbi}^{MP}(z) = 0$. Here we assume monopolistic competition.

Using this definition, total sales by brand b can be expressed as:

$$x_{nb}^{Tot}(z) = \left(\frac{\varphi_{nb}^{Tot}(z)}{P_{nb}^{brand}(z)} \right)^{\sigma_n(z)-1} \alpha_n(z) E(z) P_n(z)^{\sigma_n(z)-1} \quad (41)$$

while sales by product can be written as:

$$x_{nbi}^{MP}(z) = \left(\frac{\varphi_{nbi}^{MP}(z)}{p_{ni}} \right)^{\eta_n(z)-1} x_{nb}^{Tot}(z) P_{nb}^{brand}(z)^{\eta_n(z)-1} \quad (42)$$

In these equations, the price index by product group is defined as:

$$P_n(z) = \left[\sum_{i \in G_n} P_{nb}^{brand}(z)^{1-\sigma_n(z)} \varphi_{ni}^{Tot}(z)^{\sigma_n(z)-1} \right]^{\frac{1}{1-\sigma_n(z)}} \quad (43)$$

while the price index by brands (across products belonging to the brand) is defined as:

$$P_{nb}^{brand}(z) = \left[\sum_{i \in G_n} p_{ni}^{1-\eta_n(z)} \varphi_{nbi}^{MP}(z)^{\eta_n(z)-1} \right]^{\frac{1}{1-\eta_n(z)}} \quad (44)$$

When price elasticities $\eta_n(z)$ and $\sigma_n(z)$ (within and across brands) differ, this new definition of a brand's price index differs from traditional sales weighted price indices (e.g. Tornqvist) as they also directly depend on the number of product varieties. Let us define a price index $\bar{P}_{nb}(z)$ as a weighted average:

$$\bar{P}_{nb}(z) = \left[\frac{1}{N_{nb}} \sum_{i \in G_n} p_{ni}^{1-\eta_n(z)} \varphi_{nbi}^{MP}(z)^{\eta_n(z)-1} \right]^{\frac{1}{1-\eta_n(z)}}$$

where N_{nb} corresponds to the number of product varieties. This index only depends on a average of prices and quality, and does not directly depend on the number of product varieties, conditional on average prices and quality. On the contrary, price index $P_{nb}^{brand}(z)$ depends on N_{nb} even if prices and quality are identical across all products. Conditional on average quality and prices $\bar{P}_{nb}(z)$, total sales by brand can be written:

$$x_{nb}^{Tot}(z) = N_{nb}^{\frac{\sigma_n(z)-1}{\eta_n(z)-1}} \left(\frac{\varphi_{nb}^{Tot}(z)}{\bar{P}_{nb}(z)} \right)^{\sigma_n(z)-1} \alpha_n(z) E(z) P_n(z)^{\sigma_n(z)-1} \quad (45)$$

As shown in this equation, the number of product varieties affects whether firms sell relatively more to richer households only when $\frac{\sigma_n(z)-1}{\eta_n(z)-1}$ varies with income z . If $\frac{\sigma_n(z)-1}{\eta_n(z)-1}$ increases with income z , richer consumers tend to consume relatively more from brands with a larger number of products.

Markups and Prices for Multi-Product Firms

Markups are now determined cannibalization effects and interaction between products within the brand.

After noticing that the elasticity of the brand-level price w.r.t. product-level prices equals its market share among consumers of income z :

$$\frac{\log P_{nb}(z)}{\log p_{nbi}} = \frac{x_{nbi}(z)}{\sum_j x_{nbj}(z)}$$

and that the elasticity of the product-level sales w.r.t. brand level price index equals $\eta_n(z) - \sigma_n(z)$, we obtain that profit maximization leads to the following first-order condition associated with markups for each product i :

$$\sum_z x_{nbi}(z) - \mu_{nbi} \sum_z \eta_n(z) x_{nbi}(z) + \sum_{j,z} \left[(\eta_n(z) - \sigma_n(z)) \mu_{nbj} x_{nbj}(z) \frac{x_{nbi}(z)}{\sum_{j'} x_{nbj'}(z)} \right] = 0$$

where $\mu_{nbi} \equiv \frac{p_{nbi} - c_{nbi}}{p_{nbi}}$ denotes markup for product i and c_{nbi} refers to the marginal cost of producing good i . Let us also define $\bar{\mu}_{nb}(z) = \frac{\sum_j \mu_{nbj} x_{nbj}(z)}{\sum_j x_{nbj}(z)}$ the average markup charged by brand b on consumers of income z . Rearranging the above expression, we obtain:

$$\mu_{nbi} = \frac{\sum_z x_{nbi}(z)}{\sum_z \eta_n(z) x_{nbi}(z)} \left[1 + \frac{\sum_z (\eta_n(z) - \sigma_n(z)) \bar{\mu}_{nb}(z) x_{nbi}(z)}{\sum_z x_{nbi}(z)} \right] \quad (46)$$

or equivalently:

$$\mu_{nbi} = \frac{\sum_z x_{nbi}(z)}{\sum_z \sigma_n(z) x_{nbi}(z)} \left[1 + \frac{\sum_z (\eta_n(z) - \sigma_n(z)) (\bar{\mu}_{nb}(z) - \mu_{nbi}) x_{nbi}(z)}{\sum_z x_{nbi}(z)} \right] \quad (47)$$

In equation 46, the term $\frac{\sum_z x_{nbi}(z)}{\sum_z \eta_n(z) x_{nbi}(z)}$ reflects the markup that would be charged if each product was competing on its own, i.e. without internalizing the effect of its price on the other prices of the products of the same brand. In equation 47, the term $\frac{\sum_z x_{nbi}(z)}{\sum_z \sigma_n(z) x_{nbi}(z)}$ reflects the markup that the brand would be charging if it had only one product variety.

Two special cases are worth mentioning. First, if all products have the same share of consumers in each income group, markups would be the same as in the single-product case, i.e. $\mu_{nbi} = \frac{\sum_z x_{nbi}(z)}{\sum_z \sigma_n(z) x_{nbi}(z)}$. Second, if the difference $\eta_n(z) - \sigma_n(z)$ does not depend on income z , markups are again the same as in the single-product case. Hence, in this model, cannibalization effects arise only when the consumer base varies among products of the same brand and when the difference between the two elasticities (within and across brands) varies across consumers.

On a side note, notice that in all cases we obtain:

$$\frac{\sum_{z,i} \mu_{nbi} \sigma_n(z) x_{nbi}(z)}{\sum_{z,i} x_{nbi}(z)} = 1$$

once we take a weighted average across products. This shows that average markups are governed by the elasticity of substitution across brands rather than within brands (since brands internalize the price of each product on other products of the brand). Moreover, if $\sigma_n(z) = \sigma_n$ is homogeneous across consumers, then markups μ_{nbi} are homogeneous and equal $\frac{1}{\sigma_n}$ across all products.

Optimal Quality for Multi-Product Firms

Suppose, as in the main text, that quality $\varphi_{nb}^{Tot}(z)$ is a function of a fundamental product quality ϕ_{nb} and income-group taste for quality $\gamma_n(z)$ such that:

$$\log \varphi_{nb}^{Tot}(z) = \gamma_n(z) \log \phi_{nb}$$

Assuming that multi-product firms choose ϕ_{nb} to maximize aggregate profits:

$$\Pi = \sum_i \left[\left(1 - \frac{c_{nbi}(\phi_{nb})}{p_{nbi}} \right) \sum_i x_{nbi}(z) \right] - f_n(\phi_{nb})$$

(where $f_n(\phi_{nb}) = b_n \phi_{nb}^{\frac{1}{\beta_n}}$ are the fixed costs of quality upgrading) we obtain the following first-order condition in brand-level quality ϕ_{nb} :

$$b_n \phi_{nb}^{\frac{1}{\beta_n}} = \sum_{i,z} [\mu_{nbi}(\sigma_n(z) - 1) \gamma_n(z) x_{nbi}(z)] - \xi_n \sum_i (1 - \mu_{nbi}) x_{nbi}(z)$$

$(\sigma_n(z) - 1) \gamma_n(z)$ reflects the effect of quality upgrading on demand, while ξ_n is the effect on costs. Using our expression above for average markups (equation 47), we obtain the following expression for optimal quality that generalizes expression 12 for multi-product brands:

$$b_n \phi_{nb}^{\frac{1}{\beta_n}} = (\tilde{\gamma}_{nb}^{MP} - \xi_n) \sum_{i,z} (1 - \mu_{nbi}) x_{nbi}(z)$$

where $\tilde{\gamma}_{nb}$ is now defined at the brand level by:

$$\tilde{\gamma}_{nb}^{MP} = \frac{\sum_{i,z} \gamma_n(z) (\sigma_n(z) - 1) \mu_{nbi} x_{nbi}(z)}{\sum_{i,z} (\sigma_n(z) - 1) \mu_{nbi} x_{nbi}(z)}$$

Note that markups appear in this equation but, as described above, markups are no longer simply determined by an average of $\sigma_n(z)$ across households because of cannibalization effects and interaction between products within the brand.

4.B) Model Extension with Firm-Specific Quality Upgrading Costs

In this model extension, we allow firms upon entry to discover both an idiosyncratic productivity term a and a quality shifter ψ . This quality shifter ψ can be modeled as a cost shifter or a demand shifter: both approaches are isomorphic. Moreover, a and ψ could be correlated, or ψ could be viewed as a deterministic function of productivity a .

Here, let us model ψ as a cost shifter. More specifically, we assume that the fixed cost of quality upgrading is given by:

$$\psi^{-\frac{1}{\beta_n}} f_n(\phi) = b_n \beta_n \left(\frac{\phi}{\psi} \right)^{\frac{1}{\beta_n}}$$

For a given firm, one can see that all the previous expressions hold if we replace b_n by $\psi^{-\frac{1}{\beta_n}}$.

First, using expression 12 and substituting b_n by $\psi^{-\frac{1}{\beta_n}}$, optimal quality is now:

$$\phi(a, \psi) = \psi \cdot (X_n(a, \psi) (\tilde{\gamma}_n(a, \psi) - \xi_n) / \mu_n(a, \psi))^{\beta_n} \quad (48)$$

In this expression, it is clear that ψ can alternatively be interpreted as a taste shifter (i.e. assuming that perceived quality is multiplied by ψ for firms associated with a draw ψ).

Next, one can see that the first-order expression for prices remains the same, conditional on optimal quality φ and firm productivity a . Third, since each firm takes its productivity a and quality shifter ψ as given, all results on the uniqueness of prices and optimal quality hold, taking the distribution of a and ψ as given. Finally, econometric specifications and counterfactual simulations remain very similar in this case, given that ψ is held constant for each firm, as described below.

Estimation with Heterogeneous Quality Shifters As we estimate this new quality upgrading equation in the cross-section, one needs to account for a producer-specific term ψ_{ni} :

$$\log(p_{nit}) = \left(\beta_n - \frac{1}{\bar{\sigma}_n - 1} \right) \log(X_{nit}) - \frac{1}{N_z} \sum_z \frac{1}{\sigma_{nz} - 1} \log\left(\frac{X_{nzit}}{X_{nit}}\right) \\ + \beta_n \log((\tilde{\gamma}_{nit} - \xi_n) / \mu_{nit}) + \psi_{ni} + \eta_{nt} + \epsilon_{nit}$$

Such term is not directly observed and be positively correlated with size, $\log(X_{nit})$, leading to an upward bias in our estimated coefficient β_n in the cross section. However, this additional term ψ_{ni} disappears when we estimate these relationships in panel, as long as these firm-specific fixed costs shifters do not vary over time (or as long as the changes are not correlated with our demographic instruments for changes in scale). Given these concerns, we choose the panel specification as our baseline specification.

Counterfactuals in Model Extension with Heterogeneous Quality Shifters Here we briefly describe counterfactual equilibrium equations for the model extension with heterogeneous quality terms ψ (see next Appendix 5 for counterfactuals with the baseline model). In this extension, each firm faces an idiosyncratic quality shifter ψ (e.g. as a quality upgrading cost shifter) in addition to the firm-level productivity term a .

Quality: Quality is now determined by equation 48. However, the term ψ is held constant over time in all counterfactuals. Hence, taking ratios, we obtain exactly the same counterfactual equation as previously as a function of sales and demand parameters:

$$\frac{\phi_{n1}(a, \psi)}{\phi_{n0}(a, \psi)} = \left[\frac{(\tilde{\gamma}_{n1}(a, \psi) - \xi_n) \mu_{n1}(a, \psi)^{-1} X_{n1}(a, \psi)}{(\tilde{\gamma}_{n0}(a, \psi) - \xi_n) \mu_{n0}(a, \psi)^{-1} X_{n0}(a, \psi)} \right]^{\beta_n}$$

Sales and prices: Sales and prices depend on on the quality shifter ψ only through quality. Hence, conditional on quality, the previous equations hold. For instance, sales are given by:

$$x_n(a, \psi, z) = \phi_n(a, \psi)^{\gamma_n(z)(\sigma_n(z)-1)} p_n(a, \psi)^{1-\sigma_n(z)} \alpha_n(z) E(z) P_n^{\sigma_n(z)-1}$$

Hence, the change in sales of surviving firms equal:

$$\frac{x_{n1}(a, \psi, z)}{x_{n0}(a, \psi, z)} = \left(\frac{P_{n1}(z)}{P_{n0}(z)} \right)^{\sigma_n(z)-1} \left(\frac{\mu_{n1}(a, \psi)}{\mu_{n0}(a, \psi)} \right)^{1-\sigma_n(z)} \left(\frac{\phi_{n1}(a, \psi)}{\phi_{n0}(a, \psi)} \right)^{(\sigma_n(z)-1)(\gamma_n(z)-\xi_n)}$$

where $\mu_n(a, \psi)$ is the markup of a firm with productivity a and quality shifter ψ .

Entry and exit: Conditioning on sales, profits, exit and entry are given by the same expressions as in the baseline case.

Appendix 5: Counterfactuals and Decompositions

We use our theoretical framework to explore two types of counterfactuals. The first counterfactual changes the income tax schedule, thereby altering disposable expenditures across income groups z . The second counterfactual changes business taxes across firms.

These counterfactuals illustrate how these changes affect the demand and supply of product quality, and how they feed back into consumer inflation and real income inequality. The first part of this section describes the equilibrium conditions to solve for counterfactual outcomes. The remaining parts provide additional details on uniqueness and the decompositions of the price index changes.

5.A) Characterization of Counterfactual Equilibria

In both setups, we denote by $\phi_{n0}(a)$ and $\phi_{n1}(a)$ initial and counterfactual quality respectively, and by $x_{n0}(z, a)$ and $x_{n1}(z, a)$ initial and final sales for firm a and income group z . We denote by N_{n0} and N_{n1} the measure of firms in the baseline and counterfactual equilibrium, and we denote by $\delta_{nD}(a)$ a dummy equal to 1 if firm a survives in the counterfactual equilibrium. Finally, we denote by $P_{n0}(z)$ and $P_{n1}(z)$ the initial and counterfactual price index in product group n for income z .

In the first set of counterfactuals, changes in income taxation across consumers lead to changes in expenditures across income groups z , which we model as changes $h_n(z)$ in total sales to income group z . In the second set of counterfactuals, changes in business taxation are modeled as a change in a firm-specific sales tax $\tau_n(a)$ which depend on the rank of the firm in terms of total sales. Comparing the initial and counterfactual equilibria, we derive that the changes in firm sales, quality, entry, exit and price indices must satisfy the following five equilibrium conditions.

First, the evolution of firm sales for a given income group z depends on quality upgrading, markup and price index changes for each consumer group as well as the shocks in terms of group-specific expenditures $h_n(z)$ and shocks in sales taxes $\tau_n(a)$:

$$\frac{x_{n1}(z, a)}{x_{n0}(z, a)} = \delta_{nD}(a) h_n(z) \tau_n(a)^{1-\sigma_n} \left(\frac{P_{n1}(z)}{P_{n0}(z)} \right)^{\sigma_n(z)-1} \left(\frac{\mu_{n1}(a)}{\mu_{n0}(a)} \right)^{1-\sigma_n(z)} \left(\frac{\phi_{n1}(a)}{\phi_{n0}(a)} \right)^{(\sigma_n(z)-1)(\gamma_n(z)-\xi_n)} \quad (49)$$

In these equations, the effect of quality depends on its valuation $\gamma_n(z)$ by income group z net of the effect on the marginal cost, parameterized by ξ_n . This equation is obtained by combining equations 3 and 4:

$$\frac{x_{n1}(a, z)}{x_{n0}(a, z)} = h_n(z) \left(\frac{P_{n1}(z)}{P_{n0}(z)} \right)^{\sigma_n(z)-1} \left(\frac{\phi_{n1}(a)}{\phi_{n0}(a)} \right)^{\gamma_n(z)(\sigma_n(z)-1)} \left(\frac{p_{n1}(a)}{p_{n0}(a)} \right)^{1-\sigma_n(z)}$$

accounting for changes in prices: $p_n(a) = \frac{\phi_n(a)^{\xi_n} \mu_n(a)}{a}$ and survival $\delta_{nD}(a)$. Based on new sales $x_{n1}(z, a)$, total sales of firm a in the counterfactual equilibrium are given by $X_{n1}(a) = \int_z x_{n1}(z, a) dH(z)$.

Next, equation 12 implies that quality upgrading is determined by:

$$\frac{\phi_{n1}(a)}{\phi_{n0}(a)} = \left[\frac{(\tilde{\gamma}_{n1}(a) - \xi_n) \mu_{n1}(a)^{-1} X_{n1}(a)}{(\tilde{\gamma}_{n0}(a) - \xi_n) \mu_{n0}(a)^{-1} X_{n0}(a)} \right]^{\beta_n} \quad (50)$$

where $\tilde{\gamma}_{n0}(a)$ and $\tilde{\gamma}_{n1}(a)$ correspond to the weighted averages of $\gamma_n(z)$ among firm a 's consumers, weighting either sales in the baseline and counterfactual equilibrium respectively (see equation 13). This equation reflects how a change in the income distribution impacts firms' product quality choices, given the differences in quality valuations $\gamma_n(z)$ across consumers. It also reflects a scale effect: firms that expand the most also tend to upgrade their quality. This equation is the same

in both types of counterfactuals. In turn, new markups are determined by sales growth:

$$\mu_{n1}(a) - 1 = \frac{\int_z (x_{n1}/x_{n0}) x_{n0}(z, a) dH(z)}{\int_z (\sigma_n(z) - 1) (1 - (x_{n1}/x_{n0}) s_{n0}(z, a)) (x_{n1}/x_{n0}) x_{n0}(z, a) dH(z)} \quad (51)$$

which we obtain by combining $\mu_{n1}(a) = 1 + \frac{\int_z x_{n1}(z, a) dH(z)}{\int_z (\sigma_n(z) - 1) (1 - s_{n1}(z, a)) x_{n1}(z, a) dH(z)}$ and $\frac{s_{n1}(z, a)}{s_{n0}(z, a)} = \frac{x_{n1}(z, a)}{x_{n0}(z, a)}$.

Thirdly, we need to describe change in the price index $P_n(z)$ for each module n and income group z . Taking ratios of equation 6, and adjusting for the exit of firms in the counterfactual equilibrium, we obtain:

$$\begin{aligned} \frac{P_{n1}(z)}{P_{n0}(z)} &= \left[\frac{N_{n1} \int_a \delta_{nD}(a) \tau_n(a)^{1-\sigma_n} p_{n1}(a)^{1-\sigma_n(z)} \phi_{n1}(a)^{\gamma_n(z)(\sigma_n(z)-1)} dG(a)}{N_{n0} \int_a p_{n0}(a)^{1-\sigma_n(z)} \phi_{n0}(a)^{\gamma_n(z)(\sigma_n(z)-1)} dG(a)} \right]^{\frac{1}{1-\sigma_n(z)}} \\ &= \left[\frac{N_{n1} \int_a \delta_{nD}(a) \tau_n(a)^{1-\sigma_n} p_{n1}(a)^{1-\sigma_n(z)} \phi_{n1}(a)^{\gamma_n(z)(\sigma_n(z)-1)} \alpha_n(z) E(z) P_{n0}^{\sigma_n(z)-1} dG(a)}{N_{n0} \int_a p_{n0}(a)^{1-\sigma_n(z)} \phi_{n0}(a)^{\gamma_n(z)(\sigma_n(z)-1)} \alpha_n(z) E(z) P_{n0}^{\sigma_n(z)-1} dG(a)} \right]^{\frac{1}{1-\sigma_n(z)}} \\ &= \left[\frac{N_{n1} \int_a x_{n0}(z, a) \delta_{nD}(a) \tau_n(a)^{1-\sigma_n} \left(\frac{p_{n1}(a)}{p_{n0}(a)} \right)^{1-\sigma_n(z)} \left(\frac{\phi_{n1}(a)}{\phi_{n0}(a)} \right)^{\gamma_n(z)(\sigma_n(z)-1)} dG(a)}{N_{n0} \int_a x_{n0}(z, a) dG(a)} \right]^{\frac{1}{1-\sigma_n(z)}} \end{aligned}$$

where the second line is obtained by multiplying each line by $\alpha_n(z) E(z) P_{n0}^{\sigma_n(z)-1}$ and the third line by noticing that $p_{n0}(a)^{1-\sigma_n(z)} \phi_{n0}(a)^{\gamma_n(z)(\sigma_n(z)-1)} \alpha_n(z) E(z) P_{n0}^{\sigma_n(z)-1} = x_{n0}(a, z)$. Using the expression $p_n(a) = \frac{\phi_n(a)^{\xi_n} \mu_n(a)}{a}$ for prices, we obtain our main equation describing the change in price indices in our counterfactual equilibrium:

$$\frac{P_{n1}(z)}{P_{n0}(z)} = \left[\frac{N_{n1} \int_a x_{n0}(z, a) \delta_{nD}(a) \tau_n(a)^{1-\sigma_n} \left(\frac{\mu_{n1}(a)}{\mu_{n0}(a)} \right)^{1-\sigma_n(z)} \left(\frac{\phi_{n1}(a)}{\phi_{n0}(a)} \right)^{(\sigma_n(z)-1)(\gamma_n(z)-\xi_n)} dG(a)}{N_{n0} \int_a x_{n0}(z, a) dG(a)} \right]^{\frac{1}{1-\sigma_n(z)}} \quad (52)$$

This ratio is determined by the change in quality weighted by initial sales of each firm. It also depends on the availability of product varieties, the extent to which is a function of the price elasticity $\sigma_n(z)$. Increases in the measure of firms N_{n1} lead to a reduction in the price index, while firm exit ($\delta_{nD}(a) = 0$) leads to an increase.

The entry and exit decisions are determined in a standard way. In a Melitz-type model, free entry is such that expected profits are equal to the sunk cost of entry F_{nE} . Upon entry, firms do not know their productivity and are *ex ante* homogeneous. Firms realize their production after paying the sunk cost of entry. Here, looking at long-term outcomes, free entry implies that average profits π_{n1} (adjusting for exit) remain unchanged in the counterfactual equilibrium:

$$F_{nE} = \int_a \pi_{n0}(a) dG(a) = \int_a \delta_{nD}(a) \pi_{n1}(a) dG(a)$$

Using expression 14 for profits, this is equivalent to the following condition:

$$\begin{aligned} \int_a \left[1 - \frac{1}{\mu_{n0}(a)} - \frac{\beta_n(\tilde{\gamma}_{n0}(a) - \xi_n)}{\mu_{n0}(a)} \right] X_{n0}(a) dG_n(a) = \\ \int_a \delta_{nD}(a) \left[1 - \frac{1}{\mu_{n1}(a)} - \frac{\beta_n(\tilde{\gamma}_{n1}(a) - \xi_n)}{\mu_{n1}(a)} \right] X_{n1}(a) dG_n(a) + \int_a (1 - \delta_{nD}(a)) f_{n0} dG_n(a) \end{aligned} \quad (53)$$

The number of firms N_{n1} adjusts such that this equality holds (e.g. an increase in N_{n1} leads to a

decrease in the price index and a decrease in sales $X_{n1}(a)$, and thus a decrease in the right-hand side of the equation above).

In turn, survival ($\delta_{nD}(a)$ dummy) requires that profits are positive:

$$\left[1 - \frac{1}{\mu_{n1}(a)} - \frac{\beta_n(\tilde{\gamma}_{n1}(a) - \xi_n)}{\mu_{n1}(a)}\right] X_{n1}(a) - f_{n0} > 0 \Leftrightarrow \delta_{nD}(a) = 1 \quad (54)$$

Definition of Counterfactual Equilibrium A counterfactual equilibrium is defined by sales $x_{n1}(z, a)$, quality $\phi_{n1}(a)$, markups $\mu_{n1}(a)$, exit $\delta_{n1}(a)$, entry N_{n1} and price indices $P_n(z)$ that jointly satisfy equations 49, 50, 51, 52, 53 and 54.

Monopolistic competition Under monopolistic competition, all equilibrium conditions remain identical except equation 51. Equation 51 becomes:

$$\mu_{n1}(a) - 1 = \frac{\int_z (x_{n1}/x_{n0}) x_{n0}(z, a) dH(z)}{\int_z (\sigma_n(z) - 1) (x_{n1}/x_{n0}) x_{n0}(z, a) dH(z)}$$

Uniqueness Is the counterfactual equilibrium unique? Appendix 2 examines second-order conditions but this only shows that firms decisions are local maxima and this does not prove global uniqueness for joint decisions. We show here that, conditional on entry, exit and markups under monopolistic competition, the changes in quality, sales and price indices are uniquely determined by equations 49, 50 and 52, under the condition:

$$\beta_n(\sigma_n(z) - 1)(\gamma_n(z) - \xi_n) < 1 \quad \text{for all } z$$

Below, we further argue informally that other margins are not likely to threaten uniqueness.

The proof of (conditional) uniqueness in sales and quality is inspired by Kehoe, Levine and Romer (1992) and Kucheryavyi, Lyn and Rodriguez-Clare (2018). We consider a convex function of quality choices $\phi_{n1}(a)$ whose first order conditions coincide with our counterfactual equilibrium conditions. Specifically, taking markups $\mu_{n1}(a)$, exit $\delta_{n1}(a)$ and entry N_{n1} as given, consider the function F_n of the vector of adjusted quality choices $\tilde{\phi}_{n1} = \phi_{n1}^{\frac{1}{\beta_n}}$ across surviving firms:

$$F_n(\tilde{\phi}_{n1}) = \sum_a \int_z \frac{\beta_n(\gamma_n(z) - \xi) \mu_{n1}(a) \phi_{n0}^{-\frac{1}{\beta_n}}(a)}{\mu_{n0}(a)} x_{n0}(z, a) \tilde{\phi}_{n1}(a) dH(z) - \frac{N_{n0}}{N_{n1}} \int_z \frac{\alpha_n(z) E(z) h_n(z)}{\sigma_n(z) - 1} \log \left(\sum_a x_{n0}(a, z) \tau_n(a)^{1 - \sigma_n(z)} \left(\frac{\mu_{n1}(a)}{\mu_{n0}(a)} \right)^{1 - \sigma_n(z)} \left(\frac{\tilde{\phi}_{n1}(a)}{\phi_{n0}^{\frac{1}{\beta_n}}(a)} \right)^{\beta_n(\sigma_n(z) - 1)(\gamma_n(z) - \xi_n)} \right) dH(z) \quad (55)$$

This function is convex in the vector of quality choices $\tilde{\phi}_{n1}$ under the assumption $\beta_n(\sigma_n(z) - 1)(\gamma_n(z) - \xi_n) < 1$. To see this, notice that each term $\left(\frac{\tilde{\phi}_{n1}}{\tilde{\phi}_{n0}}\right)^{\beta_n(\sigma_n(z) - 1)(\gamma_n(z) - \xi_n)}$ is concave in $\tilde{\phi}_{n1}/\tilde{\phi}_{n0}$, and that these functions are combined with the log function, which is itself increasing and concave. The second term is therefore concave in the vector of $\tilde{\phi}_{n1}/\tilde{\phi}_{n0}$ while the first term is linear, from which we can conclude that function F_n is convex. Hence, a minimum in F_n is unique. What remains to be shown is that the derivatives are zero (first-order conditions) precisely for firm-level quality $\phi_{n1}(a)$ that satisfy the counterfactual equilibrium conditions above.

The derivatives can be expressed as:

$$\begin{aligned} \frac{\partial F_n}{\partial \log \tilde{\phi}_{n1}(a)} &= \left(\frac{\phi_{n1}(a)}{\phi_{n0}(a)} \right)^{\frac{1}{\beta_n}} \int_z \beta_n (\sigma_n(z) - 1) (\gamma_n(z) - \xi_n) x_{n0}(z, a) dH(z) \\ &\quad - \int_z \frac{\beta_n (\gamma_n(z) - \xi_n)}{\mu_{n1}(a)} x_{n0}(a, z) h_n(z) \tau_n(a)^{1 - \sigma_n(z)} \\ &\quad \times \left(\frac{\mu_{n1}(a)}{\mu_{n0}(a)} \right)^{1 - \sigma_n(z)} \left(\frac{\phi_{n1}(a)}{\phi_{n0}(a)} \right)^{(\sigma_n(z) - 1)(\gamma_n(z) - \xi_n)} \left(\frac{P_{n1}(z)}{P_{n0}(z)} \right)^{\sigma_n(z) - 1} dH(z) \end{aligned}$$

where we replace $\tilde{\phi}_{n1}$ by $\phi_{n1}^{\frac{1}{\beta_n}}$, noticing that $\frac{P_{n1}(z)}{P_{n0}(z)}$ coincides with the price index change:

$$\frac{N_{n1} \int_a x_{n0}(z, a) \delta_{nD}(a) \tau_n(a)^{1 - \sigma_n(z)} \left(\frac{\mu_{n1}(a)}{\mu_{n0}(a)} \right)^{1 - \sigma_n(z)} \left(\frac{\phi_{n1}(a)}{\phi_{n0}(a)} \right)^{(\sigma_n(z) - 1)(\gamma_n(z) - \xi_n)} dG(a)}{N_{n0} \int_a x_{n0}(z, a) dG(a)} = \left(\frac{P_{n1}(z)}{P_{n0}(z)} \right)^{1 - \sigma_n(z)}$$

The derivatives are zero when the following condition is satisfied for all firms:

$$\begin{aligned} \left(\frac{\phi_{n1}(a)}{\phi_{n0}(a)} \right)^{\frac{1}{\beta_n}} &= \\ \frac{\int_z \frac{\beta_n (\gamma_n(z) - \xi_n)}{\mu_{n1}(a)} x_{n0}(a, z) h_n(z) \tau_n(a)^{1 - \sigma_n(z)} \left(\frac{\mu_{n1}(a)}{\mu_{n0}(a)} \right)^{1 - \sigma_n(z)} \left(\frac{\phi_{n1}(a)}{\phi_{n0}(a)} \right)^{(\sigma_n(z) - 1)(\gamma_n(z) - \xi_n)} \left(\frac{P_{n1}(z)}{P_{n0}(z)} \right)^{\sigma_n(z) - 1} dH(z)}{\int_z \frac{\beta_n (\gamma_n(z) - \xi_n)}{\mu_{n0}(a)} x_{n0}(z, a) dH(z)} \end{aligned}$$

The same expression can be obtained by combining equilibrium conditions 49 and 50, after noticing that equation 50 can be rewritten as:

$$\left(\frac{\phi_{n1}(a)}{\phi_{n0}(a)} \right)^{\frac{1}{\beta_n}} = \frac{\int_z \beta_n (\gamma_n(z) - \xi) \mu_{n1}(a)^{-1} x_{n1}(z, a) dH(z)}{\int_z \beta_n (\gamma_n(z) - \xi) \mu_{n0}(a)^{-1} x_{n0}(z, a) dH(z)}$$

This proves uniqueness in the vector of quality choices across firms, conditional on their markups, entry, survival, etc. and assuming monopolistic competition. Unfortunately, this approach does not generalize to oligopoly where markups also depend on market power because the derivatives of the first order conditions are not symmetric in this case (a requirement for the existence of a maximizing function such as the function F_n above).

However, computationally we never encountered any instances of multiplicity of equilibria (or issues with convergence) within our range of parameter estimates. The simulations always converge to the same solutions and do so quickly, providing reassurance against the concern of multiplicity in our applications. This also makes sense in theory, since i) price elasticities do not vary greatly across income groups in our calibration in Section 5.1, thus avoiding interactions between quality and markup choices as discussed in Appendix 2.B; and ii) market power under oligopolistic competition acts as a negative feedback: for firms with larger market shares, choosing larger markups induces lower quality upgrading and consequently smaller market shares in our framework.

5.B) Decomposition of the Price Index Effect

The differential change in the price index for rich vs. poor households depends on a variety of channels that we can illustrate through a decomposition, taking a first-order approximation. For a given income group z , the price index change equals:

$$\begin{aligned} \frac{P_{n1}(z)}{P_{n0}(z)} &= \left[\frac{N_{n1} \int_a x_{n0}(z, a) \delta_{nD}(a) \tau_n(a)^{1-\sigma_n} \left(\frac{\mu_{n1}(a)}{\mu_{n0}(a)} \right)^{1-\sigma_n(z)} \left(\frac{\phi_{n1}(a)}{\phi_{n0}(a)} \right)^{(\sigma_n(z)-1)(\gamma_n(z)-\xi_n)} dG_n(a)}{N_{n0} \int_a x_{n0}(z, a) dG_n(a)} \right]^{\frac{1}{1-\sigma_n(z)}} \\ &= \left[\int_a s_{n1}(a, z) \left(\frac{\mu_{n1}(a)}{\mu_{n0}(a)} \right)^{1-\sigma_n(z)} \left(\frac{\phi_{n1}(a)}{\phi_{n0}(a)} \right)^{(\sigma_n(z)-1)(\gamma_n(z)-\xi_n)} dG_n(a) \right]^{\frac{1}{1-\sigma_n(z)}} \\ &\quad \times \left[\frac{N_{n1}}{N_{n0}} \int_a s_{n0}(a, z) \delta_{nD}(a) \tau_n(a)^{1-\sigma_n} dG_n(a) \right]^{\frac{1}{1-\sigma_n(z)}} \end{aligned}$$

where we denote $s_{n0}(a, z) = \frac{x_{n0}(z, a)}{\int_{a'} x_{n0}(z, a') dG_n(a')}$ and $s_{n1}(a, z) = \frac{\delta_{nD}(a) \tau_n(a)^{1-\sigma_n} x_{n0}(z, a)}{\int_{a'} \delta_{nD}(a') \tau_n(a')^{1-\sigma_n} x_{n0}(z, a') dG_n(a')}$.

Taking logs and a first-order approximation leads to:

$$\begin{aligned} \log \frac{P_{n1}(z)}{P_{n0}(z)} &= -\frac{1}{\sigma_n(z) - 1} \log \left[\int_a s_{n1}(a, z) \left(\frac{\mu_{n1}(a)}{\mu_{n0}(a)} \right)^{1-\sigma_n(z)} \left(\frac{\phi_{n1}(a)}{\phi_{n0}(a)} \right)^{(\sigma_n(z)-1)(\gamma_n(z)-\xi_n)} dG_n(a) \right] \\ &\quad - \frac{1}{\sigma_n(z) - 1} \log \left[\frac{N_{n1}}{N_{n0}} \int_a s_{n0}(a, z) \delta_{nD}(a) \tau_n(a)^{1-\sigma_n} dG_n(a) \right] \\ &\approx -(\gamma_n(z) - \xi_n) \int_a s_{n1}(a, z) \log \left(\frac{\phi_{n1}(a)}{\phi_{n0}(a)} \right) dG_n(a) + \int_a s_{n1}(a, z) \log \tau_n(a) dG_n(a) \\ &\quad + \int_a s_{n1}(a, z) \log \left(\frac{\mu_{n1}(a)}{\mu_{n0}(a)} \right) dG_n(a) - \frac{1}{\sigma_n(z) - 1} \log \left[\frac{N_{n1}}{N_{n0}} \int_a s_{n0}(a, z) \delta_{nD}(a) dG_n(a) \right] \end{aligned}$$

Next, by comparing income groups z and z_0 , we have:

$$\begin{aligned} \log \frac{P_{n1}(z)}{P_{n0}(z)} - \log \frac{P_{n1}(z_0)}{P_{n0}(z_0)} &\approx -(\gamma_n(z) - \xi_n) \int_a s_{n1}(a, z) \log \left(\frac{\phi_{n1}(a)}{\phi_{n0}(a)} \right) dG_n(a) \\ &\quad + (\gamma_n(z_0) - \xi_n) \int_a s_{n1}(a, z_0) \log \left(\frac{\phi_{n1}(a)}{\phi_{n0}(a)} \right) dG_n(a) \\ &\quad + \int_a (s_{n1}(a, z) - s_{n1}(a, z_0)) \log \tau_n(a) dG_n(a) \\ &\quad + \int_a (s_{n1}(a, z) - s_{n1}(a, z_0)) \log \left(\frac{\mu_{n1}(a)}{\mu_{n0}(a)} \right) dG_n(a) \\ &\quad - \left(\frac{1}{\sigma_n(z) - 1} - \frac{1}{\sigma_n(z_0) - 1} \right) \log \left[\frac{N_{n1}}{N_{n0}} \right] \\ &\quad - \frac{1}{\sigma_n(z) - 1} \left[\int_a s_{n0}(a, z) \delta_{nD}(a) dG_n(a) \right] \\ &\quad + \frac{1}{\sigma_n(z_0) - 1} \log \left[\int_a s_{n0}(a, z_0) \delta_{nD}(a) dG_n(a) \right] \end{aligned}$$

Using the equality $AB - A'B' = (A - A') \left(\frac{B+B'}{2} \right) + (B - B') \left(\frac{A+A'}{2} \right)$ that holds for any four numbers A, A', B and B' , we can rewrite the first two lines and the last two lines of the previous

sum:

$$\begin{aligned}
\log \frac{P_{n1}(z)}{P_{n0}(z)} - \log \frac{P_{n1}(z_0)}{P_{n0}(z_0)} &\approx \int_a (s_{n1}(a,z) - s_{n1}(a,z_0)) \log \tau_n(a) dG_n(a) \\
&- (\gamma_n(z) - \gamma_n(z_0)) \int_a \bar{s}_{n1}(a) \log \left(\frac{\phi_{n1}(a)}{\phi_{n0}(a)} \right) dG_n(a) \\
&- (\bar{\gamma}_n - \xi_n) \int_a (s_{n1}(a,z) - s_{n1}(a,z_0)) \log \left(\frac{\phi_{n1}(a)}{\phi_{n0}(a)} \right) dG_n(a) \\
&+ \int_a (s_{n1}(a,z) - s_{n1}(a,z_0)) \log \left(\frac{\mu_{n1}(a)}{\mu_{n0}(a)} \right) dG_n(a) \\
&- \left(\frac{1}{\sigma_n(z) - 1} - \frac{1}{\sigma_n(z_0) - 1} \right) \log \left[\frac{N_{n1}}{N_{n0}} \right] \\
&- \left(\frac{1}{\sigma_n(z) - 1} - \frac{1}{\sigma_n(z_0) - 1} \right) \log \left[\int_a \bar{s}_{n0}(a) \delta_{nD}(a) dG_n(a) \right] \\
&- \frac{1}{\bar{\sigma}_n - 1} \log \left[\frac{\int_a s_{n0}(a,z) \delta_{nD}(a) dG_n(a)}{\int_a s_{n0}(a,z_0) \delta_{nD}(a) dG_n(a)} \right]
\end{aligned}$$

where $\bar{s}_{n1}(a,z)$ is the average of $s_{n1}(a,z)$ and $s_{n1}(a,z_0)$, and $\frac{1}{\bar{\sigma}_n - 1}$ is the average of $\frac{1}{\sigma_n(z) - 1}$ and $\frac{1}{\sigma_n(z_0) - 1}$.

Combining lines 4 and 5 together, denoting $\bar{\delta}_{nD} = \int_a \delta_{nD}(a) \bar{s}_{n0}(a) dG(a)$, we obtain a six-term decomposition of the price index change:

$$\begin{aligned}
\log \frac{P_{n1}(z)}{P_{n0}(z)} - \log \frac{P_{n1}(z_0)}{P_{n0}(z_0)} &= \underbrace{\int_a (s_{n1}(a,z) - s_{n1}(a,z_0)) \log \tau_n(a) dG(a)}_{(0) \text{ Direct business tax effect}} \tag{56} \\
&- \underbrace{(\gamma_n(z) - \gamma_n(z_0)) \int_a \bar{s}_{n1}(a) \log \left(\frac{\phi_{n1}(a)}{\phi_{n0}(a)} \right) dG(a)}_{(1) \text{ Average quality effect}} \\
&- \underbrace{(\bar{\gamma}_n - \xi_n) \int_a (s_{n1}(a,z) - s_{n1}(a,z_0)) \log \left(\frac{\phi_{n1}(a)}{\phi_{n0}(a)} \right) dG(a)}_{(2) \text{ Asymmetric quality-adjusted cost changes}} \\
&+ \underbrace{\int_a (s_{n1}(a,z) - s_{n1}(a,z_0)) \log \left(\frac{\mu_{n1}(a)}{\mu_{n0}(a)} \right) dG(a)}_{(3) \text{ Asymmetric markup changes}} \\
&- \underbrace{\left(\frac{1}{\sigma_n(z) - 1} - \frac{1}{\sigma_n(z_0) - 1} \right) \log \left(\frac{N_{n1} \bar{\delta}_{nD}}{N_{n0}} \right)}_{(4) \text{ Love of variety}} \\
&- \underbrace{\frac{1}{\bar{\sigma}_n - 1} \log \left(\frac{\int_a s_{n0}(a,z) \delta_{nD}(a) dG_n(a)}{\int_a s_{n0}(a,z_0) \delta_{nD}(a) dG_n(a)} \right)}_{(5) \text{ Asymmetric exit effects}}
\end{aligned}$$

The first line actually only appears in the second counterfactual. It directly reflects the asymmetric effect of a business tax on consumers, holding constant all allocations on the firm or consumer side. It is non-zero if firms' market shares vary across the income distribution in a way that also covaries with the direct incidence of the tax across firms.

In both counterfactuals, the effect in term (1) is that firms on average have incentives to downgrade their product quality, which has heterogeneous effects across households depending on their preference parameters $\gamma_n(z)$ (quality upgrading benefits households with the highest $\gamma_n(z)$ relatively more).

The effect in term (2) is that the scale of production changes asymmetrically across higher and lower quality producers in both counterfactuals. With economies of scale in quality production ($\beta_n > 0$), this translates into asymmetric effects on quality and quality-adjusted prices. In turn, this second effect hurts richer households if they spend relatively more on firms with the largest decrease in scale and quality. As seen in the second term, this channel can be expressed as a covariance term between consumer-specific budget shares $s_n(a, z)$ and firms' incentives to downgrade product quality $\log\left(\frac{\phi_{n1}(a)}{\phi_{n0}(a)}\right)$.

The effect in term (3) captures the change in markups, which in our framework differ endogenously across firms as a function of their market power and the composition of consumers that they sell to. These markups can be affected asymmetrically across higher and lower quality producers. Firms who experience the largest change in scale or in the composition of their consumer base have incentives to adjust their markups the most, which can give rise to asymmetric changes in markups across consumption baskets due to uneven consumption shares of rich and poor households across the firm size distribution.

Term (4) shows that the change in the overall number of product varieties can have asymmetric impacts across households depending on their elasticity of substitution across products $\sigma_n(z)$. More product entry benefits households with higher estimated love of variety, i.e. lower $\sigma_n(z)$, and vice-versa for product exit.

In addition to differences in the love of variety, the last channel reflects the unequal effects of exit as a function of differences in consumption shares across household groups. Exiting firms tend to be the smallest firms. Since small firms tend to sell relatively more to poor consumers, exit tends to hurt poorer consumers relatively more than richer consumers (abstracting from differences in $\sigma_n(z)$). This is reflected in the sign of term (5), which depends on whether the sales-weighted survival rate is lower for income group z compared to the average.

Finally, we aggregate these terms across product modules to obtain a decomposition of the aggregate price index change for retail consumption. A standard within-between decomposition yields:

$$\begin{aligned} \log \frac{P_{G1}(z)}{P_{G0}(z)} - \log \frac{P_{G1}(z_0)}{P_{G0}(z_0)} &= \underbrace{\sum_n \left(\frac{\alpha_n(z) + \alpha_n(z_0)}{2} \right) \left(\log \frac{P_{n1}(z)}{P_{n0}(z)} - \log \frac{P_{n1}(z_0)}{P_{n0}(z_0)} \right)}_{(0+1+2+3+4+5) \text{ Within-module changes}} \quad (57) \\ &+ \underbrace{\sum_n (\alpha_n(z) - \alpha_n(z_0)) \left(\frac{\log \frac{P_{n1}(z)}{P_{n0}(z)} + \log \frac{P_{n1}(z_0)}{P_{n0}(z_0)}}{2} \right)}_{(6) \text{ Between-module changes}} \end{aligned}$$

The within term can be decomposed into the five terms described in equation 56. The between term reflects the covariance between product module-level relative price changes and the cross-module differences in consumption shares between rich and poor: this term is positive if prices tend to increase faster in product modules where households from income group z tend to spend a larger fraction of their retail expenditures relative to income group z_0 . As our analysis follows the literature on firm heterogeneity within sectors, our theory is focused on relative price changes across producers and consumers within product groups, and has little to say about price changes across sectors. Nevertheless, rich and poor households have different consumption shares across product groups (the upper-tier $\alpha_n(z)$), and even within our framework the firm size distributions and preference and technology parameters can differ across the n dimension in arbitrary ways, so that the between-module term need not be zero.

5.C) International Trade Counterfactual

Our third counterfactual illustrates the role of reducing trade costs in a setting with heterogeneous firms, as in [Melitz \(2003\)](#), in addition to heterogeneous households who source their consumption differently across the firm size distribution as observed in the scanner data. The documented empirical findings and our quantitative framework have clear implications for the distribution of the gains from trade. As in [Melitz \(2003\)](#), a decrease in trade costs induces a reallocation in which the largest firms expand through trade while less productive firms either shrink or exit. In our framework, better access to imported varieties and exit of domestic producers affect the price indices of rich and poor households asymmetrically. In addition, lower trade costs also lead to heterogeneous changes in product quality and markups across firms. Armed with our parameter estimates, we can quantify these effects on the cost of living across the income distribution.

To illustrate and quantify these forces, we introduce quality choice under two-sided heterogeneity into an otherwise standard [Melitz \(2003\)](#) model with monopolistic competition and with two symmetric countries. We simulate an increase in the openness to trade where, as is typically the case, only a fraction of the firms start exporting, and where exporters sell only a small share of their output abroad. We calibrate fixed trade costs f_X such that half of of output is produced by exporting firms. We calibrate variable trade costs τ such that export sales of exporters equal 20 percent of their output. Combining these two statistics, about 10 percent of aggregate output is traded. The counterfactual is to reduce variable trade costs from an equilibrium with no trade to the new trade equilibrium. This overall increase in trade shares is moderate. In comparison, trade over GDP has increased from 20 percent to 30 percent in the US since 1990, and other countries have seen much larger increases (since 1990, the trade-to-GDP ratios have increased from on average 40 to 60 percent across countries according to the World Development Indicators).

As for other counterfactuals, we derive that the changes in firm sales, quality, entry, exit and price indices must satisfy five equilibrium conditions respectively, and there is here an additional condition reflecting the decision to export. Equilibrium conditions describing quality and exit are identical to those in the main counterfactuals (but with markups now under monopolistic competition), whereas sales, price indices and entry have to explicitly account for trade and trade costs.

First, the evolution of firm sales for a given income group z now also depends on the export decision (in addition to quality upgrading and the price index change for each consumer group):

$$\frac{x_{n1}(z, a)}{x_{n0}(z, a)} = \delta_{nD}(a) \left(1 + \delta_n^X(a) \tau_n^{1-\sigma_n}\right) \left(\frac{P_{n1}(z)}{P_{n0}(z)}\right)^{\sigma_n(z)-1} \left(\frac{\mu_{n1}(a)}{\mu_{n0}(a)}\right)^{1-\sigma_n(z)} \left(\frac{\phi_{n1}(a)}{\phi_{n0}(a)}\right)^{(\sigma_n(z)-1)(\gamma_n(z)-\xi_n)}$$

where we denote by $\delta_n^X(a)$ an export dummy equal to one if firm a exports in the counterfactual equilibrium. Other terms are identical to those in the main counterfactuals.

Changes in the price index $P_n(z)$ for each module n and income group z now need to account for trade costs and new varieties of imported goods. Assuming symmetry between the domestic and foreign economies, this additional margin is captured by the term $(1 + \delta_n^X(a) \tau_n^{1-\sigma_n(z)})$:

$$\frac{P_{n1}(z)}{P_{n0}(z)} = \left[\frac{N_{n1} \int_a x_{n0}(z, a) \delta_{nD}(a) \left(1 + \delta_n^X(a) \tau_n^{1-\sigma_n(z)}\right) \left(\frac{\mu_{n1}(a)}{\mu_{n0}(a)}\right)^{1-\sigma_n(z)} \left(\frac{\phi_{n1}(a)}{\phi_{n0}(a)}\right)^{(\sigma_n(z)-1)(\gamma_n(z)-\xi_n)} dG(a)}{N_{n0} \int_a x_{n0}(z, a) dG(a)} \right]^{\frac{1}{1-\sigma_n(z)}}$$

The decision to export is as in [Melitz \(2003\)](#) except that the firm also has to account for its choice of quality which is itself endogenous to its export decision. Firm a decides to export if and only if its revenue gains on both the export and domestic market, exceed the fixed cost of

exporting, net of quality upgrading costs:

$$r_n^X(a, \phi_{n1}^X(a)) + r_{n1}^D(a, \phi_{n1}^X(a)) - f_n(\phi_{n1}^X(a)) - f_{nX} > r_n^D(a, \phi_{n1}^D(a)) - f_n(\phi_{n1}^D(a))$$

where $r_n^X(a, \phi_{n1}^X(a))$ denotes revenues net of variable costs on the export market (exports times $\frac{1}{\bar{\sigma}_n}$) where its quality $\phi_{n1}^X(a)$ is the optimal quality if the firm exports. The terms $r_n^D(a, \phi)$ denote revenues net of variable costs on the domestic market where its quality is the optimal quality if the firm exports (left-hand side) or if the firm does not export (right-hand side). As before, $f_n(\phi)$ denotes the fixed costs of upgrading to quality ϕ which itself depends on whether the firm exports or not.

Finally, the entry decision has to also account for profits made on the export market, as in [Melitz \(2003\)](#).

Decomposition of the Price Index Effect Having simulated the price index change, we can again decompose the overall change in various terms that reflect different channels that can potentially affect consumers differentially in this counterfactual. These differential effects rely essentially on the fact that richer households tend to purchase from larger producers.

As in the main counterfactuals, quality may change across firms. This quality increase is primarily due to a scale effect: export opportunities lead firms to expand and thus invest in quality upgrading due to economies of scale in quality production. An average increase in quality tends to benefit richer households who have the highest preferences for quality, $\gamma_n(z)$. In addition, the largest firms are the ones who become exporters and thus have the highest incentives for quality upgrading due to the larger scale of their operation. They are also the ones whose initial sales are more concentrated among richer consumers. The heterogeneity of this scale effect reinforces the effect of the average increase in product quality. Our decomposition results of the trade counterfactual (see NBER working paper version) indicate that quality upgrading account for about half of the overall effect on inflation differences.

Another effect captures the change in the overall number of product varieties, which has asymmetric impacts across households depending on their love for variety. This can be captured by a new term in the decomposition:

$$-\left(\frac{1}{\sigma_n(z) - 1} - \frac{1}{\sigma_n(z_0) - 1}\right) \log \left[\frac{N_{n1} \bar{\delta}_{nD} (1 + \bar{\delta}_X \tau_n^{1 - \bar{\sigma}_n})}{N_{n0}} \right]$$

where $\bar{\delta}_{nD}(1 + \bar{\delta}_X \tau_n^{1 - \bar{\sigma}_n})$ denotes the average of $\delta_{nD}(a)(1 + \delta_X(a) \tau_n^{1 - \sigma_n(z_0)})$ across consumers and firms. This term now combines the number of varieties that are available on the domestic market as well as new imported varieties. Quantitatively, we find that this channel account for about a quarter of the overall effect.

While the previous channel is driven by differences in $\sigma_n(z)$, finally another channel takes into account differences in consumption shares spent on new imported varieties or exiting domestic firms across rich and poor households:

$$-\frac{1}{\bar{\sigma}_n - 1} \log \left[\frac{\int_a s_{n0}(a, z) \delta_{nD}(a) (1 + \delta_X(a) \tau_n^{1 - \sigma_n(z)}) dG_n(a)}{\int_a s_{n0}(a, z_0) \delta_{nD}(a) (1 + \delta_X(a) \tau_n^{1 - \sigma_n(z_0)}) dG_n(a)} \right]$$

Due to selection into exporting, the products that are traded tend to be those consumed to a higher extent by the richest households. Access to imported varieties thus benefits richer households relatively more compared to the poor. In addition, domestic exit due to import competition is concentrated among producers whose sales are concentrated among poorer households. Quantitatively, we find that this channel account for about a quarter of the overall effect. In contrast, differential changes in markups, exit and between effects account for a small share.