

Firm Heterogeneity in Consumption Baskets: Evidence from Home and Store Scanner Data*

Benjamin Faber[†] and Thibault Fally[‡]

July 2017

Abstract

A growing literature has documented the role of firm heterogeneity within sectors in accounting 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 microdata 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 wealthier households, and that this gives rise to asymmetric effects on household price indices. These effects matter for real income inequality. We find that they amplify observed changes in nominal inequality over time, lead to a more regressive distribution of the gains from international trade, and give rise to new distributional implications of business regulations.

Keywords: Firm heterogeneity, real income inequality, household price indices, scanner data
JEL Classification: F15, F61, E31

*We would like to thank Mary Amity, Costas Arkolakis, David Atkin, 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.

[†]Department of Economics, UC Berkeley and NBER.

[‡]Department of Agricultural and Resource Economics, UC Berkeley and NBER.

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 (Acemoglu & Autor, 2011; Piketty & Saez, 2003). In this context, a recent and 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. In particular, we aim to answer three central questions:² i) to what extent do rich and poor households source their consumption baskets from different parts of the firm size distribution?; ii) what explains these differences?; and iii) what are the implications of the answers to i) and ii) for 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 large and significant differences in the weighted average firm sizes that rich and poor US households source their consumption from, and explore alternative explanations. Second, to rationalize these moments we develop a tractable quantitative model that features two-sided heterogeneity across both firms in production and consumers on the demand side, and calibrate it using the microdata to quantify the channels underlying the observed stylized facts. Third, we explore model-based general equilibrium counterfactuals to illustrate how, in a setting where households source their consumption from heterogeneous firms, economic shocks give rise to asymmetric effects on cost of living inflation across the income distribution.

At the center of the analysis lies the construction of an extremely detailed collection of microdata that, for the first time, allows 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 18 six-month periods. The combination of home and store-level scanner microdata allows us to trace the size distribution of producers of brands (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).³

The analysis proceeds in four steps. In step 1, we use the microdata 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

¹E.g. (Bloom et al., 2017; Card et al., 2013; Helpman et al., 2017). See discussion at the end of this section.

²These questions are the analogue to the existing literature on Melitz-type firm heterogeneity and inequality, namely: a) to what extent do higher and lower-income workers source their nominal earnings from different parts of the firm size distribution, b) what explains these differences, and c) what are the implications for inequality. For example, Helpman et al. (2010) and Davis & Harrigan (2011) focus on differences in wage premia across the firm size distribution for homogeneous workers. Alternatively, e.g. Harrigan & Reshef (2011) and Sampson (2014) focus on differences in skill intensity across the firm size distribution.

³The data are made available through an academic user agreement with the Kilts Center at Chicago Booth.

that these differences in weighted-average firm sizes across consumption baskets arise in a setting where the rank order of household budget shares spent on different producers within a product module is preserved across the income distribution –i.e. the largest firms command the highest budget shares for all income groups. After exploring a number of alternative explanations using the richness of the data, we interpret these stylized facts as equilibrium outcomes in a setting where both consumers and firms optimally choose their product attributes.

In step 2 we write down a tractable model that rationalizes these observed moments 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 product quality choice into a Melitz model of heterogeneous firms within sectors. These firms now operate in a setting where their choice of product quality attributes and prices endogenously affect the composition of heterogeneous consumers that shapes each firm’s market demand. Modeling optimal product choices with two-sided heterogeneity implies that shocks that affect producers differently, such as trade integration or business regulations, can feed into the consumption baskets of rich and poor households asymmetrically. Conversely, changes in the income distribution affect firms differently across the size distribution, which in turn again can affect consumption baskets asymmetrically. We use the model to derive estimation equations for the key preference and technology parameters as functions of observable moments in the home and store scanner microdata. Armed with these estimates and the raw moments from the data, we then quantify the role of different forces that underlie the firm heterogeneity across consumption baskets we document in step 1, and use our framework to explore general equilibrium counterfactuals. The remaining steps of the analysis tackle each of these in turn.

In step 3, we use the microdata to estimate the preference and technology parameters. On the consumption 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 higher-income households, but that these differences, while statistically significant, are relatively minor in terms of magnitudes. 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 product attributes that consumers evaluate as higher quality significantly 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 the existing literature, and is based on cross-sectional variation in brand quality and the scale of production. The second exploits within-firm changes in brand quality and scale over time. We regard the panel data approach as more conservative.⁴ To identify the effect of firm scale on product quality in the panel estimation, we use state-level measures of changes in brand quality on the left-hand side, and construct a shift-share instrument for national brand scale on the right-hand side that exploits pre-existing differences in brand-level sales across other US states interacted with state-level variation in average sales growth observed in other product groups. For

⁴For two reasons: first, adjustments to product quality in response to changes in scale are likely best understood as longer-term effects. Second, the cross-sectional approach can be subject to omitted variable bias that the panel data approach addresses. We further discuss this in Sections 4 and 5.

identification, we thus exploit within-state-by-product module variation in changes to the national scale of brand producers that are driven by differences in pre-existing geographical exposure to demand shocks in other states (which in turn are constructed from broad-based changes in demand that are not specific to the producer or its product group).

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 that, in turn, choose to produce at lower 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 moments from step 1 translate into statistically and economically 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. Our framework also gives rise to varying markups across the firm size distribution: because the sales of larger firms are driven to a larger extent by richer, less price-sensitive households, markups within product groups monotonically increase with firm size. Overall, we find that the calibrated model based on the estimates from step 3 can both qualitatively and quantitatively reproduce the differences in firm sizes across consumption baskets observed in step 1.

In the final step 4, we use the calibrated model to explore a number of general equilibrium counterfactuals that explore the distributional consequences of changes in nominal income inequality, trade liberalization and business regulations/taxes. In the first part, we find that increases in nominal inequality lead to an endogenous amplification in terms of real income inequality due to asymmetric knock-on effects on household price indices. We explore these effects in two counterfactuals. In the first, we simulate a hypothetical 5 percent transfer of market expenditure from the poorest to the richest 20 percent of US households. This gives rise to a 1.5-2 percentage point higher cost of living inflation in retail consumption for the poorest household quintile compared to the richest. In the second counterfactual, we quantify the implications of the change in market expenditure shares across US income quintiles observed since the 1980s. We find that this has led to a 2-2.5 percentage point higher cost of living inflation in retail consumption for the poorest quintile compared to the richest.⁵ We also find that increases in nominal income inequality amplify the extent of firm heterogeneity within sectors on the production side.

These results are driven by a number of underlying channels that we quantify. The first is that firms on average have incentives to upgrade their product quality since more of total sales are now

⁵Given that our calibration allows us to model changes in the distribution of incomes across five broad income bins, while abstracting from increases in inequality within quintiles, we consider these results to be conservative.

in the hands of households with relatively higher taste for quality. Given the estimated preference parameters, this channel decreases consumer price inflation for richer households compared to the poor. The second effect is that the scale of production changes asymmetrically across the initial firm size distribution. Given the estimated economies of scale in quality production, this reinforces the first effect in favor of richer households, who spend more of their consumption on initially higher-quality products that experience an expansion in scale and a reduction in quality-adjusted prices. The third effect is that markups are affected asymmetrically across higher and lower-quality producers due to the different extent to which the composition of their demand is affected. Finally, changes in product variety affect the price indices of rich and poor households asymmetrically. More product entry benefits richer households slightly more due to higher estimated love of variety, while the induced exit of firms is concentrated among low-quality producers, both of which tend to work in favor of relatively less cost of living inflation among higher-income households.

In our third counterfactual, we quantify the implications for the distribution of the gains from trade. We do this in an otherwise standard [Melitz \(2003\)](#) framework in which distributional effects would be zero in the absence of the price index effects that we study. We find that a 10 percentage point increase⁶ in import penetration between two symmetric countries leads to a 1.5-2.5 percentage point lower cost of living inflation in retail consumption for the richest 20 percent of US households compared to the poorest 20 percent.

These effects arise because, as in [Melitz \(2003\)](#), heterogeneous producers respond differently to trade cost shocks, but in a setting where it is now also the case that consumers source their consumption differently across the firm size distribution. Again, we decompose this total effect into several distinct channels. First, wealthier consumers benefit more from imports that are driven by the largest producers from abroad, and their price indices increase less due to the exit of less productive domestic firms compared to the poor. Richer households also benefit more from the overall increase in available variety, again due to higher estimated love of variety. Second, the trade shock induces firms on average to upgrade product quality, which benefits higher-income households more. Finally, it is the initially larger firms who become exporters and have incentives for quality upgrading due to the enlarged market. These firms also initially sell a higher proportion of their output to richer consumers, so that the covariance between the scale effect and household consumption shares further reinforces relatively lower inflation among richer consumers.

In a fourth counterfactual, we explore the distributional effects of business regulations. Following the notion that taxes and red tape increase in producer size (e.g. [Hsieh & Klenow \(2009\)](#)), we quantify the distributional implications of a 10 percent tax that applies to large firms accounting for 20 percent of sectoral production. We find that the incidence of this tax on higher cost of living is 10-20 percent stronger among the richest quintile of US households compared to the poorest 20 percent of households. Overall, the findings from the counterfactual analyses illustrate a rich new interplay of adjustments that arise in the presence of two-sided heterogeneity. These effects significantly alter the impact of economic shocks on real income inequality through usually unobserved asymmetric price index effects on consumption baskets across the income distribution.

This paper is related to the growing literature on the extent, causes and consequences of firm

⁶This is a moderate increase in trade openness that compares to about half of the average increase experienced across countries since the 1990s.

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, our paper is most closely related to existing work on the implications of firm heterogeneity for nominal income inequality (Bloom et al., 2017; Burstein & Vogel, 2015; Card et al., 2013; Davis & Harrigan, 2011; Frias et al., 2009; Helpman et al., 2017, 2010; Sampson, 2014). Our analysis using the scanner data also follows recent work Hottman et al. (2016) who use the US home scanner data to decompose Melitz-type firm heterogeneity into differences in marginal costs, product quality, markups and the number of firm varieties within a representative-agent framework on the demand side. Further in this literature, our theoretical framework builds on existing work on endogenous quality choice across heterogeneous firms (Bastos et al., 2016; Feenstra & Romalis, 2014; Johnson, 2012; Kugler & Verhoogen, 2012), and the link between trade and quality upgrading (Bustos, 2011; Dingel, 2015; Eslava et al., 2016; Verhoogen, 2008).⁷ Relative to the existing work in this area, our paper presents 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 within countries, quantifies the underlying channels, and explores a novel set of implications for real income inequality.

More broadly, our work is related to a growing empirical literature in economics that uses the Nielsen consumption scanner data (Broda & Weinstein, 2010; Handbury, 2014; Handbury & Weinstein, 2014). Most of this literature has been based on the home scanner data. More recently, Argente & Lee (2016) and Jaravel (2016) use the 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 during the Great Recession, and Jaravel (2016) to more innovation and competition in product segments consumed by richer households. In this paper, we use of the combination of the two Nielsen datasets to establish a new set of stylized facts about how Melitz-type firm heterogeneity on the producer side translates into the consumption baskets of households across the US income distribution. We then develop a theoretical framework that can explain these moments in the data, and explore the implications of our findings in general equilibrium counterfactuals.

Finally, our analysis complements existing work on the consumer price index implications of international trade. Porto (2006) combines Argentinian tariff changes under Mercosur with household expenditure shares across seven consumption sectors to simulate household inflation differences. More recently, Fajgelbaum & Khandelwal (2014) propose a quantitative framework using national accounts data on production and consumption across sectors and countries to explore heterogeneous consumer gains from trade. Atkin et al. (2016) use detailed consumption microdata from Mexico to quantify the price index implications from foreign supermarket entry. Given our

⁷Within this literature recent work by Eslava et al. (2016) is closely related as they carefully study the link between trade-induced quality upgrading and (nominal) inequality in Colombia. Our paper complements recent work in this literature by departing from its traditional focus on the role of firm heterogeneity (and quality upgrading) for nominal income inequality. We document that higher and lower-income workers can be affected differently by shocks, such as trade, not only because they source their earnings from heterogeneous firms, but also their consumption baskets.

focus on relative prices within disaggregated product groups, this paper is also close in spirit to [Faber \(2014\)](#) who uses Mexican microdata from consumption surveys, plant surveys and CPI store price surveys to estimate the effect of tariff reductions on the price of product quality in Mexican stores. More recently, [Amiti et al. \(2016\)](#) use a combination of customs and firm microdata to investigate the consequences of China’s WTO accession for US consumer price inflation, and [Cravino & Levchenko \(2016\)](#) use Mexican CPI and expenditure microdata to quantify the implications of the Peso Crisis. Relative to existing work, this paper is the first to use newly available matched home and store scanner data to trace the firm size distribution into the consumption baskets of individual households, and to propose a tractable quantitative model of two-sided heterogeneity that allows us to explore general equilibrium counterfactuals, including changes in trade costs.

The remainder of the paper proceeds as follows. Section 2 describes the data. Section 3 documents a set of stylized facts about firm heterogeneity in consumption baskets across the income distribution. Section 4 presents the theoretical framework. Section 5 presents the parameter estimation and calibration. Section 6 presents the counterfactual analysis. Section 7 concludes.

2 Data

2.1 Retail Scanner Data

We use the Retail Scanner Database collected by AC Nielsen and made available through the Kilts Center at The University of Chicago Booth School of Business. The 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 individual stores. The stores in the database vary widely in terms of formats and types: e.g. food, drug, mass merchandising, 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 Table 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.⁸ 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 size (in terms of national sales measured across on average 27,000 stores per half year) into the consumption baskets of individual households.

⁸We do not make use of Nielsen’s “Magnet” database that covers non-barcode products, such as fresh produce.

2.2 Home Scanner Data

We use the Home Scanner Database collected by AC Nielsen and also made available through the Kilts Center. AC Nielsen collects these data using 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 1 we aggregate the raw microdata to the household-by-barcode-by-half-year level. On average each six-month period covers \$109 million worth of retail sales across 59,000 individual households in more than 1000 disaggregated 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. More importantly, 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.⁹ To address potential concerns about decreasing budget shares of retail relative to other consumption with respect to nominal 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 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. We begin in Figure 1 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 differ substantially within disaggregated product groups. In this and the

⁹Per 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).

subsequent figures and tables, we define a firm as a producer of a brand within one of more than 1000 disaggregated product modules in the Nielsen data. This leads to an average of about 150 active firms within a given product module. Two possible alternatives given our data 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).

We choose the definition of firms as brands within product modules for two main conceptual 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. 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 product quality investments across firms, and it is at the level of brands within product groups that these decisions appear to be most plausible.

Figure 1 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 thousand households per half year) appears to be compressed relative to that measured using the store scanner data (27 thousand supermarkets per half year). This compression is stronger before applying the Nielsen household weights, but still clearly visible after applying the weights. There are several possible explanations. First, it could be the case that the home scanner data fail to capture a long tail of small brands that are part of total store sales but unlikely to be reported by one of the Nielsen sample households in a given half-year period. On the other hand, it could be the case that the store scanner data fail to capture a large mass of brands with predominantly average market shares due to non-participating retail chains. Third, it could also be the case that the home scanner data are subject to under-reporting by households, and that this leads to a mis-representation of the true dominance of the most popular brands: for example a household buying Coca Cola three times a week may only report the first purchase.

To further investigate which of these scenarios seem more likely, the right panel of Figure 1 plots the market share distributions for the two datasets restricting attention to brands observed in both of them. The fact that the same pattern holds in the overlapping product space suggests that the first two explanations are unlikely to account for the compression of the firm size distribution in the home scanner data relative to the store scanner data, and that problems related to household under-reporting in the home scanner could be a factor. For this reason, and the fact that the store scanner data capture more than 1000 times the amount of transactions compared to the home scanner data, we will report in the following the main new stylized fact using the firm size distributions computed from both datasets, and then choose the store scanner data as our preferred measure of brand-level national market shares.

Firm Heterogeneity Across Consumption Baskets Figure 2 depicts the main stylized fact of the paper. Pooling repeated cross-sections across 18 six-month periods, we depict percentiles of household per capita expenditure (within each half year) on the x-axis and weighted average deviations of log firm sales from the product module-by-half-year means on the y-axis.¹⁰ The underlying 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 2, 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 correspond to 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 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 2 is also robust to alternative definitions of firms in the Nielsen data: appendix Figure A.2 shows close to identical results when defining firms instead as holding companies within product groups, which is the level of aggregation that for example Hottman et al. (2016) have followed. We also replicate the relationship in Figure 2 after computing firm sizes in terms of quantities (units sold) rather than revenues in appendix Figure A.3.¹¹

What types of shopping decisions are driving these pronounced differences in weighted average firm sizes across the income distribution? In appendix Table A.1, 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.¹² 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 in any given product module whereas the poorest quintile of households have a 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).

¹⁰See the description of the home scanner data in Section 2 for discussion of percentiles in terms of per-capita expenditure and the relationship to nominal income bins in the data (also in Figure A.1). To avoid measurement error from exiting or entering households in the consumer panel, we restrict attention to households for each six-month period that we observe to make purchases in both the first and the final month of the half year.

¹¹Plotting firm sizes in terms of quantities is sometimes suggested as a check against variable markups. Three notes are in order. First, in a setting with quality differentiation, the quantity of units can be a misleading measure of output. Second, variable markups would likely push the other way: if larger firms have higher markups, and the elasticity of substitution is larger than 1 within product groups (as we find it to be in Section 5), this would imply a decrease in relative sales, ceteris paribus. Third, rather than assuming away variable markups, our theoretical framework allows for them (through both traditional and new channels), and in the estimation we quantify the heterogeneity of markups across firms and consumption baskets that underlies the observed stylized facts in this section (e.g. appendix Figure A.13 documents very small deviations in markups under both monopolistic and oligopoly competition).

¹²The 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).

We now investigate whether these observed differences in product choices are driven by a fundamental disagreement about relative product quality 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 quality-for-money across producers, such that the rank order of their budget shares is preserved across the income distribution? Appendix Figure A.5 documents that the latter appears to be the case in the data. Households seem to strongly agree on their evaluation of product quality attributes given prices as indicated by the fact 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 in Figure A.5 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.

Finally, appendix Figure A.4 explores the heterogeneity across different product groups. We estimate the relationship in Figure 2 separately 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.¹³ As depicted in appendix Figure A.4, we find that the pattern of firm size differences across consumption baskets holds across these very different product segments and is not driven by one particular type of consumer products. We also find that the stylized fact in Figure 2 holds in each of the 18 six-month periods in our data.

3.1 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. In the following, we distinguish between three different types of alternative explanations.

Data-Driven Explanations One concern is that the relationship documented in Figure 2 could in part be driven by shortcomings of the data. First, 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. Second, it could be the case that we are missing systematically different shares of 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 (but are present in the home scanner data). To address these two concerns, appendix Figure A.6 re-estimates the

¹³We 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 data sparsity in the parameter estimation.

relationship of Figure 2 after i) restricting consumption to sum to 100% for all non-generic product consumption for each household, and ii) after only including households for which we observe more than 90 percent of their total retail expenditure in both data sets. We find very similar results in these alternative specifications suggesting that shortcomings of the data are unlikely to account for the stylized fact documented in Figure 2.¹⁴

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 mis-measuring total firm sales relative to domestic-only US producers. Given existing evidence on the selection of firms into trade, as well as through the lens of the model that we use to rationalize the observed moments in the microdata in the following section, it is likely that the resulting measurement error in firm sizes is positively related to the observed US market shares in the Nielsen data. That would imply that the domestic-only data on sales somewhat understates true differences in firm sizes across consumption baskets. In our analysis, we address this data limitation in several ways. We report differences in weighted-average firm sizes across the income distribution separately for both low vs high import penetration product groups, and low vs high export share product groups.¹⁵ For both of these cuts of the data, the “low” category is defined as below median, which is equivalent to less than 10 percent import penetration or export shares. As depicted in appendix Figures A.7 and A.8, 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. Furthermore, we also address this data limitation in several robustness checks as part of our counterfactual analysis in Section 6.

Segmented Markets 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 heterogenous household preferences, could be driving the results. In appendix Figure A.9 we explore to what extent differences in household geographical location as well as differences in retail formats within locations play in accounting for Figure 2. We first re-estimate the same relationship after conditioning on county-by-half-year fixed effects when plotting the firm size deviations on the y-axis (keeping the x-axis exactly as before).¹⁶ 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 2, suggesting that differential access to producers is unlikely to be the driver.

¹⁴Further reassurance against the “missing retailers” concern is also apparent in Figure 2 that depicts very similar patterns when using 100 percent of household retail consumption as reported in the home scanner data.

¹⁵To this end, we match the Nielsen product groups to 4-digit SIC codes in 2005 US trade data. See appendix Table A.2. We measure import penetration as the share of imports in total production, and the export share as the share of exports in total production.

¹⁶De-meaning, instead, both the y and x-axis leads to almost identical point estimates.

¹⁷We condition on 79 store formats within the same county to capture potential differences in access across inner-city vs. suburbs or for example due to car ownership. Note that 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.

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 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

We document that in the medium or long run this notion seems hard to reconcile with either the raw moments in 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. \(2016\)](#), [Dingel \(2015\)](#)).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.3](#) 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 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 [3](#) and [4](#)). Furthermore, as we discuss formally in the theory and the estimation below, the fact that we use panel-data estimates to inform the model calibration for the counterfactuals –rather than relying on the cross-sectional correlations that we present here as part of the stylized facts– provides further reassurance against the concern that fixed product attributes confound the findings of the counterfactuals.

To summarize, we document large and statistically significant differences in the weighted average producer sizes that rich and poor households source their consumption from. This finding holds across product departments and all 18 six-month periods covered by the scanner data, and does not appear to be driven by shortcomings of the data such as retailer generics or non-participating retail chains, household differences in producer access across locations or store formats, or fixed product attributes that producers are born with. The finding also arises in a setting where households on average strongly agree on their ranking of value-for-money across producers: the largest firms command the highest expenditure shares within product modules across all income groups. The following section proposes a tractable theoretical framework that captures these observed moments

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 [2](#) 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 [2](#) after replacing contemporary differences in firm sales by either the firm sales of the very same brands three years before or 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 [2](#). Instead, appendix Figure [A.10](#) suggests that the estimated differences in producer sizes are practically identical.

in the microdata, and guides the empirical estimation.

4 Theoretical Framework

This section develops a tractable quantitative model that rationalizes the observed moments in the microdata and allows us to quantify the underlying channels and explore the implications for real income inequality. To this end, 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 elasticity and in their product quality evaluations. On the producer side, differently productive firms face the observed distribution of consumer preferences and optimally choose their product attributes and markups. We also extend this framework in multiple directions to be able to investigate the sensitivity of the estimation and counterfactuals in the following sections. In addition to the exposition here, Appendices 2-7 provide further 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 \(2014\)](#), 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 :

$$U = U(U_G(z), z) \tag{1}$$

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 and normalize to one the population of consumers. By allowing demand parameters for retail consumption to be a function of the outside good consumption, we introduce non-homotheticity in a reduced-form approach without imposing structure on the sign or size of the non-homotheticities.²² Utility from retail consumption is defined by:

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

where n refers to a product module in the Nielsen data and i refers to a specific brand producer within the product module.²³ The term $\varphi_{ni}(z)$ refers to the perceived quality of brand i in product module n at income level z . The term $\sigma_n(z)$ refers to the elasticity of substitution between brand

²¹E.g. [Handbury \(2014\)](#) estimates the income elasticity of retail consumption to be positive but lower than one, implying that the outside good is normal (by Engel aggregation).

²²This approach is similar to [Handbury \(2014\)](#) and [Redding & Weinstein \(2016\)](#) and follows earlier work by [McFadden & Train \(2000\)](#). In our empirical application, we work with five broad groups of consumers indexed by z that correspond to quintiles in the US income distribution. Appendix 3 provides a proof that, holding the initial consumer type z constant, changes in the retail price index provide a first-order approximation to the consumers' compensating variation in retail consumption under any arbitrary upper-tier utility function in (1). The appendix also derives the first and second-order conditions and the income elasticities for a specific upper-tier that we use to analyze counterfactuals in which the distribution of z changes endogenously. See also Footnote 32 below.

²³We show in Appendix 4 that these preferences can be derived from the aggregation of discrete-choice preferences across many agents choosing only one brand variety by product module.

varieties within each product module n at income level z . As we focus most of our attention on the 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

These preferences are common across all households but non-homothetic since utility from retail items depends on income level z (outside good consumption). An advantage of this specification of preferences is that we do not impose structure that dictates how price elasticities and quality valuations depend on income.²⁵ While building on the tractability of the standard CES structure, the preferences in (2) depart from the canonical setup by allowing for a new source for variable markups across firms: as we derive below, producers face heterogeneous market demand curves as a function of the composition of the consumers (z) they sell to. In Appendix 5, we further extend the model to feature, in addition, a traditional source for variable markups (oligopoly competition). In Section 6, this allows us to explore counterfactuals subject to alternative assumptions about market structure and the nature of variable markups across firms.

Comparing two goods i and j within the same module n , relative expenditures by consumers of income level z are then given by:

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

Equation 3 implies that we can use observable moments on income group-specific product sales in combination with unit values and demand parameters in order to estimate unobserved differences in product quality. Previous papers focusing on the supply side of quality choice assume that quality evaluations are constant across income groups (e.g. [Hottman et al. \(2016\)](#); [Kugler & Verhoogen \(2012\)](#); [Sutton \(1998\)](#)), while existing papers on heterogeneous quality choice by consumers generally assume that quality valuations depend on an intrinsic quality characteristic multiplied by income or log income ([Fajgelbaum et al., 2011](#); [Handbury, 2014](#)). The latter imposes the assumption that quality rankings across goods are preserved across income groups. Motivated by the evidence discussed in Figure A.5 above, 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 :

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

With the normalization $\int_{\Omega_z} \gamma_n(z) dz = 1$ (where Ω_z is a set of z household types)²⁶, this intrinsic

²⁴Note that we abstract from within-brand product substitution by summing up sales across potentially multiple barcodes within a given product brand by product module. Appendix 5 presents an extension of our model to multi-product firms which we also discuss below.

²⁵For instance, demand systems with a choke price can generate price elasticities that depend on income ([Arkolakis et al., 2012](#)), but offer significantly less flexibility in that relationship.

²⁶This normalization sets the simple mean of preference parameters $\gamma_n(z)$ equal to unity across a fixed set of household types z . In our empirical application, we work with 5 groups of households that are based on quintiles of the distribution of total retail expenditure per capita. As shown in Appendix 6, this normalization to unity across household types z is without loss of generality: counterfactual results would be identical for other values. Appendix 6 further discusses the choice of the fixed set Ω_z and the implications for estimation.

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:

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

This implies that changes in product prices, quality and availability across firms can have different implications for the cost of living of households across the income distribution.

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 (and brands) 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 subject to monopolistic competition.

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 increase in the quality of the good being produced. 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 quality. For tractability, we adopt a simple log-linear parameterization for incremental fixed costs:

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

Fixed costs increase with quality if $\beta_n > 0$.²⁷ Similarly, variable costs depend log-linearly on

²⁷An alternative setting would be that $\beta_n < 0$ and $b_n < 0$ (i.e. that fixed costs decrease in quality) as well as $\xi_n > \gamma_n(z)$. In that case, quality inversely relates to firm size and productivity. Our estimation indicates that $\beta_n > 0$ is the empirically relevant case, on which we focus on in our theoretical exposition.

quality, with parameter ξ_n to capture the elasticity of the cost increase to the level of 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

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. Markups are determined by the average price elasticity across income groups, and prices are given by:

$$p_n(a) = \frac{\phi(a)^{\xi_n}}{a\tilde{\rho}_n(a)} \quad (10)$$

where $\tilde{\rho}_n = \frac{\tilde{\sigma}_n(a)-1}{\tilde{\sigma}_n(a)}$ and $\tilde{\sigma}_n(a)$ is the weighted average price elasticity across consumers:

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

$x_n(z, a)$ denotes sales of firm with productivity a to consumers of income level z , which itself depends on the optimal quality of the firm. In turn, the first-order condition in ϕ characterizes optimal quality $\phi_n(a)$ for firms associated with productivity a :

$$\phi_n(a) = \left(\frac{1}{b_n} \cdot \tilde{\rho}_n(a) \cdot X_n(a) \cdot (\tilde{\gamma}_n(a) - \xi_n) \right)^{\beta_n} \quad (11)$$

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

$$\tilde{\gamma}_n(a) = \frac{\int_z \gamma_n(z) (\sigma_n(z) - 1) x_n(z, a) dH(z)}{\int_z (\sigma_n(z) - 1) x_n(z, a) dH(z)} \quad (12)$$

Optimal quality is determined by several forces that are apparent in equation 11. First, larger sales induce higher optimal quality, as reflected in the term $X_n(a)^{\beta_n}$. This is the scale effect due the 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. However, a lower elasticity of fixed costs to quality, captured by a higher β_n , induces larger scale effects and leads to a higher elasticity of optimal quality to sales and quality valuation.

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

As we prove in Appendix 2, uniqueness of equilibrium places bounds on the extent of heterogeneity in price elasticities across consumers and the size of $\beta_n > 0$ to obtain unique firm choices in quality and prices. Under these conditions, that we verify to hold in our empirical setting in Section 6, we show that market shares, quality and sales to each income group z increase monotonically with firm productivity a . In the special case where a firm sells to consumers from a single income group z , we obtain a simple expression to describe how quality varies with productivity that is similar to the representative agent case:

$$\frac{\partial \log \phi_n(a)}{\partial \log a} = \frac{\beta_n (\sigma_n(z) - 1)}{1 - \beta_n (\sigma_n(z) - 1)(\gamma_n(z) - \xi_n)} > 0 \quad (13)$$

The term $\beta_n(\sigma_n(z) - 1)(\gamma_n(z) - \xi_n)$ in the denominator corresponds to the share of revenues (net of variable costs) that are invested in quality-upgrading fixed costs $f_n(\phi)$.

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

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

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 2:

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

How $\widetilde{X}_n(z)$ varies with income (i.e. the slope of the estimated relationship in Figure 2) 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 is positively correlated 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 underlying the observed heterogeneity in firm sizes across consumption baskets along the income distribution.

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 11. In particular, the correlation between firm size and quality appearing in the first term can be expressed as:

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

In our empirical section that follows, we can use our parameter estimates in combination with moments from the microdata to quantify each of these terms and decompose the observed firm heterogeneity across the consumption baskets of rich and poor households depicted in Figure 2 into the underlying channels.

4.3 Model Extensions

Oligopoly Appendix 5 extends the model in three directions. First, as mentioned above, we deviate from the assumption of a continuum of firms, and allow for oligopolistic competition between firms as in [Hottman et al. \(2016\)](#).²⁹ In addition to the source of variable markups across firms that our model features (due to firm-specific market demand curves), this also allows for a more traditional source of heterogeneous markups across the firm size distribution. When firms internalize the effect of their decisions in prices, quantity and quality on the price index, their optimal markup depends on their market share and exceeds their monopolistic-competition markup:

$$\frac{p_n(a) - c_n(a)}{p_n(a)} = \frac{\int_z x_n(z, a) dH(z)}{\int_z [\sigma_n(z)(1 - s_n(z, a)) + s_n(z, a)] x_n(z, a) dH(z)} > \frac{1}{\tilde{\sigma}_n(a)} \quad (17)$$

where $s_{ni}(z, a)$ denotes the market share of brand a among consumers of consumer group z and firms of product module n . As described in this formula, markups are determined by the sales-weighted average of the perceived price elasticity of demand, which is now strictly smaller than the elasticity of substitution, especially when a firm has a large market share. In the appendix, we derive an alternative set of counterfactual equations (see below), and use this extension to investigate counterfactual impacts under alternative assumptions about market structure and the nature of variable markups across firms.

Multi-Product Firms Appendix 5 also presents an extension of our model to multi-product firms. As recently emphasized by [Hottman et al. \(2016\)](#), if barcode products within the same brand are not perfect substitutes then multi-product firms introduce an additional dimension of firm heterogeneity since different brands can offer different within-brand variety. In the appendix we show that, as long as the ratio of cross-brand to within-brand elasticities of substitution does

²⁹We also consider a generalized version of Cournot quantity competition following [Atkeson and Burstein \(2008\)](#), whereby firms compete in quantity, instead of both quantity and prices as in [Hottman et al. \(2016\)](#) (see Appendix 5). This second alternative framework yields very similar results, both for the cross-section of implied markups as well as for counterfactual simulation results.

not significantly differ across income groups, this additional dimension (the number of brands) does not affect firm heterogeneity across consumption baskets. In other words, even if rich and poor households significantly differ in their within-brand elasticities of substitution (i.e. different degrees of love of variety), this would not drive differences in budget shares across brands with more or less barcode products as long as the ratio of within-brand elasticities between rich and poor households is similar to their ratio of cross-brand elasticities of substitution.³⁰

Other Sources of Quality Differences Finally, we can flexibly allow for other determinants of output quality that differ across firms by adding a firm-specific quality shifter ψ that producers learn about after paying the fixed cost of entry alongside their productivity parameter a . Here we model ψ as a cost shifter, but it is isomorphic to a taste shifter (which is closer to our discussion in the final part of the previous Section 3).

Formally, we assume that firms discover (a, ψ) drawn from a joint distribution $G_n(a, \psi)$. While parameter a influences the variable cost, the fixed cost of quality upgrading in (8) above now depends on ψ and is given by:

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

Without loss of generality we can normalize $b_n = 1$. The uniqueness of equilibrium in prices and quality still holds in this case. With this extension, optimal quality is given by:

$$\phi(a, \psi) = \underbrace{\psi}_{\text{idiosyncratic}} \cdot \underbrace{(\tilde{\rho}_n(a, \psi) X_n(a, \psi) (\tilde{\rho}_n(a, \psi) - \xi_n))^{\beta_n}}_{\text{scale-dependent}} \quad (18)$$

Note that firm size X_n and other outcomes now also depend on this idiosyncratic quality shifter ψ . However, by allowing this parameter to differ freely across firms, we can capture other sources of quality differentiation that may not be directly related to firm sizes through the economies-of-scale channel that we focus on above. For example, a firm may have a low productivity draw a but a high quality shifter ψ which leads to a high quality relative to other firms of the same size.

While total firm sales X_n depend on trade and the distribution of income, and may change in our counterfactuals as we discuss below, the quality shifter ψ is held constant for a given firm. Therefore, as we show in Appendix 5 and 7, heterogeneity in this parameter does not affect the solution for counterfactual outcomes we derive below. This statement holds for a given set of preference and technology parameters. But as the above model extension shows, it could still be the case that our technology parameter estimation confounds the economies-of-scale channel (relating firm sizes to quality) with variation in other sources of quality differentiation, ψ , across firms in the cross-section. Importantly, as we discuss further in Section 5, our panel-data estimation strategy for the firm technology parameters is robust to allowing for other firm-specific determinants of product quality that may operate in addition to the economies-of-scale channel that we quantify.

³⁰Related to this result, appendix Table A.4 reports evidence suggesting that the ratio of within and cross-brand elasticities do not seem to significantly differ in the data. Moreover, we observe similar firm size differences across the income distribution, whether we define firms as brands or holding companies (Figure A.2).

4.4 Counterfactuals

Our framework naturally lends itself to quantitative estimation. In Appendix 7, we derive five equilibrium conditions that govern counterfactual changes in firm sales, quality, entry, exit and price indices. Thanks to the tractability of our framework, we can solve for counterfactual equilibria using data on initial sales $x_{n0}(z, a)$ for each firm across different consumer groups in addition to estimates of five sets of parameters: $\sigma_n(z)$, $\gamma_n(z)$, β_n , ξ_n and f_{n0} . With these moments in hand, we can directly solve for changes in quality $\frac{\phi_{n1}(a)}{\phi_{n0}(a)}$, sales $\frac{x_{n1}(z, a)}{x_{n0}(z, a)}$, the mass of firms $\frac{N_{n1}}{N_{n0}}$, firm survival $\delta_{nD}(a)$ and consumer price indices $\frac{P_{n1}(z)}{P_{n0}(z)}$ (as well as firm export decisions as discussed below). Equilibrium changes in quality can be derived by taking ratios of equation 11, changes in sales are derived from equations 3 and 10, changes in profits from equation 14, and changes in cost of living from equation 6.³¹ As described in Appendix 7, we do not require estimates of firm productivity a or initial firm quality $\phi(a)$ to conduct our counterfactual exercise. This approach follows Dekle et al. (2007) among others.

We use this framework to explore different types of counterfactuals. The first set of counterfactuals is to exogenously increase nominal income inequality across consumers, where $H_0(z)$ and $H_1(z)$ denote the initial and counterfactual cumulative distribution of z . These counterfactuals illustrate how changes in the income distribution affect the demand and supply of product quality, and how these changes feed back into consumer inflation and real income inequality. The second type of counterfactuals explores the distribution of the gains from trade in a setting where households source their consumption from different parts of the firm size distribution, as observed in the microdata. Here, we focus on a conventional Melitz (2003) framework with two symmetric countries where firms can export to an additional market by paying a fixed cost $f_{nX} > 0$ and variable iceberg trade costs $\tau_n > 1$. In absence of the asymmetric price index effects that we allow for, this conventional setup would feature not distributional implications of falling trade costs. In all counterfactuals, we quantify the differential effect on cost of living inflation comparing the top 20 percent of US households to those of the bottom 20 percent holding the initial income quintile z fixed over time.³²

Given the various sources of heterogeneity across consumers and firms, these price index effects are driven by a rich and novel interplay of adjustment channels. To guide the analysis, we derive a five-term decomposition of the effect on price indices for income group z relative to income group

³¹Combining the equations for sales and price changes in Appendix 7 yields an expression of how sales growth $\frac{x_{n1}(z, a)}{x_{n0}(z, a)}$ depends on quality upgrading $\frac{\phi_{n1}(a)}{\phi_{n0}(a)}$, while the equation for quality changes expresses how quality upgrading depends on sales growth. Conditional on entry and exit, these two relationships offer a contraction mapping that we exploit to solve the counterfactual, provided that the share of revenues invested in quality upgrading $(\beta_n(\sigma_n(z) - 1)(\gamma_n(z) - \xi_n))$ remains less than 1 for all z .

³²The choice of five consumer groups is driven by the empirical setting that the calibration of the next section will be based on. Our approach is similar to e.g. Atkin et al. (2016) and Hottman et al. (2016). While convenient for empirical tractability, the ad-hoc treatment of non-homotheticity (keeping initial z fixed) shuts down a second-order price index effect: large first-order effects of the shocks on real incomes may push some households across z group boundaries and thereby change their preference parameters as defined above. Since our empirical application allows preferences to differ across five broad income groups, it is reasonable to think that few households are shifted in this manner. Having said this, we also report results using an alternative approach in Appendix 3, in which we let z be an endogenous outcome to the counterfactual shock. Reassuringly, as we discuss in the results Section 6 below, the findings remain virtually unaffected.

z_0 for each product module n :

$$\begin{aligned}
\log \frac{P_{n1}(z)}{P_{n0}(z)} - \log \frac{P_{n1}(z_0)}{P_{n0}(z_0)} &= \underbrace{- (\bar{\gamma}_n(z) - \bar{\gamma}_n(z_0)) \int_a \bar{s}_{n1}(a) \log \left(\frac{\phi_{n1}(a)}{\phi_{n0}(a)} \right) dG(a)}_{(1) \text{ Average quality effect}} \tag{19} \\
&- \underbrace{(\bar{\gamma}_n - \bar{\xi}_n) \int_a (s_{n1}(a, z) - s_{n1}(a, z_0)) \log \left(\frac{\phi_{n1}(a)}{\phi_{n0}(a)} \right) dG(a)}_{(2) \text{ Asymmetric quality-adjusted cost changes}} \\
&- \underbrace{\int_a (s_{n1}(a, z) - s_{n1}(a, z_0)) \log \left(\frac{\tilde{\rho}_{n1}(a)}{\tilde{\rho}_{n0}(a)} \right) dG(a)}_{(3) \text{ Asymmetric markup changes}} \\
&- \underbrace{\left(\frac{1}{\sigma_n(z) - 1} - \frac{1}{\sigma_n(z_0) - 1} \right) \log \left(\frac{N_{n1} \bar{\delta}_{nD} (1 + \bar{\delta}_{nX} \tau_n^{1 - \bar{\sigma}_n})}{N_{n0}} \right)}_{(4) \text{ Love of variety}} \\
&- \underbrace{\frac{1}{\bar{\sigma}_n - 1} \log \left(\frac{\int_a s_{n0}(a, z) \delta_{nD}(a) (1 + \delta_{nX}(a) \tau_n^{1 - \sigma_n(z)}) dG_n(a)}{\int_a s_{n0}(a, z_0) \delta_{nD}(a) (1 + \delta_{nX}(a) \tau_n^{1 - \sigma_n(z_0)}) dG_n(a)} \right)}_{(5) \text{ Asymmetric import and exit effects}}
\end{aligned}$$

where $s_{n0}(a, z)$ denotes the initial market share of brand a among consumers of income z , and where $s_{n1}(a, z) = \frac{s_{n0}(a, z) \delta_{nD}(a) (1 + \delta_{nX}(a) \tau_n^{1 - \sigma_n(z)})}{\int_a s_{n0}(a, z) \delta_{nD}(a) (1 + \delta_{nX}(a) \tau_n^{1 - \sigma_n(z)})}$ in the first three terms adjusts for trade and survival (but not quality upgrading). $\bar{s}_n(a)$ refers to the average of $s_n(a, z)$, $\frac{1}{\bar{\sigma}_n - 1}$ refers to the average of $\frac{1}{\sigma_n(z) - 1}$, and $\bar{\gamma}_n$ to the average of $\gamma_n(z)$ across the two income groups. $\delta_{nX}(a)$ is a dummy indicator that denotes firms' decisions to export or not, and $\bar{\delta}_{nD} = \int_a \delta_{nD}(a) \bar{s}_{n0}(a) dG(a)$ denotes average survival rates across all firms and the two income groups.

In both types of counterfactuals, the effect in the first line of the decomposition is that firms on average have incentives to upgrade 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).³³ In the first set of counterfactuals (increases in nominal inequality), firms upgrade their quality as a larger share of their consumers are households with higher quality evaluations. In the trade counterfactual, the largest firms experience positive scale effects from trade opening, which also induces an increase in weighted average product quality in both markets.

The second effect is that the scale of production changes asymmetrically across higher and lower quality producers in both types of 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 favors richer households if they spend relatively more on firms with the largest increase 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 upgrade

³³For the sake of exposition, we approximate the first and second terms (1) and (2) by taking the average of the log instead of the log of the average. By Jensen's inequality, this leads to an underestimation of these two effects. In practice, we verify that the bias is very small.

product quality $\log\left(\frac{\phi_{n1}(a)}{\phi_{n0}(a)}\right)$.

The third effect captures the change in markups, which in our framework differ endogenously across firms as a function of 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 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.

The fourth channel 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 $\sigma_n(z)$. In the trade counterfactual, this effect combines the number of varieties that are available on the domestic market as well as new imported varieties.

In addition to differences in the love of variety, the fifth channel reflects the unequal effects of exit (in both counterfactuals) as well as access to new imported varieties (trade counterfactual) as a function of differences in consumption shares for both exiting and entering varieties. In both counterfactuals, 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. In the trade counterfactual, it is additionally the case that the market share of imported goods can differ significantly across households. Since richer households tend to buy from larger firms and since larger firms are more likely to trade in both countries, the effect of trade opening on new imported varieties tends to favor relatively richer households.

In addition to the above, we also explore a third type of counterfactuals. Following the notion that taxes and red tape increase in producer size (e.g. [Hsieh & Klenow \(2009\)](#)), we quantify the incidence of a non-uniform business tax that affects larger producers disproportionately within sectors. We derive additional equilibrium equations for this third application in Appendix 7.

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

$$\begin{aligned} \log \frac{P_{G1}(z)}{P_{G0}(z)} - \log \frac{P_{G1}(z_0)}{P_{G0}(z_0)} &= \underbrace{\sum_n \left(\frac{\alpha_n(z) + \alpha_n(z_0)}{2} \right) \left(\log \frac{P_{n1}(z)}{P_{n0}(z)} - \log \frac{P_{n1}(z_0)}{P_{n0}(z_0)} \right)}_{(1+2+3+4+5) \text{ Within-module changes}} \quad (20) \\ &+ \underbrace{\sum_n (\alpha_n(z) - \alpha_n(z_0)) \left(\frac{\log \frac{P_{n1}(z)}{P_{n0}(z)} + \log \frac{P_{n1}(z_0)}{P_{n0}(z_0)}}{2} \right)}_{(6) \text{ Between-module changes}} \end{aligned}$$

The within term can be decomposed into the five terms described in equation 19. 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 negative if prices tend to decrease 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, so that the between-module term need not be zero. For completeness, we report all six terms in the quantification of counterfactuals in Section 6.

5 Estimation

This section presents the empirical estimation. We begin by estimating the preference parameters, σ_{nz} and γ_{nz} , that combined with the microdata 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 parameter estimates, in combination with some raw moments from the scanner data, allow us to quantify the channels underlying the documented stylized facts at the end of this section, and to explore model-based counterfactuals in the final section of the paper.

5.1 Preference Parameter Estimation

Estimation Strategy We begin by estimating the elasticity of substitution σ_{nz} that we allow to vary across household income groups and product groups. From equation 3 we get the following estimation equation:

$$\Delta \log (s_{nzict}) = (1 - \sigma_{nz}) \Delta \log (p_{nict}) + \eta_{nzct} + \epsilon_{nzict} \quad (21)$$

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 CES price index term. Consistent with our CES preference specification at the level of household groups, we estimate expression 21 after aggregating consumption shares in the home scanner microdata 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} for the same county over time or within the county across household groups and modules, we cluster standard

³⁴To be consistent with our CES specification, we aggregate household purchases to the income group level as projection-factor-weighted sums to compute $\Delta \log (s_{nzict})$, and limit the sample to income group-by-county-by-half-year cells with at least 25 households per cell. 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.5, neither the decision to take mean prices (rather than medians), nor the decision to take a Tornqvist price index (rather than Laspeyres or a simple average) affects the point estimates.

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. 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: $\frac{1}{N-1} \sum_{j \neq c} \Delta \log(p_{nijt})$. As recently shown by Beraja et al. (2014), these two instruments are likely to identify potentially different local average treatment effects. The national leave-out means IV estimates the elasticity of substitution off retail chains that price their products nationally, whereas the state-level leave-out means additionally extend the complier group of the IV 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 bias in the σ_{nz} estimates, it would have to be the case that the advertisement campaign 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. 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 address potentially remaining concerns, we are also careful not to bind our counterfactual analysis in Section 6 to one particular set of point estimates. Instead, we report our findings both for our preferred baseline parameter values for σ_{nz} , as well as across alternative parameter combinations to document the sensitivity of the findings. Finally, we note that 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.

Estimation Results Panel A of Table 2 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

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

very similar and suggest an aggregated elasticity of substitution of about 2.2. These estimates are very close to existing work using barcode-level consumption data and the Hausman-type IV approach (e.g. [Hausman & Leibtag \(2007\)](#), [Handbury \(2014\)](#)). They are, however, somewhat lower than empirical work that has used the [Feenstra \(1994\)](#) approach for estimating σ_{nz} (e.g. [Broda & Weinstein \(2010\)](#), [Hottman et al. \(2016\)](#)). As a robustness exercise, we report our findings in the final section of this paper both for our baseline parameter values for σ_{nz} as well as for higher values of these parameters to document the sensitivity of the counterfactuals.

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—conditioning on differences in product mix. We choose the richest income group as our reference category that will be absorbed by the additional fixed effects. Interestingly, poorer households appear to have statistically significantly higher elasticities of substitution compared to wealthier households. In terms of magnitude, these differences are relatively minor, however. We estimate that the elasticity of substitution for the poorest two income quintiles is about 0.4 larger than that for the richest income quintile.

Panel B of [Table 2](#) then breaks up the estimates by the 8 product departments that are covered by the Nielsen data, and Panel C reports the results within 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.6](#) 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 2](#). As mentioned above, as a robustness exercise we present the quantitative analysis based on our baseline parameter estimates as well as across a range of alternative parameter combinations in order to document the sensitivity of our findings.

5.2 Estimation of Brand Quality, Quality-Adjusted Prices and Markups

Armed with estimates of σ_{nz} , [equation 3](#) allows us to use the scanner microdata to estimate product quality, $\log \phi_{ni} = \frac{1}{N_z} \sum_z \log \varphi_{nzi}$, and quality-adjusted prices, $\log \left(\frac{p_{ni}}{\phi_{ni}} \right)$, across producers of brands as well as household consumption baskets. In addition, we estimate the perceived product quality of each producer of brands across income groups z , $\log \varphi_{nzi}$. To do this, as shown in [3](#), we use an additional empirical moment from the data, product unit values, in combination with observed product sales and the estimated σ_{nz} parameters to estimate unobserved variation in product quality.

Appendix Figure A.11 depicts the distribution of mean deviations in log product unit values within product module-by-half-year cells (aggregated as consumption weighted averages across household consumption baskets) along the income distribution.³⁶ 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 income quintile.

The left panel of Figure 3 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 statistically and economically 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.12 confirms what we already noted in the stylized facts section from Figure A.5: these findings emerge in a setting where households appear to strongly agree in terms of the quality ranking of producers in their consumption baskets, but richer income households value higher quality attributes even more than poorer households. Moving from differences in product quality to quality-adjusted prices, the right panel of Figure 3 documents that the richest income quintile source their consumption at on average 10 percent lower quality-adjusted prices.

The parameter estimates for σ_{nz} in combination with the microdata on firm sales across household income groups also allow us to compute the distribution of the effective (weighted average) elasticities of substitution faced by individual producers, $\left(\tilde{\sigma}_{ni} = \frac{\sum_z \sigma_{nz} x_{nzi}}{\sum_z x_{nzi}}\right)$, across the firm size distribution that informs the distribution of firm markups. The left panel of Figure 4 presents the estimation results of $\tilde{\sigma}_{ni}$ across 18 pooled cross-sections (for fourteen half years between 2006-2014) of within-product module firm size distributions. As implied by the stylized fact in Figure 2, and the estimation results in Table 2, we find that larger firms face significantly lower price elasticities because they sell a higher share of their output to higher-income households who, in turn, have lower parameter values for σ_{nz} .

Having estimated product quality and the distribution of firm-level weighted average demand elasticities, we now proceed to estimate the final set of preference parameters, γ_{nz} , that govern the valuation of product quality characteristics across the household income distribution. From our definition of product quality in (4) and (5), we get the following estimation equation:

$$\log(\varphi_{nzit}) = \gamma_{nz} \log(\phi_{nit}) + \eta_{nzt} + \epsilon_{nzit} \quad (22)$$

where η_{nzt} are income group-by-product module-by-half-year fixed effects. To address the concern of correlated measurement errors 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.³⁷

³⁶We compute brand-level unit values as sales-weighted means across barcode transactions at the level of brands-by-half-year cells.

³⁷This differs from the previous regressions in 21 since the estimation equation is at the national level across

Appendix Table A.7 presents the estimation results across bins of household groups and product departments. In accordance with the documented raw moments in the consumption microdata, richer household groups are estimated to attach significantly higher valuations for higher quality products across each of the product departments. However, there also appear to be significant and interesting differences in the extent of this heterogeneity across different product departments. For example, among the departments with the highest difference in the taste for quality between rich and poor households are beverages, dairy products and packaged meat. On the other end, general merchandise and health and beauty care have the lowest differences in household taste for quality across income deciles.

As we do above for the firm-level parameter $\tilde{\sigma}_{ni}$, we can use the microdata on firm sales across income groups in combination with the parameter estimates reported in Table A.7 in order to compute the weighted average product quality evaluations faced by each brand producer: $\tilde{\gamma}_{ni} = \frac{\sum_z \gamma_{nz} (\sigma_{nz}-1) x_{nzi}}{\sum_z (\sigma_{nz}-1) x_{nzi}}$. The right panel in Figure 4 reports these estimation results across the firm size distribution. Following from the raw moments in the consumption microdata reported in Figure 2 and the parameter estimates in Table A.7, we find that larger producers of brands face a market demand schedule with significantly higher marginal valuations for product quality. As was the case for the left panel of that Figure, which plots the distribution of $\tilde{\sigma}_{ni}$, this is due to the fact that a larger share of their sales are driven by higher-income consumers compared to smaller firms. In addition, appendix Figure A.13 plots deviations in weighted-average brand markups embodied in consumption baskets across the income distribution. As described in the theory section and Appendix 5, we do this both under the baseline assumption of monopolistic competition and alternatively computing markups in oligopoly. In both cases, deviations in weighted-average markups across the consumption baskets of rich and poor US consumers are less than 1 percent.

5.3 Technology Parameter Estimation

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

Estimation in the Cross-Section Armed with estimates of the preference parameters $\tilde{\sigma}_{ni}$ 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. A model-consistent and intuitive way to estimate β_n is by estimating the empirical relationship between unit values and market shares within product modules. If we imposed the assumption of homogeneous consumer preferences (representative agent), we would get the following estimation equation from (3) and (11) above:

$$\log(p_{nit}) = \left(\beta_n - \frac{1}{\sigma_n - 1} \right) \log(X_{nit}) + \eta_{nt} + \epsilon_{nit} \quad (23)$$

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)

producers of brands without a county dimension. Clustering at this level yields slightly more conservative standard errors compared to potential alternatives (clustering at the level of brands, income groups-by-semesters or product modules-by-income groups).

and market shares would be governed by the slope of the demand curve $-\frac{1}{\sigma_n-1}$. Accounting for the relationship between unit values and firm scale conditional on quality differentiation, the extent to which firms of larger scale sort into producing higher product quality is then captured by the production function parameter β_n . To see this more clearly, we can re-write (23) 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 heterogeneous tastes for quality and price elasticities across consumers, that give rise to firm-specific taste-for-quality parameters and demand elasticities, $\tilde{\gamma}_{ni}$ and $\tilde{\sigma}_{ni}$ respectively, this estimation equation requires two additional correction terms. From (5) and (11) we get:

$$\log(p_{nit}) = \left(\beta_n - \frac{1}{\bar{\sigma}_n - 1} \right) \log(X_{nit}) - \frac{1}{N_z} \sum_z \frac{1}{\sigma_{nz} - 1} \log\left(\frac{X_{nzit}}{X_{nit}}\right) \quad (24)$$

$$+ \beta_n \log(\tilde{\rho}_{nit} (\tilde{\gamma}_{nit} - \xi_n)) + \eta_{nt} + \epsilon_{nit}$$

where N_z is the number of consumer groups (5 in our application), $\frac{1}{\bar{\sigma}_n-1} = \frac{1}{N_z} \sum_z \frac{1}{\sigma_{nz}-1}$, and $\tilde{\rho}_{nit} = \frac{\tilde{\sigma}_{nit}-1}{\tilde{\sigma}_{nit}}$. 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 (23), to allow for the fact that different producers may face different market demand elasticities due to differences in the composition of their customers. The second additional term captures the fact that regardless of firm scale different producers may sort into higher or lower product quality due to differences in the composition of their customer base (valuing quality more or less given prices).

For estimation, we can again re-write equation (24) as: $\log(\phi_{nit}) = \log(p_{nit}) + \frac{1}{N_z} \sum_z \frac{1}{\sigma_{nz}-1} \log(X_{nzit}) = \beta_n \log(X_{nit} \tilde{\rho}_{nit} (\tilde{\gamma}_{nit} - \xi_n)) + \eta_{nt} + \epsilon_{nit}$, following (3) and (5). Given the thousands of 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 OLS and 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 the and the minimum estimated γ_{nz} , and select the parameter combination that maximizes the goodness of fit in the IV estimation. We do not impose an ex ante assumption about the existence of economies of scale in quality production ($\beta_n > 0$).³⁸

The two main identification concerns in (24) 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 composition-adjusted firm scale $\log(X_{nit} \tilde{\rho}_{nit} (\tilde{\gamma}_{nit} - \xi_n))$ 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.³⁹

³⁸We find $\beta_n > 0$ in all cases and also verify that our estimates are consistent with the model's parameter restriction $\beta_n(\sigma_{nz} - 1)(\gamma_{nz} - \xi_n) < 1$ (ensuring that the share of revenues invested in quality upgrading remains less than 1).

³⁹Note that some of the regressors are themselves estimates from the previous section. We return to this issue

Panel Estimation Estimation equation (24) 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 estimation approach is to use within-brand variation over time. We think of this second approach as more conservative for two reasons. First, quality upgrading/downgrading by firms in response to changes in demand conditions (scale and consumer composition) are likely best understood as a longer-term effect (both in terms of changing actual quality attributes as well as making investments into brand perceptions through advertisement).

Second, to the extent that there are other firm-specific shocks to either perceived quality by consumers or the firm’s cost of producing quality, then the cross-sectional approach risks confounding economies of scale in quality production with omitted quality shifters. As discussed in Section 4 and Appendix 5, we can in theory flexibly account for such alternative determinants by letting the cost parameter b_{ni} differ across firms within product groups. In the estimation, as a low draw of b_{ni} increases both firm sales and product quality, this could bias the estimate of β_n in the cross-section. To address such concerns, we can exploit the panel dimension of the microdata.

The natural panel data approach to estimating β_n and ξ_n would be to write (24) in log changes instead of log levels on both the left and right hand sides. To exploit plausibly exogenous variation in changes in a brand’s national sales scale, one could then exploit a shift-share instrument based on pre-existing brand-level sales shares across US states interacted with average changes (leave-out means) in firm scale across states over time.

However, the estimation of the economies of scale parameter β_n would still likely be biased. To see this, imagine we helicopter-dropped a random sales shock onto a firm that does not adjust either product quality or prices: in this scenario, even though the shock to firm scale is perfectly exogenous, we would mechanically conclude that there are economies of scale in quality production ($\beta_n = \frac{1}{\sigma_n - 1} > 0$). The reason is that any demand shock that one would usually want to exploit as instrument 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 (11), we get:

$$\begin{aligned} \Delta \log(p_{nist}) = & \beta_n \Delta \log(X_{nit}) - \frac{1}{N_z} \sum_z \frac{1}{\sigma_{nz} - 1} \Delta \log(X_{nizst}) \\ & + \beta_n \Delta \log(\tilde{\rho}_{nit}(\tilde{\gamma}_{nit} - \xi_n)) + \eta_{nst} + \epsilon_{nit} \end{aligned} \quad (25)$$

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

when discussing bootstrapped standard errors as part of the quantification of counterfactuals in the final section.

time at the state level. For instance changes in firm productivity (and thus unit values on the left) conditional on product quality would be 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 (11), firm changes in product quality are a function of aggregate national firm scale and the firm’s composition of consumer taste parameters.

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 will 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. Also notice that the estimation would not confound conventional economies of scale in producing identical goods with economies of scale in the production of product quality: if marginal costs fell with larger scale –holding quality constant–, this would be fully accounted for by the conventional demand relationship between changes in firm prices on the left and changes in sales captured by the second term on the right.

For estimation, we can re-write (25) as: $\Delta \log(\phi_{nist}) = \Delta \log(p_{nist}) + \frac{1}{N_z} \sum_z \frac{1}{\sigma_{nz}-1} \Delta \log(X_{nzist}) = \beta_n \Delta \log(X_{nit} \tilde{\rho}_{nit} (\tilde{\gamma}_{nit} - \xi_n)) + \eta_{nst} + \epsilon_{nit}$. As before, this allows us to estimate the technology parameters β_n and ξ_n for each product department by estimating β_n using OLS and IV regressions across iterations of ξ_n , and selecting the best-fitting parameter combination.

The first identification concern in (25) is correlated measurement errors between the left and right hand sides. The second major 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 (25), we use leave-out mean changes in log firm sales across other states ($s' \neq s$) and computed 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 composition-adjusted firm scale ($\Delta \log(X_{nit} \tilde{\rho}_{nit} (\tilde{\gamma}_{nit} - \xi_n))$) is thus based on average changes in firm scale over time that exclude the product group of the firm as well as the state in which the measure of product quality on the left hand side is observed. The identifying assumption of this strategy is that exogenous shocks to firm scale in other regions of the US do not affect changes in state-level brand quality through other channels but firm scale.

Estimation Results Before estimating β_n and ξ_n jointly as described above, we start in Table 3 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 after addressing concerns about correlated measurement errors in 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 specification in column 6 of Table 3 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. When using plausibly exogenous variation in two-year changes in firm scale in the IV estimation, column 8 suggests that a 10 percent increase in national firm sales leads to a 5.7 percent change in brand quality.

Table 4 proceeds to the structural estimates of β_n and ξ_n . The main difference to the previous reduced form table lies in the additional inclusion of brand-level consumer compositions as well as the marginal cost parameter ξ_n as shown above in the estimation equations (24) and (25). The first panel reports the results when pooling all product groups, and reassuringly the IV point estimates of the best-fitting parameter combination of β_n and ξ_n are close to the reduced form results reported in Table 3. The second panel reports the technology parameter estimates separately for grocery and non-grocery product groups, and Appendix Table A.8 reports the estimation results separately for each product department. 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 4 for the counterfactual quantification in the following section (reporting results for both the cross-section and panel data estimates). An interesting pattern emerges from the parameter estimates: in both the cross-sectional specification and the panel data approach, the IV point estimates for the economies of scale parameter in quality production are significantly larger for non-grocery product groups (e.g. health and beauty and merchandise) compared to grocery product groups.

5.4 Quantification of Forces

Armed with the preference and technology parameter estimates, we can check whether the calibrated model quantitatively replicates our main stylized fact documented in Figure 2, and also use the calibrated model to quantify the forces underlying this observed relationship. Following expressions (15) and (16), we can de-compose the observed differences in weighted average firm sizes across consumption baskets into different sources of consumer and firm heterogeneity. In Figure 5, 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 products 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 5, 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 2 across individual households, and in Figure 5 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 products in the data, and predict income group-specific consumption shares that are now taking into account both heterogeneity in $\sigma_n(z)$ and in $\gamma_n(z)$. As shown in Figure 5, the fact that higher-income households are estimated to have significantly stronger tastes for product quality pushes in the opposite direction of 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 2.

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 2, 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 5, 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 in combination with the data to explore the implications for household price indices and real income inequality, and decompose those effects into different channels. In the first set of counterfactuals, we quantify the knock-on effects of changes in the distribution of household nominal incomes for real income inequality. In the second counterfactual exercise, we quantify the distribution of the gains from opening up to trade. In a third application, we quantify the incidence of business regulations/taxes that can affect large and small firms differently. We use these counterfactual analyses to illustrate how and to what extent, in a setting where households source their consumption from heterogeneous firms, economic shocks can give rise to new distributional implications through asymmetric effects on the price indices of

different income groups.

6.1 Counterfactuals 1 and 2: Changes in Nominal Income Inequality

Our first two counterfactuals explore the implications of changes in nominal income inequality on household price indices. Through the lens of our model, the documented empirical moments in the scanner microdata have the implication that observed changes in the distribution of nominal incomes can be magnified or attenuated through general equilibrium effects on consumer price indices. In our framework, and the data, consumers differ in their product evaluations and in their price elasticities, while firms sell to different compositions of these consumers by optimally choosing product attributes and markups. With non-homothetic preferences, changes in nominal income inequality lead to changes in the distribution of price elasticities and product tastes that firms face. Producers respond to these changes by adjusting markups, product quality choices as well as exit and entry. With heterogeneous consumers and firms, both the averages of these adjustments across producers, as well as their heterogeneity across the firm size distribution affect the price indices of rich and poor households asymmetrically.

To explore and quantify these forces, we estimate two counterfactuals in which we model the implications of changes in the nominal income distribution while holding market size fixed. In the first counterfactual, we do so by reallocating 5 percent of market sales from the poorest quintile to the richest quintile. In the second counterfactual, we simulate the implications of moving from today’s US income distribution to a more equal distribution observed in 1980. To this end, we use historical US Census data that cover the US income distribution over the period 1980-2015. Using these data, we compute the percentage changes of total market expenditure for each quintile of the distribution going from today back to 1980. These changes (going from 2015 to 1980) are +1, +2, +3, +1 and -8 percent for the lowest, 2nd lowest, median, 2nd highest and highest income quintiles respectively.⁴⁰ In other words, the richest 20 percent of US households accounted for 8 percentage points more of total incomes in 2015 relative to 1980, while the other four income quintiles have seen relatively evenly spread reductions in their shares.

We estimate both counterfactuals 18 times, based on the observed brand sales to the five income groups for each half year in the scanner data in addition to our estimates for the parameters $\sigma_n(z)$, $\gamma_n(z)$, β_n and ξ_n .⁴¹ We then solve for the counterfactual equilibrium as described in section 4.4 and Appendix 7. To describe the mechanisms in detail, we use the decomposition of the price effect described in equation (19). We also report results separately using both the cross-sectional technology parameter estimates and the panel data estimates. Following the discussion in the previous section, the panel data estimates are likely to be more conservative as they are based

⁴⁰These changes across quintiles are not very large. While inequality across quintiles has increased over this period, our exercise does not account for the increasing concentration of wealth within these broad groups (e.g. among the top 1% within the top quintile). In reference to Figure 2, accounting for such changes would significantly reinforce these results.

⁴¹For each of the 18 six-month periods, we verify that the moments in the data and estimated parameter values satisfy the uniqueness conditions discussed in Section 4.2 and Appendix 2. Note that 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), such that our fitted quality equals observed quality for each brand and income quintile. We then obtain the exact same counterfactual equations for *changes* in quality, sales and price indexes as long as we hold these adjustment terms constant.

on firm adjustments in their product quality as a function of changes in scale and consumer composition over a two year period, instead of the long-term relationship captured by the cross-section. To compute confidence intervals that account for sampling error in the parameter estimates as well as in the sales data, we bootstrap the quantification exercise 200 times for each half year. 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.⁴²

Figures 6 and 7, A.14-A.16, and Table 5 present the estimation results of the effect on price indices by income group and its decomposition for both counterfactuals. There are several findings to notice. A 5 percent reallocation of expenditures from the poorest to the richest quintile induces changes in price indices that are on average 1.7 (1.6) percentage points lower for the richest household quintile compared to the poorest when using the cross-sectional (panel data) technology parameter estimates. Switching from this hypothetical scenario, we find that the observed increase in US inequality over the past decades has led to a 2.3 (1.7) percentage point lower cost of living inflation for the richest quintile compared to the poorest using the cross-sectional (panel data) technology parameter estimates. These findings suggest that increases in nominal inequality give rise to larger increases in real income inequality due to endogenous asymmetric effects on household price indices.

Table 5 presents the six-fold decomposition of the difference between the richest and the poorest quintiles for both the cross-sectional and panel data estimates of the technology parameters. The first channel through which consumer inflation can be affected differently between rich and poor households is that weighted average product quality increases across all producers in the market place. An increase in the income share of the richest quintile leads to a higher demand for quality and thus quality upgrading by producers, which in turn benefits households with higher tastes for product quality. This effect is significantly stronger when estimated using the cross-sectional technology parameters compared to the panel data estimates, given the larger elasticity parameter β_n in the cross-sectional estimation.

The second term on the heterogeneous scale effect reinforces the first channel and corresponds to 30-40 percent of the overall effect using the cross-sectional technology parameters, and more than half for the panel data estimates. Firms at the higher end of the quality distribution experience the most positive scale effects due to the change in the composition of demand. This induces asymmetric quality upgrading and leads to changes in quality-adjusted prices due to economies of scale in the production of product quality (Table 4). The right panels in Figures 6 and 7 confirm this intuition by depicting endogenous changes in log product quality across the initial firm size distribution. On average, the largest firms upgrade their quality by several percentage points more than firms at the other end of the size distribution. Since the largest firms tend to sell relatively more to rich consumers, the richest consumers are the ones benefiting the most. Quantitatively, weighted average quality upgrading embodied in poor consumers' consumption baskets is not significantly different from zero, while the consumption baskets of the top-quintile experience a significantly positive effect on product quality upgrading on average.

As the income distribution shifts to the right, average price elasticities decrease and average

⁴²This is a parametric bootstrap (Horowitz, 2001). See for example Atkin et al. (2016) for a recent application.

markups increase. A homogeneous change, however, would affect consumers symmetrically. What our third effect captures is the heterogeneous change in markups, which affects consumers differently. We find that smaller firms initially selling more of their total sales to poorer consumers are the ones who see the largest change in their consumer base, and therefore the largest increase in markups. This larger increase in markups affects poorer consumers the most, further reinforcing the unequal changes in household price indices. This differential effect is relatively small, however, in both the cross-sectional and panel-based estimations.

Our counterfactuals allow for the number of firms to adjust with free entry, such that expected profits upon entry remain equal to the sunk entry costs. Changes in the number of firms have asymmetric impacts across households depending on their elasticity of substitution across products $\sigma_n(z)$. Our estimates indicate that richer households have slightly lower elasticities of substitution, hence higher estimated love of variety. As our inequality counterfactuals lead to additional entry in the panel data case, richer income households benefit relatively more. Using the cross-sectional technology estimates, this effect is reversed in sign, but close to zero in both cases. The reason for this difference is the higher economies of scale parameter in the cross-sectional estimates. Given that we hold total market size constant in this counterfactual, this adjustment channel is quantitatively not very important in either of the two cases for both counterfactuals.

Since exiting firms are those who tend to sell relatively more to poor consumers initially, the exit of firms affects the consumption baskets of the poor relatively more than the rich. Quantitatively, however, we find that exit has a negligible effect in both counterfactuals and using both sets of technology estimates. Since in the data very small firms are able to survive in the baseline equilibrium, only tiny producers are likely to exit in the counterfactual equilibrium leading to practically zero differential effect across consumption baskets.⁴³

The sixth and final term of the decomposition quantifies the covariance between cross-module differences in consumption shares and module-level relative price changes. Our theoretical framework is in principle silent on this covariance term. Reassuringly, this term is also close to zero in the data for both counterfactuals. These results that we discuss above hold to a very similar extent in each of the 18 six-month periods as indicated by the non-bootstrapped confidence intervals in Table 5 and Figures 6 and 7. As shown in appendix Figures A.14 and A.15, they also hold across all product departments, but the magnitudes vary significantly. Finally, as shown in appendix Figure A.16, increases in nominal income inequality appear to significantly amplify the extent of firm heterogeneity within sectors: using both cross-sectional and panel-data technology estimates, we find a positive correlation between initial firm size and growth in sales as a function of increases in inequality on the consumer side.

6.2 Counterfactuals 3 and 4: Opening to Trade and Business Regulations

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

⁴³In the main exercise, we adopt a simple strategy by taking the maximum fixed cost that would allow all firms to survive in the baseline equilibrium. The results are not sensitive to this estimation method. Alternatively, we have estimated fixed costs f_{n0} by setting $f_{n0} = 0$ or by taking the maximum fixed costs such that all but the smallest 10 firms survive in the baseline equilibrium. The estimated fixed costs f_{n0} are tiny in either case.

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.

In this counterfactual, 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).⁴⁴

Figures 8, A.17, A.18 and Table 6 present the counterfactual results. Greater openness to trade induces consumer price index changes that are on average 2.6 (1.7) percentage points lower for the richest household quintile compared to the poorest using the cross-sectional (panel data) technology parameter estimates. We can use our six-fold decomposition in equation 19 to describe the mechanisms at play.

As in the first counterfactual, weighted average quality significantly increases as depicted in the right panel of Figure 8. 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 (Table 4). This average increase in quality tends to benefit richer households who have the highest preferences for quality, $\gamma_n(z)$. This term is quantitatively important using the cross-sectional technology parameters (48 percent of total effect) and less so (22 percent) using the panel data estimates.

The second effect corresponds to a covariance term between market shares $s_n(a, z)$ and quality upgrading $\log\left(\frac{\phi_{n1}(a)}{\phi_{n0}(a)}\right)$. The largest firms are the ones who become exporters and have 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. This pattern is also illustrated in the right-hand panel of Figure 8.

The third effect (heterogeneous markup adjustments) turns out to not be quantitatively important in the trade counterfactual. The fourth effect captures the change in the overall number of product varieties, which has asymmetric impacts across households depending on their love for

⁴⁴Given that we have a discrete number of firms, it is in theory possible to face multiple equilibria depending on the sequence of firm entry into export markets. In our simulations, we follow Eaton et al. (2012) and Gaubert & Itskhoki (2016) assuming entry is sequential across firms starting with the ones who gain the most on export markets. As a robustness check, we also verify that our results scale close to 10 times if instead we calibrate trade costs such that exporters export 2 percent of their output on average rather than 20 percent.

variety. It explains 20 percent of the total effect using the cross-sectional technology estimates, and 30 percent using the panel data estimates. This effect is now larger compared to the first counterfactual because it combines the number of varieties that are available on the domestic market as well as new imported varieties. As shown in Table 6, even the relatively minor differences in price elasticities across income groups can lead to sizeable differences in the gains from new imported variety or losses from exiting domestic firms.

While the fourth channel is driven by differences in $\sigma_n(z)$, the final channel takes into account differences in consumption shares spent on new imported varieties or exiting domestic firms across rich and poor households. This channel is also quantitatively important and reinforces the pure love of variety effect. 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. The sixth term that measures the covariance between module-level relative price changes and differences in consumption shares plays a minor role, as was the case in the first two counterfactuals. Finally, as documented in appendix Figure A.18, we find that trade liberalization leads to an amplification of the extent of firm heterogeneity within sectors, and that this effect is stronger in a world with two-sided heterogeneity in both consumption and production, compared to the canonical case with one-sided heterogeneity on the producer side.

In the final counterfactual exercise that we report in the appendix, we explore the incidence of a business tax that affects larger firms disproportionately. Following the notion that taxes and red tape increase in producer size (e.g. Hsieh & Klenow (2009)), we quantify the incidence of a non-uniform tax of 10 percent that affects the largest firms within sectors accounting for 20 percent of sectoral output. Appendix 7 derives the system of equations to solve for the counterfactual equilibrium in this exercise. In addition to the adjustment channels we describe above in terms of changes in firm sizes, product quality, markups and entry and exit, there is now also a direct partial-equilibrium effect in operation in this case: large producers of brands who are initially demanded to different extents across the income distribution are affected directly by the tax, whereas this is not the case among the majority of smaller firms. The average incidence of this business tax is an approximately 2 percent price increase across all producers of brands. As we depict in appendix Figure A.19 and Table 7, we find that the incidence of this business tax is 10-20 percent more pronounced for the consumption baskets of the richest 20 percent of US households compared to the effect on the consumption baskets of the poorest 20 percent.

6.3 Robustness

In the final section, we explore the sensitivity of these findings to alternative modeling assumptions and parameter values. First, as discussed in Section 5, we explore the sensitivity of the counterfactual analysis to alternative values of the elasticities of substitution compared to our baseline estimates of $\sigma_n(z)$. Second, as discussed in Section 3, we explore the sensitivity of the trade counterfactual to the fact that we do not observe firms' foreign sales when calibrating the model to brand sales observed in the Nielsen data. Third, we document the sensitivity of the counterfactual results to alternative assumptions about market structure, allowing for an additional source of vari-

able markups under oligopoly competition. Fourth, we re-estimate counterfactual outcomes after allowing household types (z) to endogenously change as a function of the counterfactual shocks. As for the analysis above, we report these additional results using both the technology parameter estimates from the cross-sectional and the panel data estimation discussed in Section 5.

Alternative Parameter Values In line with e.g. [Handbury \(2014\)](#) and the literature in empirical IO that has used the Hausman IV approach, we find somewhat lower values of the elasticity of substitution compared to the recent trade literature. To explore the sensitivity of our counterfactual results, we thus re-estimate each of them after assuming that our baseline estimates for the vector $\sigma_n(z)$ are under-estimated. [Table 7](#) reports the results of the price index implications of each of the four counterfactual exercises above for the richest relative to the poorest 20 percent of US households across the different parameter assumptions. Reassuringly, we find that higher values of the elasticity of substitution tend to either confirm or somewhat strengthen the asymmetric effects on household price indices that we report above. For the two inequality counterfactuals and the business tax counterfactual, the asymmetric effects are strengthened, and this finding holds across the two alternative technology parameterizations. In the case of the trade counterfactual, the point estimates are close to the baseline results, with slightly higher asymmetric effects for the panel data technology estimates, and slightly lower asymmetric effects when using the cross-sectional technology parameter estimates.

Foreign Sales A remaining concern for the trade counterfactual is that the Nielsen data do not allow us to observe firm sales outside the US market, either for US exporters or for imported products. As discussed in [Section 2](#), we can explore the sensitivity of our results to this data limitation by restricting the estimation to product groups with less than 10 percent of either import shares or export shares over total domestic sales (which corresponds to falling below the median across product groups in US retail). [Table 8](#) reports the counterfactual estimation results both for our baseline counterfactual that uses all product groups and after breaking up those product groups into above and below median with respect to either import or export shares. We find that the asymmetric effects on household price indices are slightly stronger for product groups with low import or export shares compared to those above the median. These findings are in line with the empirical results in [Section 2](#), where we find that differences in firm sizes across the consumption baskets of rich and poor households are slightly reinforced when limiting attention to product groups with less trade. Both sets of findings are indicative that the data limitation present in our model calibration leads to non-classical measurement error in total firm sales, where firms with larger market shares observed in the US market also tend to have larger omitted sales in foreign markets. Reassuringly, we find minor differences relative to our baseline estimates in both [Section 2](#) and [Table 8](#) relative to our baseline estimates.

Oligopoly Competition Our baseline specification of the model follows the canonical case as in [Melitz \(2003\)](#) and assumes 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 the composition of their consumers. However, given the concentration of sales that has been documented in the retail data (e.g. [Hottman et al. \(2016\)](#)), it seems important to test the

sensitivity of the counterfactual results to allowing for an additional source of variable markups under oligopolistic competition as discussed in Section 4 and Appendix 5. To this end, we compute the four counterfactual equilibria in the model extension with oligopoly competition in prices and quantities following [Hottman et al. \(2016\)](#). Appendix 7 derives the system of counterfactual equations and appendix Table A.9 presents the results of the four counterfactuals under the baseline model and under oligopoly. Reassuringly, the point estimates are very similar. While the price index differences between rich and poor consumer groups are very slightly weakened for the first two inequality counterfactuals, the differential effect of trade liberalization and business taxes are slightly stronger than in the baseline cases. In all cases, the difference is very small.

Endogenous Changes in Outside Good Consumption (z) As discussed in the theory section (e.g. see Footnote 32), our counterfactual analysis omits a second-order adjustment channel that could potentially affect the welfare results. In particular, we evaluate the price index implications of different market shocks across rich and poor consumer groups, while holding their initial consumer type (z) fixed. But since their relative economic welfare is affected by the shock, as we document above, we compute the difference in price index changes without allowing for consumption behavior (expenditure shares) to be affected differently by that change in relative welfare. In other words, we are not allowing the endogenous change in (z) to affect the computation of the price index effects. A priori, since our empirical application allows preferences to differ across five broad income groups, it is reasonable to think that few households are shifted in this manner. But we can also investigate the robustness of our results more directly. 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 utility function between retail consumption and the outside good, and quantify the change in outside good consumption induced by price index changes in the four counterfactuals above. Reassuringly, we find that the second-order adjustment channel has a negligible effect on outside good consumption (0.3% change at most) and thus a negligible effect on the counterfactual results. In other words, the change in relative welfare due to market shocks does not significantly alter the price index implications between rich and poor households, after we allow for the differential change in welfare to also affect their choice of retail products endogenously.

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, explores the underlying channels, and quantifies the implications for real income inequality. To do so, we bring to bear detailed home and store scanner data that allow us to trace the national firm size distribution into the consumption baskets of individual households, and combine these data with a tractable quantitative model that features two-sided heterogeneity across firms in production and consumers on the demand side.

The analysis provides several findings. We document large and statistically significant differences in the weighted average firm sizes that rich and poor households source their consumption

from. We find that this pattern is mainly explained by two features of household preferences and firm technology. On the consumption side, rich and poor households on average strongly agree on their ranking of product evaluations within sectors. However, richer households value higher quality attributes significantly more compared to poorer households. On the production side, we estimate that producing higher product quality increases both the marginal and the fixed costs of production. Combined, these forces give rise to the endogenous sorting of larger, more productive firms into products that are valued relatively more by wealthier households.

These results have implications for our understanding of inequality. We find that observed changes in nominal income inequality are magnified through asymmetric general equilibrium effects on household price indices, and that the distribution of the gains from international trade becomes significantly more regressive. We also find that business regulations that affect large and small firms to different degrees give rise to new distributional implications. Underlying these findings is a rich interplay of firm adjustments in product quality, markups, exit and entry that are asymmetric across the initial firm size distribution, which in turn translate differently into the consumption baskets of households across the income distribution.

Overall, our findings suggest that firm heterogeneity affects real income inequality in more complex ways than solely through 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 product choice by both heterogeneous households and firms– into an otherwise standard Melitz framework. Empirically, the findings presented in this paper emphasize the importance of capturing asymmetric changes in price indices at a granular level of product aggregation for both the measurement of overall changes in real income inequality over time, as well as for studying the effects of policy shocks, such as trade or the introduction of regulations, on inequality.

References

- Acemoglu, D., & Autor, D. (2011). Skills, tasks and technologies: Implications for employment and earnings. *Handbook of labor economics*, 4, 1043–1171.
- Amiti, M., Dai, M., Feenstra, R. C., & Romalis, J. (2016). How did China’s WTO entry benefit US consumers? *Working Paper*.
- Argente, D., & Lee, M. (2016). Cost of living inequality during the great recession. *University of Chicago Working Paper*.
- Arkolakis, C., Costinot, A., Donaldson, D., & Rodriguez-Clare, A. (2012). The elusive pro-competitive effects of trade. *Unpublished, MIT*.
- Atkin, D., Faber, B., & Gonzalez-Navarro, M. (2016). Retail globalization and household welfare: Evidence from Mexico. *Forthcoming, Journal of Political Economy*.
- 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. (2016). Export destinations and input prices. *Columbia University Working Paper*.

- Beraja, M., Hurst, E., & Ospina, J. (2014). The regional evolution of prices and wages during the great recession. *unpublished paper*.
- Bernard, A. B., Jensen, J. B., Redding, S. J., & Schott, P. K. (2007). Firms in international trade. *The Journal of Economic Perspectives*, 105–130.
- Bloom, N., Song, J., Price, D. J., & Guvenen, F. (2017). Firming up inequality. *NBER Working Paper*.
- 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. (2015). International trade, technology, and the skill premium. *Columbia University Working Paper*.
- Bustos, P. (2011). Trade liberalization, exports, and technology upgrading: Evidence on the impact of mercosur on argentinian firms. *The American economic review*, 101(1), 304–340.
- 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.
- Cravino, J., & Levchenko, A. A. (2016). The distributional consequences of large devaluations. *University of Michigan Working Paper*.
- 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*, 351–355.
- Dingel, J. I. (2015). The determinants of quality specialization. *Chicago Booth Working Paper*.
- Eaton, J., Kortum, S. S., & Sotelo, S. (2012). International trade: Linking micro and macro. *National Bureau of Economic Research Working Paper*.
- Eslava, M., Fieler, A. C., & Xu, D. (2016). Trade, skills, and quality upgrading: A theory with evidence from colombia. *Duke University Working Paper*.
- Faber, B. (2014). Trade liberalization, the price of quality, and inequality: Evidence from mexican store prices. *UC Berkeley Working Paper*.
- Fajgelbaum, P., Grossman, G., & Helpman, E. (2011). Income distribution, product quality, and international trade. *Journal of Political Economy*, 119(4), 721–765.
- Fajgelbaum, P., & Khandelwal, A. (2014). *Measuring the unequal gains from trade* (Tech. Rep.). National Bureau of Economic Research.
- 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.
- Frias, J. A., Kaplan, D. S., & Verhoogen, E. A. (2009). Exports and wage premia: Evidence from mexican employer-employee data. *Working Paper, Columbia University*.

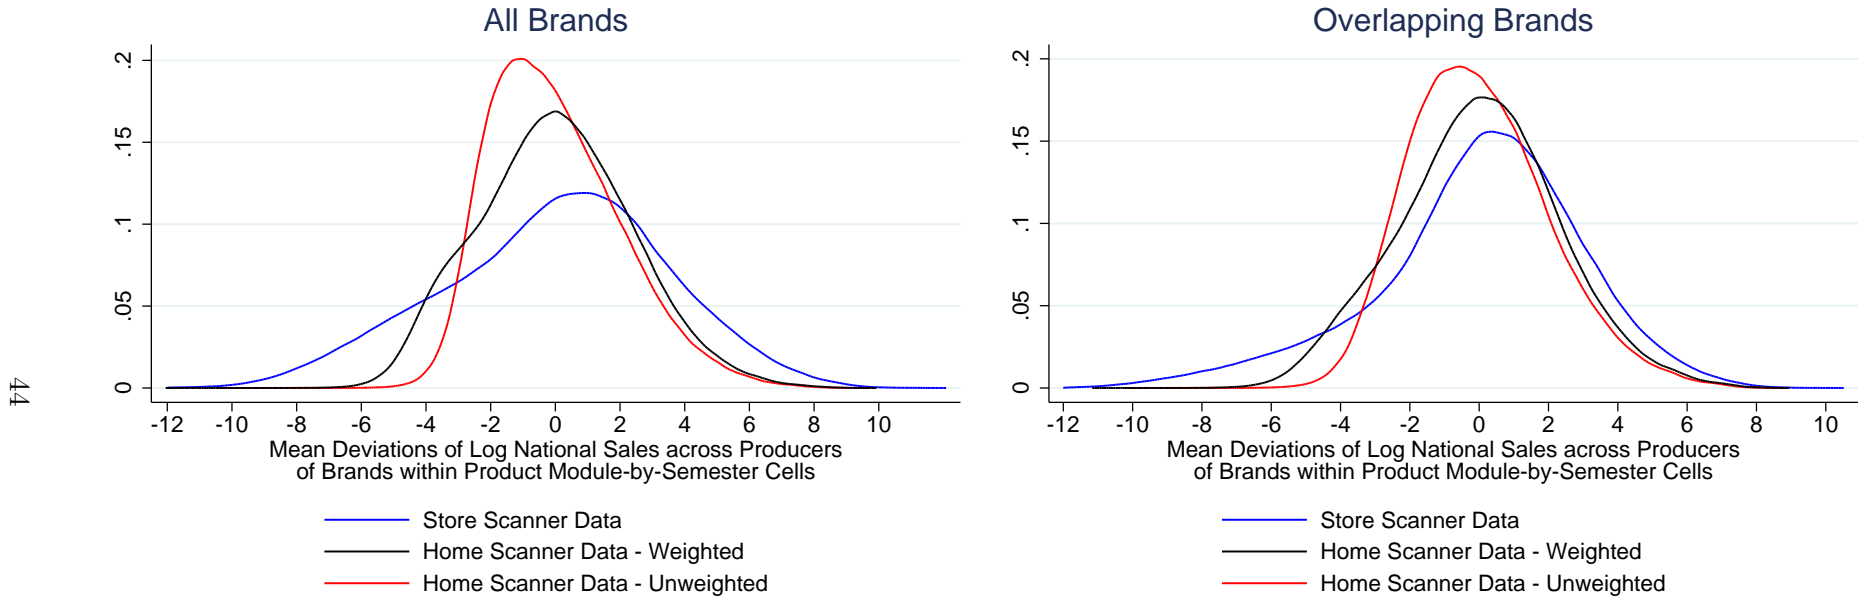
- Gaubert, C., & Itskhoki, O. (2016). Granular comparative advantage. *UC Berkeley Working Paper*.
- Handbury, J. (2014). Are poor cities cheap for everyone? Non-homotheticity and the cost of living across US cities. *Wharton Workign Paper*.
- Handbury, J., & Weinstein, D. E. (2014). Goods prices and availability in cities. *The Review of Economic Studies*, rdu033.
- Harrigan, J., & Reshef, A. (2011). *Skill biased heterogeneous firms, trade liberalization, and the skill premium* (Tech. Rep.). National Bureau of Economic Research.
- 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. *Forthcoming, Review of Economic Studies*.
- Helpman, E., Itskhoki, O., & Redding, S. (2010). Inequality and unemployment in a global economy. *Econometrica*, 78(4), 1239–1283.
- Horowitz, J. L. (2001). The bootstrap. In J. J. Heckman & E. Leamer (Eds.), *Handbook of econometrics* (Vol. 5, p. 3159 - 3228). Elsevier.
- Hottman, C., Redding, S. J., & Weinstein, D. E. (2016). What is firm heterogeneity in trade models? The role of quality, scope, markups, and cost. *Forthcoming, Quarterly Journal of Economics*.
- Hsieh, C.-T., & Klenow, P. J. (2009). Misallocation and manufacturing TFP in China and India. *Quarterly Journal of Economics*, 124(4).
- Jaravel, X. (2016). The unequal gains from product innovations. *SIEPR Working Paper*.
- 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. *Review of Economic Studies, Forthcoming*.
- 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.
- Piketty, T., & Saez, E. (2003). Income inequality in the United States 1913–1998. *The Quarterly journal of economics*, 118(1), 1–41.

- Porto, G. G. (2006). Using survey data to assess the distributional effects of trade policy. *Journal of International Economics*, 70(1), 140–160.
- Redding, S. J., & Weinstein, D. E. (2016). A unified approach to estimating demand and welfare changes. *Princeton University Working Paper*.
- Sampson, T. (2014). Selection into trade and wage inequality. *American Economic Journal: Microeconomics*, 6(3), 157–202.
- 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 (Cambridge, Mass).
- Verhoogen, E. A. (2008). Trade, quality upgrading, and wage inequality in the mexican manufacturing sector. *The Quarterly Journal of Economics*, 123(2), 489–530.

Figures and Tables


Figures

Figure 1: Firm Heterogeneity in the Home and Retail Scanner Data



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 1 provides descriptive statistics. See Section 3 for discussion.


Figure 2: Richer Households Source Their Consumption from Significantly Larger Firms



45


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 1 provides descriptive statistics. See Section 3 for discussion.

Figure 3: Distribution of Weighted Average Product Quality and Quality-Adjusted Prices across Consumption Baskets



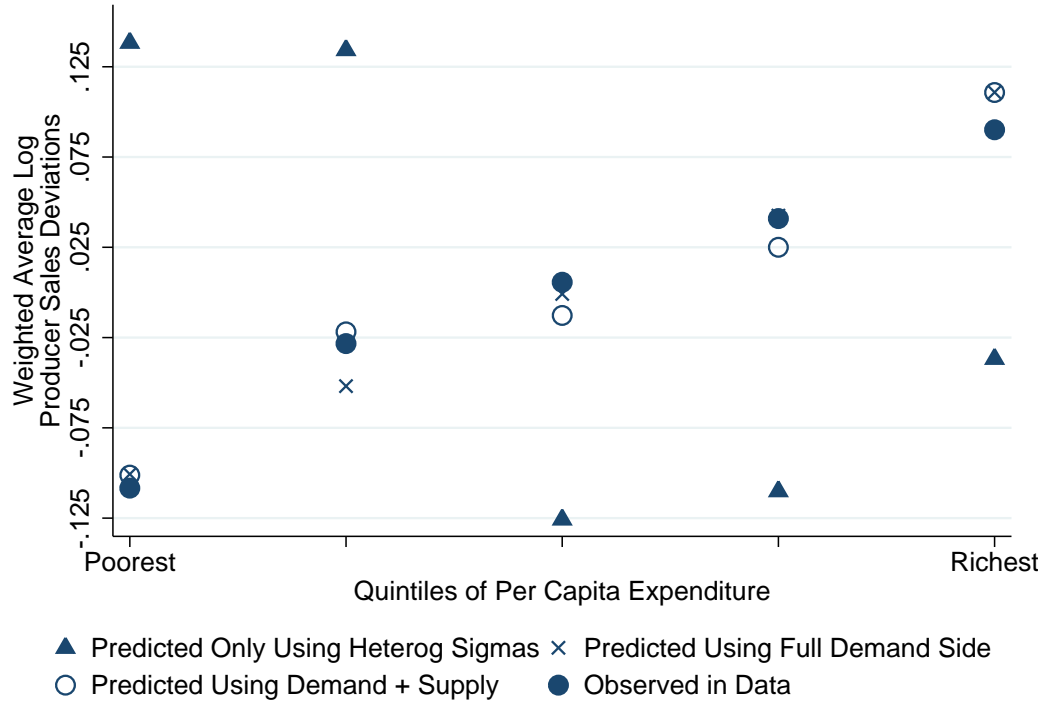
Notes: The figure depicts deviations in weighted average log brand quality embodied in the consumption baskets of on average 59,000 US households during 18 half-year periods between 2006-14. The y-axis on the left (right) displays weighted average deviations in log brand quality (quality-adjusted prices) 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 4: Producers Face Different Elasticities of Substitution and Tastes for Quality




Notes: The figure depicts deviations in the weighted average elasticities of substitution ($\tilde{\sigma}_{ni}$) and quality taste parameters ($\tilde{\gamma}_{ni}$) across the firm size distribution for 18 half-yearly cross-sections between 2006-2014. The y-axis displays deviations of these parameters 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 5: De-Composition 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. See Section 5 for discussion.

Figure 6: Counterfactual 1: Inflation Differences and Quality Upgrading due to 5 Percent Reallocation of Expenditure to Richest Group



Notes: Both graphs display confidence intervals at the 95% level. Confidence intervals in the left panel are based on robust standard errors across 18 six-month periods. See Table 5 for bootstrapped standard errors. 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. Percentiles of initial log firm sales are based on deviations of log firm sales from product-module means in the initial period. See Section 6 for discussion.


Figure 7: Counterfactual 2: Inflation Differences and Quality Upgrading due to Change in Inequality 1980-2015



50

Notes: Both graphs display confidence intervals at the 95% level. Confidence intervals in the left panel are based on robust standard errors across 18 six-month periods. See Table 5 for bootstrapped standard errors. 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. Percentiles of initial log firm sales are based on deviations of log firm sales from product-module means in the initial period. See Section 6 for discussion.

Figure 8: Counterfactual 3: Inflation Differences and Quality Upgrading due to 10% Symmetric Increase in Import Penetration



Notes: Both graphs display confidence intervals at the 95% level. Confidence intervals in the left panel are based on robust standard errors across 18 six-month periods. See Table 6 for bootstrapped standard errors. 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 for discussion.

Tables

Table 1: Descriptive Statistics

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

Table 2: Elasticities of Substitution

<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-Semester FX	✓	✓	✓	✓	✓
Brand-by-County-by-Semester 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>								
Dependent Variable: Change in Log Budget Shares	Beverages	Dairy	Dry Grocery	Frozen Foods	General Merchandise	Health and Beauty	Non-Food Grocery	Packaged Meat
(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-Semester 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>								
Dependent Variable: Change in Log Budget Shares	Beverages	Dairy	Dry Grocery	Frozen Foods	General Merchandise	Health and Beauty	Non-Food Grocery	Packaged Meat
(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-Semester 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 3: 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-Semester FX	✓	✓	✓	✓	✗	✗	✗	✗
State-by-Product Module-by-Semester 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 4: Technology Parameter Estimates

Dependent Variable:		ALL PRODUCT GROUPS							
Log Product Quality or Changes in Log Quality	Cross-Section				Panel Data				
	OLS	IV	OLS	IV	OLS	IV	OLS	IV	
Log Firm Scale or Changes in Log Firm Scale (β)	1.1092*** (0.0307)	1.1306*** (0.0305)	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		314439.78		253.09					

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)	0.9251*** (0.0118)	0.9445*** (0.0126)	0.9367*** (0.0113)	0.2857*** (0.0736)	1.5309*** (0.0844)	1.5729*** (0.0833)	1.5225*** (0.0998)	0.9438*** (0.1336)	
ξ Parameter	0.82	0.82	0.01	0.01	0.79	0.79	0.63	0.63	
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		265362.87		185.59		67033.48		104.37	

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 5: Decomposition of Inequality Counterfactuals

Difference in Consumer Inflation (Richest Quintile-Poorest Quintile)	Counterfactual 1: Hypothetical Increase in Nominal Inequality				Counterfactual 2: Changes in Income Distribution 1980 to 2015			
	Cross-Sectional Tech Estimates		Panel Data Tech Estimates		Cross-Sectional Tech Estimates		Panel Data Tech Estimates	
(1) Change in Weighted Average Product Quality	-0.906 (0.014) [0.072]	(53%)	-0.040 (0.004) [0.052]	(2%)	-1.544 (0.016) [0.122]	(67%)	-0.047 (0.007) [0.065]	(3%)
(2) Asymmetric Scale Effect	-0.756 (0.064) [0.101]	(44%)	-1.411 (0.091) [0.4]	(87%)	-0.699 (0.057) [0.063]	(31%)	-1.429 (0.094) [0.404]	(84%)
(3) Asymmetric Changes in Markups	-0.068 (0.005) [0.091]	(4%)	-0.069 (0.005) [0.101]	(4%)	-0.031 (0.002) [0.037]	(1%)	-0.033 (0.002) [0.044]	(2%)
(4) Love of Variety	0.081 (0.004) [0.43]	(-5%)	-0.069 (0.005) [0.564]	(4%)	0.101 (0.005) [0.565]	(-4%)	-0.120 (0.007) [0.743]	(7%)
(5) Asymmetric Effect of Exit	0.000 (0) [0]	(0%)	0.000 (0) [0]	(0%)	0.000 (0) [0]	(0%)	0.000 (0) [0]	(0%)
(6) Between-Group Effect	-0.055 (0.012) [0.017]	(3%)	-0.037 (0.021) [0.025]	(2%)	-0.117 (0.013) [0.022]	(5%)	-0.071 (0.017) [0.035]	(4%)
Total Effect	-1.704 (0.064) [0.418]	(100%)	-1.626 (0.101) [0.663]	(100%)	-2.289 (0.071) [0.546]	(100%)	-1.700 (0.103) [0.843]	(100%)

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

Table 6: Decomposition of Trade Counterfactual

Difference in Consumer Inflation (Richest Quintile-Poorest Quintile)	Counterfactual 3: Trade Opening			
	Cross-Sectional Tech Estimates		Panel Data Tech Estimates	
(1) Change in Weighted Average Product Quality	-1.230 (0.097) [0.062]	(48%)	-0.378 (0.029) [0.084]	(22%)
(2) Asymmetric Scale Effect	-0.332 (0.04) [0.083]	(13%)	-0.345 (0.057) [0.117]	(20%)
(3) Asymmetric Changes in Markups	0.004 (0.001) [0.002]	(0%)	0.001 (0) [0.002]	(0%)
(4) Love of Variety	-0.498 (0.07) [0.482]	(19%)	-0.498 (0.068) [0.623]	(29%)
(5) Asymmetric Effect of Exit and Imports	-0.478 (0.143) [0.061]	(18%)	-0.459 (0.219) [0.075]	(27%)
(6) Between-Group Effect	-0.054 (0.028) [0.03]	(2%)	-0.043 (0.04) [0.048]	(2%)
Total Effect	-2.588 (0.157) [0.497]	(100%)	-1.721 (0.235) [0.634]	(100%)

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

Table 7: Robustness to Alternative Parameters

	Baseline Parameters		$\sigma_n(z) + 1$		$\sigma_n(z) + 2$	
	Cross-Sectional Tech Estimates	Panel Tech Estimates	Cross-Sectional Tech Estimates	Panel Tech Estimates	Cross-Sectional Tech Estimates	Panel Tech Estimates
Counterfactual 1 (Inequality 1)	-1.704 (0.064)	-1.626 (0.101)	-2.163 (0.068)	-2.098 (0.127)	-3.591 (0.091)	-5.802 (0.275)
Counterfactual 2 (Inequality 2)	-2.289 (0.071)	-1.7 (0.103)	-2.961 (0.081)	-2.306 (0.138)	-4.665 (0.114)	-6.429 (0.283)
Counterfactual 3 (Trade)	-2.588 (0.157)	-1.721 (0.235)	-2.034 (0.123)	-1.268 (0.122)	-2.270 (0.131)	-2.522 (0.201)
Counterfactual 4 (Taxes)	0.229 (0.011)	0.306 (0.013)	0.235 (0.009)	0.349 (0.015)	0.232 (0.009)	0.365 (0.014)

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

Table 8: Robustness to Unobserved Sales


	All Product Groups	High Import Share	Low Import Share	High Export Share	Low Export Share
Cross-Sectional Tech Parameter Estimates	-2.588 (0.157)	-2.464 (0.285)	-2.607 (0.167)	-2.494 (0.284)	-2.581 (0.269)
Panel Data Tech Parameter Estimates	-1.721 (0.235)	-1.679 (0.229)	-1.767 (0.176)	-1.663 (0.507)	-1.581 (0.427)

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

Appendix

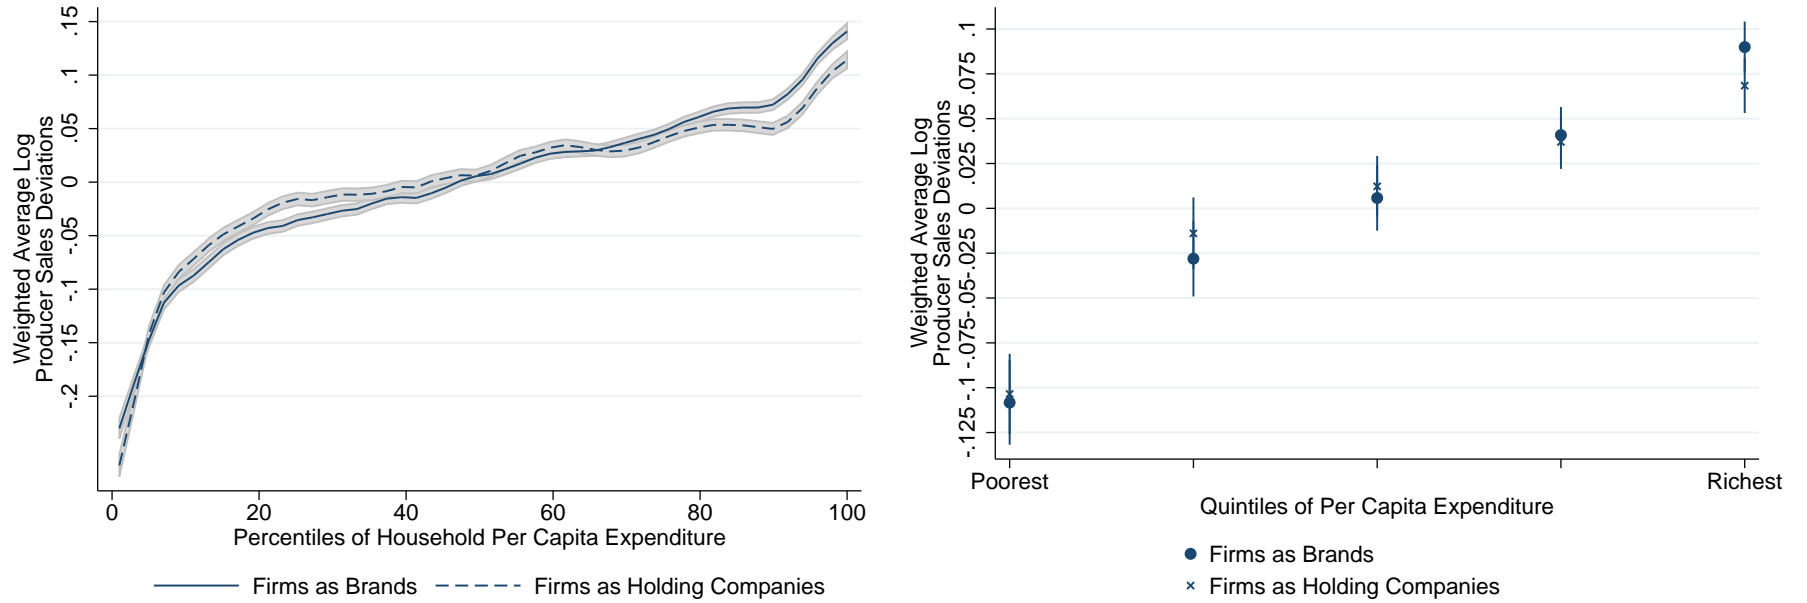
Appendix 1: Additional Figures and Tables

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 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 1 provides descriptive statistics. See Section 3 for discussion.

Figure A.3: Firm Heterogeneity Across Consumption Baskets - Firm Size in Terms of Quantities




Notes: The graph replicates the right panel of Figure 2 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.

Table A.1: 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	BISELL 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.4: Heterogeneity across Product Departments



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

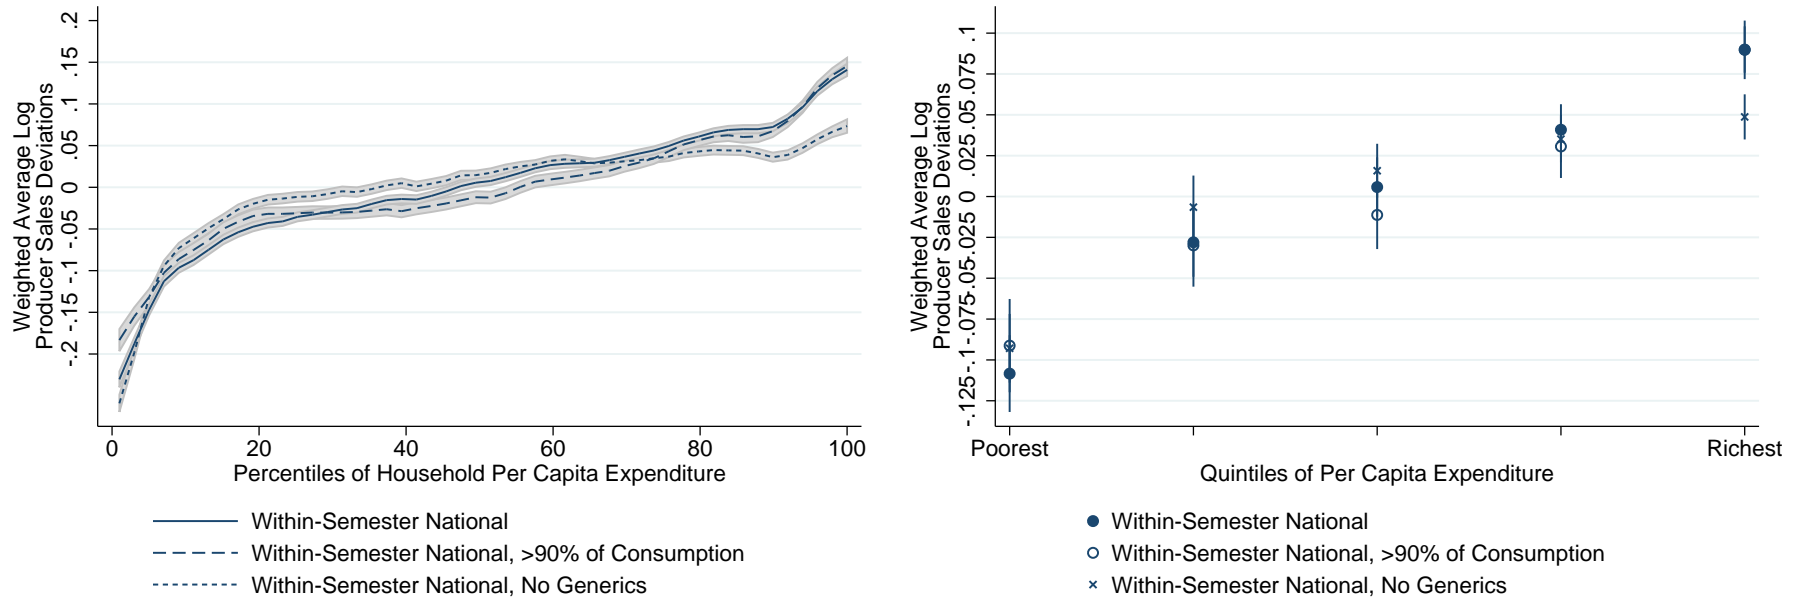
Figure A.5: Households on Average Agree on Relative Product Quality Evaluations



9

Notes: The figure depicts the relationship between income group-specific budget shares spent across producers within more than 1000 product modules (y-axis), and total market shares 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.6: The Role of Generic Retailer Brands and Non-Participating Store Chains




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.

Table A.2: 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, TOPPINGS	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, TOPPINGS	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			

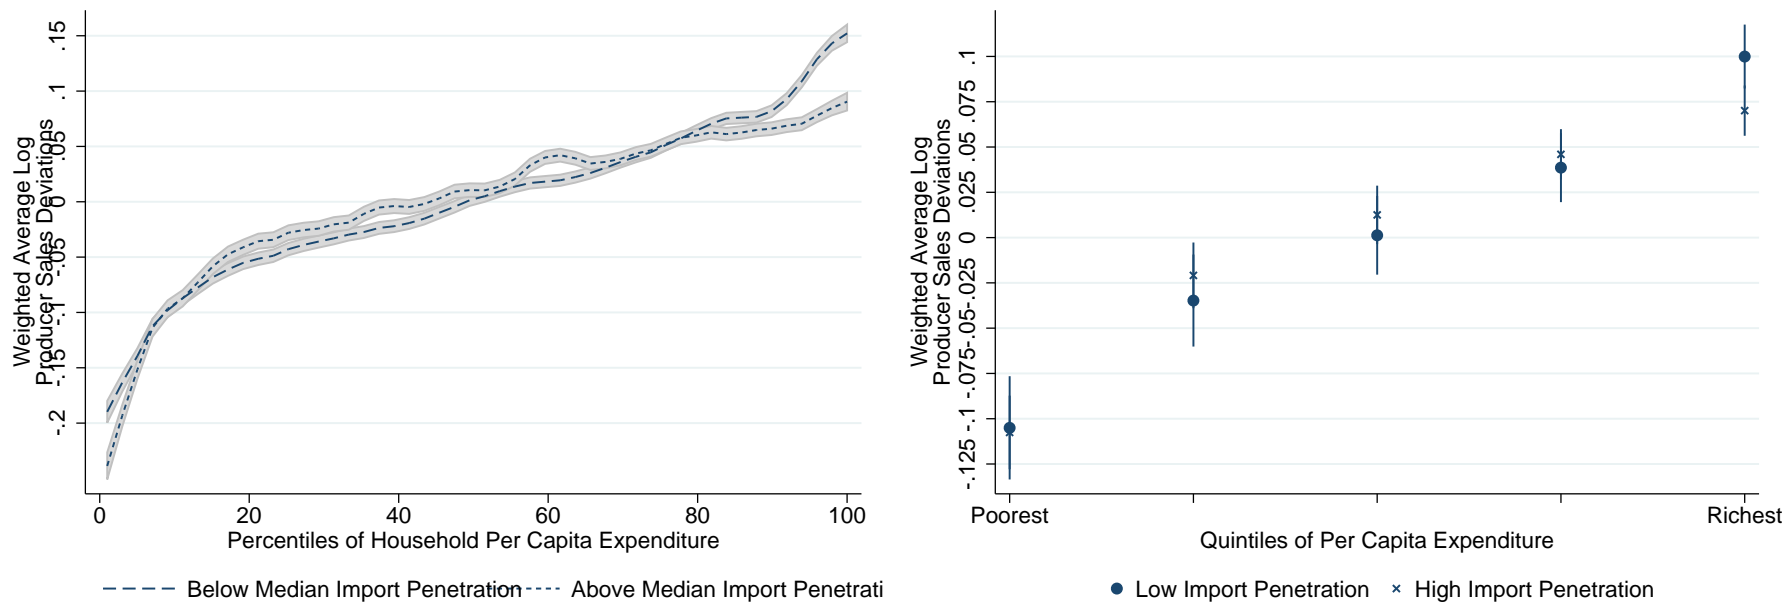
Notes: See Section 3 for discussion. Based on US trade data for 2005 at the 4-digit SIC product level.

Figure A.7: Firm Heterogeneity Across Consumption Baskets - Robustness to Export Shares



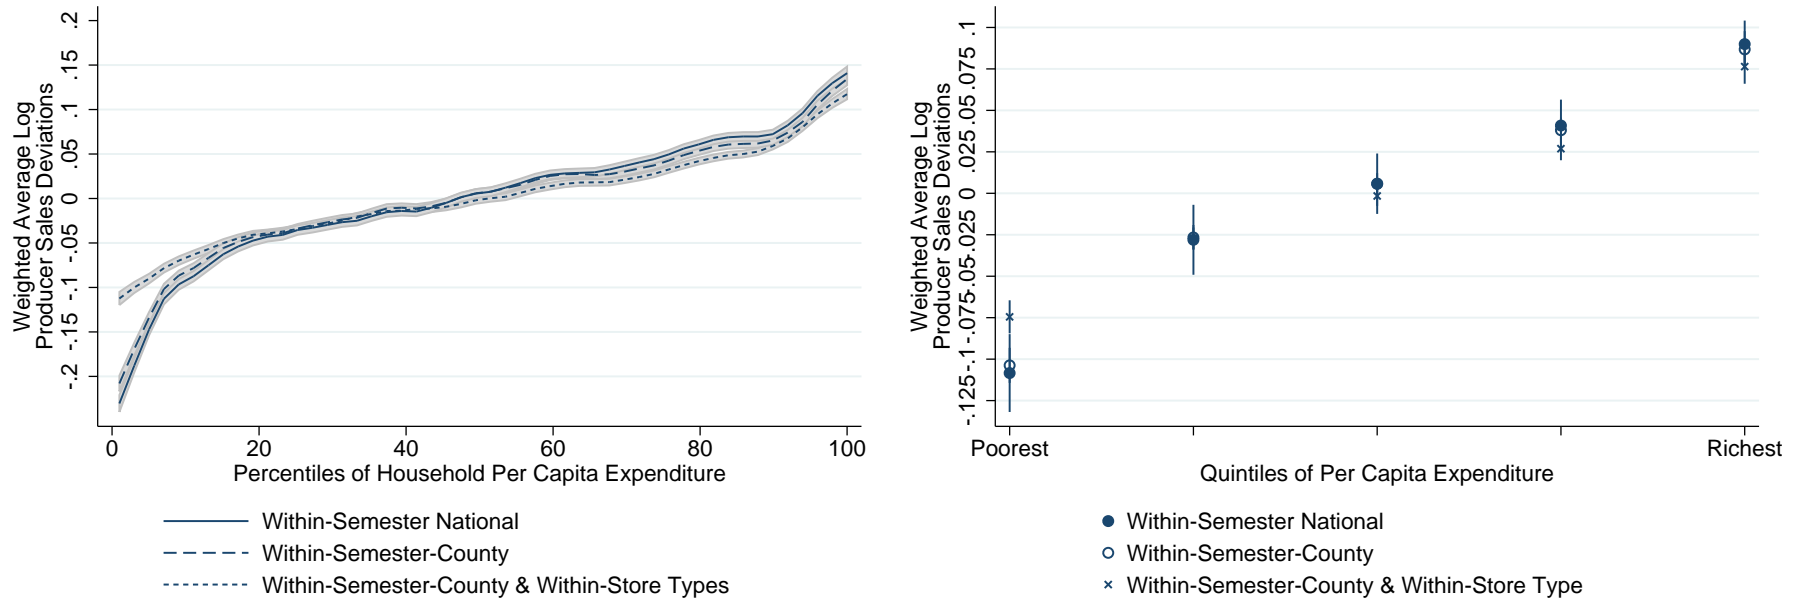
Notes: The fitted relationship in the left graph corresponds to a local polynomial regression. Standard errors in both graphs are clustered at the county level, and the displayed confidence intervals are at the 95% level. See Section 3 for discussion.

Figure A.8: 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.

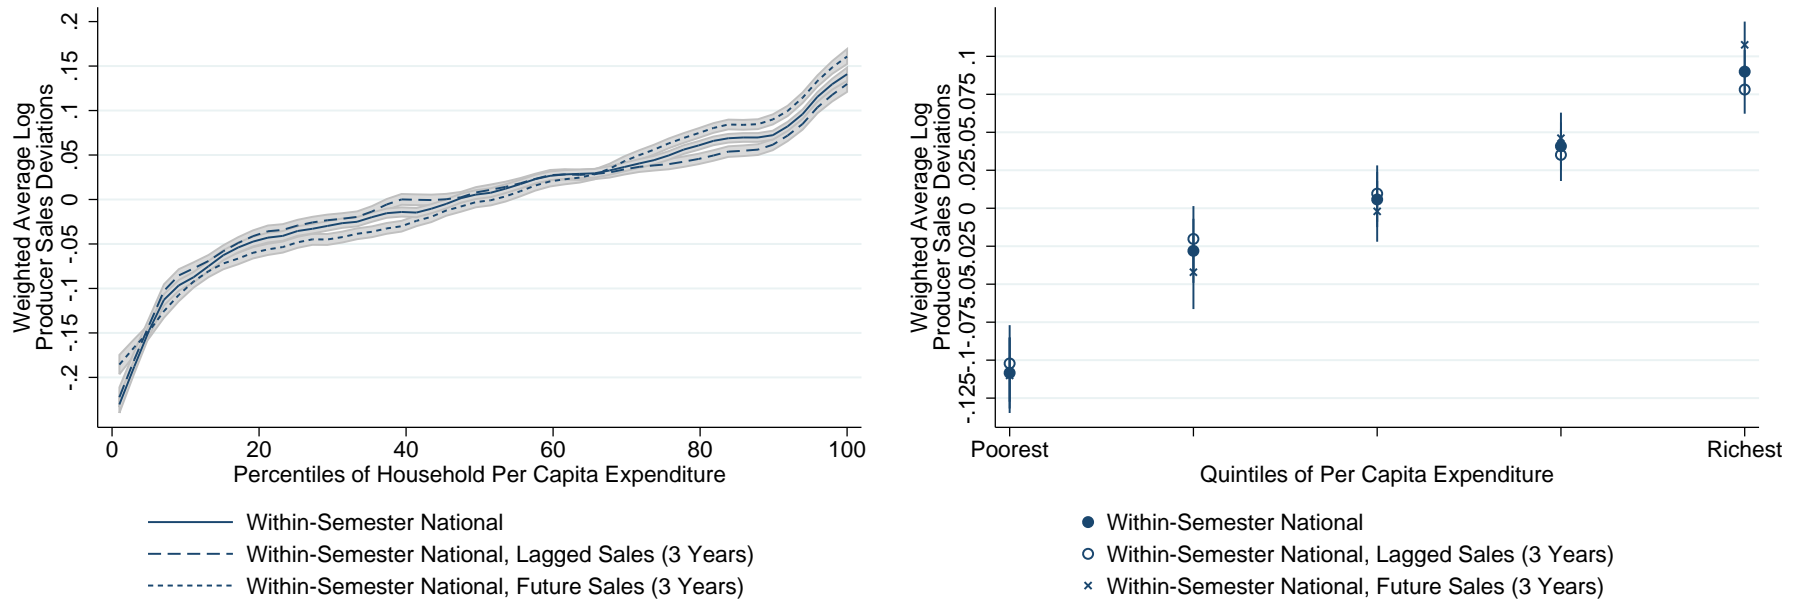
Figure A.9: The Role of Differential Access to Producers



11

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.10: The Role of Temporary Taste Shocks that Differ across Rich and Poor Households



12

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 1 provides descriptive statistics. See Section 3 for discussion.

Table A.3: 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.4: Income Group Ratios of Within and Cross-Brand Elasticities of Substitution

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-Semester FX	✓	✗
Quintile-by-Module-by-Brand-by-County-by-Semester 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.5: Alternative Specifications for Estimating the Elasticity of Substitution

<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-Semester 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
<hr/>								
<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-Semester 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
<hr/>								
<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-Semester 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.6: Full Cross of 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-Semester 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.

Figure A.11: 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.12: Households Agree on Product Quality Evaluations (But Rich Households Value Quality Relatively More)



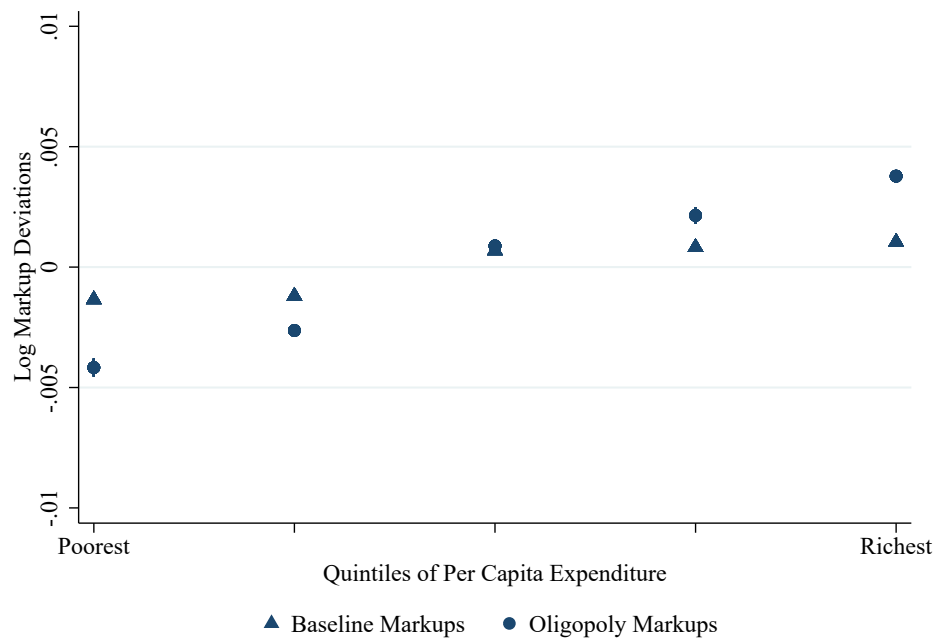
Notes: The figure depicts the relationship between deviations in log brand quality or quality-adjusted prices and deviations in log firm total sales for on average more than 150,000 producers of brands during 18 half-year periods between 2006-14. We estimate brand-level quality and quality-adjusted prices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

Table A.7: 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-Semester 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-Semester 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-Semester 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-Semester 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-Semester 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-Semester 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-Semester 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-Semester 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-Semester 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.

Figure A.13: Heterogeneity in Markups Across Consumption Baskets



Notes: The figure plots mean deviations in log markups within product modules across the consumption baskets of per-capita expenditure quintiles. Baseline markups are computed as a function of brand-level consumer composition as derived in the theory section and depicted in Figure 4. Oligopoly markups are computed following Hottman et al., (2016) as described in the theory section and Appendix 5. See Section 5 for discussion.

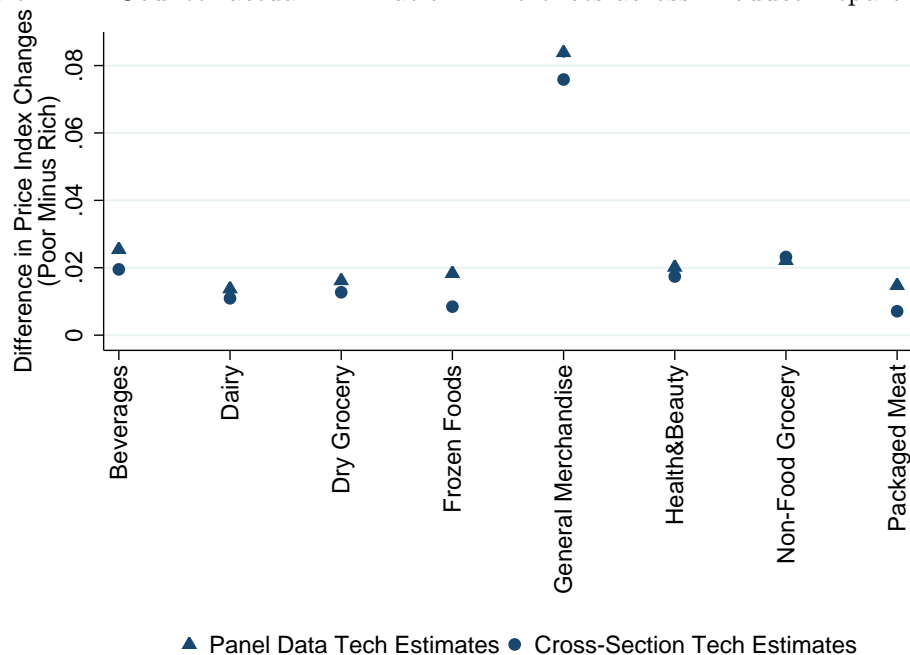
Table A.8: Technology Parameter Estimates

Dependent Variable:	BEVERAGES				DAIRY				DRY GROCERY				FROZEN FOODS			
Log Product Quality or Changes in Log Quality	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)	1.036*** (0.0212)	1.0587*** (0.026)	1.0559*** (0.0292)	-0.5982 (0.9736)	1.4821*** (0.0052)	1.5164*** (0.0076)	1.3725*** (0.0543)	0.2999 (0.3461)	0.7981*** (0.0026)	0.8211*** (0.0031)	0.8593*** (0.0073)	0.3985*** (0.0647)	0.8379*** (0.0034)	0.8479*** (0.0042)	0.815*** (0.0122)	0.1536 (0.1647)
ξ Parameter	0.76	0.76	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	50233.75		4.77		30308.12		20.37		135994.29		144.95		21625.45		27.28	

Dependent Variable:	GENERAL MERCHANDISE				HEALTH & BEAUTY CARE				NON-FOOD GROCERY				PACKAGED MEAT			
Log Product Quality or Changes in Log Quality	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)	0.5262*** (0.005)	0.5352*** (0.0057)	0.5179*** (0.0096)	0.2086*** (0.0793)	2.206*** (0.0073)	2.2555*** (0.0083)	2.156*** (0.0156)	1.6097*** (0.1415)	0.9718*** (0.0071)	0.9737*** (0.0072)	1.0147*** (0.0125)	0.2254 (0.3491)	0.8564*** (0.0063)	0.8713*** (0.0059)	0.7791*** (0.029)	-0.0256 (0.3647)
ξ Parameter	0.94	0.94	0.94	0.94	0.96	0.96	0.96	0.96	0.42	0.42	0.13	0.13	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	29440.55		46.6		36888.86		57.51		40152.76		10.5		40796.93		12.31	

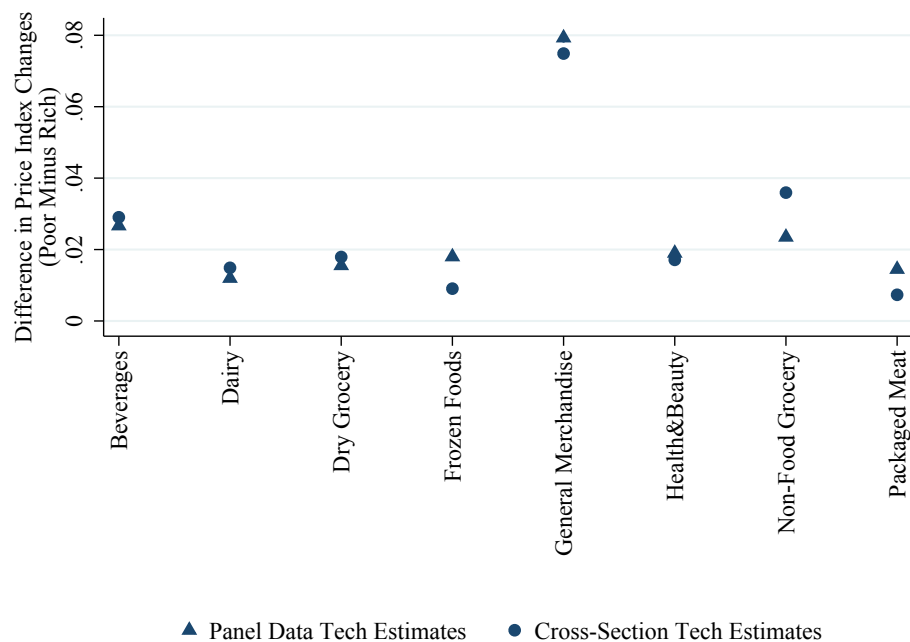
Notes: See Section 5 for discussion. Standard errors are in parentheses below point estimates and clustered at the level of product modules. ***, **, * indicate 1, 5 and 10 percent confidence levels.

Figure A.14: Counterfactual 1: Inflation Differences across Product Departments




Notes: The graph displays confidence intervals at the 95% level that are based on robust standard errors across 18 six-month periods. See Section 6 for discussion.

Figure A.15: Counterfactual 2: Inflation Differences across Product Departments




Notes: The graph displays confidence intervals at the 95% level that are based on robust standard errors across 18 six-month periods. See Section 6 for discussion.

Figure A.16: Counterfactual 1: Amplification of Firm Heterogeneity



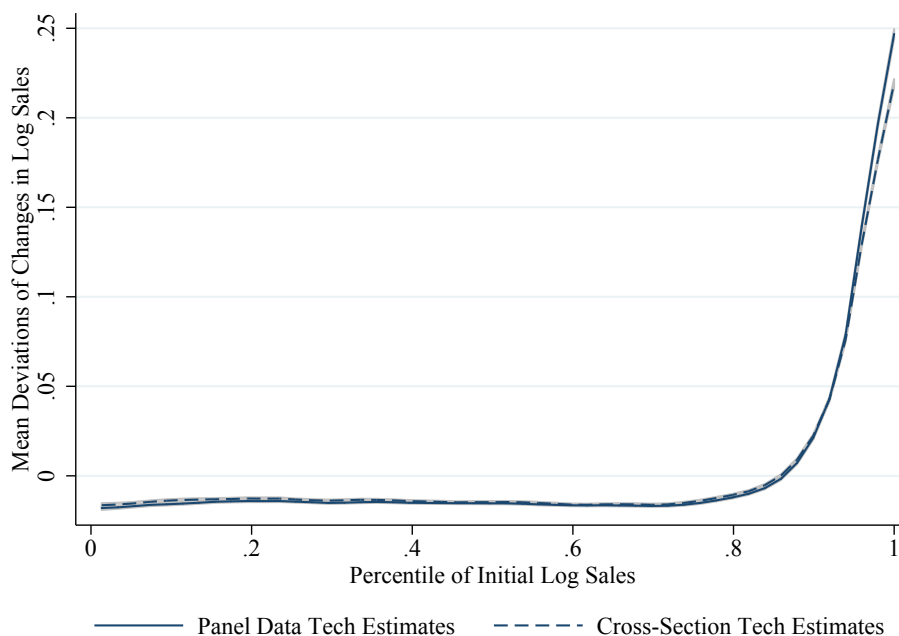
Notes: The graph plots counterfactual changes across producers of brands within product modules and within 18 six-month periods. The 95% confidence intervals are based on standard errors that are clustered at the level of product modules. See Section 6 for discussion.

Figure A.17: Counterfactual 3: Inflation Differences across Product Departments




Notes: The graph displays confidence intervals at the 95% level that are based on robust standard errors across 18 six-month periods. See Section 6 for discussion.

Figure A.18: Counterfactual 3: Amplification of Firm Heterogeneity



Notes: The graph plots counterfactual changes across producers of brands within product modules and within 18 six-month periods. The 95% confidence intervals are based on standard errors that are clustered at the level of product modules. See Section 6 for discussion.

Figure A.19: Counterfactual 4: Difference in the Incidence of Large-Firm Business Tax



Notes: The average incidence of the tax is approximately 1 percent higher inflation. The difference in cost of living inflation between the richest and the poorest household quintiles is 0.3 percent in absolute terms using the panel-data estimates of the technology parameters (see Table 7). In relative terms, this is incidence is about 18 percent higher for the richest quintile compared to the poorest. The graph displays confidence intervals at the 95% level that are based on robust standard errors across 18 six-month periods. See Section 6 for discussion.

Table A.9: Robustness to Oligopoly Competition

	Baseline (Monopolistic Competition)		Oligopoly	
	Cross-Sectional Tech Estimates	Panel Tech Estimates	Cross-Sectional Tech Estimates	Panel Tech Estimates
Counterfactual 1 (Inequality 1)	-1.704 (0.064)	-1.626 (0.101)	-1.478 (0.046)	-1.416 (0.085)
Counterfactual 2 (Inequality 2)	-2.289 (0.071)	-1.7 (0.103)	2.050 (0.055)	1.488 (0.089)
Counterfactual 3 (Trade)	-2.588 (0.157)	-1.721 (0.235)	-2.897 (0.185)	-2.198 (0.216)
Counterfactual 4 (Taxes)	0.229 (0.011)	0.306 (0.013)	0.235 (0.009)	0.349 (0.015)

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

Appendix 2: Mathematical Appendix - Sales, Markups and Quality

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:

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

with the demand shifter defined as $A_n(z) = \alpha_n(z) E(z) P_n(z)^{\sigma_n(z)-1}$ and $E(z)$ referring to total retail expenditure by consumer of income group z .

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

$$\begin{aligned} 0 = \frac{\partial \pi_n}{\partial \log \mu} &= - \left(1 - \frac{1}{\mu}\right) \int_z (\sigma_n(z) - 1) A_n(z) a^{\sigma_n(z)-1} \mu^{1-\sigma_n(z)} \phi^{(\sigma_n(z)-1)(\gamma_n(z)-\xi_n)} dH(z) \\ &\quad + \frac{1}{\mu} \int_z A_n(z) a^{\sigma_n(z)-1} \mu^{1-\sigma_n(z)} \phi^{(\sigma_n(z)-1)(\gamma_n(z)-\xi_n)} dH(z) \\ &= - \int_z (\sigma_n(z) - 1) A_n(z) a^{\sigma_n(z)-1} \mu^{1-\sigma_n(z)} \phi^{(\sigma_n(z)-1)(\gamma_n(z)-\xi_n)} dH(z) \\ &\quad + \frac{1}{\mu} \int_z \sigma_n(z) A_n(z) a^{\sigma_n(z)-1} \mu^{1-\sigma_n(z)} \phi^{(\sigma_n(z)-1)(\gamma_n(z)-\xi_n)} dH(z) \end{aligned}$$

Hence optimal markups satisfy:

$$\mu = \frac{\int_z \sigma_n(z) x_n(a, z, \mu, \phi) dH(z)}{\int_z (\sigma_n(z) - 1) x_n(a, z, \mu, \phi) dH(z)}$$

We obtain the expression in the text by defining:

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

and:

$$\tilde{\rho}_n(a) \equiv \frac{\int_z (\sigma_n(z) - 1) x_n(a, z) dH(z)}{\int_z \sigma_n(z) x_n(a, z) dH(z)} = \frac{\tilde{\sigma}_n(a) - 1}{\tilde{\sigma}_n(a)}$$

where $x_n(a, z)$ refers to sales of firm a .

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

$$0 = \frac{\partial \pi_n}{\partial \log \phi} \tag{26}$$

$$= \left(1 - \frac{1}{\mu}\right) \int_z (\sigma_n(z) - 1) (\gamma_n(z) - \xi_n) A_n(z) a^{\sigma_n(z)-1} \mu^{1-\sigma_n(z)} \phi^{(\sigma_n(z)-1)(\gamma_n(z)-\xi_n)-1} dH(z) - b_n \phi^{\frac{1}{\beta_n}-1} \tag{27}$$

$$= \left(1 - \frac{1}{\mu}\right) \int_z (\sigma_n(z) - 1)(\gamma_n(z) - \xi_n) x_n(z, a) dH(z) - b_n \phi^{\frac{1}{\beta_n}} \quad (28)$$

With $\mu = \frac{1}{\tilde{\rho}_n(a)} = \frac{\tilde{\sigma}_n(a)}{\tilde{\sigma}_n(a) - 1}$ and with $\tilde{\gamma}_n(a)$ defined as:

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

we obtain the expression in the text for optimal quality:

$$\phi_n(a) = \left(\frac{1}{b_n} \tilde{\rho}_n(a) X_n(a) (\tilde{\gamma}_n(a) - \xi_n) \right)^{\beta_n}$$

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 first derivative of profits is:

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

Hence the second derivative equals:

$$\begin{aligned} \frac{\partial^2 \pi}{\partial \log \mu^2} &= \int_z (\sigma_n(z) - 1)^2 A_n(z) a^{\sigma_n(z) - 1} \mu^{1 - \sigma_n(z)} \phi^{(\sigma_n(z) - 1)(\gamma_n(z) - \xi_n)} dH(z) \\ &\quad - \frac{1}{\mu} \int_z \sigma_n(z)^2 A_n(z) a^{\sigma_n(z) - 1} \mu^{1 - \sigma_n(z)} \phi^{(\sigma_n(z) - 1)(\gamma_n(z) - \xi_n)} dH(z) \\ &= \int_z (\sigma_n(z) - 1)^2 x_n(z, a) - \frac{1}{\mu} \int_z \sigma_n(z)^2 x_n(z, a) dH(z) \end{aligned}$$

The second order condition in markups is satisfied if:

$$\int_z (\sigma_n(z) - 1)^2 x_n(z, a) dH(z) < \frac{1}{\mu} \int_z \sigma_n(z)^2 x_n(z, a) dH(z)$$

where μ satisfies the first-order condition:

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

Hence, we need to show that:

$$\frac{\int_z \sigma_n(z)^2 x_n(a, z) dH(z)}{\int_z \sigma_n(z) x_n(a, z) dH(z)} - \frac{\int_z (\sigma_n(z) - 1)^2 x_n(a, z) dH(z)}{\int_z (\sigma_n(z) - 1) x_n(a, z) dH(z)} > 0$$

which is successively equivalent to:

$$\begin{aligned} & \left(\int_z \sigma_n(z)^2 x_n(a, z) dH(z) \right) \left(\int_z (\sigma_n(z) - 1) x_n(a, z) dH(z) \right) \\ & > \left(\int_z \sigma_n(z) x_n(a, z) dH(z) \right) \left(\int_z (\sigma_n(z) - 1)^2 x_n(a, z) dH(z) \right) \\ \Leftrightarrow & \left(\int_z \sigma_n(z)^2 x_n(a, z) dH(z) \right) \left(\int_z (\sigma_n(z) - 1) x_n(a, z) dH(z) \right) \\ & > \left(\int_z \sigma_n(z) x_n(a, z) dH(z) \right) \left[\int_z (\sigma_n(z)^2 - 2\sigma_n(z) + 1) x_n(a, z) dH(z) \right] \\ \Leftrightarrow & - \left(\int_z \sigma_n(z)^2 x_n(a, z) dH(z) \right) \left(\int_z x_n(a, z) dH(z) \right) \\ & > \left(\int_z \sigma_n(z) x_n(a, z) dH(z) \right) \left[\int_z (-2\sigma_n(z) + 1) x_n(a, z) dH(z) \right] \\ \Leftrightarrow & \frac{\int_z (2\sigma_n(z) - 1) x_n(a, z) dH(z)}{\int_z x_n(a, z) dH(z)} > \frac{\int_z \sigma_n(z)^2 x_n(a, z) dH(z)}{\int_z \sigma_n(z) x_n(a, z) dH(z)} \end{aligned}$$

In this last inequality, the left-hand side is larger than $2 \min_z \sigma_n(z) - 1$ while the right-hand side is not larger than $\max_z \sigma_n(z)$. With $\min_z \sigma_n(z) > \frac{\max_z \sigma_n(z) + 1}{2}$, this inequality is always satisfied. Note also that this inequality is always satisfied when $\sigma_n(z)$ is identical across income groups and larger than unity.

Second-Order Condition in Quality: Using equation (27), the second derivative in quality ϕ is:

$$\begin{aligned} \frac{\partial^2 \pi_n}{\partial \log \phi^2} &= \left(1 - \frac{1}{\mu} \right) \int_z (\sigma_n(z) - 1)^2 (\gamma_n(z) - \xi_n)^2 A_n(z) a^{\sigma_n(z) - 1} \mu^{1 - \sigma_n(z)} \phi^{(\sigma_n(z) - 1)(\gamma_n(z) - \xi_n)} dH(z) - \frac{b_n}{\beta_n} \phi^{\frac{1}{\beta_n}} \\ &= \left(1 - \frac{1}{\mu} \right) \int_z (\sigma_n(z) - 1)^2 (\gamma_n(z) - \xi_n)^2 x_n(z, a) dH(z) - \frac{b_n}{\beta_n} \phi^{\frac{1}{\beta_n}} \end{aligned}$$

This second derivative is negative when β_n is small enough. More specifically, when quality satisfies the first order condition, this second derivative is negative when:

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

This ensures that the first-order condition yields a unique level of quality for a given markup. A sufficient condition is that $\beta_n (\sigma_n(z) - 1)(\gamma_n(z) - \xi_n) < 1$ for all income groups z . The 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 $x_n(z, a)$ across income groups.

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

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

In addition to the second-order condition in markups, the Hessian is definite negative only 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, this can be rewritten:

$$\begin{aligned} &\left[\frac{1}{\mu} \int_z \sigma_n(z)^2 x_n(z, a) - \int_z (\sigma_n(z) - 1)^2 x_n(z, a) \right] \times \left[\frac{b_n}{\beta_n} \phi^{\frac{1}{\beta_n}} - \left(1 - \frac{1}{\mu}\right) \int_z (\sigma_n(z) - 1)^2 (\gamma_n(z) - \xi_n)^2 x_n(z, a) \right] \\ &> \left[- \left(1 - \frac{1}{\mu}\right) \int_z (\sigma_n(z) - 1)^2 (\gamma_n(z) - \xi_n) x_n(z, a) + \frac{1}{\mu} \int_z (\sigma_n(z) - 1) (\gamma_n(z) - \xi_n) x_n(z, a) \right]^2 \end{aligned}$$

where μ and ϕ satisfy the first order conditions described above:

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

and:

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

This inequality is equivalent to:

$$\begin{aligned} &\frac{b_n}{\beta_n} \phi^{\frac{1}{\beta_n}} - \left(1 - \frac{1}{\mu}\right) \int_z (\sigma_n(z) - 1)^2 (\gamma_n(z) - \xi_n)^2 x_n(z, a) \\ &> \frac{\left[- \left(1 - \frac{1}{\mu}\right) \int_z (\sigma_n(z) - 1)^2 (\gamma_n(z) - \xi_n) x_n(z, a) + \frac{1}{\mu} \int_z (\sigma_n(z) - 1) (\gamma_n(z) - \xi_n) x_n(z, a) \right]^2}{\frac{1}{\mu} \int_z \sigma_n(z)^2 x_n(z, a) - \int_z (\sigma_n(z) - 1)^2 x_n(z, a)} \end{aligned}$$

Again, this inequality holds when β_n is not too large.

2.C) Quality and Sales Increase with Productivity

We show below that quality and sales increase with productivity, i.e. $\frac{d \log \phi_n(a)}{d \log a}$ and $\frac{d \log X_n(a)}{d \log a}$ are positive, as long as β_n is small enough (to avoid offsetting feedback effects) and as long as there is not too much heterogeneity in $\sigma_n(z)$. A sufficient condition on the heterogeneity across $\sigma_n(z)$ is:

$$\min_z \sigma_n(z) > \frac{\max_z \sigma_n(z) + 1}{2}$$

which is the same as to ensure the second-order condition described above. In a first step, we show that markups do not increase as fast as productivity.

Markups and Productivity: First, the log-derivative of sales w.r.t. productivity a can be written as:

$$\frac{d \log x_n(a, z)}{d \log a} = (\sigma_n(z) - 1) \left(1 - \frac{d \log \mu}{d \log a}\right) + (\gamma_n(z) - \xi_n) (\sigma_n(z) - 1) \frac{d \log \phi}{d \log a}$$

With the optimal markup $\mu = \frac{\int_z \sigma_n(z) x_n(a, z) dH(z)}{\int_z (\sigma_n(z) - 1) x_n(a, z) dH(z)}$, we obtain:

$$\frac{d \log \mu}{d \log a} = \left(1 - \frac{d \log \mu}{d \log a}\right) B_1 + \frac{d \log \phi}{d \log a} B_2$$

with:

$$B_1 = \frac{\int_z \sigma_n(z) (\sigma_n(z) - 1) x_n(a, z) dH(z)}{\int_z \sigma_n(z) x_n(a, z) dH(z)} - \frac{\int_z (\sigma_n(z) - 1)^2 x_n(a, z) dH(z)}{\int_z (\sigma_n(z) - 1) x_n(a, z) dH(z)}$$

and:

$$B_2 = \frac{\int_z (\gamma_n(z) - \xi_n) (\sigma_n(z) - 1) \sigma_n(z) x_n(a, z) dH(z)}{\int_z \sigma_n(z) x_n(a, z) dH(z)} - \frac{\int_z \sigma_n(z) (\gamma_n(z) - \xi_n) (\sigma_n(z) - 1)^2 x_n(a, z) dH(z)}{\int_z (\sigma_n(z) - 1) x_n(a, z) dH(z)}$$

Rearranging, we obtain:

$$\frac{d \log \mu}{d \log a} = \frac{B_1}{1 + B_1} + \frac{d \log \phi}{d \log a} \frac{B_2}{1 + B_1}$$

First, one can see that B_1 is negative:

$$\begin{aligned} B_1 &= \frac{\int_z \sigma_n(z) (\sigma_n(z) - 1) x_n(a, z) dH(z)}{\int_z \sigma_n(z) x_n(a, z) dH(z)} - \frac{\int_z (\sigma_n(z) - 1)^2 x_n(a, z) dH(z)}{\int_z (\sigma_n(z) - 1) x_n(a, z) dH(z)} \\ &= \frac{\int_z \frac{\sigma_n(z)}{\sigma_n(z) - 1} (\sigma_n(z) - 1)^2 x_n(a, z) dH(z)}{\int_z \frac{\sigma_n(z)}{\sigma_n(z) - 1} (\sigma_n(z) - 1) x_n(a, z) dH(z)} - \frac{\int_z (\sigma_n(z) - 1)^2 x_n(a, z) dH(z)}{\int_z (\sigma_n(z) - 1) x_n(a, z) dH(z)} \end{aligned}$$

The two terms are weighted averages of $\sigma_n(z) - 1$ and are identical except for the additional weight $\frac{\sigma_n(z)}{\sigma_n(z) - 1}$ on the left hand side. Since $\frac{\sigma_n(z)}{\sigma_n(z) - 1}$ is monotonically decreasing with $\sigma_n(z) - 1$, one can conclude that the left term is smaller than the right term (Chebyshev's order inequality).

Next, one can see that $B_1 > -1$:

$$\begin{aligned} B_1 > -1 &\Leftrightarrow \frac{\int_z \sigma_n(z) (\sigma_n(z) - 1) x_n(a, z) dH(z)}{\int_z \sigma_n(z) x_n(a, z) dH(z)} - \frac{\int_z (\sigma_n(z) - 1)^2 x_n(a, z) dH(z)}{\int_z (\sigma_n(z) - 1) x_n(a, z) dH(z)} > -1 \\ &\Leftrightarrow \frac{\int_z \sigma_n(z)^2 x_n(a, z) dH(z)}{\int_z \sigma_n(z) x_n(a, z) dH(z)} - \frac{\int_z (\sigma_n(z) - 1)^2 x_n(a, z) dH(z)}{\int_z (\sigma_n(z) - 1) x_n(a, z) dH(z)} > 0 \end{aligned}$$

This condition is equivalent to the second-order condition in markups. With $\min_z \sigma_n(z) > \frac{\max_z \sigma_n(z) + 1}{2}$, this inequality is always satisfied.

With $B_1 > -1$, we obtain that $\frac{B_1}{1+B_1}$ is negative. Hence, $\frac{d \log \mu}{d \log a}$ is smaller than a constant term times $\frac{d \log \phi}{d \log a}$. When β_n is small enough, $\frac{d \log \phi}{d \log a}$ is not too large and one can ensure that $\frac{d \log \mu}{d \log a}$ is smaller than one.

Quality and Productivity: Using the first-order condition in optimal quality:

$$b_n \phi^{\frac{1}{\beta_n}} = \frac{1}{\tilde{\sigma}_n(a)} \int_z (\gamma_n(z) - \xi_n) (\sigma_n(z) - 1) x_n(a, z) dH(z)$$

we obtain:

$$\frac{1}{\beta_n} \frac{d \log \phi_n(a)}{d \log a} = \left(1 - \frac{d \log \mu}{d \log a}\right) B_3 + \frac{d \log \phi}{d \log a} B_4$$

where:

$$B_3 = \frac{\int_z (\sigma_n(z) - 1) x_n(a, z) dH(z)}{\int_z x_n(a, z) dH(z)} - \frac{\int_z \sigma_n(z) (\sigma_n(z) - 1) x_n(a, z) dH(z)}{\int_z \sigma_n(z) x_n(a, z) dH(z)} + \frac{\int_z (\gamma_n(z) - \xi_n) (\sigma_n(z) - 1)^2 x_n(a, z) dH(z)}{\int_z (\gamma_n(z) - \xi_n) (\sigma_n(z) - 1) x_n(a, z) dH(z)}$$

and:

$$B_4 = \frac{\int_z (\gamma_n(z) - \xi_n)(\sigma_n(z) - 1)x_n(a, z)dH(z)}{\int_z x_n(a, z)dH(z)} - \frac{\int_z \sigma_n(z)(\gamma_n(z) - \xi_n)(\sigma_n(z) - 1)x_n(a, z)dH(z)}{\int_z \sigma_n(z)x_n(a, z)dH(z)} + \frac{\int_z (\gamma_n(z) - \xi_n)^2(\sigma_n(z) - 1)^2x_n(a, z)dH(z)}{\int_z (\gamma_n(z) - \xi_n)(\sigma_n(z) - 1)x_n(a, z)dH(z)}$$

Rearranging, we obtain:

$$\frac{d \log \phi_n(a)}{d \log a} = \left(1 - \frac{d \log \mu}{d \log a}\right) \frac{\beta_n B_3}{1 - \beta_n B_4}$$

We have already shown that $\frac{d \log \mu}{d \log a} < 1$ provided that β_n is not too large and $\min_z \sigma_n(z) > \frac{\max_z \sigma_n(z) + 1}{2}$. Hence $\frac{d \log \phi}{d \log a}$ has the same sign as B_3 . In turn, B_3 is positive as long as:

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

The left-hand side is at least larger than $\min_z \sigma_n(z) - 1$, while the right-hand side the difference between two weighted averages of $\sigma_n(z) - 1$. This difference cannot exceed $\max_z \sigma_n(z) - \min_z \sigma_n(z)$. Hence, the condition $\min_z \sigma_n(z) - 1 > \max_z \sigma_n(z) - \min_z \sigma_n(z)$ is sufficient to ensure that B_3 is positive (this is the same condition to ensure that the second-order condition in markups is satisfied).

Note also that, as can be seen in the denominator, an important restriction is that $\beta_n B_4$ be smaller than one. When preference parameters $\gamma_n(z)$ and $\sigma_n(z)$ are homogeneous across consumers, this is equivalent to imposing that $\beta_n(\gamma_n - \xi_n)(\sigma_n - 1) < 1$. We do not have an well-defined equilibrium with heterogeneous quality choices if this condition is not satisfied.

Sales and Productivity: Finally, using:

$$\frac{d \log x_n(a, z)}{d \log a} = (\sigma_n(z) - 1) \left(1 - \frac{d \log \mu}{d \log a}\right) + (\gamma_n(z) - \xi_n)(\sigma_n(z) - 1) \frac{d \log \phi}{d \log a}$$

one can see that sales increase with productivity a for all consumers as long as $\frac{d \log \mu}{d \log a} < 1$ and $\frac{d \log \phi}{d \log a} > 0$. As shown in the last two subsections, these two conditions are satisfied when β_n is small enough and when the second-order condition in markups is satisfied ($\min_z \sigma_n(z) - 1 > \max_z \sigma_n(z) - \min_z \sigma_n(z)$ is a sufficient condition). Defining firm size as total sales, this also implies that firm size increases with productivity.

Homogeneous Consumers: Here we examine how quality depends on productivity a , focusing on the particular case where firm a sells to only one income group z_0 . In this case, we have:

$$b_n \phi^{\frac{1}{\beta_n}} = \rho_n(z_0)(\gamma_n(z_0) - \xi_n)x_n(a, z_0)$$

Note that the elasticity of $x_n(a, z_0)$ w.r.t a is $\sigma_n(z_0) - 1$ and the elasticity w.r.t to ϕ_n is $(\sigma_n(z_0) - 1)(\gamma_n(z_0) - \xi_n)$. Differentiating, this leads to:

$$\frac{1}{\beta_n} \frac{d \log \phi}{d \log a} = (\sigma_n(z_0) - 1) + \frac{d \log \phi}{d \log a} (\sigma_n(z_0) - 1)(\gamma_n(z_0) - \xi_n)$$

and thus:

$$\frac{d \log \phi_n(a)}{d \log a} = \frac{\beta_n(\sigma_n(z_0) - 1)}{1 - \beta_n(\sigma_n(z_0) - 1)(\gamma_n(z_0) - \xi_n)}$$

In turn, the total elasticity of sales w.r.t productivity a is the same as for ϕ , divided by β_n :

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

Note that this elasticity is larger than the elasticity $\sigma_n(z_0) - 1$ when quality is fixed and exogenous.

2.D) Other Expressions for Sales and Profits

Profits (Equation 14): As shown above:

$$\phi_n(a) = \left(\frac{1}{b_n} \cdot \tilde{\rho}_n(a) \cdot X_n(a) \cdot (\tilde{\gamma}_n(a) - \xi_n) \right)^{\beta_n}$$

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

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

This implies that fixed costs spent on quality upgrading equal:

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

Given that variable costs correspond to a share $\tilde{\rho}_n(a) = 1 - \frac{1}{\tilde{\sigma}_n(a)}$ of total sales, we obtain that profits equal:

$$\pi_n(a) = \frac{1}{\tilde{\sigma}_n(a)} (1 - \beta_n (\tilde{\gamma}_n(a) - \xi_n) (\tilde{\sigma}_n(a) - 1)) X_n(a) - f_{0n}$$

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

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

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 2 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 indexes) 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:

$$\begin{aligned}
\frac{\partial \log \widetilde{X}_n(z)}{\partial z} &= \frac{\partial \gamma_n}{\partial z} (\sigma_n(z) - 1) \left[\frac{\int_a x_n(z, a) (\log X_n(a)) (\log \phi_n(a)) dG_n(a)}{\int_a x_n(z, a) dG_n(a)} \right. \\
&\quad \left. - \left(\frac{\int_a x_n(z, a) \log X_n(a) dG_n(a)}{\int_a x_n(z, a) dG_n(a)} \right) \left(\frac{\int_a x_n(z, a) \log \phi_n(a) dG_n(a)}{\int_a x_n(z, a) dG_n(a)} \right) \right] \\
&\quad - \frac{\partial \sigma_z}{\partial z} \cdot \left[\frac{\int_a x_n(z, a) (\log X_n(a)) (\log(p_n(a)/\phi_n(a)^{\gamma_n(z)}))}{\int_a x_n(z, a) dG_n(a)} \right. \\
&\quad \left. - \left(\frac{\int_a x_n(z, a) \log X_n(a) dG_n(a)}{\int_a x_n(z, a) dG_n(a)} \right) \left(\frac{\int_a x_n(z, a) \log(p_n(a)/\phi_n(a)^{\gamma_n(z)}) dG_n(a)}{\int_a x_n(z, a) dG_n(a)} \right) \right]
\end{aligned}$$

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

Estimation Equation for β_n and ξ_n (Equation 24): 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} :

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

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

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

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

Appendix 3: Changes in 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 is given by:

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

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

Proof: 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 indexes across households in our counterfactuals may have induced a change in the consumption of the outside good z , holding nominal income constant. In our first three counterfactuals, retail prices tend to decrease faster for rich relative to poor households, which may affect the consumption of outside goods z . In the fourth counterfactual (business taxes) the opposite is the case. 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 utility takes the form:

$$U = A(z) U_G(z)^{\alpha_G} \tag{29}$$

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 $A(z) = z^{1-\alpha}$. 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 indexes $P_G(z)$. 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 29. Yet, a change in the patterns of prices across z 's still influences the consumption of the outside good z . Moreover, equation 29 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:

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

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:

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

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 an infinitesimally small $d \log w$. Differentiating the first-order condition, 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 Suppose that $\varepsilon_P(z) = \frac{zP'_G(z)}{P_G(z)}$ increases by $d\varepsilon_P$. How does this shift affect z ?

Taking the derivative, as a first-order approximation, we obtain:

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

Hence:

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

This term is negative: if retail shopping becomes relatively more expensive for higher z 's, 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:

$$\frac{d \log z}{d\varepsilon_P} = - \frac{\eta_w s_G^2}{1 - s_G}$$

where $s_G = \frac{w-z}{w}$ denotes the share of retail in consumer expenditures.

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

$$\Delta \Delta \log P = [\log P_1(z_1) - \log P_1(z_0)] - [\log P_0(z_1) - \log P_0(z_0)]$$

comparing different quintiles of consumers, z_1 and z_0 , across different sets of prices, $P_1(\cdot)$ and $P_0(\cdot)$ (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):

$$\varepsilon_P(z_1) - \varepsilon_P(z_0) = \frac{zP'_1(z)}{P_1(z)} - \frac{zP'_0(z)}{P_0(z)} \approx \frac{\Delta\Delta \log P}{\log z_0 - \log z_1}$$

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:

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

Numerical Application: Using the following estimates:

- $\eta_z \approx 2$ (upper bound, to be conservative)
- $\Delta\Delta \log P \approx -0.03$ (maximum across all counterfactuals)
- $\log z_0 - \log z_1 \approx \log 100,000 - \log 10,000 \approx 2.3$
- Share of retail shopping $s_G \approx 0.3$

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

$$d \log z \approx 2 \times 0.13 \times 0.03 / 2.3 \approx 0.003$$

Hence, quantitatively, the changes in z induced by the price changes in our counterfactuals are very small and negligible relative to the changes in nominal income distribution in the first set of counterfactuals. They would also have a negligible impact on sales, quality and prices relative to the changes obtained in our counterfactual shocks.

Appendix 4: Equivalent Discrete-Choice Model

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:

$$U_{Gz} = \prod_n \left[\sum_{i \in G_n} (q_{zni} \varphi_{zni})^{\frac{\sigma_{nz}-1}{\sigma_{nz}}} \right]^{\alpha_{nz} \cdot \frac{\sigma_{nz}}{\sigma_{nz}-1}} \quad (30)$$

Instead, suppose that individual j from income group z has utility:

$$U_{jz} = \sum_n \alpha_{nz} \max_{i \in G_n, q_{jzni}} [\log q_{jzni} + \log \varphi_{zni} + \mu_{nz} \epsilon_{jzni}] \quad (31)$$

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

$$\sum_n y_{jzn} \leq E_z$$

$$\sum_{i \in G_n} q_{jzni} p_{ni} \leq y_{jzn}$$

where E_z refers to total income allocated to grocery shopping for consumers of income group z . In expression 31 above, $\log \varphi_{zni}$ is a quality shifter associated with product z in module n that is specific to income group z . In turn, the last term $\mu_{nz} \epsilon_{jzni}$ 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 $\{y_{jzn}\}_n$ of expenditures in each module n , the good i^* being chosen maximizes:

$$i^* = \operatorname{argmax}_{i \in G_n} [\log y_{jzn} - \log p_{ni} + \log \varphi_{zni} + \mu_{nz} \epsilon_{jzni}]$$

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

$$i^* = \operatorname{argmax}_{i \in G_n} [-\log p_{ni} + \log \varphi_{zni} + \mu_{nz} \epsilon_{jzni}] \quad (32)$$

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

$$\begin{aligned} U_{jz} &= \max_{\{y_{jzn}\}} \left\{ \sum_n \alpha_{nz} \max_{i \in G_n} [\log y_{jzn} - \log p_{ni} + \log \varphi_{zni} + \mu_{nz} \epsilon_{jzni}] \right\} \\ &= \max_{\{y_{jzn}\}} \left\{ \sum_n \alpha_{nz} \log y_{jzn} \right\} + \sum_n \alpha_{nz} \max_{i \in G_n} [-\log p_{ni} + \log \varphi_{zni} + \mu_{nz} \epsilon_{jzni}] \end{aligned}$$

which leads to y_{jzn} being equal to a fraction α_{nz} of income E_z spent on grocery shopping (for consumers in income group z):

$$y_{jzn} = \alpha_{nz} E_z$$

Note that this independence property does not hold in Handbury (2013). Handbury (2013) assumes an elasticity of substitution different from unity across product modules n , which implies that the amount spent on each product model depends on the set of specific shocks ϵ_{jzni} of each consumer j . This would render the discrete-choice version of Handbury (2013) analytically intractable.

Using this property and additional assumptions on the distribution of shocks ϵ_{jzni} , 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_{jzni} is i.i.d. and drawn from a Gumbel distribution (type-II extreme value distribution) as in Anderson et al (1987). Equation 32 implies that a share:

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

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

$$x_{zni} = \frac{\left(\frac{\varphi_{zni}}{p_{ni}}\right)^{\sigma_{nz}-1}}{\sum_{i' \in G_n} \left(\frac{\varphi_{zni'}}{p_{ni'}}\right)^{\sigma_{nz}-1}} \alpha_{nz} E_z$$

where $\sigma_{nz} = 1 + \frac{1}{\mu_{nz}}$ denotes the elasticity of substitution between goods i on aggregate for

consumers of income group z . This shows that utility described in equation 31 is exactly equivalent to the consumption patterns obtained with the preferences described in equation 30 above and equation 2 in the main text.

Appendix 5: Model Extensions

This appendix extends the model in three dimensions that we present in subsections 5.A, 5.B and 5.C in the following.

5.A) Oligopolistic Competition

As before (Appendix 2), 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:

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

Now, we also allow firms to take into account the effect of their own price on the aggregate price index for the product module, $P_n(z)$. Note that the elasticity of $P_n(z)$ w.r.t μ is equal to the market share of the firm among consumers of income z , which we denote by $s_n(a, z, \mu, \phi)$:

$$\frac{\partial \log P_n(z)}{\partial \log \mu} = s_n(a, z, \mu, \phi) = \frac{x_n(a, z, \mu, \phi)}{\alpha_n(z) E(z)}$$

Hence, for markups μ , the first-order condition becomes:

$$\begin{aligned} 0 = \frac{\partial \pi_n}{\partial \log \mu} &= \left(1 - \frac{1}{\mu}\right) \int_z (\sigma_n(z) - 1) x_n(a, z, \mu, \phi) \frac{\partial \log P_n(z)}{\partial \log \mu} dH(z) \\ &\quad - \left(1 - \frac{1}{\mu}\right) \int_z (\sigma_n(z) - 1) x_n(a, z, \mu, \phi) dH(z) \\ &\quad + \frac{1}{\mu} \int_z x_n(a, z, \mu, \phi) dH(z) \end{aligned}$$

It follows that optimal markups now satisfy:

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

and thus we obtain the expression in the text:

$$\frac{p - c}{p} = 1 - \frac{1}{\mu} = \frac{\int_z x_n(a, z, \mu, \phi) dH(z)}{\int_z (\sigma_n(z) (1 - s_n(a, z, \mu, \phi)) + s_n(a, z, \mu, \phi)) x_n(a, z, \mu, \phi) dH(z)} \equiv \frac{1}{\bar{\sigma}_n^{MP}(a)}$$

As in Hottman et al. (2016), markups equal to the inverse of a weighted average of the elasticity of substitution between varieties within the industry, $\sigma_n(z)$, and the upper-tier elasticity of substitution, here set to unity (Cobb-Douglas). In addition to Hottman et al. (2016), here we have

heterogeneous elasticities of substitution that firms face across consumers with different incomes, but we still obtain a relatively simple formula.

For quality ϕ , the first order condition also differs from the baseline model since firms may internalize the effect of their quality choice on the aggregate price index. This effect is again determined by market share:

$$\frac{\partial \log P_n(z)}{\partial \log \phi} = (\gamma_n(z) - \xi_n) s_n(a, z, \mu, \phi)$$

We obtain the following first-order condition in quality:

$$\begin{aligned} 0 &= \frac{\partial \pi_n}{\partial \log \phi} \\ &= \left(1 - \frac{1}{\mu}\right) \int_z (\sigma_n(z) - 1)(\gamma_n(z) - \xi_n)(1 - s_n)x_n dH(z) - b_n \phi^{\frac{1}{\beta_n}} \end{aligned}$$

With $1 - \frac{1}{\mu} = \frac{\int_z x_n(a, z, \mu, \phi) dH(z)}{\int_z (\sigma_n(z)(1 - s_n) + s_n)x_n dH(z)}$, we obtain optimal quality.

Defining $\tilde{\gamma}_n^{MP}(a)$ as:

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

and $\tilde{\rho}_n^{MP}(a)$ as:

$$\tilde{\rho}_n^{MP}(a) = \frac{\tilde{\sigma}_n^{MP}(a) - 1}{\tilde{\sigma}_n^{MP}(a)} = \frac{\int_z (\sigma_n(z) - 1)(1 - s_n)x_n dH(z)}{\int_z (\sigma_n(z)(1 - s_n) + s_n)x_n dH(z)},$$

we obtain:

$$\phi_n(a) = \left(\frac{1}{b_n} \tilde{\rho}_n^{MP}(a) X_n(a) \left(\tilde{\gamma}_n^{MP}(a) - \xi_n \right) \right)^{\beta_n}.$$

5.B) Multi-Product Firms

Firm Heterogeneity Across Consumption Baskets

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: $\sum_i \log \varphi_{nbi}^{MP}(z) = 0$.

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

$$x_{nb}^{Tot}(z) = \left(\frac{\varphi_{nb}^{Tot}(z)}{P_{nb}^{brand}(z)} \right)^{\sigma_n(z)-1} \alpha_n(z) E(z) P_n(z) \sigma_n(z)-1 \quad (33)$$

while sales by product can be written as:

$$x_{nbi}^{MP}(z) = \left(\frac{\varphi_{nbi}^{MP}(z)}{p_{ni}} \right)^{\eta_n(z)-1} x_{nb}^{Tot}(z) P_{nb}^{brand}(z) \eta_n(z)-1 \quad (34)$$

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

$$P_n(z) = \left[\sum_{i \in G_n} P_{nb}^{brand}(z)^{1-\sigma_n(z)} \varphi_{ni}^{Tot}(z)^{\sigma_n(z)-1} \right]^{\frac{1}{1-\sigma_n(z)}} \quad (35)$$

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

$$P_{nb}^{brand}(z) = \left[\sum_{i \in G_n} p_{ni}^{1-\eta_n(z)} \varphi_{nbi}^{MP}(z)^{\eta_n(z)-1} \right]^{\frac{1}{1-\eta_n(z)}} \quad (36)$$

When price elasticities $\eta_n(z)$ and $\sigma_n(z)$ (within and across brands) differ, this new definition of a brand's price index differ from traditional sales weighted price indexes (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 does not depend on the number of product varieties. 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:

$$x_{nb}^{Tot}(z) = N_{nb}^{\frac{\sigma_n(z)-1}{\eta_n(z)-1}} \left(\frac{\varphi_{nb}^{Tot}(z)}{\bar{P}_{nb}(z)} \right)^{\sigma_n(z)-1} \alpha_n(z) E(z) P_n(z)^{\sigma_n(z)-1} \quad (37)$$

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 no longer simply determined by a sales-weighted average of price elasticities because of 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:

$$\mu_{nbi} = \frac{\sum_z x_{nbi}(z)}{\sum_z \eta_n(z) x_{nbi}(z)} \left[1 + \frac{\sum_z (\eta_n(z) - \sigma_n(z)) \bar{\mu}_{nb}(z) x_{nbi}(z)}{\sum_z x_{nbi}(z)} \right] \quad (38)$$

or equivalently:

$$\mu_{nbi} = \frac{\sum_z x_{nbi}(z)}{\sum_z \sigma_n(z) x_{nbi}(z)} \left[1 + \frac{\sum_z (\eta_n(z) - \sigma_n(z)) (\bar{\mu}_{nb}(z) - \mu_{nbi}) x_{nbi}(z)}{\sum_z x_{nbi}(z)} \right] \quad (39)$$

In equation 38, 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 39, 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 $f_n(\phi_{nb}) = b_n \phi_{bn}^{\frac{1}{\beta_n}}$ are the fixed costs of quality upgrading) we obtain the following first-order condition in brand-level quality ϕ_{nb} :

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

$(\sigma_n(z) - 1) \gamma_n(z)$ reflects the effect of quality upgrading on demand, while ξ_n is the effect on costs. Using our expression above for average markups (equation 39), we obtain the following expression

for optimal quality that generalizes expression 11 for multi-product brands:

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

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.

5.C) Other Determinants of Product Quality across Firms

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.

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 11 and substituting b_n by $\psi^{-\frac{1}{\beta_n}}$, one can see that optimal quality is now:

$$\phi(a, \psi) = \psi \cdot (\tilde{\rho}_n(a, \psi) X_n(a, \psi) (\tilde{\rho}_n(a, \psi) - \xi_n))^{\beta_n}$$

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, counterfactual simulations remain very similar in this case, given that ψ is held constant for each firm. We refer to Appendix 7.E for details on the description of the counterfactual equilibrium conditions.

Appendix 6: 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.

Note, however, that we could have instead defined these quintiles across all years (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.

In the remainder of this appendix, 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:

$$\log \phi_{ni} = \frac{1}{\lambda} \int_{\Omega_z} \log \varphi_{ni}(z) dz \quad (40)$$

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

$$\frac{x_{ni}(z)}{x_{nj}(z)} = \left(\frac{\phi_{ni}}{\phi_{nj}} \right)^{\gamma_n(z)(\sigma_n(z)-1)} \left(\frac{p_{ni}}{p_{nj}} \right)^{1-\sigma_n(z)} \quad (41)$$

$$P_n(z) = \left[\sum_{i \in G_n} p_{ni}^{1-\sigma_n(z)} \phi_{ni}^{\gamma_n(z)(\sigma_n(z)-1)} \right]^{\frac{1}{1-\sigma_n(z)}} \quad (42)$$

$$\phi_n(a) = \left(\frac{1}{b_n} \cdot \tilde{\rho}_n(a) \cdot X_n(a) \cdot (\tilde{\gamma}_n(a) - \xi_n) \right)^{\beta_n} \quad (43)$$

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 indexes 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:

$$\gamma'_n(z) = \gamma_n(z)/\lambda \quad (44)$$

$$\beta'_n = \beta_n \lambda \quad (45)$$

$$b'_n = b_n/\lambda \quad (46)$$

$$\xi'_n = \xi_n/\lambda \quad (47)$$

$$\phi'_{ni} = (\phi_{ni})^\lambda \quad (48)$$

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 41-43 hold. For instance, quality satisfies:

$$\phi'_n(a) = \phi_n(a)^\lambda = \left(\frac{1}{b_n} \cdot \tilde{\rho}_n(a) \cdot X_n(a) \cdot (\tilde{\gamma}_n(a) - \xi_n) \right)^{\lambda \beta_n} = \left(\frac{1}{b'_n} \cdot \tilde{\rho}_n(a) \cdot X_n(a) \cdot (\tilde{\gamma}'_n(a) - \xi'_n) \right)^{\beta'_n} \quad (49)$$

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

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

$$= \frac{1}{\bar{\sigma}_n(a)} \left[\int_z (1 - \beta'_n(\gamma'_n(z) - \xi'_n)(\sigma_n(z) - 1)) x_n(a, z) dH(z) \right] - f_{0n}$$

- The estimation of supply side parameters also remains the same since they are directly obtained from equation 49.
- Counterfactual equilibrium conditions (see below) also remain the same. For instance, for quality upgrading and sales, we have:

$$\frac{\phi'_{n1}(a)}{\phi'_{n0}(a)} = \left(\frac{\phi_{n1}(a)}{\phi_{n0}(a)} \right)^\lambda = \left[\frac{\int_z (\gamma_n(z) - \xi_n) x_{n1}(z, a) dH(z)}{\int_z (\gamma_n(z) - \xi_n) x_{n0}(z, a) dH(z)} \right]^{\lambda \beta_n} = \left[\frac{\int_z (\gamma'_n(z) - \xi'_n) x_{n1}(z, a) dH(z)}{\int_z (\gamma'_n(z) - \xi'_n) x_{n0}(z, a) dH(z)} \right]^{\beta'_n} \quad (50)$$

Appendix 7: Counterfactuals and Decompositions

We use our theoretical framework to explore two types of counterfactuals. The first set of counterfactuals is to exogenously increase nominal income inequality across consumers. These counterfactuals illustrate how changes in the income distribution $H(z)$ affect the demand and supply of product quality, and how these changes feed back into consumer inflation and real income inequality. Our second set of counterfactuals explores the gains from trade in a setting with heterogeneous firms where households source their consumption differently across the firm size distribution, as observed in our microdata. Here, we focus on a conventional [Melitz \(2003\)](#) framework with two symmetric countries where firms can export to an additional market by paying a fixed cost $f_{nX} > 0$ and variable iceberg trade costs $\tau_n > 1$. The first part of this section describes the 5 equilibrium conditions to solve for counterfactual outcomes. The remaining part provides additional details on the decompositions of the price index changes.

7.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, where we model changes in nominal inequality, we denote the initial cumulative distribution of z by $H_0(z)$ and we denote by $H_1(z)$ the counterfactual income distribution. In the second counterfactual where we introduce fixed trade costs f_{nX} and iceberg trade costs τ_n , we denote by $\delta_n^X(a)$ an export dummy equal to one if firm a exports in the counterfactual equilibrium. In all equations below, $\delta_n^X(a)$ is implicitly equal to zero for the first set of counterfactuals. 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. The trade counterfactual has an additional condition reflecting the decision to export.

First, the evolution of firm sales for a given income group z depends on quality upgrading and the price index change for each consumer group:

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

where $\tilde{\rho}_n(a)$ corresponds to a weighted average of $\rho_n(z)$ among firm a 's consumers weighting by either sales in the baseline equilibrium ($\tilde{\rho}_{n0}$) or sales in the counterfactual equilibrium ($\tilde{\rho}_{n1}$). 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:

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

accounting for changes in prices: $p_n(a) = \frac{\phi_n(a)\xi_n}{a\tilde{\rho}_n(a)}$ and adjustments in exit and exporter status. Note that the export dummy $\delta_n^X(a)$ is equal to zero for all firms in the first set of counterfactuals where there is no change in trade costs. Based on initial sales $x_{n0}(z, a)$ and the new distribution of income $H_1(z)$, total sales of firm a in the counterfactual equilibrium are then given by $X_{n1}(a) = \int_z x_{n1}(z, a) dH_1(z)$.

Next, equation 11 implies that quality upgrading is determined by:

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

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. 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.⁴⁵ This equation is the same in both types of counterfactuals.

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 export status and exit of firms in the counterfactual equilibrium, we obtain:

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

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 $p_n(a) = \frac{\phi_n(a)\xi_n}{a\tilde{\rho}_n(a)}$ for prices, we obtain our main equation describing the change in price indexes in our counterfactual equilibrium:

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

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. Moreover, one needs to account for the imports of

⁴⁵Conditional on entry and exit, these first two equilibrium relationships offer a contraction mapping that we exploit to solve the counterfactual, provided that $\beta_n(\sigma_n(z) - 1)(\gamma_n(z) - \xi_n)$ is strictly smaller than unity for all z .

new product varieties in the trade counterfactual. Assuming symmetry between the domestic and foreign economies, this additional margin is captured by the term $(1 + \delta_n^X(a)\tau_n^{1-\sigma_n(z)})$.

The entry, exit and export 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:

$$\int_a \frac{1}{\tilde{\sigma}_{n0}(a)} [1 - \beta_n (\tilde{\sigma}_{n0}(a) - 1) (\tilde{\gamma}_{n0}(a) - \xi_n)] X_{n0}(a) dG_n(a) = \int_a \frac{1}{\tilde{\sigma}_{n1}(a)} \delta_{nD}(a) [1 - \beta_n (\tilde{\sigma}_{n1}(a) - 1) (\tilde{\gamma}_{n1}(a) - \xi_n)] X_{n1}(a) dG_n(a) + \int_a (1 - \delta_{nD}(a)) f_{n0} dG_n(a)$$

The number of firms N_{n1} adjusts such that this equality holds.

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

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

In the trade counterfactual, 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:

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

where $r_n^X(a, \phi_{n1}^X(a))$ denotes revenues net of variable costs on the export market (exports times $\frac{1}{\tilde{\sigma}_n}$) 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.

7.B) Decomposition of the Price Index Effect

For a given income group z , the price index change equals:

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

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

Taking logs and a first-order approximation leads to:

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

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

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

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:

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

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 $\bar{\delta}_{nD} = \int_a \delta_{nD}(a) \bar{s}_{n0}(a) dG(a)$ and $\bar{\delta}_{nD} (1 + \bar{\delta}_X \tau_n^{1-\bar{\sigma}_n})$

the average of $\delta_{nD}(a)(1 + \delta_X(a)\tau_n^{1-\sigma_n(z_0)})$ across consumers and firms, we obtain the five-term decomposition described in the text.

7.C) Effect of a Non-Uniform Tax across Firms

In this counterfactual we examine the heterogeneous effect across consumers of a tax imposed heterogeneously across firms. If each firm held its markup and quality constant, the differential effect of such a tax across consumers would correspond to:

$$\log \frac{P_{n1}(z)}{P_{n0}(z)} - \log \frac{P_{n1}(z_0)}{P_{n0}(z_0)} = - \int_a (s_{n1}(a,z) - s_{n1}(a,z_0)) \log(1 + t_n(a)) dG_n(a)$$

comparing consumers of income group z and z_0 , where $s_{n1}(a,z)$ is the market share of firm a within income group z and $t_n(a)$ is the tax applied to firm a .

This direct partial-equilibrium effect however does not account for general-equilibrium adjustments. In what follows, we describe how sales, quality and prices adjust to these heterogeneous taxes.

Sales: Relative to the baseline equilibrium, prices equal:

$$p_n(a) = (1 + t_n(a)) \frac{\phi_n(a)^{\xi_n}}{a\tilde{\rho}_n(a)}$$

where t refers to the tax applied to firm a .

Firm sales equal:

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

Hence, taking ratios:

$$\frac{x_{n1}(z, a)}{x_{n0}(z, a)} = \left(\frac{P_{n1}(z)}{P_{n0}(z)} \right)^{\sigma_n(z)-1} \left(\frac{\tilde{\rho}_{n1}(a)}{\tilde{\rho}_{n0}(a)} \right)^{\sigma_n(z)-1} \left(\frac{\phi_{n1}(a)}{\phi_{n0}(a)} \right)^{(\sigma_n(z)-1)(\gamma_n(z)-\xi_n)} (1 + t_n(a))^{1-\sigma_n(z)}$$

(for surviving firms).

Quality: For optimal quality, we obtain the same equation:

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

Price index: As previously, we have:

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

Using the expression $p_n(a) = (1 + t_n(a)) \frac{\phi_n(a)^{\xi_n}}{a\tilde{\rho}_n(a)}$ for prices, we get:

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

Entry and exit: Equations are the same as for the first counterfactual as described above.

Decomposition: We can use the same decomposition as for previous counterfactuals, with one small change. Given that prices not only depend on quality and markups but also taxes, we add a seventh decomposition term that captures the direct effect of taxes (first line of the following decomposition) and corresponds to the partial-equilibrium effect described earlier:

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

7.D) Counterfactuals in Model Extension with Oligopoly Markups

Prices: Denote:

$$\varepsilon_n^{MP}(a, z) = \sigma_n(z)(1 - s_n(a, z)) + s_n(a, z)$$

where $s_n(a, z)$ denotes the market share of firm a among consumers of income z . This leads to the following weighted average of the price elasticity for brand i :

$$\tilde{\sigma}_{ni}^{MP} = \frac{\int_z \varepsilon_{ni}(z) x_{ni}(z) dH(z)}{\int_z x_{ni}(z) dH(z)} = \frac{\int_z [\sigma_n(z)(1 - s_n(a, z)) + s_n(a, z)] x_{ni}(z) dH(z)}{\int_z x_{ni}(z) dH(z)}$$

and the following markups (see Appendix 5.A):

$$\frac{p_{ni}}{c_{ni}} = \frac{\tilde{\sigma}_{ni}^{MP}}{\tilde{\sigma}_{ni}^{MP} - 1} \equiv \frac{1}{\tilde{\rho}_{ni}^{MP}}$$

Sales: Sales equal:

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

Hence, with these new prices $p_n(a) = \frac{\phi_n(a)^{\xi_n}}{a \tilde{\rho}_n^{MP}(a)}$, after taking ratios we obtain:

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

(for surviving firms).

Quality: With these new markups, we have shown in Appendix 5.A that optimal quality satisfies:

$$\begin{aligned}\phi_n(a) &= \left[\frac{1}{b_n \tilde{\sigma}_n^{MP}(a)} \int_z (\gamma_n(z) - \xi_n) (\sigma_n(z) - 1) (1 - s_n(a, z)) x_n(a, z) dH(z) \right]^{\beta_n} \\ &= \left[\frac{1}{b_n} \cdot \tilde{\rho}_n^{MP}(a) \cdot X_n(a) \cdot \left(\tilde{\gamma}_n^{MP}(a) - \xi_n \right) \right]^{\beta_n}\end{aligned}$$

with $\tilde{\gamma}_n^{MP}(a)$ is a weighted average of $\gamma_n(z)$ defined as:

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

Hence, quality upgrading is determined by:

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

where both $\tilde{\gamma}_n(a)$ and $\tilde{\rho}_n(a)$ correspond to weighted averages of $\gamma_n(z)$ and $\rho_n(z)$ among firm a 's consumers, weighting by either sales in the baseline equilibrium ($\tilde{\gamma}_{n0}(a)$ and $\tilde{\rho}_{n0}(a)$) or sales in the counterfactual equilibrium ($\tilde{\gamma}_{n1}(a)$ and $\tilde{\rho}_{n1}(a)$).

Entry and Exit: We treat entry and exit the same as in the baseline case, except that profits are now given by:

$$\pi_n^{MP}(a) = \frac{1 - \beta_n (\gamma_n^{MP}(z) - \xi_n) (\sigma_n^{MP}(z) - 1)}{\tilde{\sigma}_n^{MP}(a)} X_n(a) - f_{0n}$$

Consumer Price Indexes: Hence, taking ratios of prices and a weighted average across consumers, the change in the price index for income group z is now given by:

$$\frac{P_{n1}(z)}{P_{n0}(z)} = \left[\frac{N_{n1} \int_a x_{n0}(z, a) \delta_{nD}(a) \left(\frac{\tilde{\rho}_{n1}^{MP}(a)}{\tilde{\rho}_{n0}^{MP}(a)} \right)^{\sigma_n(z)-1} \left(\frac{\phi_{n1}(a)}{\phi_{n0}(a)} \right)^{(\sigma_n(z)-1)(\gamma_n(z)-\xi_n)} dG(a)}{N_{n0} \int_a x_{n0}(z, a) dG(a)} \right]^{\frac{1}{1-\sigma_n(z)}}$$

7.E) Counterfactuals in Model Extension with Heterogeneous Quality Shifters

Here we briefly describe counterfactual equilibrium equations for the model extension with heterogeneous quality terms ψ described in Section 4.3 and Appendix 5.C. 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 18. 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:

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

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:

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

where $\tilde{\rho}_n(a, \psi)$ is itself the sales-weighted average of $\rho_n(z)$ across consumers of firm with productivity a and quality shifter ψ .

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

$$\pi_n(a, \psi) = \frac{1 - \beta_n (\gamma_n(z) - \xi_n)(\sigma_n(z) - 1)}{\tilde{\sigma}_n(a)} X_n(a) - f_{0n}$$

It follows that entry and exit equilibrium conditions also remain unchanged.