

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

Benjamin Faber[†] and Thibault Fally[‡]

February 2016

PRELIMINARY AND INCOMPLETE

Abstract

A growing literature has emphasized the role of Melitz-type firm heterogeneity within sectors in accounting for nominal income inequality. This paper explores the implications of firm heterogeneity for household price indices across the income distribution. Using detailed matched US home and store scanner microdata that allow us to trace the firm size distribution into the consumption baskets of individual households, we present evidence that richer US households source their consumption from on average significantly larger producers of brands within disaggregated product groups compared to poorer US households. We use the microdata to explore alternative explanations, write down a quantitative framework that rationalizes the observed moments, and estimate its parameters to quantify model-based counterfactuals. Our central findings are 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. We find that these price index effects significantly amplify observed changes in nominal income inequality over time, and that they lead to a significantly more regressive distribution of the gains from international trade.

Keywords: Firm heterogeneity, real income inequality, household price indices, scanner data
JEL Classification: D22, F61, O51

*We would like to thank Cecile Gaubert, James Harrigan, Thierry Mayer, Volker Nocke, Andres Rodriguez-Clare, Nicolas Schultz, Eric Verhoogen, Jon Vogel and participants at multiple seminars and conferences for helpful comments and discussions. Dmitri Koustas provided outstanding research assistance.

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

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

1 Introduction

Over the past four decades, nominal 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). A growing literature has emphasized the role of Melitz-type firm heterogeneity within sectors in accounting for the observed changes in nominal wage inequality in a variety of different empirical settings, ranging from the US (Bloom et al., 2015) and Germany (Card et al., 2013) to Brazil (Helpman et al., 2012) and Mexico (Frias et al., 2009). Theoretically, the literature on firm heterogeneity and inequality has focused on the extent to which higher and lower-income workers source their earnings from different parts of the firm size distribution, and the implications thereof.¹

Rather than focusing on nominal incomes, this paper sets out to explore the implications of firm heterogeneity for household price indices across the income distribution. We aim to contribute to our understanding of three central questions: i) To what extent is it the case that rich and poor households source their consumption baskets from different parts of the firm size distribution?; ii) What explains the observed differences in weighted average firm sizes embodied in the consumption baskets of rich and poor households?; and iii) What are the implications of the answers to i) and ii) for real income inequality?

In answering these questions, the paper makes two main contributions to the existing literature. 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 within highly disaggregated product groups, and explore alternative explanations. Second, we write down a quantitative model of heterogeneous firms and consumers that rationalizes the observed moments, and estimate its parameters using the microdata to quantify the underlying channels and explore model-based counterfactuals.

The analysis presents three central findings. First, we document that the richest 20 (10) percent of US households source their consumption from on average more than 25 (30) percent larger producers of brands within highly disaggregated product groups compared to the poorest 20 (10) percent of US households. Second, we find that this is due to the endogenous sorting of larger, more productive firms into catering to the taste of wealthier households. This is mainly driven by two features of household preferences and firm technologies that we estimate empirically. On the consumption side, we find that rich and poor households on average strongly agree on their ranking of quality evaluations across products, but that higher income households value higher quality attributes significantly more. On the production side, producing attributes that households evaluate as higher quality increases both the marginal as well as fixed costs of production. In combination, these two features give rise to the endogenous sorting of larger, more productive firms into products that are valued relatively more by wealthier households. Our third finding is that these results have important implications for inequality due to asymmetric effects on household price indices. We find that these price index effects significantly amplify observed changes in nominal income

¹For example Helpman et al. (2010) and Davis & Harrigan (2011) focus on differences in wage premia across the firm size distribution for homogeneous workers. On the other hand, e.g. Harrigan & Reshef (2011) and Sampson (2014) focus on differences in skill intensity across the firm size distribution. See related literature at the end of this section for further discussion.

inequality over time, and that they lead to a significantly more regressive distribution of the gains from international trade. Underlying these results, we quantify a rich and novel interplay of microeconomic channels including asymmetric firm adjustments to product quality, markups, entry and exit across the firm size distribution that affect household consumption baskets differently.

At the center of the analysis lies the construction of an extremely rich collection of microdata that allows us to trace the firm size distribution into the consumption baskets of rich and poor households across the income distribution. We combine a dataset of 270 million consumer transactions when aggregated to the household-by-store-by-barcode-by-semester level from the AC Nielsen US Home Scanner data over the period 2006-2012, with a dataset of 9.3 billion store transactions when aggregated to the store-by-barcode-by-semester level from the AC Nielsen US Retail Scanner data covering the same 14 semesters. The combination of both 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 25,000 retail establishments each semester) into the consumption baskets of on average 58,000 individual households within more than 1000 disaggregated retail product modules (such as carbonated drinks, shampoos, pain killers, desktop printers or microwaves).²

The analysis proceeds in four steps. In Step 1, we use the scanner microdata to document a new set of stylized facts. We estimate large and statistically significant differences in the weighted average firm sizes that rich and poor households source their consumption from.³ This finding holds across all product departments covered by the Nielsen data and for all years in the dataset. We also document that these differences in firm sizes across consumption baskets arise in a setting where the rank order of household budget shares spent across different producers is preserved between rich and poor households –i.e. the largest firms command the highest budget shares for all income groups. After exploring a number of alternative data-driven or mechanical explanations using the microdata, we interpret these stylized facts as equilibrium outcomes in a setting where both consumers and firms optimally choose their product attributes over time.

In Step 2 we write down a 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 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 with ex ante heterogeneous firms within sectors. We use the model to derive estimation equations for the key preference and technology parameters as a function of observable moments in the home and store scanner microdata. Armed with these parameter estimates in addition to the raw moments in our microdata, we can quantify the role of different channels, and use our framework to explore general equilibrium counterfactuals. The remaining two 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

²The Nielsen data are made available through an academic user agreement with the Kilts Center at Chicago Booth.

³Household nominal incomes in the Nielsen data are only reported with 2-year lags and in broad brackets. As we focus on quintiles of the household income distribution, we also verify in Appendix Figure A.1 that our measure of expenditure per capita is monotonically increasing in reported nominal incomes.

and in terms of their valuation of product quality attributes. We find that poorer households have statistically significantly higher price elasticities relative to higher income households, but that these differences are relatively minor in terms of magnitudes. We also find that while households on average agree on the ranking of quality evaluations across producers given prices, richer households value higher quality attributes significantly more. On the production side, we estimate that producing attributes that all households evaluate as higher quality significantly increases both the marginal as well as the fixed costs of production.

To estimate the technology parameters, we follow two different estimation strategies. The first is based on cross-sectional variation in brand quality and firm scale, and the second exploits observed changes in brand quality and scale over time. Given that firm adjustments to product quality in response to changes in demand (scale and consumer composition) are likely best understood as a longer-term effect, we think of the panel data approach as more conservative. To identify the effect of firm scale on quality upgrading, we use state-level measures of changes in brand quality on the left hand side, and exploit plausibly exogenous changes in firm scale among other US regions on the right hand side. In particular, we construct a shift-share instrument for changes in national firm scale that uses variation in pre-existing brand-level sales shares across other US states interacted with state-level variation in average sales growth among other product groups.

These estimates give rise to two opposing forces that determine both firm sizes across consumption baskets and the sorting of larger firms across quality attributes. On the one hand, larger firms offer lower quality-adjusted prices, which increases the share of their sales coming from more price elastic lower income consumers. Since these consumers value quality less, this channel, *ceteris paribus*, leads poorer households to source their consumption from on average larger firms producing 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 that is valued relatively more by wealthier households. Empirically, we find that this second channel by far outweighs 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 parameter estimates, we find that the observed moments from step 1 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 (10) percent of US households source their consumption from on average more than 35 (40) percent higher quality producers compared to the poorest 20 (10) percent of households. At the same time, we find that the richest income quintile (decile) source their consumption at on average 20 (25) percent lower quality adjusted prices. We also find that markups endogenously vary across the firm size distribution: Because the sales of larger firms are driven to a larger extent by richer households who are less price sensitive, markups within product groups monotonically increase with firm size.

In the final Step 4, we use the parameterized model in combination with raw moments from the microdata to explore a number of general equilibrium counterfactuals. In our first counterfactual, we find that an exogenous increase in nominal income inequality leads to a significant endogenous amplification in terms of real income inequality due to asymmetric general equilibrium effects on household price indices. In particular, we find that a 5 percent transfer of market expenditure

from the poorest to the richest household quintile gives rise to a 2.4-3 percentage point lower cost of living inflation for the richest quintile compared to the poorest.

This amplification is driven by a rich interplay of underlying channels. The first is that firms on average have incentives to upgrade their product quality since more of total sales are now in the hands of households with higher quality evaluations. Given the estimated preference parameters on relative tastes for quality, this channel significantly decreases consumer inflation for richer households compared to the poor. The second effect is that the scale of production changes asymmetrically across higher and lower quality producers. 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 firms with lower changes in quality adjusted prices compared to the poor. 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 is concentrated among low quality producers, which again tends to work in favor of relatively less cost of living inflation among higher income households.

In our second counterfactual, we quantify the distribution of the gains from trade. We find that a 10 percent increase in import penetration between two symmetric countries leads to a 2.7-3.6 percentage point lower cost of living inflation for the richest 20 percent of US households compared to the poorest 20 percent. This effect arises because, as in [Melitz \(2003\)](#), heterogeneous producers respond differently to trade cost shocks, but in a setting where it is 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 richer income households more than the poor. 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. Overall, these findings illustrate a number of new adjustment channels that in both counterfactuals significantly amplify changes in observed nominal income inequality due to asymmetric price index effects across the income distribution.

This paper is related to the large and growing literature on the extent, causes and consequences of firm heterogeneity within sectors that has spanned different fields in economics, including international trade ([Bernard et al., 2007](#); [Melitz, 2003](#)), industrial organization ([Bartelsman et al., 2013](#)), macroeconomics ([Hsieh & Klenow, 2009](#)), development ([Peters, 2013](#)), labor economics ([Card et al., 2013](#)) and management ([Bloom & Van Reenen, 2007](#)). Within this literature, our paper is most closely related to existing work on the implications of firm heterogeneity for nominal wage inequality ([Bloom et al., 2015](#); [Burststein & Vogel, 2015](#); [Card et al., 2013](#); [Frias et al., 2009](#); [Helpman et al., 2012, 2010](#); [Sampson, 2014](#); [Verhoogen, 2008](#)). Relative to existing work

in this area, this paper presents empirical evidence that the widely documented presence of firm heterogeneity within sectors translates asymmetrically into the consumption baskets of rich and poor households across the income distribution, quantifies the underlying channels and explores the implications for real income inequality.

The paper is also closely related to recent work by [Hottman et al. \(2014\)](#) who use AC Nielsen’s US home scanner data to empirically decompose Melitz-type firm heterogeneity into differences in marginal costs, product quality, markups and the number of firm varieties. The paper is also related to recent work on endogenous quality choice across heterogeneous firms ([Johnson, 2012](#); [Kugler & Verhoogen, 2012](#); [Mandel, 2010](#); [Sutton, 1998](#)) as well as the literature on the linkage between trade and quality upgrading ([Bustos, 2011](#); [Verhoogen, 2008](#)). Relative to existing work, this paper departs from the representative agent assumption on the consumer side, and explores the implications that arise from firms’ profit maximizing product and markup choices in a setting where firm heterogeneity on the producer side interacts with household heterogeneity on the consumer side.

The paper also relates to existing work on the household 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 simulate heterogeneous consumer gains from trade. [Atkin et al. \(2015\)](#) use detailed microdata from Mexico to quantify the price index implications from foreign supermarket entry. Given our focus on relative prices within disaggregated product groups, this paper is closest in spirit to [Faber \(2014\)](#) who uses microdata from consumption surveys, plant surveys and CPI price surveys to estimate the effect of tariff reductions on the price of product quality in Mexican stores. Relative to existing work, this paper is the first to use newly available matched home and store scanner data to empirically trace the firm size distribution into the consumption baskets of individual households, and to combine this analysis with a quantitative framework that allows us to explore general equilibrium counterfactuals, including changes in trade costs.

Finally, the paper relates to the growing empirical literature using Nielsen consumption scanner data in economics ([Broda & Weinstein, 2010](#); [Handbury, 2014](#); [Handbury & Weinstein, 2014](#); [Hottman et al., 2014](#)). Most of this literature has relied on the Nielsen home scanner data covering approximately 58 thousand households in recent years. More recently, for example [Beraja et al. \(2014\)](#) have also used the store-level retail scanner data to estimate local price indices across US States. This paper leverages the combination of the two Nielsen datasets. This allows us to trace the national market shares of producers of brands across on average 25 thousand retail establishments located in more than 2500 US counties within disaggregated product groups into individual household consumption baskets.

The remainder of the paper is structured as follows. Section 2 describes the scanner microdata used in the estimations. 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 estimation of the preference and technology parameters, and uses these estimates to quantify the distribution of firm quality and quality adjusted prices across firms and

consumption baskets. Section 6 uses the model in combination with the microdata to explore counterfactuals. 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 2012. When a retail chain agrees to share their data, all of their stores enter the database. As a result, the database includes roughly 45,000 individual stores. The stores in the database vary in terms of the channel they represent: 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-semester level. On average each semester covers \$110 billion worth of retail sales across 25,000 individual stores in more than 1000 disaggregated product modules, 2500 US counties and across more than 700,000 barcodes belonging to 170,000 producers of brands. As described in more detail in the following section, we use these data in combination with the home scanner data described below in order to trace the distribution of firm sizes (in terms of national sales measured across on average 25k stores per semester) into the consumption baskets of individual households.

2.2 Home Scanner Data

We use the Home Scanner Database collected by AC Nielsen and made available through the Kilts Center at The University of Chicago Booth School of Business. 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-semester level. On average each semester covers \$105 million worth of retail sales across 58,000 individual households in more than 1000 disaggregated product modules, 2600 US counties and across more than 500,000 barcodes belonging to 180,000 producers of brands. One shortcoming of the home scanner dataset

is that nominal household incomes are measured inaccurately. 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 semester into percentiles of total retail expenditure per capita.⁴ 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 consumption microdatasets. 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 unfortunately only cover a tiny 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 semester. This paper combines national sales by product from the store scanner data, with the detailed information on individual household consumption shares in the home scanner data for the empirical analysis.

3 Stylized Facts

This section draws on the combination of 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 the home and store scanner microdata, 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 subsequent figures and tables, we define a firm as a producer of a unique brand within one of more than 1000 disaggregated product modules in the Nielsen data. This leads to an average number of firms active within a given product module of about 150. Two possible alternatives given our data would be to define a firm as a barcode product (leading to an average number of 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 reasons. Our objective is to define a producer within any 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 thousands of brands produced in hundreds of different establishments. The definition of firms at the barcode level would be problematic for the opposite reason, because the same establishment produces for example different pack sizes of the same product that are marked by different barcodes. For these reasons,

⁴Per 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).

we argue that 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 58 thousand households per semester) appears to be compressed relative to that measured using the store scanner data (25 thousand supermarkets per semester). Interestingly, 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 happen to be not picked up by any of the Nielsen sample households in a given half year period. Second, 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 dominance of the most popular brands: for example a household buying Coca Cola five 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 do not 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 are more likely. 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.

Figure 2 depicts the main stylized fact of the paper. Pooling repeated cross-sections across 14 semesters, we depict percentiles of household per capita expenditure (within each semester) on the x-axis and weighted average deviations of log firm sales from the product module-by-semester means on the y-axis. The underlying weights correspond to each household's retail consumption shares across all brands in all product modules consumed during the semester. 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 25 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 more than 30 percent when comparing the richest and poorest 10 percent of households.

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, across all of the listed product modules it is the case that 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. 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).

Figure 3 explores the heterogeneity of this pattern 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. As depicted in Figure 3, we find that the pattern of firm size differences across consumption baskets holds in each of these different product segments.

Finally, in Figure 4 we ask whether the 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 on average agree on their relative evaluations of quality-for-money across producers so that the rank order of their budget shares –facing the same relative prices– is preserved across the income distribution? Figure 4 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 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 4 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.

3.1 Alternative Explanations

One natural interpretation of these stylized facts is that they arise as equilibrium outcomes in a setting where both heterogeneous households and firms choose the product attributes they consume or produce. 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 stories.

Data-Driven Explanations One might be worried 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 underestimate 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 some important retail chains (notably Walmart) 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, Figure 5 re-estimates the 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. Regarding the “missing retailers” concern, we should also note that Figure 2 depicts very similar patterns when using 100 percent of household retail consumption paired with the national sales distribution computed using the home scanner data. While reassuring against the concern of missing retailers, the compression of the firm size distribution in the home scanner data also makes this comparison less than perfect, however.

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 heterogeneous household preferences, could be driving the results. In Figure 6 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-semester fixed effects when plotting the firm size deviations on the y-axis (keeping the x-axis exactly as before).⁵ 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).⁶ We find a very similar relationship compared to Figure 2, suggesting that differential access to producers is unlikely to be the driver.

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).⁷

⁵De-meaning both y and x-axis leads to almost identical point estimates.

⁶We 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.

⁷This 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

We document that in the medium or long run this notion seems hard to reconcile with either the raw moments in our data or the existing literature on endogenous quality choice by firms. A large 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. [Verhoogen \(2008\)](#), [Bastos et al. \(2014\)](#), [Dingel \(2015\)](#)). Another body of literature in support of this is of course the vast marketing literature on firm strategies using advertising to affect brand perceptions over time (e.g. [Keller et al. \(2011\)](#)). Furthermore, the scanner data suggest that it is a pervasive feature that producers of brands alter the physical characteristics and/or presentation of their products over time. Appendix Table [A.2](#) documents that each semester 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 varieties on offer by the same brand –clearly suggesting that producers are indeed capable of choosing their product attributes as a function of market conditions. In support of these descriptive moments, we also provide more direct empirical evidence on brand-level quality upgrading over time as part of our technology parameter estimation in Section [5](#).

However, it could still be the case that our 14 repeated cross-sections (semesters) depicted in Figure [2](#) are actually capturing the result of short-term taste shocks 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, what Figure [7](#) suggests is that the estimated differences in weighted average producer sizes are practically identical after replacing the measures of relative firm size by three-year lags or three-year leads for the same producers. These results indicate that the observed relationship in Figure [2](#) arises in a setting where both heterogeneous households and firms can endogenously choose their product attributes as equilibrium outcomes.

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 for all product departments 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 temporary taste shocks that differ across rich and poor households. The finding also arises in a setting where households on average appear to 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 theoretical framework that captures these observed moments in the microdata, and guides the empirical estimation.

differentiation.

4 Theoretical Framework

This section proposes a quantitative model that rationalizes the observed moments in the microdata and allows us to explore the implications for real income inequality. To this end, we introduce two important features into an otherwise standard Melitz model of ex ante 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 evaluations. On the producer side, differently productive firms face the observed distribution of consumer preferences and optimally choose their product attributes and markups.

4.1 Model Setup

4.1.1 Household Preferences

The economy consists of two broad sectors: retail shopping 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 :

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

For the sake of exposition, we do not explicitly specify the allocation of expenditures in retail vs. non-retail items but we assume that the consumption of the outside good is normal.⁸ 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.⁹ Utility from retail shopping is defined by:

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

where n refers to a product module in the Nielsen data and i refers to a specific brand producer within the product module.¹⁰ 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 varieties within each product module n at income level z . As we focus most of our attention to 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).¹¹

⁸[Handbury \(2014\)](#) estimates the income elasticity of retail consumption to be significantly positive (only slightly lower than one).

⁹We prefer to let the data speak to these relationships. Note that it would in principle be straight forward to microfound a positive relationship between the taste for product quality and household real incomes, by introducing a complementarity between outside good consumption and product quality within retail product groups (similar to for example [Fajgelbaum et al. \(2011\)](#)).

¹⁰We show in Appendix B that these preferences are equivalent to the aggregation of discrete-choice preferences across many agents choosing only one brand variety by product module.

¹¹Note 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 D presents an extension of our model to

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 the preferences specified above is that we do not impose any structure that dictates how price elasticities and quality valuations depend on income.¹²

Comparing two goods i and j within the same module n , 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 is constant across income groups (e.g. [Hottman et al. \(2014\)](#); [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. 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 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 $\int_z \gamma(z) dH(z) = 1$ (where $H(z)$ refers to the cumulative distribution of z across households), this intrinsic quality term also corresponds to the democratic average quality evaluation across households:

$$\log \phi_{ni} = \int_z \log \varphi_{ni}(z) dH(z). \quad (5)$$

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 can have different implications for the cost of living of households across the income distribution.

multi-product firms which we also discuss below.

¹²For 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.

4.1.2 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_E , as in Melitz (2003). For the remainder for Section 4, 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_E and discover their productivity a . Second, each entrepreneur decides at which level of quality to produce, or exit. Third, production occurs and markets clear.

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 the fixed cost of production to increase in the quality of the good being produced. This 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 production. 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, assuming that $\beta_n > 0$. Similarly, we let variable costs depend log-linearly on quality, with an elasticity ξ_n to capture the elasticity of the cost increase to the level of quality.¹³

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

We impose the restriction that ξ_n is smaller than the minimum quality evaluation $\gamma_n(z)$ in order to insure positive quality levels in equilibrium, as described 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:

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

¹³There is no need for a constant term as it would be isomorphic to a common productivity shifter after redefining $G_n(a)$.

$x_n(z, a)$ 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 to the fixed costs related to producing at higher quality. If we compare two firms with the same customer base, the larger one would more profitably invest in upgrading quality. Second, optimal quality depends on how much the 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.

When a firm sells to consumers from a single income group z , we obtain a simple expression to describe how quality varies with productivity:

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

Note that equilibrium requires $\beta_n (\sigma_n(z) - 1) (\gamma_n(z) - \xi_n) < 1$. Profits are then given by:

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

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

4.2.1 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, one can 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 how firm size correlates with quality and quality adjusted prices. The first term 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 tends to be negative (lower quality-adjust prices lead to larger sales when $\sigma_n(z) > 1$). If, in turn, 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 impose any assumption on the supply 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 quality satisfies equation 11. In particular, the correlation between firm size and quality appearing in the first term can be expressed as:

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

In our empirical section, we can then use our estimates to quantify each of these terms in order to de-compose the observed degree of firm heterogeneity across the consumption baskets of rich and poor households depicted in Figure 2.

4.2.2 Extension with Multi-Product Firms

Appendix D presents an extension of our model to multi-product firms. As recently emphasized by [Hottman et al. \(2014\)](#), 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 formally, that as long as the ratio of cross-brand to within-brand elasticities of substitution does not significantly differ across income groups, this additional dimension 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 lead to 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 the same ratio of cross-brand elasticities of substitution. Related to this insight, Appendix Table A.3 reports empirical evidence suggesting that the ratio of within-brand elasticities of substitution between rich and poor households does not significantly differ from the estimated ratio of cross-brand elasticities of substitution.

4.3 Counterfactual Equilibria

We use the model to explore two types of counterfactuals. The first counterfactual is to exogenously increase nominal income inequality by reallocating expenditure from the poorest to richest income quintile. This counterfactual illustrates 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. Our second counterfactual 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. We focus on a simple Melitz (2003) framework with two symmetric countries where firms can export to an additional market by paying a fixed cost $f_X > 0$ and variable iceberg trade costs $\tau > 1$.

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 counterfactual where we alter the income distribution, we denote the initial cumulative distribution of z (indexing nominal income levels) 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_X and iceberg trade costs τ , 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 counterfactual.

Comparing the initial and counterfactual equilibria, we find that the changes in firm sales, quality, entry, exit and price indices must satisfy the following five equilibrium conditions. The second counterfactual has an additional condition reflecting the decision to export.

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

$$\frac{x_{n1}(z, a)}{x_{n0}(z, a)} = \delta_{nD}(a) \left(1 + \delta_n^X(a)\tau^{1-\sigma_n}\right) \left(\frac{P_{n1}(z)}{P_{n0}(z)}\right)^{\sigma_n(z)-1} \left(\frac{\tilde{\rho}_{n1}(a)}{\tilde{\rho}_{n0}(a)}\right)^{\sigma_n(z)-1} \left(\frac{\phi_{n1}(a)}{\phi_{n0}(a)}\right)^{(\sigma_n(z)-1)(\gamma_n(z)-\xi_n)} \quad (16)$$

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 equation 16, 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 . Effects of prices on sales also vary across consumers

depending on their price elasticity $\sigma_n(z)$. Note that the export dummy $\delta_n^X(a)$ is equal to zero for all firms in the first counterfactual where there is no change in trade costs. Based on initial sales $x_{n0}(z, a)$ and the new distribution of income $H_1(z)$ (which differs from the baseline distribution in the first counterfactual), 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:

$$\frac{\phi_{n1}(a)}{\phi_{n0}(a)} = \left[\frac{(\tilde{\gamma}_{n1}(a) - \xi_n) \tilde{\rho}_{n1}(a) X_{n1}(a)}{(\tilde{\gamma}_{n0}(a) - \xi_n) \tilde{\rho}_{n0}(a) X_{n0}(a)} \right]^{\beta_n} \quad (17)$$

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. Equation 17 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 (highest counterfactual sales X_{n1}) also tend to upgrade their quality. This equation is the same in both counterfactuals.

Thirdly, the change in the price index $P_n(z)$ for each module n and income group z is determined by the change in quality weighted by initial sales of each firm:

$$\frac{P_{n1}(z)}{P_{n0}(z)} = \left[\frac{N_{n1} \int_a x_{n0}(z, a) \delta_{nD}(a) \left(1 + \delta_n^X(a) \tau^{1-\sigma_n(z)}\right) \left(\frac{\tilde{\rho}_{n1}(a)}{\tilde{\rho}_{n0}(a)}\right)^{\sigma_n(z)-1} \left(\frac{\phi_{n1}(a)}{\phi_{n0}(a)}\right)^{(\sigma_n(z)-1)(\gamma_n(z)-\xi_n)} dG(a)}{N_{n0} \int_a x_{n0}(z, a) dG(a)} \right]^{\frac{1}{1-\sigma_n(z)}} \quad (18)$$

It also also depends on the availability of product varieties, the extent to which depends on 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 new product varieties in the second counterfactual. Assuming symmetry between the domestic and foreign economies, this additional margin is captured by the term $(1 + \delta_n^X(a) \tau^{1-\sigma_n(z)})$.

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* homogenous. 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:

$$\int_a \frac{1}{\tilde{\sigma}_{n0}(a)} [1 - \beta_n (\tilde{\sigma}_{n0}(a) - 1) (\tilde{\gamma}_{n0}(a) - \xi_n)] X_{n0}(a) dG_n(a) = \int_a \frac{1}{\tilde{\sigma}_{n1}(a)} \delta_{nD}(a) [1 - \beta_n (\tilde{\sigma}_{n1}(a) - 1) (\tilde{\gamma}_{n1}(a) - \xi_n)] X_{n1}(a) dG_n(a) + \int_a (1 - \delta_{nD}(a)) f_{n0} dG_n(a) \quad (19)$$

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

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

$$\frac{1}{\tilde{\sigma}_{n1}(a)} [1 - \beta_n (\tilde{\sigma}_{n1}(a) - 1) (\tilde{\gamma}_{n1}(a) - \xi_n)] X_{n1}(a) - f_{n0} > 0 \Leftrightarrow \delta_{nD}(a) = 1 \quad (20)$$

In the second 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:

$$r_n^X(a, \phi_{n1}^X(a)) + r_{n1}^D(a, \phi_{n1}^X(a)) - f_n(\phi_{n1}^X(a)) - f_X > r_n^D(a, \phi_{n1}^D(a)) - f_n(\phi_{n1}^D(a)) \quad (21)$$

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.

Solution and Decomposition

Our framework naturally lends itself to quantitative estimation. As we can see from equations 16-21, our counterfactual estimation requires data on initial sales $x_{n0}(z, a)$ in addition to estimates of five sets of parameters: $\sigma_n(z), \gamma_n(z), \beta_n, \xi_n$ and f_{n0} . With these parameter values in hand, we can directly solve these equations for the relative changes in quality $\frac{\phi_{n1}(a)}{\phi_{n0}(a)}$, sales $\frac{x_{n1}(z, a)}{x_{n0}(z, a)}$, mass of firms $\frac{N_{n1}}{N_{n0}}$, survival $\delta_{nD}(a)$, export $\delta_{nX}(a)$ and price indices $\frac{P_{n1}(z)}{P_{n0}(z)}$. Note that we do not require estimates of firm productivity a or firm quality $\phi(a)$ to conduct our counterfactual exercise. This approach follows Dekle et al. (2007) among others.¹⁴

Our primary objective is to quantify the effect of either changes in the income distribution or opening to trade on differences in price indices across households. 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 propose a five-term decomposition of the effect on price indices for income group z relative to income group z_0 :

$$\begin{aligned} \log \frac{P_{n1}(z)}{P_{n0}(z)} - \log \frac{P_{n1}(z_0)}{P_{n0}(z_0)} &= \underbrace{- (\gamma_n(z) - \gamma_n(z_0)) \int_a \bar{s}_{n1}(a) \log \left(\frac{\phi_{n1}(a)}{\phi_{n0}(a)} \right) dG(a)}_{(1) \text{ Average quality effect}} \quad (22) \\ &\quad - \underbrace{(\tilde{\gamma}_n - \xi_n) \int_a (s_{n1}(a, z) - s_{n1}(a, z_0)) \log \left(\frac{\phi_{n1}(a)}{\phi_{n0}(a)} \right) dG(a)}_{(2) \text{ Asymmetric quality-adjusted cost changes}} \\ &\quad - \underbrace{\int_a (s_{n1}(a, z) - s_{n1}(a, z_0)) \log \left(\frac{\tilde{\rho}_{n1}(a)}{\tilde{\rho}_{n0}(a)} \right) dG(a)}_{(3) \text{ Asymmetric markup changes}} \end{aligned}$$

¹⁴Combining equations 16 and 18 describes how sales growth $\frac{x_{n1}(z, a)}{x_{n0}(z, a)}$ depend on quality upgrading $\frac{\phi_{n1}(a)}{\phi_{n0}(a)}$, while equation 17 describes 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 $\beta_n(\sigma_n(z) - 1)(\gamma_n(z) - \xi_n)$ is strictly smaller than unity for all z .

$$\begin{aligned}
& - \underbrace{\left(\frac{1}{\sigma_n(z)-1} - \frac{1}{\sigma_n(z_0)-1} \right) \log \left(\frac{N_{n1} \bar{\delta}_{nD} (1 + \bar{\delta}_X \tau^{1-\bar{\sigma}_n})}{N_{n0}} \right)}_{(4) \text{ Love of variety}} \\
& - \underbrace{\frac{1}{\bar{\sigma}_n-1} \log \left(\frac{\int_a s_{n0}(a,z) \delta_{nD}(a) (1 + \delta_X(a) \tau^{1-\sigma_n(z)}) dG_n(a)}{\int_a s_{n0}(a,z_0) \delta_{nD}(a) (1 + \delta_X(a) \tau^{1-\sigma_n(z_0)}) dG_n(a)} \right)}_{(5) \text{ Asymmetric import and exit effects}}
\end{aligned}$$

where $s_{n0}(a,z)$ denotes the initial market share of brand a among consumers of income z , and where $s_{n1}(a,z) = \frac{s_{n0}(a,z) \delta_{nD}(a) (1 + \delta_X(a) \tau^{1-\sigma_n(z)})}{\int_a s_{n0}(a,z) \delta_{nD}(a) (1 + \delta_X(a) \tau^{1-\sigma_n(z)})}$ 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. $\bar{\delta}_{nD} = \int_a \delta_{nD}(a) \bar{s}_{n0}(a) dG(a)$ denotes average survival rates across all firms and the two income groups.

In both counterfactuals, the first underlying channel is that firms on average have incentives to upgrade their product quality, which has heterogeneous effects across households depending on their preference parameters $\gamma_n(z)$.¹⁵ In the first counterfactual, firms upgrade their quality as a larger share of their consumers are households with higher quality evaluations. In the second counterfactual, the largest firms experience positive scale effects from trade opening, which also induces an increase in weighted average product quality.

The second effect is that the scale of production changes asymmetrically across higher and lower quality producers in both counterfactuals. Given the estimated economies of scale in quality production, this translates into asymmetric effects on quality and quality-adjusted prices. In turn, this favors richer households if they spend relatively more on firms with the largest increase in scale. This channel can be expressed as a covariance term between market shares $s_n(a,z)$ and quality upgrading $\log \left(\frac{\phi_{n1}(a)}{\phi_{n0}(a)} \right)$.

The third effect captures the change in markups, which depends on the composition of demand and 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 second (trade) counterfactual, this effect combines the number of varieties that are available on the domestic market as well as new imported varieties.

Finally, the fifth channel reflects the unequal effects of exit and import penetration. 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

¹⁵For 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.

(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 second (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.

5 Estimation

This section presents the empirical estimation. We begin by estimating the preference parameters, σ_{nz} and γ_{zn} , that combined with the microdata allow us to quantify the distribution of product quality and quality adjusted prices 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, 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:

$$\Delta \ln(s_{znict}) = (1 - \sigma_{nz}) \Delta \ln(p_{nict}) + \eta_{znct} + \epsilon_{znict} \quad (23)$$

where as before z , n and i denote household groups, product modules and brands. c and t indicate US counties and 14 semesters (13 changes), and s_{znict} are budget shares within product module n . η_{znict} are household group-by-product module-by-county-by-semester fixed effects that capture the CES price index term. Consistent with our CES preference specification at the level of household groups, we estimate expression 23 after aggregating consumption shares in the home scanner microdata for the period 2006-2012 to the level of household quintile-by-county-by-module-by-semester.¹⁶ To address concerns about autocorrelation in the error term ϵ_{znict} for the same county over time or within the county across household groups and modules, we cluster standard errors at the county level.¹⁷

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

¹⁶We aggregate household-level sales as projection-factor-weighted sums to compute $\Delta \ln(s_{znict})$. To be consistent with our CES specification and limit the bias due to zeroes in observed consumption, we restrict estimations to income group-by-county-by-semester cells with at least 50 households per cell. To compute brand-level log price changes we first compute projection-factor-weighted price means for each barcode-by-county-by-semester cell, and then compute $\Delta \ln(p_{nict})$ as a brand-level Tornqvist price index across all barcodes belonging to the same brand. As reported in Appendix Table A.2, 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.

¹⁷Alternatively clustering at the level of product modules, county-by-income group, county-by-semester or county-by-product module lead to slightly smaller standard errors.

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 $\Delta \ln(p_{nict})$ with either national or state-level leave-out mean price changes: $\frac{1}{N-1} \sum_{j \neq c} \Delta \ln(p_{nijt})$. 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.¹⁸

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 will report our findings both for our preferred baseline parameter values for σ_{nz} , as well as across a range of alternative parameter combinations to document the sensitivity of the findings. Finally, it is important to notice that the key empirical moment in our welfare quantification does not rely on the average level of σ_{nz} , but on its observed heterogeneity across different income groups. And while it is of course a possibility that some of the discussed endogeneity concerns may affect rich and poor households to differently, such concerns would require somewhat more elaborate stories compared to the traditional simultaneity bias in demand estimation.

Estimation Results In the estimations reported in [Table 2](#), we allow for sigma heterogeneity across both product departments as well as household income groups. Panel A shows the pooled estimation results across all household and product groups. In support of the IV strategy, we find that the point estimates change from slightly positive in the OLS specification to negative and statistically significant in both IV estimations as well as the joint IV column. The estimates from the two different instruments are very similar and suggest an aggregated elasticity of substitution

¹⁸It is not for lack of imagination that our main specification follows the Hausman IV tradition: e.g. we have attempted using brand-specific variation in local store openings (in many different forms) as an instrument for local price changes, but found that the first stage in these regressions was insufficient for a valid IV approach.

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, significantly lower than empirical work that has used the Feenstra (1994) approach for estimating (σ_{nz}) (e.g. Broda & Weinstein (2010), Hottman et al. (2014)). As a robustness exercise, we will report our findings in the final section of this paper both for our baseline parameter values for (σ_{nz}) as well as across a range of alternative parameter combinations 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 income quintile is about 0.5 larger than that for the richest income quintile. In terms of patterns across the income distribution, it appears that the clearest difference is between the bottom two income quintiles, which have roughly identical estimates, relative to the top three income quintiles as the difference to the top quintile in the reference category becomes significantly smaller for the median income quintile and the 2nd richest.

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 16 (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 the number of observations starts to become sparse 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, Table A.3 in the Appendix reports the results when estimating 40 (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 all central findings of this paper both based on our baseline parameter estimates as well as across a wide range of alternative parameter combinations in order to document the sensitivity of our findings.

5.2 Estimation of Brand Quality and Quality-Adjusted Prices

Armed with estimates of σ_{nz} , equation 3 allows us to use the scanner microdata to estimate product quality, $\ln\phi_{ni} = \int_z \ln\varphi_{ni}(z)dH(z)$, and quality adjusted prices, $\ln\left(\frac{p_{ni}}{\phi_{ni}}\right)$, across producers of brands as well as household consumption baskets. These empirical relationships are also directly

linked to the decomposition of the estimated slope in Figure 2 that we derive in expression 15.

As shown in 3, the key additional empirical moment are product unit values that we use in addition to observed product sales and the estimated σ_{nz} parameters to estimate unobserved variation in product quality. Figure 8 depicts the distribution of mean deviations in log product unit values within product module-by-semester cells (aggregated as consumption weighted averages across household consumption baskets) along the income distribution.¹⁹ The richest quintile of US households source their consumption from firms that have on average 15 percent higher unit values within product modules compared to the poorest income quintile.

Figure 9 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 (10) percent of US households source their consumption from on average more than 35 (40) percent higher quality producers compared to the poorest 20 (10) percent of households. Appendix Figure A.4 confirms what we already noted in the stylized facts section from Figure 4: 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, Figure 10 documents that the richest income quintile (decile) source their consumption at on average 20 (25) 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, $\left(\tilde{\sigma}_{ni} = \frac{\int_z \sigma_n(z) x_n(z,i) dH(z)}{\int_z x_n(z,i) dH(z)}\right)$, across the firm size distribution that informs the distribution of firm markups. The left panel of Figure 11 presents the estimation results of $\tilde{\sigma}_{ni}$ across 14 pooled cross-sections (for fourteen semesters between 2006-2012) 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:

$$\ln(\varphi_{znit}) = \gamma_{zn} \ln(\phi_{nit}) + \eta_{znt} + \epsilon_{znit} \quad (24)$$

where η_{znt} are income group-by-product module-by-semester 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 $\ln(\phi_{nit})$ with two semester lagged values of product quality.

¹⁹We compute brand-level unit values as sales-weighted means across barcodes and stores in case of multiple observations at the level of brand-by-household-by-semester cells.

Table 3 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 3 in order to compute the weighted average product quality evaluations faced by each brand producer: $\tilde{\gamma}_{ni} = \frac{\int_z \gamma_n(z) (\sigma_n(z)-1) x_n(z,i) dH(z)}{\int_z (\sigma_n(z)-1) x_n(z,i) dH(z)}$. The right panel in Figure 11 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 3, 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.

5.3 Technology Parameter Estimation

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:

$$\ln(p_{nit}) = \left(\beta_n - \frac{1}{\sigma_n - 1} \right) \ln(X_{nit}) + \eta_{nt} + \epsilon_{nit} \quad (25)$$

where η_{nt} are product module-by-semester fixed effects. Intuitively, if brands were of the same quality then the relationship between unit values (that would be identical to prices in this case) and market shares would be governed by the slope of the demand curve $-\frac{1}{\sigma_n-1}$. 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 (25) with product quality on the left hand side: $\ln(\phi_{nit}) = \beta_n \ln(X_{nit}) + \eta_{nt} + \epsilon_{nit}$, where following (3) and (5) $\ln(\phi_{nit}) = \ln(p_{nit}) + \frac{1}{\sigma_n-1} \ln(X_{nit})$. 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:

$$\begin{aligned} \ln(p_{nit}) = & \left(\beta_n - \frac{1}{\bar{\sigma}_n - 1} \right) \ln(X_{nit}) - \frac{1}{N_z} \sum_z \frac{1}{\sigma_{nz} - 1} \ln\left(\frac{X_{nizt}}{X_{nit}}\right) \\ & + \beta_n \ln(\tilde{\rho}_{nit}(\tilde{\gamma}_{nit} - \xi_n)) + \eta_{nt} + \epsilon_{nit} \end{aligned} \quad (26)$$

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 $\left(-\frac{1}{\sigma_n - 1} \ln(X_{nit})\right)$ in equation (25), 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 (26) as: $\ln(\phi_{nit}) = \ln(p_{nit}) + \frac{1}{N_z} \sum_z \frac{1}{\sigma_{nz} - 1} \ln(X_{nizt}) = \beta_n \ln(X_{nit} \tilde{\rho}_{nit}(\tilde{\gamma}_{nit} - \xi_n)) + \eta_{nt} + \epsilon_{nit}$, 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 $\tilde{\gamma}_{nit}$, and select the parameter combination that maximizes the goodness of fit while satisfying the model's parameter restriction $\beta_n(\tilde{\sigma}_{ni} - 1)(\tilde{\gamma}_{ni} - \xi_n) < 1$ (i.e. that the share of revenues re-invested in quality upgrading is less than 1).²⁰

The two main identification concerns in (26) 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 $\ln(X_{nit} \tilde{\rho}_{nit}(\tilde{\gamma}_{nit} - \xi_n))$ with two-semester lags. To address concerns about autocorrelation in the error term, we cluster the standard errors at the level of product modules as before.²¹

Panel Estimation Estimation equation (26) 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, because quality upgrading/downgrading by firms in response to changes in demand conditions (scale and consumer composition) are likely best understood as a

²⁰Given that the R-squared has no useful interpretation using IV's, we use the minimization of the root mean squared error of the reduced form regression (instrument on the right hand side).

²¹At the moment we do not adjust standard errors for the fact that some of the regressors are themselves estimates from the previous section. In future versions, we bootstrap the estimation procedure across the different steps.

longer-term effect (both in terms of changing actual quality attributes as well as making investments into brand perceptions through advertisement).

The natural panel data approach to estimating β_n and ξ_n would be to write (26) 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:

$$\begin{aligned} \Delta \ln(p_{nist}) = & \beta_n \Delta \ln(X_{nit}) - \frac{1}{N_z} \sum_z \frac{1}{\sigma_{nz} - 1} \Delta \ln(X_{nizst}) \\ & + \beta_n \Delta \ln(\tilde{\rho}_{nit}(\tilde{\gamma}_{nit} - \xi_n)) + \eta_{nst} + \epsilon_{nit} \end{aligned} \quad (27)$$

where subscript s indexes US states, η_{nst} are state-by-product module-by-semester fixed effects, and Δ indicates a two-year change (3 changes in our database starting from the first semester in 2006). As before, the second term on the right captures the demand-side relationship between sales and product unit values conditional on product quality, but this time at the state level. For instance changes in firm productivity (and thus unit values on the left) conditional on product quality 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 (27) as: $\Delta \ln(\phi_{nist}) = \Delta \ln(p_{nist}) + \frac{1}{N_z} \sum_z \frac{1}{\sigma_{nz}-1} \Delta \ln(X_{nizst}) = \beta_n \Delta \ln(X_{nit} \tilde{\rho}_{nit} (\tilde{\gamma}_{nit} - \xi_n)) + \eta_{nst} + \epsilon_{nit}$. 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 (27) 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 (27), we exploit 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 ($\Delta \ln(X_{nit} \tilde{\rho}_{nit} (\tilde{\gamma}_{nit} - \xi_n))$) 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 4 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 4 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 percent change in brand quality.

Table 5 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 brand-level consumer compositions as well as the marginal cost parameter ξ_n as shown above in the estimation equations (26) and (27). 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 4. The second panel reports the technology parameter estimates separately for grocery and non-grocery product groups, and Appendix Table A.4 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 5 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.

6 Counterfactuals

In this section, we use the model in combination with the microdata to explore the implications for household price indices and real income inequality, and decompose those effects into different channels. In the first counterfactual, we quantify the implications of an exogenous change in the distribution of household nominal incomes for real income inequality. In the second counterfactual, we quantify the distribution of the gains from opening up to trade.

6.1 Counterfactual 1: Changes in Nominal Income Inequality

Our first counterfactual explores 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, 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.

The first counterfactual is to exogenously increase nominal income inequality while holding market size fixed. We do so by reallocating 5 percent of market sales from the poorest quintile to the richest quintile.²² Using initial sales and our estimates for parameters $\sigma_n(z)$, $\gamma_n(z)$, β_n and ξ_n , we solve for the counterfactual equilibrium as described in section 4.3. To describe the

²²This is a simple and transparent way to increase nominal income inequality, but since we implicitly decrease the number of poor consumers by more than we increase the number of rich consumers, it also increases the mean of household income. One concern is that average product quality only increases because of this. An alternative approach would have been to keep average income constant and increase the variance. Notice, however, that what matters for firm decisions is the sales-weighted average of consumer incomes (and thus taste-for-quality parameters and price elasticities), which increases in both approaches. We prefer our current approach for its simplicity.

mechanisms in detail, we use the decomposition of the price effect described in equation 22. 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 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.

Figures 12 and A.5 and Table 6 present the estimation results for the total effect on price indices by income group and its decomposition. Several findings emerge. A 5 percent reallocation of expenditures from the poorest to the richest quintile induces changes in price indices that are on average 2.9 (2.4) percentage points lower for the richest household quintile compared to the poorest when using the cross-sectional (panel data) technology parameter estimates. An exogenous increase in nominal inequality leads to a significantly larger increase in real income inequality once we take into account endogenous asymmetric effects on household price indices. The first two columns of Table 6 present the five-term 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. Interestingly, this effect is significantly stronger when estimated using the cross-sectional technology parameters compared to the panel data estimates. This is intuitive, as the economies of scale parameter is significantly higher in the cross-sectional estimation. The right panel in Figure 12 confirms this intuition by depicting endogenous changes in log product quality across the initial firm size distribution.

The second term on the heterogeneous scale effect reinforces the first channel and corresponds to half 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 5). Once again this pattern is also illustrated in the right panel of Figure 12. 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 markups increase. A homogenous 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 counterfactual allows 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 counterfactual leads 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. 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.

Finally, 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, we find, however, that exit has a negligible effect. 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.²³

These results hold to a very similar extent in each of the 14 semesters as indicated by the depicted confidence intervals in the figures. Finally, as shown in Figure A.5, they also hold across all product departments, but the magnitudes vary significantly.

6.2 Counterfactual 2: Opening to Trade

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

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 (adding their domestic and export sales). We calibrate variable trade costs τ such that export sales of exporters equal 20 percent of their output. Combining these two statistics, about 10% 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

²³In 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.

ratios have increased from 40 percent to 60 percent on average across countries).

Figures 13 and A.6 and Table 6 present the counterfactual results. Greater openness to trade induces consumer price index changes that are on average 3.6 (2.7) 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 five-term decomposition in equation 22 to describe the mechanisms at play.

As in the first counterfactual, weighted average quality significantly increases as depicted in the right panel of Figure 13. 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 5). 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 (40 percent of total effect) and less so (17 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 thus reinforces the effect of the average increase in product quality. This pattern is also illustrated in the right-hand panel of Figure 13.

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 variety. It explains more than one half of the total effect in both quantifications. 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 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.

6.3 Robustness Checks across Parameter Ranges

[Work in progress.]

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 rich and poor households, explores the underlying channels, and quantifies the implications for real income inequality. To do so, we bring to bear newly available matched home and retail 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 quantitative model of endogenous product quality choice by heterogeneous firms and households.

The analysis provides several new 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 outcome is mainly driven 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 two features give rise to the endogenous sorting of larger, more productive firms into products that are valued relatively more by wealthier households.

These results have a new set of implications for 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 trade becomes significantly more regressive due to asymmetric effects on household price indices. Underlying these findings is a rich interplay of firm adjustments to product quality, markups, exit and entry that are asymmetric across the pre-existing distribution of firm sizes, which in turn translate differently into the consumption baskets of rich and poor households.

Our findings suggest that firm heterogeneity affects real income inequality in more complex ways than solely through nominal earnings, which have been the focus of the existing literature. This insight arises after introducing a very basic set of features that we observe in the data—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 estimating the partial effects of policy shocks on inequality.

References

- Acemoglu, D., & Autor, D. (2011). Skills, tasks and technologies: Implications for employment and earnings. *Handbook of labor economics*, 4, 1043–1171.
- 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. (2015). *Retail globalization and household welfare: Evidence from Mexico* (Tech. Rep.). National Bureau of Economic Research.

- 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. (2014). *Export destinations and input prices* (Tech. Rep.). National Bureau of Economic Research.
- 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. (2015). 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 mimeo*.
- 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.
- Davis, D. R., & Harrigan, J. (2011). Good jobs, bad jobs, and trade liberalization. *Journal of International Economics*, 84(1), 26–36.
- Deaton, A. (2003). Prices and poverty in India, 1987-2000. *Economic and Political Weekly*, 362–368.
- Dekle, R., Eaton, J., & Kortum, S. (2007). Unbalanced trade. *The American Economic Review*, 351–355.
- Diewert, W. (1979). *The economic theory of index numbers: A survey*. Department of Economics, University of British Columbia.
- Dingel, J. I. (2015). The determinants of quality specialization. *mimeo Chicago Booth*.
- Faber, B. (2014). Trade liberalization, the price of quality, and inequality: Evidence from mexican store prices. *UC Berkeley mimeo*.
- 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.
- Frias, J. A., Kaplan, D. S., & Verhoogen, E. A. (2009). Exports and wage premia: Evidence from mexican employer-employee data. *Working Paper, Columbia University*.

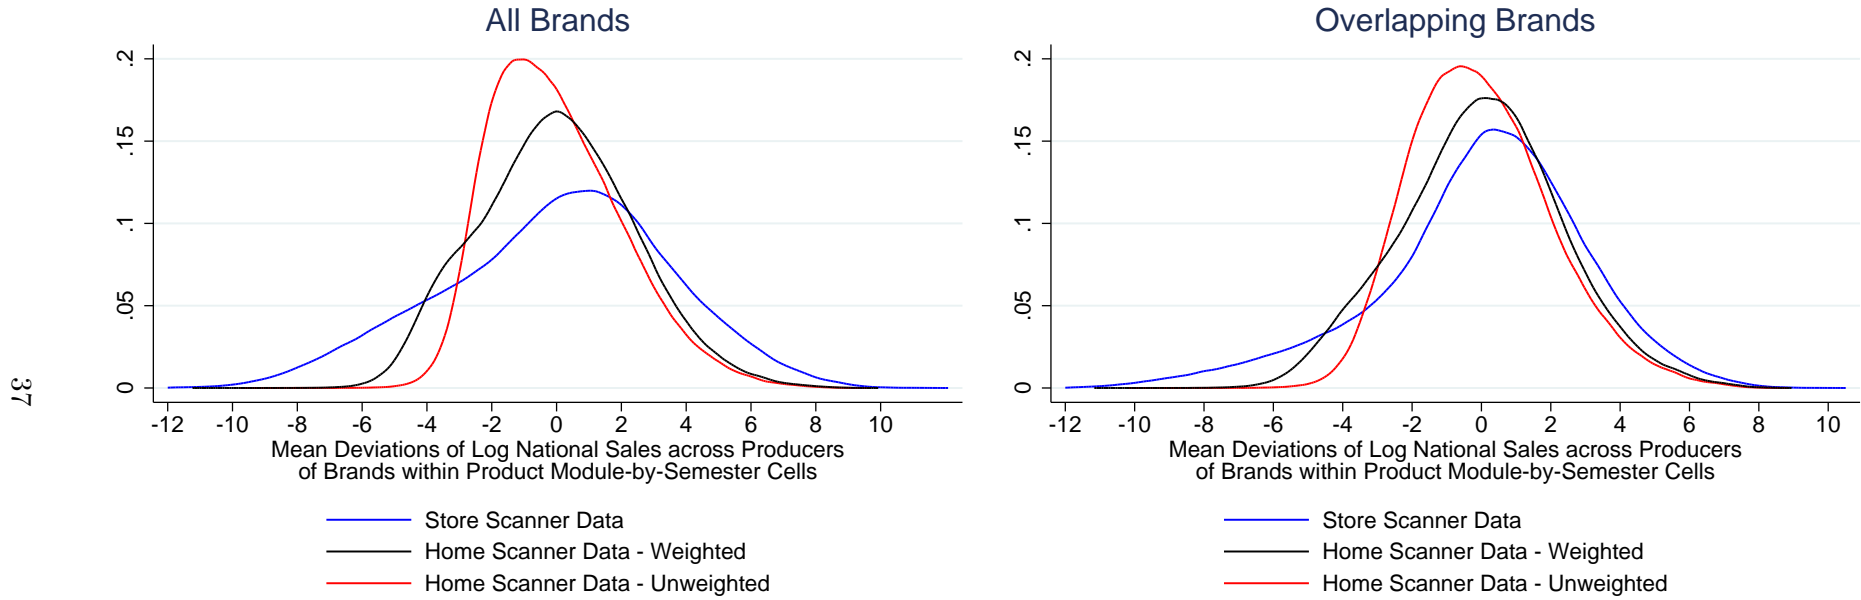
- Handbury, J. (2014). Are poor cities cheap for everyone? Non-homotheticity and the cost of living across US cities. *Wharton mimeo*.
- Handbury, J., & Weinstein, D. E. (2014). Goods prices and availability in cities. *The Review of Economic Studies*, rdu033.
- Harrigan, J., & Reshef, A. (2011). *Skill 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. (2003). Sources of bias and solutions to bias in the consumer price index. *Journal of Economic Perspectives*, 23–44.
- 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. (2012). *Trade and inequality: From theory to estimation* (Tech. Rep.). National Bureau of Economic Research.
- Helpman, E., Itskhoki, O., & Redding, S. (2010). Inequality and unemployment in a global economy. *Econometrica*, 78(4), 1239–1283.
- Hottman, C., Redding, S. J., & Weinstein, D. E. (2014). *What is firm heterogeneity in trade models? the role of quality, scope, markups, and cost* (Tech. Rep.). National Bureau of Economic Research.
- Hsieh, C.-T., & Klenow, P. J. (2009). Misallocation and manufacturing tfp in china and india. *Quarterly Journal of Economics*, 124(4).
- 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.
- Khandelwal, A. (2010). The long and short (of) quality ladders. *Review of Economic Studies*, 77(4), 1450–1476.
- Kugler, M., & Verhoogen, E. (2012). Prices, plant size, and product quality. *Review of Economic Studies*, *Forthcoming*.
- Mandel, B. (2010). Heterogeneous firms and import quality: Evidence from transaction-level prices. *International Finance Discussion Papers*.
- Melitz, M. (2003). The impact of trade on intra-industry reallocations and aggregate industry productivity. *Econometrica*, 71(6), 1695–1725.
- Melitz, M., Helpman, E., Marin, D., Verdier, T., & Costantini, J. (2007). The dynamics of firm-level adjustment to trade liberalization. *The Organization of Firms in a Global Economy*.
- Muellbauer, J. (1974). Prices and inequality: The United Kingdom experience. *The Economic Journal*, 84(333), 32–55.
- 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.
- Saez, E., & Zucman, G. (2014). Wealth inequality in the united states since 1913: Evidence from capitalized income tax data [Working Paper]. (20625).
- 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.

8 Figures and Tables


Figures

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



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.


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



38


The figure depicts deviations in weighted average log firm sales embodied in the consumption baskets of on average 58 thousand US households during 14 half year periods between 2006-12. 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 more than 150,000 brand producers. In the first step, we calculate brand-level deviations from mean log national sales within product module-by-semester cells from either the home or the store-level scanner data. 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 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. The fitted relationships in the left graph corresponds 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.

Figure 3: Heterogeneity across Product Departments




The figure depicts deviations in weighted average log firm sales embodied in the consumption baskets of on average 58 thousand US households during 14 half year periods between 2006-12. The y-axis displays weighted average deviations in log producer sales computed within more than 1000 product modules where the weights are household expenditure shares across more than 150,000 brand producers. National firm size deviations are based on the store scanner data. These firm size deviations are depicted separately for consumption in eight product departments that are defined by Nielsen as indicated in the figure. The x-axis displays national percentiles of per capita total household retail expenditure per half year period. The fitted relationship in the left graph corresponds to a local polynomial regression.

Figure 4: Households on Average Strongly Agree on Relative Product Quality Evaluations




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 58 thousand US households during 14 half year periods between 2006-12. 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.

Figure 5: The Role of Generic Retailer Brands and Non-Participating Store Chains



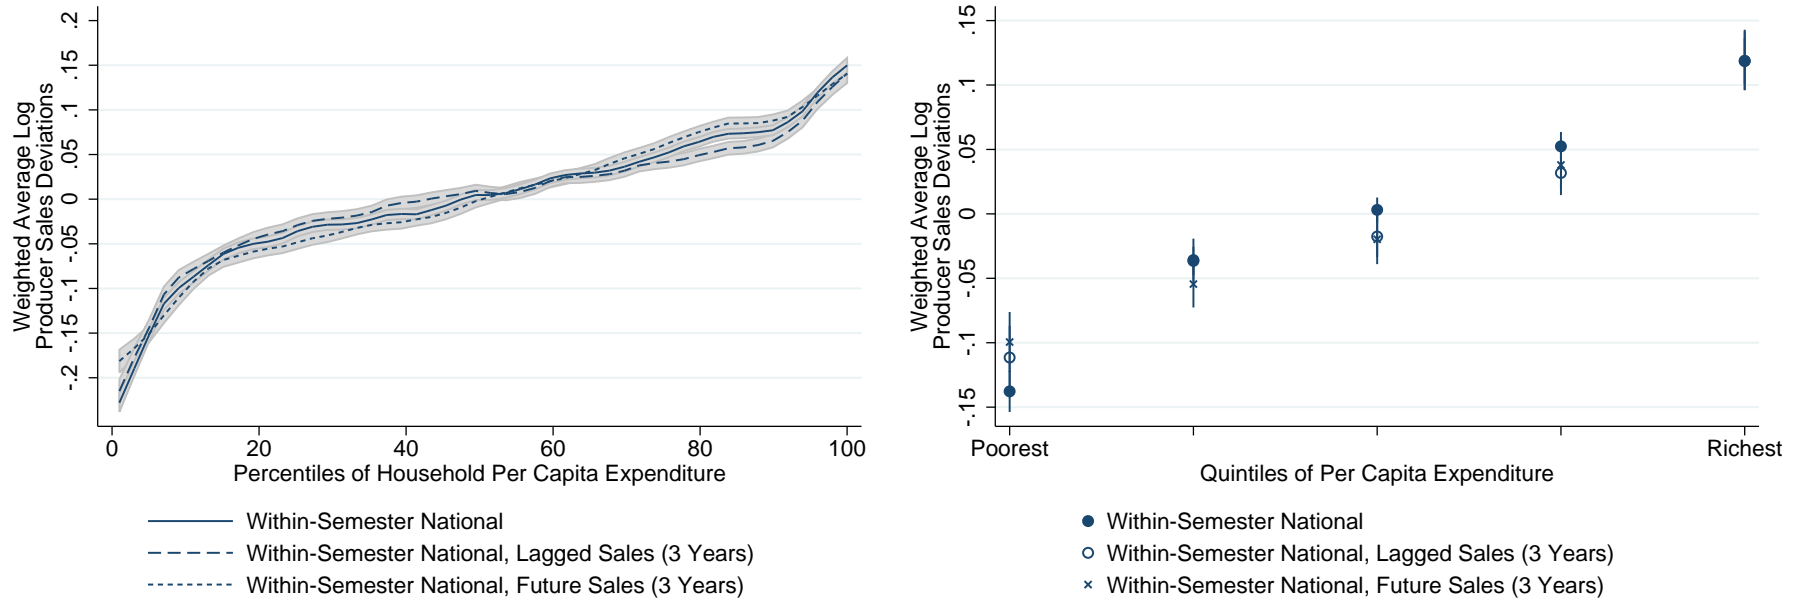
The figure depicts deviations in weighted average log firm sales embodied in the consumption baskets of on average 58 thousand US households during 14 half year periods between 2006-12. The y-axis in both graphs displays weighted average deviations in log producer sales within more than 1000 product modules where the weights are household expenditure shares across more than 150,000 brand producers. National firm size deviations are based on the store scanner data. These firm size deviations are depicted across consumption baskets conditional on semester 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. 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.

Figure 6: The Role of Differential Access to Producers




The figure depicts deviations in weighted average log firm sales embodied in the consumption baskets of on average 58 thousand US households during 14 half year periods between 2006-12. 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 more than 150,000 brand producers. National firm size deviations are based on the store scanner data. These firm size deviations are depicted across consumption baskets i) conditional on semester fixed effects, ii) conditional on semester-by-county fixed effects, and iii) conditional on semester-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. 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.

Figure 7: The Role of Temporary Taste Shocks that Differ across Rich and Poor Households



The figure depicts deviations in weighted average log firm sales embodied in the consumption baskets of on average 58 thousand US households during 14 half year periods between 2006-12. 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 more than 150,000 brand producers. National firm size deviations are based on the store scanner data. These firm size deviations are depicted across consumption baskets conditional on semester 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. 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.


Figure 8: Distribution of Weighted Average Unit Values across Consumption Baskets



44

The figure depicts deviations in weighted average log firm unit values embodied in the consumption baskets of on average 58 thousand US households during 14 half year periods between 2006-12. 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 more than 150,000 brand producers. In the first step, we calculate brand-level deviations from mean log national unit values within product module-by-semester 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. 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.

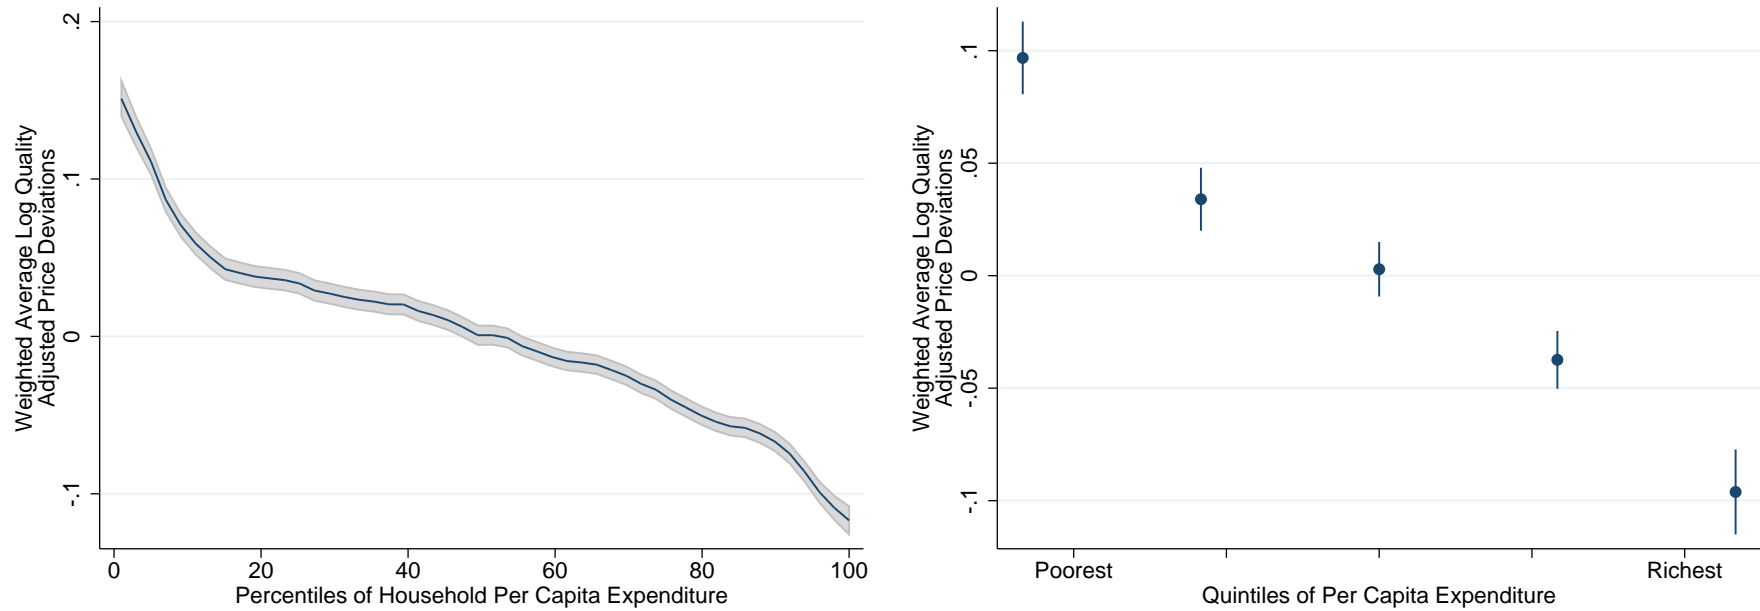
Figure 9: Distribution of Weighted Average Product Quality across Consumption Baskets



45


The figure depicts deviations in weighted average log brand quality embodied in the consumption baskets of on average 58 thousand US households during 14 half year periods between 2006-12. The y-axis in both graphs displays weighted average deviations in log producer (brand) quality within more than 1000 product modules where the weights are household expenditure shares across more than 150,000 brand producers. The x-axis displays national percentiles of per capita total household retail expenditure per half year period. 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.

Figure 10: Distribution of Weighted Average Quality-Adjusted Prices across Consumption Baskets



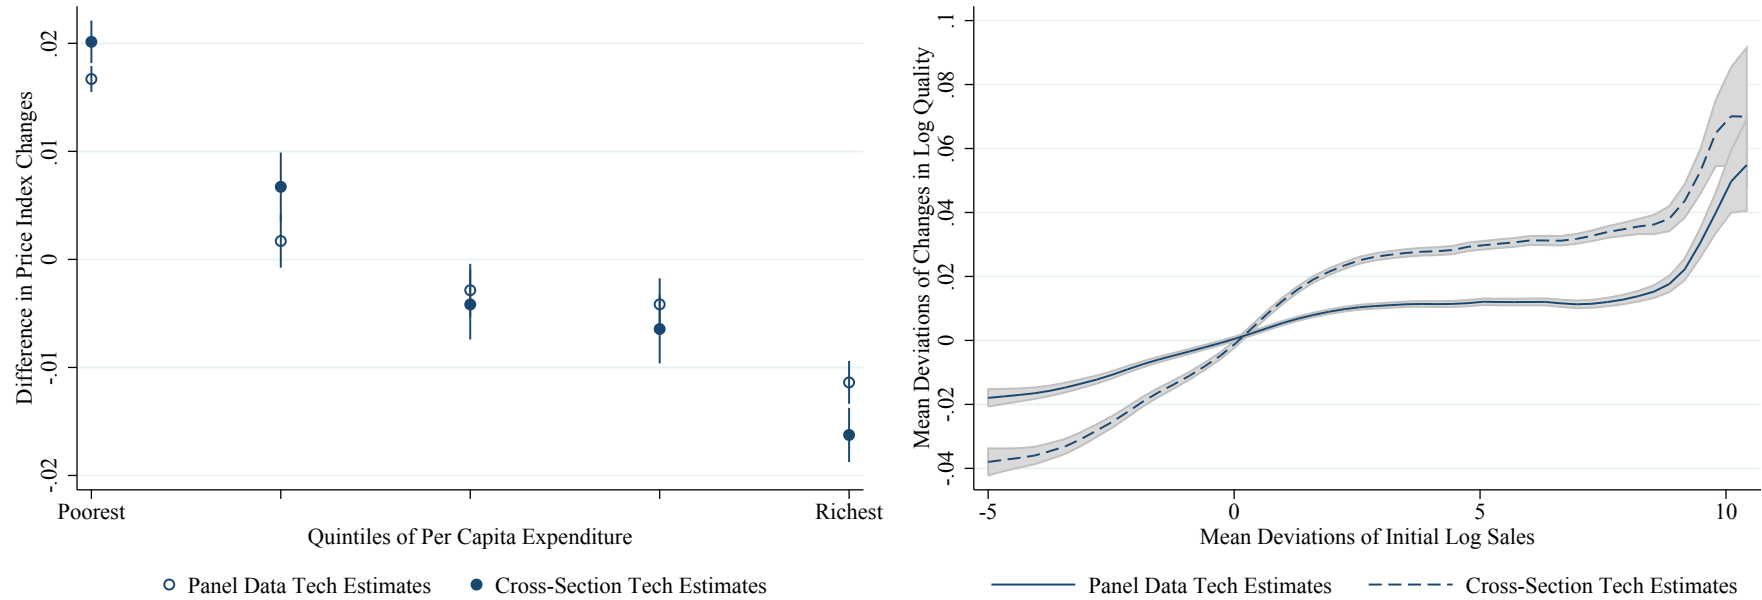
The figure depicts deviations in weighted average log quality adjusted prices embodied in the consumption baskets of on average 58 thousand US households during 14 half year periods between 2006-12. The y-axis in both graphs displays weighted average deviations in log producer (brand) quality adjusted prices within more than 1000 product modules where the weights are household expenditure shares across more than 150,000 brand producers. The x-axis displays national percentiles of per capita total household retail expenditure per half year period. 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.

Figure 11: Producers Face Different Elasticities of Substitution and Tastes for Quality



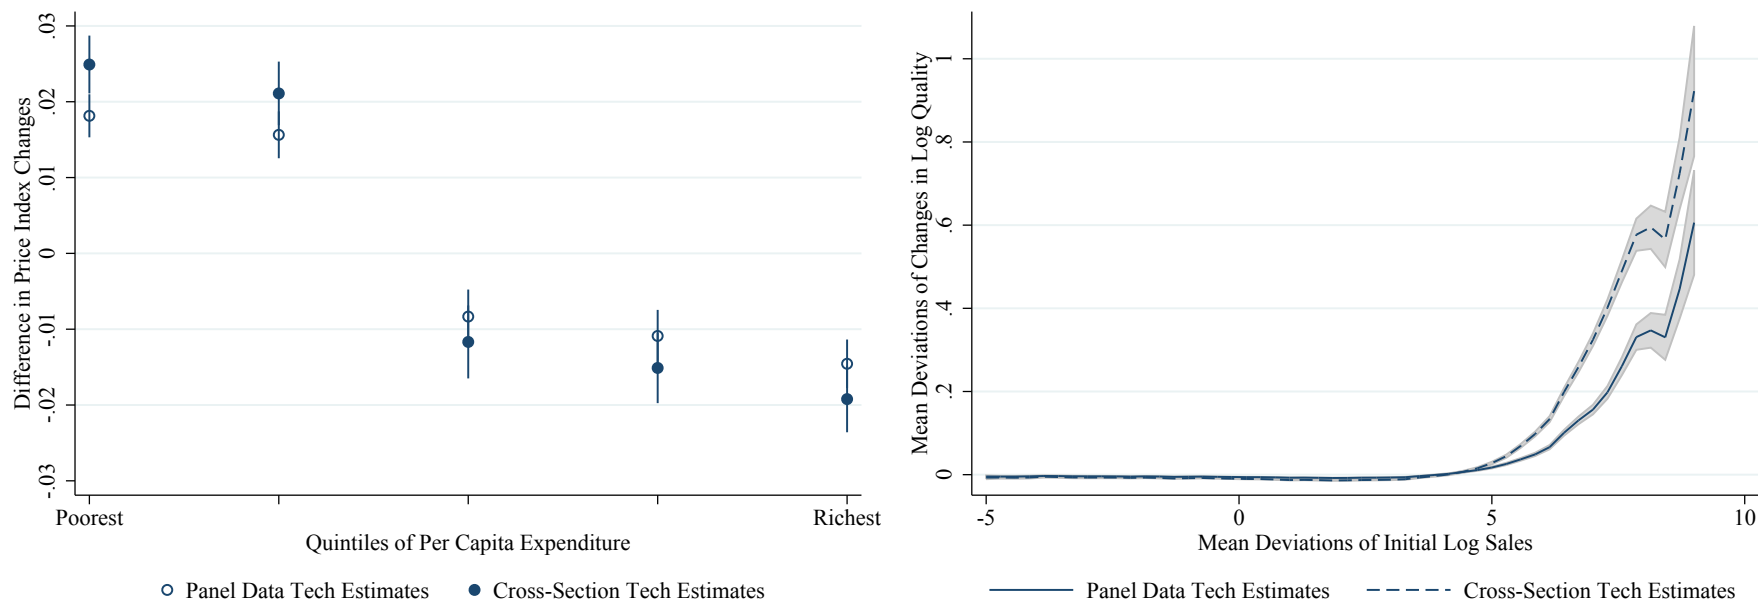
The figure depicts deviations in the weighted average elasticities of substitution (σ tilde) and quality taste parameters (γ tilde) across the firm size distribution for 14 semester cross-sections between 2006-2012. The y-axis displays de-measured values of the parameters within product module-by-semester cells. The x-axis displays de-measured log firm sales at the same level. 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.

Figure 12: Counterfactual 1: Inflation Differences and Quality Upgrading due to 5 Percent Reallocation of Expenditure to Richest Group



The figure depicts mean deviations in household retail price index changes for on average 58 thousand US households during 14 half year periods between 2006-12. The estimated price index changes correspond to the counterfactual where 5 percent of total market sales are reallocated from the poorest household income group to the richest as discussed in Section 6. Both graphs display confidence intervals at the 95% level.

Figure 13: Counterfactual 2: Inflation Differences and Quality Upgrading due to 10% Symmetric Increase in Import Penetration



The figure depicts mean deviations in household retail price index changes for on average 58 thousand US households during 14 half year periods between 2006-12. The estimated price index changes correspond to the second counterfactual discussed in Section 6. The graph displays confidence intervals at the 95% level.. Both graphs display confidence intervals at the 95% level.

Tables

Table 1: Descriptive Statistics

	Home Scanner Data		Retail Scanner Data
Number of Semesters 2006-12	14	Number of Semesters 2006-12	14
Number of Observations (Summed Up to Household-Semester-Barcode-Retailer)	267,477,828	Number of Observations (Summed Up to Store-Semester-Barcode)	9,324,297,920
Number of Households per Semester	58,041	Number of Stores per Semester	24,882
Number of Product Groups per Semester	1,089	Number of Product Groups per Semester	1,086
Number of Brands per Semester	182,996	Number of Brands per Semester	171,695
Number of Barcodes per Semester	583,131	Number of Barcodes per Semester	722,685
Number of Retailers per Semester	783	Number of Retailers per Semester	100
Number of Counties per Semester	2,662	Number of Counties per Semester	2,482
Total Sales per Semester (Using Projection Weights)	105,737,356 (208,530,458,605)	Total Sales per Semester	110,839,293,906

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	0.150*** (0.0368)	-1.163*** (0.0545)	-1.137*** (0.0490)	-1.183*** (0.0440)				
(1- σ) Poorest Quintile (Relative to Richest)					-0.514*** (0.121)			
(1- σ) 2nd Poorest Quintile (Relative to Richest)					-0.552*** (0.103)			
(1- σ) Median Quintile (Relative to Richest)					-0.173 (0.104)			
(1- σ) 2nd Richest Quintile (Relative to Richest)					-0.279*** (0.0817)			
Quintile-by-Module-by-County-by-Semester FX	✓	✓	✓	✓	✓			
Brand-by-County-by-Semester FX	✗	✗	✗	✗	✓			
Observations	4,804,155	4,804,155	3,980,418	3,980,418	3,980,418			
First Stage F-Stat		723.0	176.0	348.7	139.0			
<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	-1.075*** (0.191)	-0.701*** (0.104)	-1.338*** (0.0631)	-1.444*** (0.0830)	-2.226*** (0.249)	-0.596*** (0.121)	-1.004*** (0.138)	-1.263*** (0.162)
Quintile-by-Module-by-County-by-Semester FX	✓	✓	✓	✓	✓	✓	✓	✓
Observations	304,797	347,321	1,909,138	423,660	98,279	352,567	433,769	110,837
First Stage F-Stat	147.7	421.3	290.3	66.27	178.4	123.7	554.2	53.54
<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	-1.751*** (0.296)	-0.920*** (0.173)	-1.480*** (0.154)	-1.552*** (0.219)	-2.252*** (0.372)	-0.579 (0.352)	-1.328*** (0.367)	-1.715*** (0.281)
(1- σ) Median and Above Quintiles	-0.868*** (0.184)	-0.633*** (0.0952)	-1.301*** (0.0607)	-1.422*** (0.0766)	-2.233*** (0.254)	-0.579*** (0.176)	-0.932*** (0.119)	-1.136*** (0.212)
Quintile-by-Module-by-County-by-Semester FX	✓	✓	✓	✓	✓	✓	✓	✓
Observations	304,797	347,321	1,909,138	423,660	98,279	352,567	433,769	110,837
First Stage F-Stat	139.0	347.5	254.1	50.17	131.4	109.4	298.0	37.68

Table 3: 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	0.883*** (0.00761)	0.882*** (0.00763)	0.916*** (0.00849)	0.919*** (0.00835)	1.062*** (0.00485)	1.059*** (0.00457)	1.073*** (0.00561)	1.070*** (0.00542)	1.066*** (0.00640)	1.069*** (0.00673)
	Product Module-by-Semester FX	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓
	Observations	1,424,389	986,452	1,424,389	986,452	1,424,389	986,452	1,424,389	986,452	1,424,389	986,452
	Number of Product Module Clusters	1046	1027	1046	1027	1046	1027	1046	1027	1046	1027
BEVERAGES	Log Average Brand Sales	0.625*** (0.00963)	0.632*** (0.00973)	0.661*** (0.0118)	0.681*** (0.0124)	1.226*** (0.0102)	1.213*** (0.0106)	1.256*** (0.00774)	1.243*** (0.00626)	1.232*** (0.00933)	1.231*** (0.00917)
	Product Module-by-Semester FX	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓
	Observations	134,847	89,228	134,847	89,228	134,847	89,228	134,847	89,228	134,847	89,228
	Number of Product Module Clusters	69	68	69	68	69	68	69	68	69	68
DAIRY	Log Average Brand Sales	0.761*** (0.00412)	0.766*** (0.00439)	0.780*** (0.00251)	0.788*** (0.00310)	1.146*** (0.00276)	1.138*** (0.00310)	1.155*** (0.00292)	1.147*** (0.00294)	1.159*** (0.00450)	1.161*** (0.00487)
	Product Module-by-Semester FX	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓
	Observations	90,887	67,701	90,887	67,701	90,887	67,701	90,887	67,701	90,887	67,701
	Number of Product Module Clusters	46	45	46	45	46	45	46	45	46	45
DRY GROCERY	Log Average Brand Sales	0.887*** (0.00243)	0.888*** (0.00251)	0.919*** (0.00128)	0.923*** (0.00135)	1.055*** (0.00109)	1.050*** (0.00120)	1.067*** (0.00145)	1.064*** (0.00145)	1.071*** (0.00201)	1.075*** (0.00239)
	Product Module-by-Semester FX	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓
	Observations	554,129	394,266	554,129	394,266	554,129	394,266	554,129	394,266	554,129	394,266
	Number of Product Module Clusters	398	391	398	391	398	391	398	391	398	391
FROZEN FOODS	Log Average Brand Sales	0.919*** (0.00407)	0.918*** (0.00604)	0.950*** (0.00313)	0.950*** (0.00362)	1.037*** (0.00280)	1.032*** (0.00350)	1.045*** (0.00267)	1.045*** (0.00301)	1.050*** (0.00378)	1.056*** (0.00496)
	Product Module-by-Semester FX	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓
	Observations	98,077	69,798	98,077	69,798	98,077	69,798	98,077	69,798	98,077	69,798
	Number of Product Module Clusters	78	76	78	76	78	76	78	76	78	76
GENERAL MERCHANDISE	Log Average Brand Sales	0.955*** (0.00250)	0.955*** (0.00320)	0.990*** (0.00178)	0.991*** (0.00218)	1.013*** (0.00160)	1.009*** (0.00150)	1.017*** (0.00157)	1.018*** (0.00214)	1.025*** (0.00232)	1.027*** (0.00248)
	Product Module-by-Semester FX	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓
	Observations	152,478	97,502	152,478	97,502	152,478	97,502	152,478	97,502	152,478	97,502
	Number of Product Module Clusters	143	139	143	139	143	139	143	139	143	139
HEALTH & BEAUTY CARE	Log Average Brand Sales	0.964*** (0.00216)	0.964*** (0.00295)	1.002*** (0.00360)	1.004*** (0.00418)	1.009*** (0.00325)	1.008*** (0.00301)	1.018*** (0.00208)	1.016*** (0.00250)	1.007*** (0.00523)	1.008*** (0.00599)
	Product Module-by-Semester FX	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓
	Observations	217,735	144,697	217,735	144,697	217,735	144,697	217,735	144,697	217,735	144,697
	Number of Product Module Clusters	173	172	173	172	173	172	173	172	173	172
NON-FOOD GROCERY	Log Average Brand Sales	0.775*** (0.00190)	0.780*** (0.00242)	0.798*** (0.00259)	0.805*** (0.00314)	1.138*** (0.00294)	1.132*** (0.00278)	1.149*** (0.00197)	1.143*** (0.00230)	1.141*** (0.00409)	1.139*** (0.00468)
	Product Module-by-Semester FX	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓
	Observations	147,186	101,640	147,186	101,640	147,186	101,640	147,186	101,640	147,186	101,640
	Number of Product Module Clusters	128	125	128	125	128	125	128	125	128	125
PACKAGED MEAT	Log Average Brand Sales	0.742*** (0.00451)	0.752*** (0.00453)	0.756*** (0.00251)	0.765*** (0.00176)	1.156*** (0.00701)	1.145*** (0.00750)	1.171*** (0.00450)	1.166*** (0.00472)	1.174*** (0.00568)	1.171*** (0.00701)
	Product Module-by-Semester FX	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓
	Observations	29,050	21,620	29,050	21,620	29,050	21,620	29,050	21,620	29,050	21,620
	Number of Product Module Clusters	11	11	11	11	11	11	11	11	11	11

Table 4: 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	0.0307*** (0.00343)	0.0285*** (0.00392)	1.094*** (0.0258)	1.109*** (0.0255)				
Δ Log National Firm Sales					0.0380*** (0.00341)	0.0730*** (0.0160)	1.092*** (0.0347)	0.530*** (0.0637)
Product Module-by-Semester FX	✓	✓	✓	✓	✗	✗	✗	✗
State-by-Product Module-by-Semester FX	✗	✗	✗	✗	✓	✓	✓	✓
Observations	986,467	986,467	986,467	986,467	1,309,744	1,309,744	1,309,744	1,309,744
Number of Product Module Clusters	1029	1029	1029	1029	998	998	998	998
First Stage F-Stat		280781		280781		192.4		192.4

Table 5: Technology Parameter Estimates

Dependent Variable:		ALL PRODUCT GROUPS			
Log Product Quality or Changes in Log Quality	Cross-Section		Panel Data		
	OLS	IV	OLS	IV	
Log Firm Scale or Changes in Log Firm Scale (β)	1.0837*** (0.0254)	1.1027*** (0.0252)	1.0996*** (0.0364)	0.6189*** (0.0999)	
ξ Parameter	0.63	0.63	0.63	0.63	
Observations	986,467	986,467	664,882	664,882	
Number of Clusters	1,029	1,029	930	930	
First Stage F-Stat		228890.53		78.53	

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)	0.9498*** (0.013)	0.9668*** (0.0138)	0.9567*** (0.0149)	0.2285** (0.0965)	1.3848*** (0.0695)	1.4214*** (0.0683)	1.3784*** (0.083)	0.8866*** (0.1284)	
ξ Parameter	0.63	0.63	0.01	0.01	0.74	0.74	0.56	0.56	
Observations	744,262	744,262	770,324	770,324	242,205	242,205	291,036	291,036	
Number of Clusters	717	717	685	685	312	312	256	256	
First Stage F-Stat		177226.02		133.91		56087.69		79.86	

Table 6: Decomposition of Counterfactuals

Difference in Consumer Inflation (Richest Quintile - Poorest Quintile)	Counterfactual 1: Increase in Nominal Inequality				Counterfactual 2: Trade Opening			
	Cross-Sectional Tech Estimates		Panel Data Tech Estimates		Cross-Sectional Tech Estimates		Panel Data Tech Estimates	
(1) Change in Weighted Average Product Quality	-1.475	(50%)	-0.027	(1%)	-1.438	(40%)	-0.468	(17%)
	(0.027)		(0.004)		(0.05)		(0.021)	
(2) Asymmetric Scale Effect	-1.451	(49%)	-1.771	(75%)	-0.127	(3%)	-0.141	(5%)
	(0.122)		(0.154)		(0.063)		(0.069)	
(3) Asymmetric Changes in Markups	-0.169	(6%)	-0.178	(8%)	0.018	(0%)	0.001	(0%)
	(0.014)		(0.014)		(0.002)		(0.001)	
(4) Love of Variety	0.155	(-5%)	-0.371	(16%)	-1.626	(45%)	-1.609	(59%)
	(0.013)		(0.012)		(0.036)		(0.033)	
(5) Asymmetric Effect of Exit and Imports	0.000	(0%)	0.000	(0%)	-0.460	(13%)	-0.488	(18%)
	(0.000)		(0.000)		(0.063)		(0.062)	
Total Effect	-2.939	(100%)	-2.346	(100%)	-3.633	(100%)	-2.705	(100%)
	(0.119)		(0.155)		(0.151)		(0.127)	

9 Online Appendix

Appendix A: Mathematical Appendix

Prices and markups (Equation 10): Firms choose prices p to maximize profits. As a function of prices (holding quality as given), profits can be written as:

$$\pi_n = (p - c) \int_z \phi^{\gamma_n(z)(\sigma_n(z)-1)} p^{-\sigma_n(z)} A_n(z) dH(z) - f_n(\phi) - f_{n0}$$

where $c = \frac{c_n(\phi)}{a}$ denotes the marginal cost and the integral term refers to total quantities, summing across consumers, $A_n(z)$ is an income-group specific demand shifter corresponding to $A_n(z) = \alpha_n(z)E(z)P_n(z)^{\sigma_n(z)-1}$, and $E(z)$ refers to retail expenditures for consumers with income z .

The first order condition in prices leads to:

$$\int_z \phi^{\gamma_n(z)(\sigma_n(z)-1)} p^{-\sigma_n(z)} A_n(z) dH(z) = \left(\frac{p-c}{p} \right) \int_z \sigma_n(z) \phi^{\gamma_n(z)(\sigma_n(z)-1)} p^{-\sigma_n(z)} A_n(z) dH(z)$$

Using the expression for sales to consumers z : $x_n(z, a) = \phi^{\gamma_n(z)(\sigma_n(z)-1)} p(a)^{1-\sigma_n(z)} A_n(z)$, we obtain:

$$\int_z x_n(z, a) dH(z) = \left(\frac{p-c}{p} \right) \int_z \sigma_n(z) x_n(z, a) dH(z)$$

which leads to the following markups:

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

where, again, the marginal cost c is given by $c = \frac{c_n(\phi)}{a}$ as a function of ability a and quality ϕ .

Optimal quality (Equation 11): Assuming that firms choose quality ϕ and prices p jointly to maximize profits. As a function of prices and quality, profits can be written as:

$$\pi_n = \left(p - \frac{c_n(\phi)}{a} \right) \int_z \phi^{\gamma_n(z)(\sigma_n(z)-1)} p^{-\sigma_n(z)} A_n(z) dH(z) - f_n(\phi) - f_{n0}$$

where $f_n(\phi) = b_n \phi^{\frac{1}{\beta_n}}$ are the fixed costs of quality upgrading, and where the product $\phi^{\gamma_n(z)(\sigma_n(z)-1)} p^{-\sigma_n(z)} A_n(z)$ corresponds to quantities sold to consumers of income z with prices p and quality ϕ . Looking at the first order condition in quality (in log), we obtain:

$$\begin{aligned} \phi f'_n(\phi) &= \left(p - \frac{c_n(\phi)}{a} \right) \int_z \gamma_n(z)(\sigma_n(z) - 1) \phi^{\gamma_n(z)(\sigma_n(z)-1)} p^{-\sigma_n(z)} A_n(z) dH(z) \\ &\quad - \frac{\xi_n c_n(\phi)}{a} \int_z \phi^{\gamma_n(z)(\sigma_n(z)-1)} p^{-\sigma_n(z)} A_n(z) dH(z) \end{aligned}$$

where ξ_n is the elasticity of the marginal cost w.r.t quality ϕ . Using the expression for sales to consumers z : $x_n(a, z) = \phi^{\gamma_n(z)(\sigma_n(z)-1)} p^{1-\sigma_n(z)} A_n(z)$, we obtain:

$$\phi f'_n(\phi) = \left(1 - \frac{c_n(\phi)}{ap} \right) \int_z \gamma_n(z)(\sigma_n(z) - 1) x_n(a, z) dH(z) - \frac{\xi_n c_n(\phi)}{ap} \int_z x_n(a, z) dH(z)$$

However, as showed earlier, prices at equilibrium are such that:

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

Replacing $\frac{c_n(\phi)}{ap}$ by $\frac{\int_z (\sigma_n(z)-1)x_n(a, z) dH(z)}{\int_z \sigma_n(z) x_n(a, z) dH(z)}$ and rearranging, we obtain:

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

With $\phi f'_n(\phi) = b_n \phi^{\frac{1}{\beta_n}}$, with $\tilde{\rho}_n(a) = \frac{\tilde{\sigma}_n(a)-1}{\tilde{\sigma}_n(a)}$ 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:

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

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 $\gamma_n(z)$:

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

Derivative of quality w.r.t. a (Equation 13) with homogenous 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.

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 26): 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 B: 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:

$$U_{Gz} = \prod_n \left[\sum_{i \in G_n} (q_{zni} \varphi_{zni})^{\frac{\sigma_{nz} - 1}{\sigma_{nz}}} \right]^{\alpha_{nz} \cdot \frac{\sigma_{nz}}{\sigma_{nz} - 1}} \quad (28)$$

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

$$U_{jz} = \sum_n \alpha_{nz} \max_{i \in G_n, q_{jzni}} [\log q_{jzni} + \log \varphi_{zni} + \mu_{nz} \epsilon_{jzni}] \quad (29)$$

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 29 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:

$$i^* = \operatorname{argmax}_{i \in G_n} [-\log p_{ni} + \log \varphi_{zni} + \mu_{nz} \epsilon_{jzni}] \quad (30)$$

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 renders the discrete-choice version of Handbury (2013) analytically untractable.

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 30 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 29 is exactly equivalent to the consumption patterns obtained with the preferences described in equation 28 above and equation 2 in the main text.

Appendix C: Counterfactuals and Decompositions

1) Counterfactual 1: Equilibrium

Sales: By combining equations 3 and 11, we obtain that firm sales satisfy:

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

Prices, in turn, equal:

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

Hence, taking ratios:

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

(for surviving firms).

Quality: Optimal quality in each equilibrium is given by equation 11. Taking ratios, we obtain the expression in the text for optimal quality:

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

where 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 $\tilde{\rho}_{n1}(a)$).

Price index: In equilibrium, it is given by:

$$P_n(z) = \left[N_n \int_a p_n(a)^{1-\sigma_n(z)} \phi_n(a)^{\gamma_n(z)(\sigma_n(z)-1)} dG_n(a) \right]^{\frac{1}{1-\sigma_n(z)}}$$

Taking ratios, and adjusting for the exit of firms in the counterfactual equilibrium, we obtain:

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

where the second line is obtained by multiplying each line by $\alpha_n(z) E(z) P_{n0}^{\sigma_n(z)-1}$. Noticing that $p_{n0}(a)^{1-\sigma_n(z)} \phi_{n0}(a)^{\gamma_n(z)(\sigma_n(z)-1)} \alpha_n(z) E(z) P_{n0}^{\sigma_n(z)-1} = x_{n0}(a, z)$, we obtain:

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

Using the expression $p_n(a) = \frac{\phi_n(a)^{\xi_n}}{a\tilde{\rho}_n(a)}$ for prices, we obtain the expression in the text:

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

Entry: At equilibrium, free entry is such that expected profits are equal to the sunk cost of entry F_{nE} , which 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 above for profits, this is equivalent to expression in the text:

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

Exit: Survival ($\delta_{nD}(a)$ dummy) requires that profits are positive:

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

2) Counterfactual 1: Decompositions

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

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

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

Taking logs, and then a first-order approximation, we obtain:

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

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

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

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

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

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

Denoting $\bar{\delta}_{nD} = \int_a \delta_{nD}(a) \bar{s}_{n0}(a) dG(a)$ and combining lines 4 and 5 together, we obtain the five-term decomposition described in the text.

3) Counterfactual 2: Equilibrium

Sales: Same as in counterfactual 1 except that we now have to add export sales. With trade costs equal to τ , the increase in sales for exporters (dummy $\delta_n^X(a) = 1$) is given by $(1 + \tau^{1 - \sigma_n(z)})$, which yields the following change in sales:

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

Quality: same expression as in counterfactual 1.

Export decision: Let $r_n^X(a, \phi)$ denotes revenues net of variable costs on the export market (exports times $\frac{1}{\sigma_n}$) depending on its quality ϕ . Let $r_n^D(a, \phi)$ denote revenues net of variable costs on the domestic market. As before, $f_n(\phi)$ denotes the fixed costs of upgrading to quality ϕ which itself depends on whether the firm exports or not.

Profits earned by firm a if it does not export are then:

$$r_n^D(a, \phi_n^D(a)) - f_n(\phi_n^D(a))$$

where $\phi_n^D(a)$ denotes optimal quality when it does not export. Concersely, profits earned by firm a if it exports are equal to:

$$r_n^X(a, \phi_n^X(a)) + r_n^D(a, \phi_n^X(a)) - f_n(\phi_n^X(a)) - f_X$$

where $\phi_n^X(a)$ denotes optimal quality when it exports. In general, exporters can produce at higher quality since they have a larger size: $\phi_n^X(a) > \phi_n^D(a)$. Hence, this difference in optimal quality may influence the export decision, and it is important to take this potential quality upgrading decision into account in our counterfactual.

Exit decision: Similarly, the firm exits if $r_n^D(a, \phi_n^D(a)) - f_n(\phi_n^D(a)) - f_{n0} < 0$.

Entry decision: We model entry the same way as in the previous counterfactual, by imposing average profits (adjusting for exit) to remain constant.

Price index: As before, we can take the ratios of the price indexes in the counterfactual and baseline equilibria, and multiply the numerator and denominator by $\alpha_n(z)E(z)P_{n0}^{\sigma_n(z)-1}$ to express the price index change as a function of $x_{n0}(a, z)$ and the changes in $\varphi_n(a)$, $x_n(a, z)$ and $\tilde{\rho}_n(a)$.

We also multiply by $(1 + \delta_n^X(a)\tau^{1-\sigma_n(z)})$ all varieties that are traded since they are now available to both domestic and foreign consumers (hence, symmetrically, domestic consumers can access both foreign and domestic varieties of firms of productivity a).

4) Counterfactual 2: Decompositions

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

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

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

Like for the first counterfactuals, taking logs and a first-order approximation leads to:

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

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

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

Like before, using the equality $AB - A'B' = (A - A') \left(\frac{B+B'}{2} \right) + (B - B') \left(\frac{A+A'}{2} \right)$, we can rewrite the first two lines as well as the last two lines of the previous sum:

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

where $\bar{s}_{n1}(a, z)$ is the average of $s_{n1}(a, z)$ and $s_{n1}(a, z_0)$, and $s_{n1}(a, z)$ is now constructed as:

$$s_{n1}(a, z) = \frac{s_{n0}(a, z) \delta_{nD}(a) (1 + \delta_X(a) \tau^{1-\sigma_n(z)})}{\int_a s_{n0}(a, z) \delta_{nD}(a) (1 + \delta_X(a) \tau^{1-\sigma_n(z)})}$$

Also, the term $\bar{\delta}_{nD}(1 + \bar{\delta}_X \tau^{1-\bar{\sigma}_n})$ corresponds to the average of $\delta_{nD}(a)(1 + \delta_X(a) \tau^{1-\sigma_n(z_0)})$ across consumers and firms.

Appendix D: Extension with Multi-Product Firms

1) Heterogeneity in 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, 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:

$$x_{nb}^{Tot}(z) = \left(\frac{\varphi_{nb}^{Tot}(z)}{P_{nb}^{brand}(z)} \right)^{\sigma_n(z)-1} \alpha_n(z) E(z) P_n(z)^{\sigma_n(z)-1} \quad (31)$$

while sales by product can be written as:

$$x_{nbi}^{MP}(z) = \left(\frac{\varphi_{nbi}^{MP}(z)}{p_{ni}} \right)^{\eta_n(z)-1} x_{nb}^{Tot}(z) P_{nb}^{brand}(z)^{\eta_n(z)-1} \quad (32)$$

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

$$P_n(z) = \left[\sum_{i \in G_n} P_{nb}^{brand}(z)^{1-\sigma_n(z)} \varphi_{nbi}^{Tot}(z)^{\sigma_n(z)-1} \right]^{\frac{1}{1-\sigma_n(z)}} \quad (33)$$

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

$$P_{nb}^{brand}(z) = \left[\sum_{i \in G_n} p_{ni}^{1-\eta_n(z)} \varphi_{nbi}^{MP}(z)^{\eta_n(z)-1} \right]^{\frac{1}{1-\eta_n(z)}} \quad (34)$$

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:

$$x_{nb}^{Tot}(z) = N_{nb}^{\frac{\sigma_n(z)-1}{\eta_n(z)-1}} \left(\frac{\varphi_{nb}^{Tot}(z)}{\bar{P}_{nb}(z)} \right)^{\sigma_n(z)-1} \alpha_n(z) E(z) P_n(z) \sigma_n(z)^{-1} \quad (35)$$

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.

2) 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:

$$\mu_{nbi} = \frac{\sum_z x_{nbi}(z)}{\sum_z \eta_n(z) x_{nbi}(z)} \left[1 + \frac{\sum_z (\eta_n(z) - \sigma_n(z)) \bar{\mu}_{nb}(z) x_{nbi}(z)}{\sum_z x_{nbi}(z)} \right] \quad (36)$$

or equivalently:

$$\mu_{nbi} = \frac{\sum_z x_{nbi}(z)}{\sum_z \sigma_n(z) x_{nbi}(z)} \left[1 + \frac{\sum_z (\eta_n(z) - \sigma_n(z)) (\bar{\mu}_{nb}(z) - \mu_{nbi}) x_{nbi}(z)}{\sum_z x_{nbi}(z)} \right] \quad (37)$$

In equation 36, 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 37, 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 homogenous across consumers, then markups μ_{nbi} are homogeneous and equal $\frac{1}{\sigma_n}$ across all products.

3) 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 37), 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.

Appendix E: Additional Figures and Tables

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



The figure depicts the relationship between our measure of log expenditure per capita and reported nominal income brackets two years before across fourteen semester cross-sections between 2006-2012. The y-axis displays within-semester 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 semester. Standard errors in both graphs are clustered at the county level, and the displayed confidence intervals are at the 95% level.

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	BISSELL POWER FORCE	0.065	2.084
HEALTH & BEAUTY CARE	PAIN REMEDIES - HEADACHE	ADVIL	GENERIC STORE BRAND	0.078	0.086
HEALTH & BEAUTY CARE	SANITARY NAPKINS	ALWAYS MX PD/WG ULTR THN OVRNT	GENERIC STORE BRAND	0.030	1.591
HEALTH & BEAUTY CARE	SHAMPOO-AEROSOL/ LIQUID/ LOTION/ POWDER	PANTENE PRO-V	ALBERTO VO5	0.109	1.444
NON-FOOD GROCERY	CIGARS	HAV-A-TAMPA	POM POM OPERAS	0.023	0.375
NON-FOOD GROCERY	DETERGENTS - HEAVY DUTY - LIQUID	TIDE - H-D LIQ	PUREX - H-D LIQ	0.283	0.779
NON-FOOD GROCERY	SOAP - BAR	DOVE	DIAL	0.221	0.772
PACKAGED MEAT	BACON-REFRIGERATED	OSCAR MAYER	BAR S	0.214	0.961
PACKAGED MEAT	BRATWURST & KNOCKWURST	JOHNSONVILLE	KLEMENT'S	0.678	0.141
PACKAGED MEAT	FRANKS-COCKTAIL-REFRIGERATED	HILLSHIRE FARM	CAROLINA PRIDE	0.388	0.243

Figure A.2: 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	-

Figure A.3: 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	Cross-Brand Both IVs	Within-Brand Both IVs
(1- σ) All Households	-1.183*** (0.0440)	-1.223*** (0.0653)		
(1- σ) Below Median Quintiles			-1.376*** (0.0815)	-1.303*** (0.0859)
(1- σ) Median and Above Quintiles			-1.133*** (0.0444)	-1.203*** (0.0688)
Quintile-by-Module-by-County-by-Semester FX	✓	✗	✓	✗
Quintile-by-Module-by-Brand-by-County-by-Semester FX	✗	✓	✗	✓
Observations	3,980,418	7,222,751	3,980,418	7,222,751
First Stage F-Stat	348.7	492.4	312.5	410.8
Estimate of Ratio of σ 's (Poor/Rich)			1.114 (0.0388)	1.046 (0.0378)
95% Confidence Interval of Ratio			[1.0378, 1.19]	[0.972, 1.12]


Table A.2: 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	0.150*** (0.0368)	-1.163*** (0.0545)	-1.137*** (0.0490)	-1.183*** (0.0440)	0.0780*** (0.0278)	-1.138*** (0.0553)	-1.060*** (0.0430)	-1.138*** (0.0411)
Quintile-by-Module-by-County-by-Semester FX	✓	✓	✓	✓	✓	✓	✓	✓
Observations	4,804,155	4,804,155	3,980,418	3,980,418	4,804,155	4,804,155	3,980,418	3,980,418
First Stage F-Stat		723.0	176.0	348.7		763.1	166.6	370.9
<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	0.156*** (0.0359)	-1.000*** (0.0559)	-1.040*** (0.0552)	-1.046*** (0.0483)	0.0791*** (0.0270)	-1.014*** (0.0566)	-0.996*** (0.0481)	-1.038*** (0.0459)
Quintile-by-Module-by-County-by-Semester FX	✓	✓	✓	✓	✓	✓	✓	✓
Observations	4,804,155	4,804,155	3,980,418	3,980,418	4,804,155	4,804,155	3,980,418	3,980,418
First Stage F-Stat		681.0	174.2	304.9		761.6	167.2	357.1
<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	0.159*** (0.0361)	-1.175*** (0.0558)	-1.150*** (0.0528)	-1.196*** (0.0460)	0.0912*** (0.0278)	-1.158*** (0.0565)	-1.079*** (0.0467)	-1.158*** (0.0440)
Quintile-by-Module-by-County-by-Semester FX	✓	✓	✓	✓	✓	✓	✓	✓
Observations	4,804,155	4,804,155	3,980,418	3,980,418	4,804,155	4,804,155	3,980,418	3,980,418
First Stage F-Stat		638.7	166.6	293.1		633.9	159.4	285.8

Table A.3: 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	-1.285*** (0.358)	-1.001*** (0.172)	-1.593*** (0.225)	-1.613*** (0.407)	-1.908 (1.265)	-1.104** (0.476)	-0.968** (0.407)	-1.906*** (0.476)
(1- σ) 2nd Poorest Quintile	-2.055*** (0.384)	-0.872*** (0.291)	-1.417*** (0.145)	-1.525*** (0.265)	-2.347*** (0.387)	-0.0706 (0.524)	-1.464*** (0.395)	-1.610*** (0.365)
(1- σ) Median Quintile	-0.540* (0.279)	-0.423** (0.176)	-1.381*** (0.0921)	-1.520*** (0.219)	-1.253 (0.796)	-0.648 (0.607)	-0.422 (0.285)	-0.844** (0.376)
(1- σ) 2nd Richest Quintile	-1.050*** (0.317)	-0.829*** (0.171)	-1.316*** (0.0844)	-1.322*** (0.212)	-3.262*** (0.327)	-0.464*** (0.131)	-1.116*** (0.182)	-0.991*** (0.357)
(1- σ) Richest Quintile	-0.909*** (0.208)	-0.599*** (0.106)	-1.246*** (0.0828)	-1.445*** (0.182)	-2.006*** (0.410)	-0.612* (0.341)	-1.042*** (0.175)	-1.444*** (0.264)
Quintile-by-Module-by-County-by-Semester FX	✓	✓	✓	✓	✓	✓	✓	✓
Observations	304,797	347,321	1,909,138	423,660	98,279	352,567	433,769	110,837
First Stage F-Stat	139.0	347.5	254.1	50.17	131.4	109.4	298.0	37.68

Figure A.4: Households Agree on Product Quality Evaluations (But Rich Households Value Quality Relatively More)



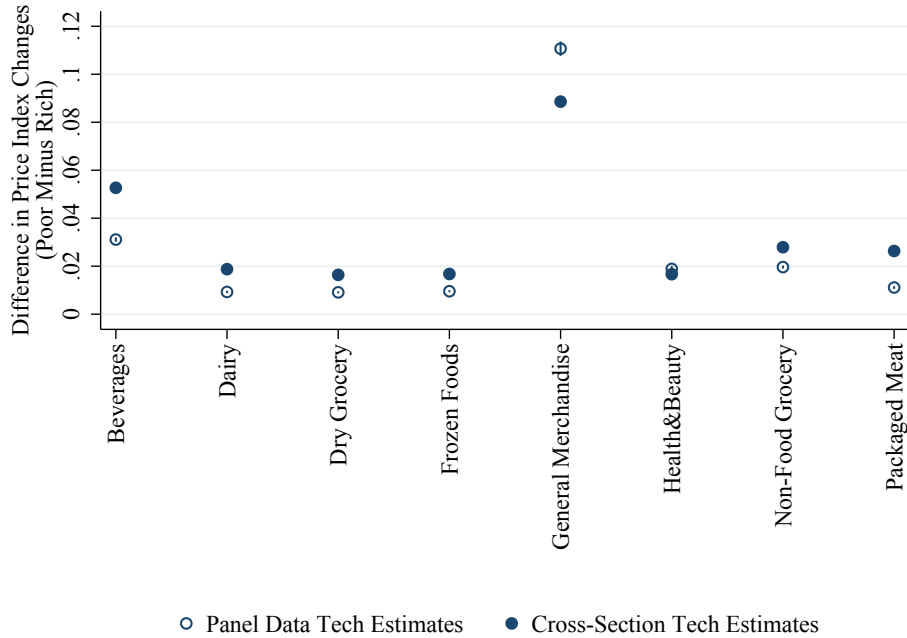
The figure depicts the relationship between deviations in log brand quality or quality adjusted prices and deviations in log firm total sales for on average more than 150,000 producers of brands during 14 half year periods between 2006-12. We estimate brand-level quality and quality adjusted prices as evaluated by each quintile of total household per capita expenditure as discussed in Sections 4 and 5.

Table A.4: Technology Parameter Estimates

Dependent Variable: Log Product Quality or Changes in Log Quality	BEVERAGES				DAIRY				DRY GROCERY				FROZEN FOODS			
	Cross-Section		Panel Data		Cross-Section		Panel Data		Cross-Section		Panel Data		Cross-Section		Panel Data	
	OLS	IV	OLS	IV	OLS	IV	OLS	IV	OLS	IV	OLS	IV	OLS	IV	OLS	IV
Log Firm Scale or Changes in Log Firm Scale (β)	1.0323*** (0.0216)	1.0797*** (0.0278)	1.0928*** (0.0267)	-0.5623 (0.9029)	1.5531*** (0.0089)	1.567*** (0.0102)	1.3506*** (0.0624)	0.2127 (0.4624)	0.8287*** (0.0029)	0.8364*** (0.0032)	0.855*** (0.0071)	0.3898*** (0.0736)	0.7621*** (0.0037)	0.7716*** (0.0047)	0.7415*** (0.013)	0.0598 (0.2073)
ξ Parameter	0.63	0.63	0.01	0.01	0.51	0.51	0.76	0.76	0.60	0.60	0.01	0.01	0.63	0.63	0.86	0.86
Observations	89,229	89,229	75,174	75,174	67,701	67,701	50,259	50,259	394,273	394,273	427,413	427,413	69,799	69,799	68,654	68,654
Number of Clusters	68	68	64	64	45	45	43	43	392	392	380	380	76	76	73	73
First Stage F-Stat	24102.14		4.83		30113.71		13.41		110350.95		110.15		19125.2		20.1	


Dependent Variable: Log Product Quality or Changes in Log Quality	GENERAL MERCHANDISE				HEALTH & BEAUTY CARE				NON-FOOD GROCERY				PACKAGED MEAT			
	Cross-Section		Panel Data		Cross-Section		Panel Data		Cross-Section		Panel Data		Cross-Section		Panel Data	
	OLS	IV	OLS	IV	OLS	IV	OLS	IV	OLS	IV	OLS	IV	OLS	IV	OLS	IV
Log Firm Scale or Changes in Log Firm Scale (β)	0.5696*** (0.005)	0.5722*** (0.0057)	0.5678*** (0.0112)	0.2046** (0.101)	1.9852*** (0.0074)	1.996*** (0.0071)	1.8943*** (0.0147)	1.4752*** (0.1413)	1.0657*** (0.0077)	1.0689*** (0.0078)	1.1167*** (0.0136)	0.0367 (0.5569)	0.8744*** (0.0062)	0.8866*** (0.0057)	0.8004*** (0.0197)	-0.0885 (0.4914)
ξ Parameter	0.91	0.91	0.86	0.86	0.27	0.27	0.96	0.96	0.48	0.48	0.01	0.01	0.55	0.55	0.01	0.01
Observations	97,506	97,506	95,054	95,054	144,699	144,699	195,982	195,982	101,640	101,640	129,844	129,844	21,620	21,620	18,980	18,980
Number of Clusters	140	140	102	102	172	172	154	154	125	125	115	115	11	11	10	10
First Stage F-Stat	26208.16		32.22		39273.84		49.94		45276.84		6.77		25994.86		7.21	

Figure A.5: Counterfactual 1: Inflation Differences across Product Departments



The figure depicts mean deviations in household retail price index changes for on average 58 thousand US households during 14 half year periods between 2006-12. The estimated price index changes correspond to the counterfactual where 5 percent of total market sales are reallocated from the poorest household income group to the richest as discussed in Section 6. Both graphs display confidence intervals at the 95% level.

Figure A.6: Counterfactual 2: Inflation Differences across Product Departments



The figure depicts mean deviations in household retail price index changes for on average 58 thousand US households during 14 half year periods between 2006-12. The estimated price index changes correspond to the third counterfactual discussed in Section 6. Both graphs display confidence intervals at the 95% level.