NBER WORKING PAPER SERIES

# ARE POOR CITIES CHEAP FOR EVERYONE? NON-HOMOTHETICITY AND THE COST OF LIVING ACROSS U.S. CITIES

Jessie Handbury

Working Paper 26574
http://www.nber.org/papers/w26574

NATIONAL BUREAU OF ECONOMIC RESEARCH
1050 Massachusetts Avenue
Cambridge, MA 02138
December 2019

This paper was the main chapter of my Ph.D. dissertation. I am grateful to my advisors David Weinstein and Kate Ho for invaluable advice and discussions. I thank Songyuan Ding, Yue Cao, and Bo Cheng for outstanding excellent research assistance. I also thank David Bieri, Chris Conlon, Don Davis, Jonathan Dingel, Gilles Duranton, Galina Hale, Jean-Francois Houde, Amit Khandelwal, Joan Monras, Aviv Nevo, Eleonora Patacchini, Mike Riordan, Molly Schnell, Katja Seim, Ina Simonovska, Todd Sinai, Daniel Sturm, Eric Verhoogen, Jonathan Vogel, and Maisy Wong and several presentation attendees for valuable discussions and comments. I gratefully acknowledge financial support from the Program of Economic Research at Columbia and the Research Sponsors' Program of the Wharton Zell/Lurie Real Estate Center. Researcher's own analyses are calculated (or derived) based in part on data from The Nielsen Company (US), LLC and marketing databases provided through the Nielsen Datasets at the Kilts Center for Marketing Data Center at The University of Chicago Booth School of Business. The conclusions drawn from the Nielsen data are those of the researcher and do not reflect the views of Nielsen, nor do they necessarily reflect the views of the National Bureau of Economic Research.. Nielsen is not responsible for, had no role in, and was not involved in analyzing and preparing the results reported herein. All errors are my own.

Are Poor Cities Cheap for Everyone? Non-Homotheticity and the Cost of Living Across U.S. Cities
Jessie Handbury
NBER Working Paper No. 26574
December 2019
JEL No. F11,R11,R22,R32,L81

**ABSTRACT**

This paper shows that the products and prices offered in markets are correlated with local income-specific tastes. To quantify the welfare impact of this variation, I calculate local price indexes micro-founded by a model of non-homothetic demand over thousands of grocery products. These indexes reveal large differences in how wealthy and poor households perceive the choice sets available in wealthy and poor cities. Relative to low-income households, high-income households enjoy 40 percent higher utility per dollar expenditure in wealthy cities, relative to poor cities. Similar patterns are observed across stores in different neighborhoods. Most of this variation is explained by differences in the product assortment offered, rather than the relative prices charged, by chains that operate in different markets.

Jessie Handbury
The Wharton School
University of Pennsylvania
1463 Steinberg-Dietrich Hall
Philadelphia, PA 19104
and NBER
handbury@wharton.upenn.edu

# 1 Introduction

It is well known that prices and product variety vary systematically across space: high-end goods are more available in rich neighborhoods than poor ones. Yet the cost-of-living indexes that economists employ to account for these spatial price differences aggregate prices using the same expenditure weights for all consumers, implicitly assuming that tastes do not vary with income.[1,2] Under this assumption, a high-income Washington D.C. resident would be indifferent between the set of goods available in their local stores and the set available in a city with less than half the per capita income, like Detroit. In reality, preferences are non-homothetic.[3] This paper is the first to study the implications of non-homotheticity for spatial price indexes.[4]

I first document how availability and prices of grocery products varies with local income across U.S. cities as well as across neighborhoods within these cities. To measure the implications of these spatial availability and pricing patterns for the welfare of consumers at different income levels, I next develop a model of non-homothetic demand. I estimate the model with a combination of data describing the aggregate sales of different products in a sample of stores across the U.S. and the purchases of individual households in those stores. I use the estimated model to construct price indexes that summarize how households at different income levels value the prices and products available to them in different geographic markets. Finally, I characterize how and why the price level varies across cities and neighborhoods in the U.S. differently for consumers at different income levels. This analysis yields three sets of novel results.

First, stores favor high-income consumers more in wealthy locations than in poor ones through both their pricing and product offerings. Stores in wealthier cities offer products representing a greater share of the high-income consumption bundle than the low-income consumption bundle. Stores in wealthier cities also charge relatively less for the high-income consumption bundle than the low-income one, conditional on availability. The same patterns are observed across stores in different neighborhoods of the same city.

[1]Albouy (2009) and Moretti (2013), for example, use the ACCRA indexes to calculate real tax burdens and income inequality accounting for intra-national price variation, while Deaton (2010) and Almas (2012) use homothetic indexes based on the Penn World Table data to calculate poverty thresholds and real income inequality adjusting for international price variation.

[2]Notable exceptions include Deaton and Dupriez (2011a) who calculate country-specific poverty thresholds based on purchasing power parity deflators that reflect the consumption patterns of the global poor, and Li (2012) who uses income-specific price indexes to measure the difference in the potential welfare gains from variety for high- relative to low-expenditure households moving from rural to urban areas in India.

[3]Deaton and Muellbauer (1980) note that homotheticity is consistently violated in cross-sectional household studies. More direct evidence includes Bils and Klenow (2001), Broda et al. (2009), Faber and Fally (2017), Hottman and Monarch (2018), and Jaimovich et al. (2019) for the U.S., Faber (2014) and Cravino and Levchenko (2017) for Mexico, and Li (2012) for India.

[4]The importance of recognizing the these non-homotheticities in regional price indexes was recognized over 50 years ago in Snyder (1956). Related concurrent work has considered the impact of non-homotheticity in demand for housing in measuring real income inequality (Albouy et al., 2016).

Second, these differences in availability and pricing matter for consumers. Income-specific spatial price indexes reveal large differences in how high- and low-income households perceive the prices and variety available in different U.S. cities. Once you account for income-specific tastes, markets that are relatively expensive for poor households can be instead relatively cheap for the wealthy. For example, a low-income household earning \$25,000 a year faces approximately 2 percent higher grocery costs in Bridgeport, CT, with per capita income \$50,000, relative to El Paso, TX, with per capita income below \$20,000. But the same is not true for high-income households earning \$200,000 a year whose grocery costs are 47 percent lower in Bridgeport than in El Paso. On average, high-income households perceive the difference in grocery costs between high- and low-income cities to be 40 percentage points larger than low-income households.

Third, I show that the differences in relative grocery costs across cities are driven more by cross-city variation in product variety than by variation in prices. Higher income households find groceries cheaper in wealthier cities primarily because more varieties of the high-quality products that high-income consumers prefer to consume are available in these locations. These high-quality products are sold at lower unit prices relative to low-quality products in wealthy cities, but these price differences only explain a small portion of the gap between the grocery costs perceived by high- and low-income households across wealthy and poor cities. This result points towards a second short-coming of conventional price indexes, which compare only the prices of common goods, and not variety differences, across locations.[5] Even if they are non-homothetic, price indexes that do not account for differences in product availability will fail to capture any of the true cost-of-living differences for wealthy, relative to poor, consumers.

I also study how store-level price indexes vary across and within cities. I find that higher income households face relatively lower price indexes in stores located in higher income neighborhoods, even within the same CBSA. In fact, the cross-CBSA variation in income-specific price indexes is strongest between stores located in above median neighborhoods within each CBSA. Thus, within-city sorting can maximize a wealthy consumer's variety gains from living in a wealthy city, and mitigate the relative losses for a poor consumer. I also use these indexes to better understand why variety varies across and within cities. Here I find that the variation in variety offerings across CBSAs and neighborhoods is entirely driven by variation in the local mix of retail chains. There is no systematic variation in the price indexes high- and low-income households face across stores belonging to the same retail chain.

The main methodological challenge I overcome in this paper is to summarize the costs

[5] Handbury and Weinstein (2014) find a huge amount of variation in availability of grocery varieties across U.S. cities and show that conventional price indexes underestimate the correlation between city size and the grocery price level, for a homothetic representative consumer, by about a third. Variety differences play a much larger role here, explaining all of the positive correlation between city income and the differences in the grocery price levels faced by wealthy, relative to poor, consumers.

that consumers face across multiple differentiated product categories in a way that parsimoniously accounts for the non-homothetic tastes demonstrated in household behavior. To do this, I build income-specific price indexes. A major reason why existing regional price indexes do not take non-homotheticities into account is that the single-sector models used to identify non-homotheticities in micro studies do not lend themselves to aggregation. I nest a variant of these micro models into a utility framework that represents non-homothetic preferences across many sectors of differentiated products. The starting point for these price indexes is the log-logit/constant elasticity of substitution (CES) family of utility functions. Log-logit sub-utility functions govern how consumers allocate expenditures between products within product categories, while a CES superstructure governs the substitutability of products across different categories. The key feature of this structure is that it can be aggregated in such a way that one could also express aggregate demands for goods as if they had been derived from a representative (non-homothetic) household.[6] This provides a way of bridging the gap between the micro data that I use to identify parameters and an aggregate non-homothetic price index that can be used to compare price levels across locations.

The model nests two forms of non-homotheticity and is structured in a way that enables me to test for their relative importance in explaining the differences between the purchases of high- and low-income consumers. The elasticity of demand with respect to price and product quality depends on the consumer's expenditure on a composite of non-grocery products which I assume to be normal. The intuition here is that, if high-income households spend more on cars, schooling, and mortgages, for example, then they have a greater willingness to pay for their own ideal product variety or for products that are ranked as high quality by all consumers. These are the most common ways in which international economists hypothesize that non-homotheticities might matter.[7] Where previous papers have verified each of these channels of non-homotheticity independently, this is the first to test their empirical relevance concurrently and to assess their relative importance in explaining consumer behavior. My results demonstrate the salience of non-homothetic demand for quality in U.S. grocery consumption. I compare three different models of non-homotheticity: a specification in which the taste for quality rises with income, a specification in which high-income households are less price sensitive, and a

[6]The log-logit and CES are linked mathematically such that the CES-nested log-logit utility framework yields the same aggregate outcomes as a nested-CES utility function. The origins of this result are Anderson et al. (1987), whose proof is extended to models that account for product quality in Verhoogen (2008). This link has also been explored recently in Hortaçsu and Joo (2015) who present a generalized version of the demand system developed here that allows for tastes for product quality to vary with both observed and unobserved consumer attributes.

[7]Hummels and Lugovskyy (2009) and Simonovska (2010), for example, are based on the idea that substitution elasticities vary systematically with income, while Fajgelbaum et al. (2011) and Faber (2012) model non-homotheticities as a changing taste for quality. There are other reasons that demand may vary with income, related to demand for variety (Li (2012)) and shopping behavior (Aguiar and Hurst (2005)). These do not appear to be the primary factors driving differences in the purchases of high- and low-income households in this dataset and are, therefore, not included in the model.

specification in which both factors play a role. I find that the specification that allows for non-homothetic demand for quality alone explains the differences between the purchases of rich and poor households most parsimoniously.[8]

The main contribution of this paper is to provide the first direct evidence of income-specific tastes for local consumption amenities. A recent urban economics literature hypothesizes that these tastes may help explain spatial disparities in income and skill observed across U.S. cities: high-skill, high-income workers co-locate because they enjoy more utility from certain endogenous local amenities than low-skill, low-income consumers.[9] Previous empirical support of this theory relies on spatial equilibrium models that assume people are perfectly mobile. Diamond (2016), for example, infers changes to skill-biased amenities as those which reconcile changes in housing price and wage data with the changing skill composition of U.S. cities between 1980 and 2000.[10] I instead measure these skill-specific amenities directly, providing cross-sectional evidence that non-housing price indexes are correlated with local incomes in such a way that might encourage further skill-biased agglomeration.

In particular, I show that product variety is skewed towards the income-specific tastes of local consumers. This result is consistent with the theory that, in markets with increasing returns and demand heterogeneity, differentiated product firms cater to local tastes generating "preference externalities" or "home market effects." Fajgelbaum et al. (2011), for example, show theoretically that high-income consumers with non-homothetic preferences enjoy greater consumption utility when living in high-income countries. Like Waldfogel (2003), I provide evidence suggesting that the mechanism behind these effects is local distributors catering to local tastes. My main contribution here, however, is to demonstrate the economic significance of these externalities by measuring their impact on consumer costs. My results showing that these preference externalities are mediated by chain-level pricing and product assortment decisions corroborate a growing literature on these decisions (DellaVigna and Gentzkow (2019)) and the role that they play in generating cross-city variation in aggregate variety (Hottman (2014)).[11,12]

[8]Faber and Fally (2017) estimate the same demand system non-parametrically using only the household-level data and also find that the differences in price elasticities across income quintiles are small relative to the cross-quintile differences in the elasticities of demand for quality.

[9]Glaeser et al. (2001) brought attention to the role of high-skilled workers as consumers of urban amenities. Diamond (2016) and Couture and Handbury (2017) thereafter studied the differential role that access to amenities play in explaining the changing location choices of the college and non-college educated across and within U.S. cities over recent decades.

[10]Black et al. (2009) show that returns to education are inversely related to housing prices and posit that non-homothetic preferences reconcile these results with a spatial equilibrium model. Non-homothetic preferences for endogenous private amenities, such as the distribution of high-quality retail goods, are one factor that makes high house price cities less expensive for the rich.

[11]The observed distribution of product availability is also consistent with a comparative advantage story and my analysis does not identify this story from the preference externalities. More recent work by Dingel (2016) shows that the specialization of high-income counties in exporting high-quality products is explained as much by home-market demand as by differences in factor usage and endowments.

[12]Complementary work finds variation in inflation across income groups. The BLS has a long tradition of

These results have mixed implications for the question of how to account for cost-of-living differences across locations when measuring welfare. Standard homothetic price indexes implicitly ignore that households with different incomes have different tastes and, therefore, may perceive these relative costs differently. I find that these cost differences are large in the context of non-durable goods. If similar group-specific externalities are at play in other non-tradable sectors (such as housing, non-tradable services, and durables), it may be necessary to account for income-specific tastes when measuring relative real incomes and expenditures of households at opposite ends of the income distribution. Such adjustments may, for example, have implications for the recent findings on how ignoring intra-national price variation biases measures of real income inequality (Moretti, 2013; Albouy et al., 2016) and the geographic distribution of real tax expenditures in the U.S. (Albouy, 2009).[13] Finally, these results suggest that it may also be worth revisiting whether to use homothetic price indexes to account for location-specific costs when calculating poverty thresholds or entitlement payments (*e.g.*, Slesnick (2002), Deaton (2010), and Ziliak (2011)).[14]

# 2 Data

The analysis in this paper is based on detailed store sales and household purchase data, both from Nielsen.[15] I use the store sales data to infer the set of products and prices available in U.S. cities and the household purchase data to identify how consumers at different income levels value these products and prices. These two Nielsen datasets are available from 2006 onward. I analyze data from a single year, 2012, so as to abstract from dynamics in both the product set and tastes. I complement the 2012 Nielsen data with 5-year 2010-2014 average of tract-,

using confidential survey data to construct inflation indexes that use income-specific expenditure weights (see, e.g., Snyder (1961); Kokoski (1987); Jorgenson et al. (1989); Garner et al. (1996); Cage et al. (2002)). More recent papers apply a method developed by Broda and Romalis (2009) to calculate income-specific exact price indexes for the U.S. with the same household purchase data used here (Argente and Lee, 2017; Jaravel, 2018). On the structural side, Albouy et al. (2016) quantify a model of non-homothetic housing demand to show that the poor have been disproportionately impacted by rising relative rents in the U.S., and Atkin et al. (2019) use an AIDS model to calculate aggregate income-specific inflation rates for Indian households.

[13]Results here suggest that we underestimate the relative tax burden of high-income consumers (Albouy, 2009) and overestimate welfare inequality between individuals with and without college degrees (Moretti, 2013) if when we do not account for the fact that the wealthy and college-educated are more likely to live in cities with high house prices. The results here suggest that income-biased preference externalities might offset the expensiveness of wealthy cities for high-income households.

[14]By contrast, Deaton and Dupriez (2011b) find that re-weighting the International Comparison Project (ICP)'s country-level purchasing power parity (PPP) indexes to reflect the consumption patterns of the world's poor does not change the indexes or, therefore, poverty counts dramatically. These indexes are not adjusted for spatial variation in product availability, however, so do not account for the variety differences that drive the price index differentials documented here.

[15]These data are provided by the Kilts-Nielsen Data Center at the University of Chicago Booth School of Business.

county-, and CBSA-level population and income data from the American Community Survey (ACS) to measure how prices and product availability co-vary with local wealth across cities and neighborhoods.[16] In what follows, I describe the structure of each Nielsen dataset and the key variables I draw from them. Further details are available in Appendix A.

The Nielsen store-level (RMS) data contains a panel of weekly sales and quantities by Universal Product Code (UPC) collected by point-of-sale systems in over 30,000 participating retailers across the U.S., along with the county in which each store is located. I complement the RMS data with the Nielsen household-level (HMS) data, which contains information on all bar-coded product purchases made by a panel of over 100,000 households in markets across the United States. Each household in this sample was provided with a bar-code scanner and instructed to collect information such as the UPC, the value and quantity, the date, and the name, location, and type of store for every purchase they made. Nielsen also surveys each household to collect information on, among other things, income, household size, and residential 5-digit zip code.

The RMS data is collected in an automated process so it is less prone to measurement error than the HMS household survey data. As such, the RMS data is better-suited for the construction of non-linear sales share moments I use to identify price elasticity and quality parameters common to all households. The HMS data, meanwhile, provides a detailed picture of the products selected by households at different income levels in the same store and is useful for documenting differences in purchases by income level, controlling for their choice set, and estimating the parameters that generate these differences in the model.

The HMS data also allows me to obtain a more precise estimate of household income in the neighborhoods surrounding each store. I measure the income distribution in a store's vicinity with a distance-weighted average of the income distributions observed in the Census tracts within 30km of the centroid of the modal residential zip code of Nielsen panelists that report shopping at that store over all available years (2006 through 2017).

### Product Definitions

Nielsen categorizes UPCs into categories called "modules" and provides a UPC's brand, size (including units), and container count.[17] Within each module, I aggregate UPCs into a classification that I call a "product." A product is defined as the set of UPCs within a module with the same brand. For example, in the module "SOFT DRINKS - CARBONATED", there are 104 UPCs that refer to drinks sold under the brand "COCA-COLA R" (R stands for regular, as

[16]The ACS data are sourced from NHGIS (Manson, Schroeder, Van Riper, and Ruggles, Manson et al.).

[17]The container count is equal to one when each container of the product is sold individually and greater than one when multiple containers of the good are sold in a multi-pack.

opposed to diet). These UPCs belong to the same product.[18]

Table 1 shows how UPCs are distributed across products and modules in the sample used to estimate demand. This sample has been cleaned in various ways. To ensure that differences in container sizes or multi-packs do not mechanically generate spurious differences in prices in my sample, I define prices on a per unit basis throughout the paper, using the modal unit definition for each module. So, I limit my attention to products whose container size is expressed in the modal units for their module and exclude modules whose modal container size is either not expressed in meaningful units (e.g., counts instead of weights or volume) or in the same units for at least 75% of UPCs. To avoid differences in product quality that could be correlated with store amenities or neighborhood income, I exclude random weight items.[19] To control for data recording errors, I drop any market (store-month) in which I observe a UPC sold at a unit price greater than three times or less than a third of the median unit price paid per unit of any UPC within the same product or module categorization. For computational reasons, I put products whose average positive sales shares across store-month markets fall below the 60th percentile into an outside product and drop sales from any markets that sell less than two non-outside products. Finally, for identification purposes, I limit my attention to modules that have some overlap between the product-store-month RMS store sales data and the HMS household purchase data and to products that are sold in 5 or more of the remaining markets. The cleaned data contains approximately 200,000 UPCs categorized into 22,655 products across 530 product modules. Almost three quarters of these products are purchased by households in the HMS data. The median numbers of products and UPCs per module are 28 and 119, respectively.

The utility function presented below assumes that, conditional on price, consumers do not differentiate between UPCs in the same product. The assumption might be violated in cases where different UPCs that I have defined to be the same product are differentiated by their packaging or flavor. To check the extent to which consumers differentiate between UPCs within product categories, I compared the coefficient of variation for the average unit price paid for each UPC with the coefficient of variation for the average unit price paid for the set of UPCs with the same product categorization. The median coefficient of variation of unit values across UPCs in a given module is 0.506, only slightly higher than the median coefficient of variation of unit values across products in a given module at 0.498, and the two statistics are highly correlated across modules ($\rho = 0.96$). This indicates that there is little variation in the prices charged for UPCs within the same product.

[18]The analysis abstracts from other product characteristics, such as container, flavor, size, and whether the product was sold in a multi-pack or not. Differentiating between products along these dimensions leads to many products with sales shares too low to allow for the matrix inversions required in the estimation procedure.

[19]The quality of random weight items, such as fruit, vegetables, and deli meats, varies over time as the produce loses its freshness and it is likely that stores set prices to reflect this. This potential inter-temporal correlation between their unobserved quality of random weight products and their prices would introduce biases in the price elasticities estimated below.

Table 1: Summary Statistics for the Nielsen Data Used in Estimation

| Data: | RMS (Store) | | | | | | | HMS (HH) |
|---|---|---|---|---|---|---|---|---|
| | Total | Count Per Module | | | Count Per Product | | | Total |
| | Count | Min | Median | Max | Min | Median | Max | Count |
| Modules | 530 | - | - | - | - | - | - | 530 |
| Products | 22,655 | 2 | 28 | 627 | - | - | - | 15,615 |
| UPCs | 203,049 | 2 | 119 | 6,595 | 1 | 6 | 1,347 | 106,118 |

Notes: This table shows the distribution of UPCs across product and module categories in the Nielsen RMS store sales and HMS household purchase data used for estimation. This estimation sample has been cleaned from the raw Nielsen data as described in Section 2 of the paper. A product is defined as the set of UPCs within a module with the same brand. The table does not include the "outside" product (into which 60 percent of products are allocated, in the base specification).

### Household Income

The Nielsen HMS data is uniquely suited for estimating how consumers at different income levels value products because it links detailed information on household purchases to information on their reported annual income and demographics. Nielsen classifies households into 16 brackets of reported income. For my analysis, I exclude households with reported incomes below $11,000 and/or missing demographic data. I convert household income to a continuous variable equal to the mid-point of the income range represented by their Nielsen income category and an income of $150,000 to the households in the "above $100,000" income category. I then adjust income for household size using a square-root equivalence scale.[20]

Figure 1 shows the distribution of the resulting size-adjusted household income across the households considered in the analysis below. The bulk of the distribution is between $10,000 and $80,000, which seems reasonable given that the per capita incomes of the cities represented in the sample ranges from approximately $30,000 to $60,000.[21]

### City-Level Product and Price Availability

I infer the products and prices available in CBSAs in 2012 with those that I observe in the sales of local outlets of Nielsen participating retailers in that year. Not all stores participate in the RMS sample, so I likely observe only a sub-set of the products available in each city. This sample might not be representative, so my measure of product availability and prices will be

[20]This simple rule of thumb has been employed by the OECD Income Distribution Database (IDD) since 2012 (http://www.oecd.org/els/soc/IDD-ToR.pdf).

[21]Nielsen under-samples low-income households and, to a lesser degree, high-income households (see Appendix Figures A.2), but has positive weight of households at most income levels – up to the top-code – which, combined with functional form assumptions, allows for the calculation of price indexes at all points along the income distribution.

Figure 1: Distribution of Size-Adjusted Household Income

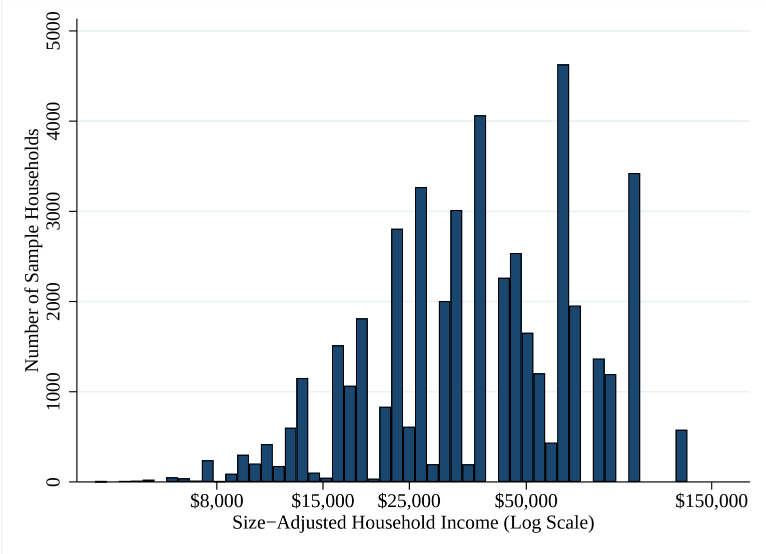


Notes: Plot depicts the number of households with a purchasing record in the 2012 Nielsen HMS data with non-missing demographic information and reported income above $11,000. Household income is adjusted for size by dividing by the square root of the number of household members.

subject to biases related to the number and type of stores sampled in each city.[22] To deal with these potential biases, I infer CBSA-level product availability and pricing using the sales of randomly-selected sub-samples of stores from each city. For the main analysis, I use products and unit prices represented in the sales of 50 randomly-selected stores, limiting my attention to 125 cities with 50 or more retailers in the RMS sample.[23]

# 3 Stylized Facts

This section draws on the Nielsen HMS and RMS data described above to document two stylized facts. Taken together, these facts demonstrate the empirical patterns behind the main results of the paper. The first also serves to motivate the theoretical framework presented in Section 4 below.

[22]This data limitation is common to all work that builds spatial price indexes from micro data. Handbury and Weinstein (2014) show how homothetic non-parametric cross-city price indexes can be adjusted to account for potential sample size biases. Unfortunately the parametric price index methodology used in this paper does not allow for such adjustments.

[23]Appendix A.4 lists this data for each of the 125 cities considered in this analysis and shows how the stores are distributed across these locations. Appendix Figure A.3 shows that the Nielsen participating retailer sample is over-weighted towards stores in higher-income neighborhoods, relative to the distribution of grocery stores in the County Business Patterns zip-level data, but only to a small degree.

## 3.1 High-Income Households Purchase Different, More Expensive, Products than Low-Income Households

Figure 2 shows that high-income households pay more than low-income households for the same type of products. The level of each circle shows how much more households in each Nielsen income category pay per unit for products within a module than households in the lowest income category, earning between $10,000 and $12,000. These relative prices are measured in a regression of log unit price paid against income category dummies and module fixed effects, controlling for other demographics with dummies for household size, marital status, race, Hispanic origin, and male and female head-of-household education and age. There is a distinct upward slope, with households in the upper-most income category paying approximately 17 percent more for products in the same module than households in the lowest income category. This could be either because high-income households are paying more for the same products within a module or because they are purchasing different, more expensive products. The following result suggests that the latter effect dominates.

Figure 2: Average Log Price Paid by Household Income Category

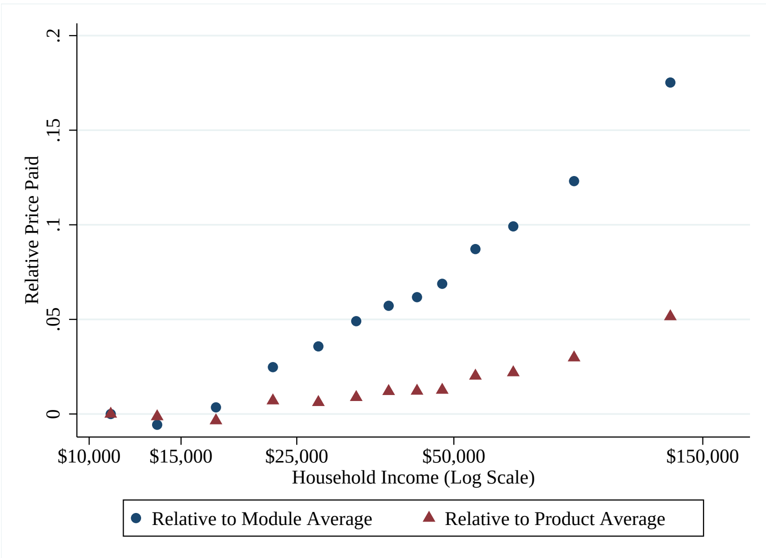


Notes: This figure plots the average unit price paid by Nielsen household panelists at different income levels relative to the unit price paid by all households for either the same product or products in the same module. Relative price paid is the coefficient on a household income dummy in a regression of the log unit price paid by a household for a product in a month on module or product fixed effects and demographic controls. The relative price paid by each household income category is plotted against the mid-point of the bounds of the reported incomes for that category for all but the highest "income greater than $100,000" category, whose relative price paid is plotted at $130,000.

The level of each triangle in Figure 2 shows how much more households in each Nielsen income category pay for the same product, relative to households in the lowest income category, measured in the same regression as described above but with product, instead of module, fixed effects. The slope of the log unit price paid controlling for product fixed effects is positive but much smaller than the slope of the log unit paid only controlling for module fixed effects. High-income households do pay more for the same products but, consistent with Broda et al. (2009), most of this gradient is explained by the fact that they are buying different products that are sold at higher prices to all consumers.

## 3.2 Stores in Wealthier Markets Offer More Products that are Purchased by High-Income than by Low-Income Households at Slightly Lower Relative Prices

Above, I showed that high-income households purchase different, more expensive, products to low-income households. Here I infer the extent to which high- and low-income households favor each product from the collective product-level expenditures of HMS panelists in each size-adjusted income decile. To assess the favorability of the product variety offered in a given market towards high- or low-income tastes, I then calculate how much more of the top income decile's consumption bundle is represented in the set of products available in a market than of the bottom income decile's consumption bundle. Then, to study whether the prices charged by stores in a market favor the tastes of high- or low-income households, I calculate the difference in the weighted average relative price charged for products in the top income decile's bundle from the weighted average relative price charged for products in the bottom income decile's bundle, using each income decile's expenditures as respective weights.[24]

Figure 3 shows how the difference in these availability and price indexes between the top and bottom income decile varies across CBSAs with different per capita incomes. The upper plot in Figure 3 shows a statistically-significant correlation between the city wealth and product availability: the consumption opportunities in high-income cities are skewed towards those products that are consumed more heavily by high-income consumers relative to those consumed more heavily by low-income consumers. For example, around 1.2 percentage points more of the top income decile's expenditure share than that of the bottom income decile is represented in the sample for the wealthiest city, Bridgeport-Stamford-Norwalk, CT (BRI), while 1 percentage point less is represented in the sample for the poorest city, El Paso, Texas (ELP). To put these differences into context, the mean CBSA has products representing 40 percent each of the top and bottom decile's expenditures, with a standard deviation across CBSAs of 1.9 and 1.4, respectively.

The lower plot in Figure 3 replicates this analysis looking instead at how the gap in the average relative price faced by high- and low-income households for the products they consume more of varies across CBSAs with different per capita income. The plot shows a noisier rela-

[24] Specifically, I calculate the price level in a market for a given income decile as the weighted average log ratio of the unit price a product is sold at in that market and the unit price it is sold at nationally, where products are weighted by the value of purchases of that product by the respective income decile of Nielsen household panelists.

Figure 3: Difference in the Availability and Relative Price of High-Income and Low-Income Baskets Across CBSAs

a. Availability

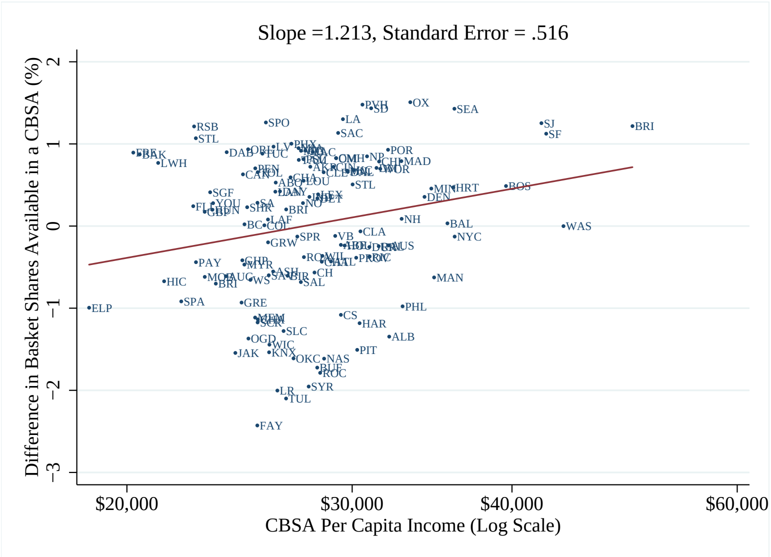


b. Relative Price


Notes: Figure a. plots CBSA-level data for the difference between the expenditure shares of high-income Nielsen HMS panelists represented in the CBSA product set and the expenditure share of low-income panelists represented in that product set against CBSA per capita income. The panelist expenditure shares are calculated for 2012 and are CBSA-specific, in that they exclude the expenditures of any panelists residing in the CBSA whose availability is being measured. Figure b. plots CBSA-level data for the difference between the average price level faced by consumers in the top income decile and the average price level faced by households in the bottom income decile against CBSA per capita income. The price level in each CBSA for a given income decile is calculated as the weighted average log of the ratio between the price a product is sold for in a CBSA relative to the price that product is sold at in the national sample where weights are defined as the value of the purchases of that product made by households in the respective income decile in the Nielsen household-level panel. Panelists are defined as high- (or low-) income if their size-adjusted income falls in the top (bottom) decile of panelist incomes. The products available and prices charged in each CBSA are defined as the set of products sold and average unit prices charged in a random sample of 50 Nielsen stores in a given CBSA in 2012. The plots show the mean availablility share and price indexes calculated in 100 bootstrap iterations of this sampling procedure. CBSA income is household income adjusted for size using a square-root equivalence scale. The marker labels for each CBSA are acronyms linked to the full CBSA name in Appendix A.4.

tionship. Stores in high-income CBSAs tend to charge less for the products that high-income households purchase more of (relative to low-income households) than stores in low-income CBSAs, but the slope is smaller (-0.9 relative to the standard deviation of the price levels across CBSA for the top and bottom deciles of 2.1 and 2.2, respectively).

In summary, there are large systematic differences between the products available in wealthy and poor cities that are correlated with the purchase behavior of high- and low-income households. Stores in wealthy cities also charge relatively less for products that the top income decile's consumption basket than the bottom income decile's consumption basket, but this correlation explains less of the variation in prices than the corresponding correlation in relative availability with local income explains of the variation in that variable.

Table 2 replicates this analysis comparing the products available and price charged across individual grocery stores, rather than across CBSAs.[25] Panel A compares availability patterns across stores. In column [1], we see that, in aggregate, stores in higher-income neighborhoods offer more of the products high-income households purchase more of. These availability patterns are stronger looking across stores within the same CBSA, in column [3], than across stores in CBSAs with different per capita incomes, in column [5]. In all three cases, the availability patterns are less than half as large when looking across stores in the same retail chain. The patterns in price levels, shown in Panel B, are similar, also favoring high-income consumers in higher-income neighborhoods and CBSAs, with less variation looking within chain than across chains. The only exception here is that the relative price charged for products that high-income consumers favor is less correlated with local income across stores in different neighborhoods of the same CBSA (column [3]) than across neighborhoods both within and across CBSAs (column [1]). Consistent with evidence of chain-level pricing in DellaVigna and Gentzkow (2019), this correlation falls almost to zero when looking within chain and CBSA (column [4]). In summary, the spatial differences in product availability and prices documented in this paper can be attributed primarily to variation in store location and product distribution patterns across chains, and less to variation in product distribution patterns across stores within the same chain.

In the structural analysis below, I will revisit these stylized facts in the context of a non-homothetic demand system. This analysis will serve to characterize which products are preferred by high-income households and quantify how much high-income households gain from the relative abundance and low prices of these products available in wealthy cities and neighborhoods across the U.S..

[25] In the store-level analysis, I compare pricing and availability across different grocery stores (listed in the Nielsen data as in the "food" channel), dropping mass merchandisers, drug, and convenience stores, which may exhibit different relative pricing and availability patterns.

Table 2: Difference in the Availability and Relative Price of High-Income and Low-Income Baskets Across Stores

Panel A: Availability

| Dependent Variable: Difference in Basket Shares (%) | | | | | | |
|---|---|---|---|---|---|---|
| | [1] | [2] | [3] | [4] | [5] | [6] |
| Ln(Local Per Capita Income) | 2.12*** | 0.70*** | 2.47*** | 1.07*** | | |
| | (0.053) | (0.036) | (0.088) | (0.042) | | |
| Ln(CBSA Per Capita Income) | | | | | 1.87*** | 0.49*** |
| | | | | | (0.44) | (0.18) |
| CBSA Fixed Effects | No | No | Yes | Yes | No | No |
| Chain Fixed Effects | No | Yes | No | Yes | No | Yes |
| Number of CBSAs | - | - | - | - | 691 | 691 |
| Observations | 9,019 | 9,019 | 8,849 | 8,849 | 9,019 | 9,019 |
| adj. $R^2$ | 0.15 | 0.79 | 0.56 | 0.89 | 0.08 | 0.78 |

Panel B: Relative Price

| Dependent Variable: Difference in Hedonic Price Index (%) | | | | | | |
|---|---|---|---|---|---|---|
| | [1] | [2] | [3] | [4] | [5] | [6] |
| Ln(Local Per Capita Income) | -1.30*** | -0.36*** | -0.58*** | -0.067* | | |
| | (0.029) | (0.024) | (0.055) | (0.036) | | |
| Ln(CBSA Per Capita Income) | | | | | -1.44*** | -0.46*** |
| | | | | | (0.23) | (0.12) |
| CBSA Fixed Effects | No | No | Yes | Yes | No | No |
| Chain Fixed Effects | No | Yes | No | Yes | No | Yes |
| Number of CBSAs | - | - | - | - | 691 | 691 |
| Observations | 9,019 | 9,019 | 8,849 | 8,849 | 9,019 | 9,019 |
| adj. $R^2$ | 0.18 | 0.72 | 0.51 | 0.79 | 0.14 | 0.72 |

Notes: *** p< 0.01, ** p<0.05, * p<0.10; ; standard errors, clustered by store in columns 1 through 4 and by CBSA in columns 5 and 6, are in parentheses. The table reports the results of fixed-effect regressions. In the Panel A, the dependent variable is the difference between the share of the high-income Nielsen HMS panelist expenditures represented in the set of products sold by a store in 2012 and the share of low-income panelist expenditures represented in that same product set. In Panel B, the dependent variable is the difference between the average price level faced by consumers in the top income decile and the average price level faced by households in the bottom income decile against local per capita income. The price level in each store for a given income decile is calculated as the weighted average ratio between the price a product is sold for in a store relative to the price that product is sold at in the national sample where weights are defined as the value of the purchases of that product made by households in the respective income decile in the Nielsen household-level panel. In each column, this dependent variable is regressed against the log per capita income of the neighborhood (in columns 1 through 4) or CBSA (in columns 5 and 6) where the store is located, as well as chain fixed effects in columns 2, 4, and 6. The number of observations decreases when introducing CBSA fixed effects because not all stores are located in CBSAs.

# 4 Model

This section introduces the demand system I use to study why high-income households purchase different products to low-income households and at different prices. This framework also forms the basis of the price indexes that summarize how high- and low-income households value the prices and products available to them in different markets.

## 4.1 Notation

Figure 4 shows how consumers choose to allocate expenditures. At the upper-most level, a consumer $i$ spends $W$ on a set of grocery products, denoted $\mathbf{G}$, and $Z$ on a set of other goods, denoted **Z**, subject to the budget constraint $W + Z \leq Y_i$. I do not explicitly model this upper-level expenditure allocation decision, but it is crucial in one respect: preferences over grocery products are non-homothetic because they depend on aggregate non-grocery expenditures.[26] This is generically the case if optimal non-grocery expenditures are normal.[27]

Figure 4: Consumer Choices

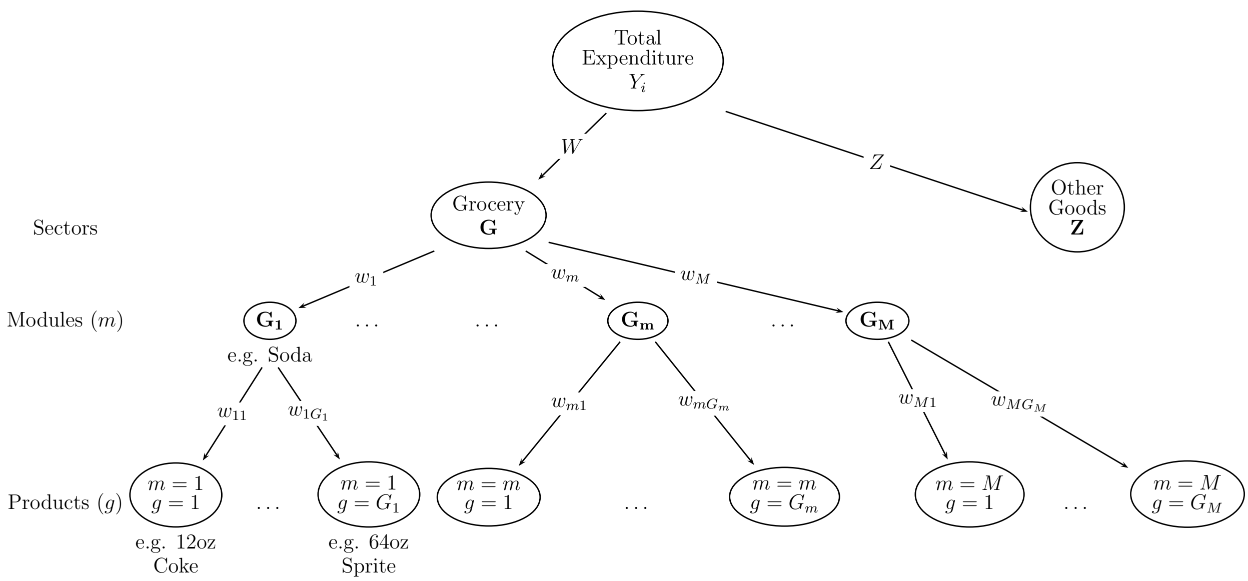


This paper focuses on the choices that consumers make within the grocery sector; that is, how consumers allocate their grocery expenditure $W$ between product modules, $\mathbf{M} = \{1, ..., M\}$, and their module expenditure $w_m$ between the varieties of grocery products in module $m$, $\mathbf{G_m} = \{1, ..., G_m\}$, for each module $m$. A consumer chooses to spend some $w_{mg}$ on

[26] Formally, preferences cannot depend on expenditures, so $Z$ is rather an aggregate of non-grocery consumption.

[27] In Appendix C.1, I solve for an implicit restriction on utility and prices under which the optimal non-grocery expenditure, $Z_i^*$, will be increasing in income. I cannot show that this restriction holds generally, but am instead able to show that it holds in the data.

each product $g$ in module $m$, purchasing $q_{mg} = w_{mg}/p_{mg}$ units of the product at a unit price $p_{mg}$. I denote the set of observed grocery prices and purchase quantities for module $m$ as $\mathbb{P}_m = \{p_{mg}\}_{g\in\mathbf{G_m}}$ and $\mathbb{Q}_m = \{q_{mg}\}_{g\in\mathbf{G_m}}$, respectively. $\mathbb{P}$ and $\mathbb{Q}$ are the unions of these price and quantity sets over all modules. A consumer's across-module and within-module expenditure allocation decisions are linked by the fact that they cannot allocate more than their total module expenditure, $w_m$, between products $g \in \mathbf{G_m}$; that is, $\sum_{g\in\mathbf{G_m}} w_{mg} = w_m$.

## 4.2 Consumption Utility

I model consumer demand for the products in $\mathbf{G}$ using a combination of CES and log-logit preferences. A consumer $i$'s utility from grocery consumption, conditional on their non-grocery expenditure $Z$, is a CES aggregate over consumer-specific module-level utilities:

$$(1) \qquad U_{iG}(\mathbb{Q}, Z) = \left\{ \sum_{m\in M} u_{im}\left(\mathbb{Q}_m, Z\right)^{\frac{\sigma-1}{\sigma}} \right\}^{\frac{\sigma}{\sigma-1}}$$

where $\sigma > 1$ is the elasticity of substitution between modules.

Consumer $i$'s utility from consumption in module $m$, conditional on their non-grocery expenditure $Z$, is equal to the sum of their consumer-specific product-level utilities:

$$(2) \qquad u_{im}\left(\mathbb{Q}_m, Z\right) = \sum_{g\in\mathbf{G_m}} u_{img}(\mathbb{Q}_m, Z)$$

where consumer $i$'s utility from consuming $q_{mg}$ of product $g$ in module $m$, conditional on their non-grocery expenditure $Z$, is defined as:

$$(3) \qquad u_{img}(Z) = q_{mg} \exp(\gamma_m(Z)\beta_{mg} + \mu_m(Z)\varepsilon_{img})$$

where $\beta_{mg}$ is the quality of product $g$ in module $m$ and $\varepsilon_{img}$ is the idiosyncratic utility of consumer $i$ from product $g$ in module $m$ drawn from a type I extreme value distribution. $\gamma_m(Z)$ and $\mu_m(Z) > 0$ are weights that govern the extent to which consumers with non-grocery expenditure $Z$ care about product quality and their idiosyncratic utility draws.[28]

### 4.2.1 Functional Forms

Before proceeding, it is worth making three observations about the general functional forms assumed above. First, the CES utility function governing the cross-module substitution patterns

[28]The log-logit utility function defined in equations (2) and (3) is a generalization of a utility function used by Auer (2010) to theoretically derive the effects of consumer heterogeneity on trade patterns and the welfare gains from trade.

implies that consumers will optimally consume a positive amount in each module. In the data for 2012, the typical household buys products in around one third of sample modules. This purchase behavior could reflect that households are, on average, consuming small quantities of products in some modules and, therefore, purchase the product so infrequently that we do not observe a purchase over the time period that they are in the sample.[29]

Second, the assumption that module utility is additive in product utilities that themselves are proportional to random draws from a continuous (type I extreme value) distribution implies that households allocate all of their module expenditure to a single product (the product that maximizes their marginal utility from expenditure, $\exp(\gamma_m(Z)\beta_{mg} + \mu_m(Z)\varepsilon_{img})/p_{mg})$. This matches the discrete-continuous behavior observed in the data: conditional on purchasing any products in a module in a month, households typically only purchase on product.[30]

Finally, the log-logit function governing preferences within modules yields the same Marshallian demand function for a set of consumers as the nested-CES utility function for a representative consumer with non-grocery expenditure $Z$ and an elasticity of substitution between products equal to one plus the inverse of the idiosyncratic utility draw weight, i.e., $\sigma_m(Z) = 1 + 1/\mu(Z)$. This link provides a natural analytic approximation for the relative utility that consumers with the discrete-continuous preferences described above face across markets offering different choice sets. The log-logit functional form also implies that, conditional on non-grocery expenditure, preferences are weakly-separable between modules. I exploit these features in the empirical strategy presented in Section 5.1 below.

#### 4.2.2 Non-Homotheticities

Consumers get utility from consuming quantity $q_{mg}$ of a product $g$, scaled up by the exponents of the quality of the product, $\beta_{mg}$, and their idiosyncratic utility draw for the product, $\varepsilon_{img}$. Preferences will be non-homothetic when at least one of the weights on these scalars, $\gamma_m(Z)$ or $\mu_m(Z)$, varies with non-grocery expenditure and, as discussed above, this expenditure varies with income. In order to interpret how these weights vary with income empirically, I make further functional form assumptions.

[29]In this scenario, households make purchases in all modules in expectation. The moments used to estimate the model parameters are based on individual household product selections within modules, conditional on their making a purchase in a given module, and expected store sales, i.e., the purchases of many households that shop in a store. The fact that some households do not purchase products in certain modules during a given period will be reflected in the fact that these modules have low within-store sales shares, and explained by the fact that the products in these modules are, on average, either more expensive or lower quality, relative to products in other modules. Models that reflect these more realistic cross-module consumption patterns, either by accounting for dynamic purchase behavior (see, *e.g.*, Hendel (1999); Dube (2004)) or explicitly modeling consumer's discrete-continuous preferences over modules (see, *e.g.*, Song and Chintagunta (2007); Pinjari and Bhat (2010)), would be difficult to estimate given the dimensions of the problem that this paper addresses.

[30]That is, the median household-month-module level observation with strictly positive expenditure has just one product purchased.

I interpret $\gamma_m(Z)$ to be the valuation for product quality, $\beta_{mg}$, for product $g$ in module $m$ shared by consumers with non-grocery expenditure $Z$. I assume that $\gamma_m(Z)$ is log-linear in $Z$ with a module specific slope, $\gamma_m$, such that:

$$\gamma_m(Z) = 1 + \gamma_m \ln(Z) \tag{4}$$

A consumer's valuation for product quality in module $m$ is increasing in $Z$ when $\gamma_m > 0$.

I employ a revealed preference approach to estimate the product quality $\beta_{mg}$ parameters as the average willingness to pay for product $g$ in module $m$ across all consumers. The idea here is that product $g$ in module $m$ is estimated as having high quality, $\beta_{mg}$, relative to that of another product $\tilde{g}$ in the same module $m$, $\beta_{m\tilde{g}}$, when a set of consumers facing the same price for both products spends a higher share of their expenditure on product $g$ than on product $\tilde{g}$. All consumers agree on this distribution of product qualities but, for $\gamma_m > 0$, consumers who spend more on non-grocery items place a greater weight on product quality, relative to quantity, in selecting which product to purchase in a module. Since $Z$ is normal, a positive $\gamma_m$ implies that high-income consumers spend a disproportionate amount of their module expenditures on higher quality products, relative to low-income consumers.

This form of non-homotheticity is common in the international trade literature where, for example, Fajgelbaum et al. (2011) show the theoretical implications of non-homothetic demand with a model that allows for complementarities between product quality and expenditure on a non-differentiated outside good. These complementarities imply that the elasticity of demand for quality is increasing with income, as in Hallak (2006) and Feenstra and Romalis (2012), who calculate cross-country price indexes similar to those estimated below.

The within-module utility function defined in equations (2) and (3) is also non-homothetic through the weight, $\mu_m(Z)$, on the idiosyncratic utility, $\varepsilon_{img}$. These idiosyncratic utility weights govern the dis-utility from consuming products that are horizontally differentiated from the consumer's ideal type of product, or the extent to which consumers find the available products substitutable with their ideal. I assume that the inverse of the idiosyncratic utility draw weight for module $m$ is log linear in non-grocery expenditures:

$$\frac{1}{\mu_m(Z)} = \sigma_m(Z) - 1 \equiv \alpha_m^0 + \alpha_m^1 \ln(Z) \tag{5}$$

where recall that $\sigma_m(Z)$ reflects the elasticity of substitution between products in module $m$ for a representative consumer with non-grocery expenditure $Z$. For $\alpha_m^1 < 0$ , $\sigma_m(Z)$ decreases with $Z$ such that consumers with high non-grocery expenditures find the available products less substitutable with each other and their ideal product and will, therefore, have a higher willingness to pay for the product closest to their ideal than consumers with low non-grocery expenditures. That is, for $Z$ normal, $\alpha_m^1 < 0$ implies that consumers' elasticity of substitution between products within a module and their tendency to switch between products in response to relative price changes is decreasing in consumer income.[31]

This form of non-homothetic price sensitivity is also similar to those used in recent international trade models. Hummels and Lugovskyy (2009), for example, develop a Lancaster ideal variety utility function where the dis-utility from distance between a product and a consumer's ideal type is an increasing function of their consumption quantity $q^\gamma$ for $\gamma \in [0, 1]$. This weight implies an income-specific price elasticity in a similar manner to the idiosyncratic utility weights, $\mu_m(Z)$, above.[32]

## 4.3 Individual Utility Maximization Problem

The grocery utility function defined in equations (1)-(3) is specific to the individual through a consumer's income, their non-grocery expenditure, and their idiosyncratic utility draws. I assume that consumers draw an idiosyncratic utility $\varepsilon_{img}$ for each product $g \in \mathbf{G}$ prior to making their purchase decision. Consumers then solve for their optimal grocery consumption bundle for a given non-grocery expenditure level $Z$ by maximizing grocery utility subject to budget and non-negativity constraints:

$$\sum_{m \in \mathbf{M}} \sum_{g \in \mathbf{G_m}} p_{mg} q_{mg} \leq Y_i - Z \quad \text{and} \quad q_{mg} \geq 0 \; \forall mg \in \mathbf{G} \tag{6}$$

The solution to this problem is a vector of optimal product selections (one for each module), $\mathbf{g_i^*}(Z) = (g_{i1}^*(Z), ..., g_{iM}^*(Z))$, and module-level expenditures, $\mathbf{w_i^*}(Z) = (w_{i1}^*(Z), ..., w_{iM}^*(Z))$. These are derived in Appendix C.2 to be:

$$g_{im}^*(Z) = \arg \max_{g \in \mathbf{G_m}} \left( \gamma_m(Z) \beta_{mg} + \mu_m(Z) \varepsilon_{img} \right) / p_{mg} \tag{7}$$

[31]The role that the $\alpha$ parameters that govern the substitution elasticities play in determining the elasticity of demand with respect to price is clearly demonstrated in the expenditure share equations presented in Section 5.3 below.

[32]Macro-economists have found alternative models to be empirically relevant for explaining differences in the prices paid by high- and low-income households. These models appear to be less relevant in the Nielsen data, so it is unlikely that ignoring them biases the aggregate estimates found below. The cross-income differences in search costs and shopping behavior explored in Simonovska (2010) could, in theory, enable low-income households to mitigate the high prices in wealthy cities at a lower cost than high-income households. However, Figure 2 shows that the cross-income differences in prices paid for identical items purchased in different stores or at different sale/non-sale periods are relatively small compared to the unit expenditure differences attributable to the fact that high- and low-income consumers are buying entirely different products. I also find no evidence that high-income consumers purchase more varieties of bar-coded products than low-income consumers, as would be the case in a hierarchic demand model like that use to explain Indian household consumption in Li (2012) or the translated additive-log utility function used in Simonovska (2010).

and

$$w_{im}^{*}(Z) \quad = \quad (Y_i - Z)\frac{\left(\max\limits_{g\in\mathbf{G_m}} \left(\gamma_m(Z)\beta_{mg} + \mu_m(Z)\varepsilon_{img}\right)/p_{mg}\right)^{\sigma-1}}{P(\mathbb{P}, Z, \varepsilon_{\mathbf{i}})^{1-\sigma}} \tag{8}$$

where $P(\mathbb{P}, Z, \varepsilon_i)$ is a CES price index over the grocery products that a consumer $i$ optimally consumes in each module:

$$P(\mathbb{P}, Z, \varepsilon_{\mathbf{i}}) = \left[\sum_{m\in\mathbf{M}} \left(\max_{g\in\mathbf{G_m}} \left(\gamma_m(Z)\beta_{mg} + \mu_m(Z)\varepsilon_{img}\right)/p_{mg}\right)^{\sigma-1}\right]^{\frac{1}{1-\sigma}} \tag{9}$$

Plugging these optimal product choices and module expenditures into the direct utility function defined in equations (1)-(3), I obtain the indirect utility of consumer $i$ from grocery consumption in a market offering prices and products summarized in the vector $\mathbb{P}$:

$$V(\mathbb{P}, Y_i, Z, \varepsilon_{\mathbf{i}}) = \frac{(Y_i - Z)}{P(\mathbb{P}, Z, \varepsilon_{\mathbf{i}})} \tag{10}$$

# 5 Empirical Strategy

A key goal of this paper is to characterize how consumers at different income levels value the different products and prices available to them across different markets in the U.S.. In this section, I derive the income- and city-specific price indexes I use to measure this variation. These indexes require two key components: vectors of the prices that provide comparable representations of the prices and product variety available in different U.S. cities, and estimates for model parameters that govern consumer's perceptions of these price vectors. In the remainder of the section, I describe how I use the Nielsen data to obtain each of these components.

## 5.1 Measuring Relative Utility Across Markets

Section 4.3 above solved for the indirect utility of a consumer from grocery consumption in a generic market offering a vector of prices $\mathbb{P}$. This paper seeks to compare the utility consumers get from the prices and products available to them in different markets, so I now introduce a market subscript to equation (10), writing the indirect utility of a consumer $i$ in market $t$ as

$$V(\mathbb{P}_t, Y_i, Z_{it}, \varepsilon_{\mathbf{i}}) = \frac{(Y_i - Z_{it})}{P(\mathbb{P}_t, Z_{it}, \varepsilon_{\mathbf{i}})} \tag{11}$$

where the set of prices and products available to household $i$, $\mathbb{P}_t = \{p_{mgt}\}_{g\in\mathbf{G}_t}$, and their optimal non-grocery expenditures, $Z_{it}$, are both allowed to vary across markets.

This indirect utility function is consumer-specific in three ways: it depends on a consumer's income, $Y_i$, on their optimal non-grocery expenditures, $Z_{it}$, and on their idiosyncratic utility draws, $\varepsilon_{\mathbf{i}}$. To study the systematic variation in utility across consumers earning different incomes, I abstract from any variation in non-grocery expenditures $Z_{it}$ and/or idiosyncratic utility draws $\varepsilon_{\mathbf{i}}$ that is uncorrelated with income. The idiosyncratic utility $\varepsilon_{\mathbf{i}}$ draws are, by definition, uncorrelated with consumer income $Y_i$. The most direct way to abstract from this random variation would be to take the expectation of the indirect utility defined in equation (11) over the idiosyncratic draws. Unfortunately, there is no analytic solution to this problem, and numerical solutions are computationally intensive. Instead, I approximate the relative utility of households at a given income level across different markets with the relative utility of an income-specific representative consumer at the same income across the same markets. The representative consumer's utility from consuming a grocery bundle $\mathbb{Q}$ is a nested-CES function conditional on their non-grocery expenditure $Z$ defined as:

$$(12) \qquad U_G^{CES}(\mathbb{Q}, Z) = \left\{ \sum_{m\in M} \left[ \sum_{g\in\mathbf{G_m}} [q_{mg} \exp(\beta_{mg}\gamma_m(Z))]^{\frac{\sigma_m(Z)-1}{\sigma_m(Z)}} \right]^{\left(\frac{\sigma_m(Z)}{\sigma_m(Z)-1}\right)\left(\frac{\sigma-1}{\sigma}\right)} \right\}^{\frac{\sigma}{\sigma-1}},$$

where $q_{mg}$, $\beta_{mg}$, $\gamma_m(Z)$, $\sigma_m(Z)$, and $\sigma$ take the same definitions as in the log-logit utility function presented in Section 4 above.[33] The indirect utility of this representative consumer from income $Y_i$ and prices and products $\mathbb{P}_t$, $V^{CES}(\mathbb{P}_t, Y_i)$, takes a similar form to the indirect utility of the idiosyncratic consumer provided in equation (11) above. It can also be expressed as the ratio of the consumer's grocery expenditure to a price index that summarizes the consumer's marginal utility from expenditure given the prices and products available in the market:

$$(13) \qquad V^{CES}(\mathbb{P}_t, Y_i, Z_{it}) = \frac{(Y_i - Z_{it})}{P^{CES}(\mathbb{P}_t, Z_{it})},$$

where

$$P^{CES}(\mathbb{P}_t, Z_{it}) = \left[ \sum_{m\in\mathbf{M}} \left( \left[ \sum_{g\in\mathbf{G_{mt}}} \left( \frac{p_{mgt}}{\exp(\beta_{mg}\gamma_m(Z_{it}))} \right)^{(1-\sigma_m(Z_{it}))} \right]^{\frac{1-\sigma}{1-\sigma_m(Z_{it})}} \right) \right]^{\frac{1}{1-\sigma}}$$

for $p_{mgt}$ equal to the unit price at which product $g$ in module $m$ is sold in market $t$.

To summarize this indirect utility function across households so that it varies with $i$ only

[33] In Appendix C.3, I show that this income-specific, nested, asymmetric CES utility function yields identical within-grocery budget shares as the CES-nested log-logit utility function that I estimate.

through income, $Y_i$, I approximate household non-grocery expenditures by assuming that non-grocery expenditures, $Z_{it}$, vary only with household income, $Y_i$, such that $Z_{it} = Z(Y_i)$.[34] Under this assumption, we can express the consumer's indirect utility as a function of market prices, $\mathbb{P}_t$, and consumer income, $Y_i$ alone:

$$V^{CES}(\mathbb{P}_t, Y_i) = \frac{(Y_i - Z(Y_i))}{P^{CES}(\mathbb{P}_t, Z(Y_i))}, \tag{14}$$

where

(15)

$$P^{CES}(\mathbb{P}_t, Z(Y_i)) = \left[ \sum_{m \in \mathbf{M}} \left( \left[ \sum_{g \in \mathbf{G_{mt}}} \left( \frac{p_{mgt}}{\exp(\beta_{mg}\gamma_m(Z(Y_i)))} \right)^{(1-\sigma_m(Z(Y_i)))} \right]^{\frac{1-\sigma}{1-\sigma_m(Z(Y_i))}} \right) \right]^{\frac{1}{1-\sigma}}$$

In particular, a consumer's relative indirect utility across two markets $t$ and $t'$ is equal to the inverse of the relative price indexes they face across the same markets:

$$\frac{V(\mathbb{P}_t, Y_i)}{V(\mathbb{P}_{t'}, Y_i)} = \frac{P^{CES}(\mathbb{P}_{t'}, Z(Y_i))}{P^{CES}(\mathbb{P}_t, Z(Y_i))} \tag{16}$$

That is, the magnitude of the price index a consumer with income $Y_i$ faces in market $t$ relative to the price index they face in market $t'$ indicates how much lower (or higher) the consumer's grocery utility is in market $t$ relative to market $t'$. The remainder of this section outlines how I obtain the two key inputs for these price indexes: market-specific price vectors and demand parameters.

Note that this approach to measuring income-specific spatial price indexes is different from the approach that Broda and Romalis (2009) developed to calculate income-specific inflation with the same Nielsen household-level data. Broda and Romalis (2009), and subsequent papers by Argente and Lee (2017) and Jaravel (2018), use the Feenstra (1994) methodology to calculate price indexes that are exact to a nested-CES utility function similar to the one above, but with two key differences. The Broda and Romalis (2009) approach is more restrictive in that the authors do not allow the substitution elasticities, $\sigma_m$ in the framework above, to vary with income. It is, however, more flexible in implicitly allowing for households at different income levels to have entirely different revealed preferences ($\beta_{mg}$s) for products. In the model presented here, households agree on the qualities of products and only the willingness to pay for quality

[34]Theoretically, this assumption could be violated since consumers at each income level may optimally choose different aggregate expenditure allocations across cities to suit the different grocery and non-grocery prices they face in these locations. Empirically, however, I observe that the relationship between non-grocery expenditures and income is surprisingly consistent across cities. Appendix Figure A.6 demonstrates that households earning higher incomes spend a smaller share of their income on grocery products. Within income groups, however, the average grocery expenditure share does not vary much across cities and, in particular, it does not vary systematically with city income.

varies with household income. The additional structure imposed on the relationship between perceived quality and income in this paper, as well as in more recent work by Feenstra and Romalis (2014), provides a clearer economic interpretation for the cross-income differences in the relative costs measured here relative to those measured in Broda and Romalis (2009).[35]

## 5.2 Inferring Prices and Product Availability

The first input to the price index defined in equation (15) is a market-specific price vector, $\mathbb{P}_t$, representing the set of prices and products available to consumers in a market $t$. I calculate price indexes comparing grocery costs across two types of markets in 2012: CBSAs and stores. I proxy for the set of prices and products available to consumers in each CBSA in 2012 using the set of products and unit prices represented in the 2012 sales of a random sample of the RMS participating retailers located in that CBSA, as described in Section 2 above. I proxy for the prices and products available to consumers in individual grocery stores in 2012 using the set of products and unit prices observed in the sales of each establishment in 2012.[36]

## 5.3 Parameter Estimation

The second set of inputs into the price index defined in equation (15) are model parameters that characterize how consumers value the products and prices available to them in a market, and how this valuation varies with consumer income. I denote this set of parameters using a vector $\theta$ defined as

$$\theta = \{(\theta_1, \ldots, \theta_M), \sigma\}$$

where $\theta_m = \{\alpha_m^0, \alpha_m^1, \beta_{m1}, \ldots, \beta_{mG_m}, \gamma_m\}$.

I estimate these parameters in two stages. The first stage identifies the parameters that govern the relative shares households spend on different products within each module. The second identifies the parameters that govern the relative shares households spend on different modules. In the first estimation stage, I follow Berry et al. (2004) in a GMM procedure that fits two sets of predicted moments to their data analogs: (1) store-level product sales shares and (2) the covariance of the prices and estimated qualities of the products purchased by each household with household income. In the second estimation stage, I fit only store-level module sales shares. This section describes the moment conditions and the identification of the parameters for each estimation stage.

[35]The Feenstra and Romalis (2014) approach is similar to mine in that the authors estimate the parameters of the underlying utility function and use these estimates to adjust prices for product quality. While the resulting price indexes are not income-specific, they are based on a utility function that is non-homothetic in demand for quality in the same way as the utility function presented above.

[36]The cross-store comparisons are again between stores in one channel – grocery or “food” in the RMS data – dropping mass merchandise, drug, and convenience stores.

### 5.3.1 Within-Module Estimation Methodology

**Moments** Given the distributional assumption on $\varepsilon_{img}$, the conditional probability of purchasing product $g$ in module $m$ for a household with non-grocery expenditure $Z_i$ and facing a vector of prices $\mathbb{P}$ takes the familiar multinomial logit form:

$$P_{mg}(Z_i, \mathbb{P}, \theta_m) = \frac{\exp\left[\alpha_{im}\left(\gamma_{im}\beta_{mg} - \ln p_{mg}\right)\right]}{\sum\limits_{g' \in \mathbf{G_m}} \left(\exp\left[\alpha_{im}\left(\gamma_{im}\beta_{mg'} - \ln p_{mg'}\right)\right]\right)} \tag{17}$$

where $\alpha_{im} = (\alpha_m^0 + \alpha_m^1 \ln Z_i)$ and $\gamma_{im} = (1 + \gamma_m \ln Z_i)$.

The first set of moments fits predicted store-level product sales shares to the shares observed in the RMS data. To calculate the store-level sales shares, I adjust the standard purchase probability expressed in equation (17) to reflect time-varying store-specific pricing and promotion activity:

$$P_{mg}(Z_i, \mathbb{P}_{st}, \theta_m, \xi_{st}) = \frac{\exp\left[\alpha_{im}\left(\gamma_{im}\beta_{mgst} - \ln p_{mgst}\right)\right]}{\sum\limits_{g' \in \mathbf{G_{mst}}} \left(\exp\left[\alpha_{im}\left(\gamma_{im}\beta_{mg'st} - \ln p_{mg'st}\right)\right]\right)}$$

where $\beta_{mgst} = \beta_{mg} + \xi_{mgst}$ and $\xi_{mgst}$ is a transitory taste shock for product $g$ in store $s$ at time $t$, demeaned from the fixed product quality parameter, $\beta_{mg}$. The fixed product quality parameter refers to characteristics of the product that are common across stores and over time, such as physical characteristics of the product itself and national recognition of the product's brand. The transitory taste shock is associated with local brand tastes, non-price promotions, and relative prominence of the product in the store's layout. In this stage of estimation, the product quality and the transitory taste shock will be identified for all but one product in each module, so will be estimated relative to the taste shock for the outside product (the set of products with average positive sales shares below the 60th percentile for all products).

The predicted sales of product $g$ in module $m$ at store $s$ at time $t$ is obtained by aggregating individual choice probabilities over the units purchased by customers at each non-grocery expenditure level:

$$Q_{mgst}(\theta_m; \mathbb{P}_{st}) = \int \frac{\exp\left[\alpha_{im}\left(\gamma_{im}\beta_{mgst} - \ln p_{mgst}\right)\right]}{\sum\limits_{g' \in \mathbf{G_{mst}}} \exp\left[\alpha_{im}\left(\gamma_{im}\beta_{mg'st} - \ln p_{mg'st}\right)\right]} dF(Z_i|s,t) \tag{18}$$

where $F(Z_i|s,t)$ is the distribution of non-grocery expenditures over all customers $i$ of store $s$ at time $t$ weighted by the number of module-$m$ units each purchases.

The first set of within-module moment conditions is expressed over the transitory component of the unobserved product quality

$$\bar{g}^1(\theta_m) = \frac{1}{n_m} \sum_{mg,s,t} g^1_{mgst}(\theta_m) = \frac{1}{n_m} \sum_{mg,s,t} \tilde{\xi}_{mgst}(\mathbf{X}_m; \theta_{1m}) \tilde{\mathbf{W}}_{mgst}$$

where $n_m$ is the number of (store-product-time) observations and $\mathbf{W}_{mgst}$ is a vector of pre-determined variables including product fixed effects and instrumental variables outlined in the identification discussion below. The tilde denotes that a variable has been differenced from the respective value for the outside product in each module, $\bar{g}_m$, e.g., $\tilde{\xi}_{mgst}(\mathbf{X}_m; \theta_{1m}) = \xi_{mgst}(\mathbf{X}_m; \theta_{1m}) - \xi_{m\bar{g}_m st}(\mathbf{X}_m; \theta_{1m})$.

The second and third set of within-module moment conditions respectively compare the covariance between the relative quality and unit value of the products purchased by households and their non-grocery expenditure to that predicted by the model. Following Berry et al. (2004), I fit the model's predictions for the uncentered covariance of quality and price with household non-grocery expenditure, i.e., $E(x_{mg}Z)$ for $x_{mgst} \in \left\{\tilde{\beta}_{mg}, \tilde{p}_{mgst}\right\}$, to that observed in the HMS data.

The quality-covariance moments are obtained from the difference between the average non-grocery expenditure of Nielsen panelists who purchase each product $g$ in store $s$ at time $t$ and the average non-grocery expenditure predicted by the model for households that purchase product $g$ in store $s$ at time $t$. If $y = mg$ denotes that a household purchases a unit of product $g$ in module $m$, $i_{mg}$ denote one of the $N_{mg}$ units purchased by sample households, and $N_m = \sum_{g \in \mathbf{G}_m} N_{mg}$, the quality-co variance moments are:

$$\bar{g}^2(\theta_m) \approx \frac{1}{N_m} \sum_{mg} N_{mg} \beta_{mg} \left\{ \frac{1}{N_{mg}} \sum_{i_{mg}=1}^{n_{mg}} Z_{i_{mg}} - E[Z|y = mg, \theta_m] \right\}$$

I calculate $E[Z|y = mg, \theta]$ by first transforming it into an expression that depends on the model's predicted choice probabilities for each unit purchased:

$$E[Z|y = mg, \theta_m] = \frac{\int \int ZP(y = mg|Z, \theta_m, y = mst) F(Z|m, s, t) G(s, t|y = m)}{\int Pr(y = mg, |\theta_m, y = m) G(s, t|y = m)}$$

where $F(Z|m, s, t)$ is now the distribution of non-grocery expenditures of the households observed to be purchasing units of module-$m$ products in store $s$ at time $t$, weighted by units purchased, and $G(s, t|y = m)$ is the distribution of these purchases across stores and time periods. In practice, I calculate

$$E[Z|y = mg, \theta_m] \quad = \quad \frac{\frac{1}{N_m} \sum_i Z_i P_{mg}(Z_i, \mathbb{P}_{st}, \theta_m, \xi_{st})}{\frac{1}{N_m} \sum_i P_{mg}(Z_i, \mathbb{P}_{st}, \theta_m, \xi_{st})}$$

where $N_m = \sum_{mg} N_{mg}$ is the total number of units sold and $i$ indexes each unit purchased by a household $i$ with non-grocery expenditure $Z_i$. This assumes that households receive an independent taste shock for each unit they purchase. $P_{mg}(Z_i, \mathbb{P}_{st}, \theta_m, \xi_{st})$ is defined above in equation (17).

The price-covariance moments compare the covariance between the relative unit price paid by households for their selection and their non-grocery expenditure to that predicted by the model:

$$\bar{g}^3(\theta_m) \approx \frac{1}{N_m}\sum_i \left(Z_i - \bar{Z}\right)\sum_{s,t}\left(\tilde{p}_{imst} - E[\tilde{p}_{imst}|\theta_m] - \frac{1}{N_m}\sum_{i,s,t}\left(\tilde{p}_{imst} - E[\tilde{p}_{imst}|\theta_m]\right)\right)$$

where $\bar{Z} = \frac{1}{N_m}\sum_i \bar{Z}_i$ is the unit-weighted mean non-grocery expenditure of sample households. The relative unit price paid by a household $i$ in module $m$ in store $s$ at time $t$ is defined as the difference between the unit price charge by the store for product household $i$ selected from the weighted average unit price charged by the store for products in that module: $\tilde{p}_{imst} = (p_{imgst} - \bar{p}_{mst})$, where $\bar{p}_{mst} = \sum_{g\in\mathbf{G}_{mst}} w_{mgst}p_{mgst}$ and $w_{mgst} = s_{mg}/\sum_{g\in\mathbf{G}_{mst}} s_{mg}$ is the product sales weight taken from the store-level data. I calculate the predicted relative unit price paid by household $i$ in module $m$ in store $s$ and time $t$, as

$$E[\tilde{p}_{imst}|\theta_m] = \sum_{g\in\mathbf{G}_{mst}} \tilde{p}_{mgst}P_{mg}(Z_i, \mathbb{P}_{st}, \theta_m, \xi_{st})$$

**Estimation Procedure** The three moment conditions defined above identify all of the module-specific parameters, $\theta_m$, except for the quality parameter $\beta_{m\bar{g}_m}$ of the outside product $\bar{g}_m$ in each module. I denote this set of parameters by $\theta_1 = \{\theta_{1m}\}_{m\in\mathbf{M}}$ where

$$\theta_{1m} = \left\{\alpha_m^0, \alpha_m^1, \gamma_m, \left\{\tilde{\beta}_{mg}\right\}_{g\in\mathbf{G}_m}\right\}$$

for each module $m \in \mathbf{M}$ and tildes continue to denote that a variable has been differenced from the respective value for the outside product in each module, $\bar{g}_m$ (e.g., $\tilde{\beta}_{mg} = \tilde{\beta}_{mg} - \tilde{\beta}_{m\bar{g}_m}$).

The $\theta_1$ parameters are estimated in separate non-linear GMM procedures that minimize a quadratic function over the moment conditions $\{\bar{g}^1(\theta_m), \bar{g}^2(\theta_m), \bar{g}^3(\theta_m)\}$ for each module $m$. I use the nested fixed-point algorithm proposed by Berry et al. (1995) to obtain the relative product quality parameters, $\left\{\tilde{\beta}_{mg}\right\}_{g\in\mathbf{G}_m}$, as a function of the three non-linear parameters for each module, $\theta_{1m}^{NL} = \{\alpha_m^0, \alpha_m^1, \gamma_m\}$. Given a guess of $\theta_{1m}^{NL}$, I first invert the share equation for the relative product quality shocks, $\tilde{\beta}_{mgst}(\theta_{1m}^{NL}) = \beta_{mgst}(\theta_{1m}^{NL}) - \beta_{m\bar{g}_m st}(\theta_{1m}^{NL})$, that solve a system of non-linear equations equating predicted and observed demand at each store in each time period. I project $\tilde{\beta}_{mgst}(\theta_{1m}^{NL})$ on product dummies to obtain estimates for relative product quality $\tilde{\beta}_{mg}(\theta_{1m}^{NL})$. The residuals provide estimates for the transitory shocks, $\tilde{\xi}_{mgst}(\theta_{1m}^{NL}) = \tilde{\beta}_{mgst}(\theta_{1m}^{NL}) - \tilde{\beta}_{mg}(\theta_{1m}^{NL})$. Both of these terms are used to calculate the moment conditions $\{\bar{g}^1(\theta_m), \bar{g}^2(\theta_m), \bar{g}^3(\theta_m)\}$ and, in turn, the objective function that I minimize over the remaining parameters, $\theta_{1m}^{NL} = \{\alpha_m^0, \alpha_m^1, \gamma_m\}$. Details on this full procedure can be found in Appendix D.2.1.

I proxy non-grocery expenditure, $Z$, with household income, $Y$.[37] I assume that the distribution of units purchased in each module across income levels $F(Y|m,s,t)$ is constant over time within a store $s$. I estimate the store-level income distribution $F(Y|s)$ as an average of the income distributions of tracts whose centroids fall within 30km of the store zip centroid (estimated to be the modal residential zip code of the store's customers observed in the Nielsen HMS data). [38] In the baseline estimation, I assume a degenerate distribution for the clientele of each store.

**Identification** The store-level moments serve primarily to identify the mean price elasticity, $\alpha_m^0$ and product quality, $\beta_{mg}$, parameters. Conditional on product quality, the base price sensitivity $\alpha_m^0$ parameter is identified by the extent to which relative within-store sales shares co-vary with the components of relative price variation captured by the price instruments, described in more detail below. Relative product quality, $\tilde{\beta}_{mg} = \beta_{mg} - \beta_{m\bar{g}_m}$, is identified by variation in the average within-store sales shares of each product $g$, relative to the sales share of the outside product $\bar{g}_m$, conditional on price. The idea here is that, if products with two different products sell at the same price, but product A has a higher average relative market share across all store-months than product B, then product A will be assigned a higher quality parameter relative to the base good for that module.[39]

The household moments serve to identify the non-homothetic parameters, $\alpha_m^1$ and $\gamma_m$. The $\alpha_m^1$ parameter that governs how the price sensitivity varies with income is identified primarily by the covariance between the prices of products purchased by households and their income. Like $\alpha_m^1$, the quality-income gradient $\gamma_m$ parameter that governs how demand for quality varies with income are primarily identified by the covariance between the estimated quality of products

[37] In a slight abuse of notation, I will denote the coefficients on log income using the same notation used to denote the coefficients on log non-grocery expenditure in defining the moment conditions above. These new coefficients are in fact approximations of the original coefficient multiplied by the elasticity of non-grocery expenditure with respect to household income.

[38] See the Appendix A.2 for more details on this procedure.

[39] Variation in the quality of the outside product across store-months may bias the relative quality estimates that, in practice, are calculated as the mean of store-month-specific quality shocks that rationalize the relative sales shares on that product relative to the outside product given the non-linear parameter estimates, across the store-months in which the product is sold; i.e., $\hat{\tilde{\beta}}_{mg} = \frac{1}{N_g}\sum_{st} \tilde{\beta}_{mgst}(\hat{\theta}_{1m}^{NL})$ where $\tilde{\beta}_{mgst}(\hat{\theta}_{1m}^{NL}) = \beta_{mgst}(\hat{\theta}_{1m}^{NL}) - \beta_{m\bar{g}_m st}(\hat{\theta}_{1m}^{NL})$. I discuss these errors in more detail in Section 6.4.2, where I find them to be small in magnitude and not correlated with the spending patterns of high- or low-income households in such a way that would yield biases in other parameter estimates.

purchased by households and their income.

**Price Instruments** The store-level moments are based on the assumption that $\mathbb{E}[\tilde{\xi}_{mg}(\theta_{1m}^{NL})\tilde{\mathbf{W}}_{mg}^{1}] = 0$ for a set of instruments $\mathbf{W}^1$. These instruments include a set of brand dummies, price instruments, and interaction terms between these sets of variables and a set moments of the income distribution in the vicinity of the store.[40] As noted above, the tildes in the moment equation indicate that these errors and instruments are differenced from the outside product within each market to control, among other things, for market-level variation in the quality of the outside product. The set of brand dummies includes one dummy for each brand except this base product $\bar{g}_m$. To reduce the dimension of the estimation data, I conduct principal components analysis on this final set of instruments and use components that together explain over 99 percent of the variation of the data.[41]

I do not use prices as instruments because they might be correlated with the transient product-market-specific taste shocks, $\xi_{mg}(\theta_{1m}^{NL})$. I instrument for the price charged by a store for a given product with the contemporaneous price charged for the same product by stores that belong to the same retail chain but are located in different Demographic Market Areas (geographic market areas defined by Nielsen, which are roughly akin to MSAs). This "same chain-other city" instrument, also employed in DellaVigna and Gentzkow (2019), relies on similar relevance and exogeneity arguments as in Hausman et al. (1994) and Nevo (2001).

For relevance, I rely on cross-product inter-temporal and across-chain variation in the prices charged by chains, driven by the timing of chain-level sales or changes in wholesale pricing arrangements. Recall that the data is differenced from the outside product within market and and implicitly from the product mean, by the inclusion of the product fixed effects. Even after controlling for market and product fixed effects, there is sufficient variation in the instrument to provide a strong first stage, with F-statistics above 100 in all modules, and above 500 in 99% of modules.[42]

For exogeneity, cross-product variation in retail chain-level pricing cannot be correlated with changes in relative product tastes in a market. Such a correlation could arise, for example, if prices adjust in response to changes in the tastes of a retail chain's national customer base. A chain might, for example, lower the frequency of promotional sales for a product or re-negotiate a wholesale price agreement in response to declining national demand for that product. Though

[40]Specifically, the average, the average squared, and the 5th, 10th, 25th, 50th, 75th, 90th, and 95th percentiles of the income distribution.

[41]The principal components IV reduces the scale of the optimization problem with minimal sacrifice to identifying variation, noting that linear combinations of valid instruments remain valid instruments – c.f. Bai and Ng (2010). The exact number of principal components used based on Winkelried and Smith (2011)'s retention rule with $\delta = -2$. This retains instruments explaining over 99 percent of the variation in the instruments while reducing the number of instruments by 90-95 percent.

[42]See Appendix Figures A.7 and A.8.

I am unable to test this exclusion restriction directly, I can – for a subset of my data – construct an instrument that is plausibly uncorrelated with national demand shocks by residualizing my baseline "same chain-other city" instrument from the average contemporaneous price charged for the same product by other retail chains in different DMAs. I use this alternate instrument to test the validity of my base instrument in the sub-sample of products over which the residualized instrument is non-missing – i.e., products sold in multiple chains in multiple DMAs.

First, I run a GMM distance test comparing the J-statistics from the model estimated using both "same chain-other city" instruments to the J-statistics from the model estimated using only the residualized version. In most modules, I fail to reject the null that the base instrument is exogenous.

Then, I show that the price elasticity estimates using the baseline and the residualized instruments are comparable. Both instruments similarly remove negative biases in the price coefficient relative to an "OLS" specification that uses the endogenous observed price as the instrument. The price coefficients estimated using the base instrument are slightly lower than those estimated using the residualized version, but the difference is small with respective medians of 1.34 and 1.69.[43] In Section 6.4.1 below, I show that the main index results are robust to increases in the mean price coefficient of this, or even larger, magnitudes.

### 5.3.2 Across-Module Estimation Procedure

**Moments** The remaining parameters include the cross-module substitution parameter, $\sigma$, and the quality of the base product in each module, $\beta_{m\bar{g}_m}$, for all modules $m \in \mathbf{M}$, except for the base module $\bar{m}$.[44] I denote this set of parameters by $\theta_2$:

$$\theta_2 = \left\{\sigma, \{\beta_{m\bar{g}_m}\}_{m\in\mathbf{M},m\neq\bar{m}}\right\}$$

To estimate these parameters, I use a single set of moments that fit the predicted store-level module sales shares observed in the Nielsen RMS data to those predicted by the model.

The expected log expenditure share in module $m$ relative to $\bar{m}$ for a group of households with the same non-grocery expenditure, $Z_i$, facing a common vector of grocery prices, $\mathbb{P}$, is derived in Appendix C.2. Adjusting this expression to reflect time-varying store-specific pricing and promotion activity yields:

$$\mathbb{E}_\varepsilon\left[\ln s_{imst} - \ln s_{i\bar{m}st}\right] = (\sigma - 1)\ln \tilde{V}_m(Z_i, \mathbb{P}_{mst}, \mathbb{P}_{\bar{m}st}) \tag{19}$$

where $\tilde{V}_{mst}(Z_i, \mathbb{P}_{mst}, \mathbb{P}_{\bar{m}st}) = V_{mst}(Z_i, \mathbb{P}_{mst})/V_{\bar{m}st}(Z_i, \mathbb{P}_{\bar{m}st})$. $V_{mst}(Z_i, \mathbb{P}_{mst})$ is a CES-style

[43] See Appendix Figure A.11 for the full distributions.

[44] I normalize the fixed quality of the base product in the base module (unpopped popcorn), $\beta_{\bar{m}\bar{g}_{\bar{m}}}$, to equal zero.

index over price-adjusted product qualities:

$$(20)\qquad V_m(Z_i, \mathbb{P}_{mst}) = \left[ \sum_{g \in \mathbf{G_m}} \left( \frac{\exp(\gamma_{im}\beta_{mgst})}{p_{mgst}} \right)^{-\alpha_{im}} \right]^{\frac{1}{-\alpha_{im}}}$$

Note that the inclusive value is a function of the parameters estimated in both the first and second stage, i.e., $\theta_1$ and $\theta_2$. To see this recall that $\alpha_{im} = (\alpha_m^0 + \alpha_m^1 \ln Z_i)$ and $\gamma_{im} = (1 + \gamma_m \ln Z_i)$ and each store-specific product quality shock, $\beta_{mgst}$, is the the sum of $(\beta_{mgst} - \beta_{m\bar{g}_m st})$, estimated in stage 1, and an unknown base product quality shock, $\beta_{m\bar{g}_m st}$. We can express the inclusive value function as the product of the base product quality parameter, $\beta_{m\bar{g}_m st}$, to be estimated in the second stage and an inclusive value function calculated using only elements of $\theta_{1m}$ estimated in the first stage:

$$V_m(Z_i, \mathbb{P}_{mst}) = \exp(\gamma_{im}\beta_{m\bar{g}_m st}) V_{1m}(Z_i, \mathbb{P}_{mst})$$

where

$$(21)\qquad V_{1m}(Z_i, \mathbb{P}_{mst}) = \left[ \sum_{g \in \mathbf{G_m}} \left( \frac{\exp(\gamma_{im}\tilde{\beta}_{mgst})}{p_{mgst}} \right)^{-\alpha_{im}} \right]^{\frac{1}{-\alpha_{im}}}$$

and $\tilde{\beta}_{mgst} = \beta_{mgst} - \beta_{m\bar{g}_m st}$. Under the normalization that $\beta_{\bar{m}\bar{g}_{\bar{m}} st} = 0$ for all $s, t$, and using the decomposition of the inclusive value function above, we can now rewrite equation (19) as:

$$(22)\qquad \mathbb{E}_\varepsilon \left[ \ln s_{imst} - \ln s_{i\bar{m}st} \right] = (\sigma - 1) \left( \gamma_{im}\beta_{m\bar{g}_m st} + \ln \tilde{V}_{1mst}(Z_i, \mathbb{P}_{mst}, \mathbb{P}_{\bar{m}st}) \right)$$

where $\ln \tilde{V}_{1mst}(Z_i, \mathbb{P}_{mst}, \mathbb{P}_{\bar{m}st}) = \ln V_{1m}(Z_i, \mathbb{P}_{mst}) - \ln V_{1\bar{m}}(Z_i, \mathbb{P}_{\bar{m}st})$.

The predicted log expenditure share of module $m$ relative to module $\bar{m}$ at store $s$ at time $t$ is obtained by aggregating $i$-specific expected relative shares over the units purchased by customers at each non-grocery expenditure level:

$$(23)\qquad \mathbb{E}_z \left[ \mathbb{E}_\varepsilon \left[ \ln s_{imst} - \ln s_{i\bar{m}st} \right] \right] = \beta_{m\bar{g}_m st} (\sigma - 1) \overline{\gamma}_{mst} + (\sigma - 1) \overline{\tilde{v}}_{mst}$$

where $\overline{\gamma}_{mst} = \int \gamma_{im} dF(Z|s,t)$ and $\overline{\tilde{v}}_{mst} = \int \ln \tilde{V}_{1mst}(Z_i, \mathbb{P}_{mst}, \mathbb{P}_{\bar{m}st}) dF(Z|s,t)$ can be calculated using price data and parameter estimates for $\theta_1$ obtained in stage 1 above.

The moment equation is then defined as:

$$\bar{h}(\theta_2) = \frac{1}{n} \sum_{m,s,t} h_{mst}(\theta_2) = \frac{1}{n} \sum_{m,s,t} u_{mst}(\mathbf{X}; \hat{\theta}_1, \theta_2) \mathbf{W}_{mst}$$

where $n$ is the number of (store-module-month) observations; $\mathbf{W}_{mst}$ includes the average store-level quality coefficient $\overline{\gamma}_{mst}$ interacted with module fixed effects and an instrument for the average relative inclusive value for the module, $\overline{\tilde{v}}_{mst}$, described below; and $u_{mst}$ denotes the difference between the observed log relative module shares between modules $m$ and $\bar{m}$ in store $s$ at time $t$ and their predicted values, i.e.,

$$u_{mst}(\mathbf{X};\hat{\theta}_1,\theta_2) = \ln\left(s_{mst}/s_{\bar{m}st}\right) - \beta_{m\bar{g}_m}\left(\sigma-1\right)\overline{\gamma}_{mst}(\hat{\theta}_1) - \left(\sigma-1\right)\overline{\tilde{v}}_{mst}(\hat{\theta}_1) \tag{24}$$

Identification of $\sigma$ and $\beta_{m\bar{g}_m}$ relies on the assumption that the errors in the model predicted shares ($u_{mst}$) are orthogonal from $\mathbf{W}_{mst}$. The $u_{mst}$ errors can be broken into two components, $u_{mst} = u^1_{mst} + u^2_{mst}$. The first, $u^1_{mst} = (\sigma-1)\left(\beta_{m\bar{g}_m}\left(\overline{\gamma}_{mst} - \overline{\gamma}_{mst}(\hat{\theta}_1)\right) + \overline{\tilde{v}}_{mst} - \overline{\tilde{v}}_{mst}(\hat{\theta}_1)\right)$, reflect errors in the first stage estimates, while the second, $u^2_{mst} = \xi_{m\bar{g}_m st}\left(\sigma-1\right)\overline{\gamma}_{mst}$ for $\xi_{m\bar{g}_m st} = \beta_{m\bar{g}_m st} - \beta_{m\bar{g}_m}$, reflect the transitory components of the product-store-time taste shocks that are not estimated directly. To deal with the endogeneity of prices with respect to these transitory taste shocks, I instrument for the average inclusive value, $\overline{v}_{mst}$, using a data analog calculated with the same contemporaneous chain-specific national cost shock instruments that are used in the module-level estimation in place of store-time-specific price data.

The $\sigma$ substitution elasticity parameter is identified by the extent to which relative module shares react to national chain-specific cost shocks for each module. Recall that the relative inclusive value, $\overline{\tilde{v}}_{mst}$, is scaled up or down by the quality of the base product, $\bar{g}_m$, in a module $m$ relative to the quality of the base product, $\bar{g}_{\bar{m}}$, in the base module $\bar{m}$, unpopped popcorn (a product type sold in most stores), which is normalized to equal zero. Any difference between the expenditure share of module $m$ relative to popcorn and what would be expected given the relative inclusive value of the two modules and the $\sigma$ estimate will identify the quality of the base product in the module, $\beta_{m\bar{g}_m}$, scaled by the market average taste for quality, $\overline{\gamma}_{mst}$. Together with the relative product quality estimates from the first stage of estimation, $\beta_{mg} - \beta_{m\bar{g}_m}$, the base product quality estimates define the quality of each product in the dataset relative to the quality of the base product in the base module.

This estimation procedure yields consistent estimates for $\theta_2$, but the variance-covariance matrix of these parameters will be biased due to the presence of the first-step estimates for $\theta_1$ in the $u^1$ component of the error. I adjust this variance-covariance matrix to account for the errors from the first stage of the estimation following the GMM analog of the Murphy and Topel (1985) procedure outlined in Newey and McFadden (1994).[45] The adjusted variance-covariance matrix yields consistent standard errors for the $\theta_2$ estimates.

[45] Appendix D.2 details how these adjustments are calculated.

# 6 Results

## 6.1 Parameter Estimates

I estimated the model under four sets of parameter restrictions. These restrictions allow preferences to vary with income through the demand elasticities with respect to both quality and price, through only one of these channels, or through neither of these channels, in which case the model is homothetic.

Table 3 summarizes the estimates for the module-level parameters in each of these four models over the 400-500 modules where the optimization procedure reached internal solutions.[46]

Columns [1] through [3] summarize the parameter estimates for the unrestricted version of the model. Column [1] reports the parameter that governs the substitution elasticity of a consumer with the mean log income level in the sample for each module, $\hat{\alpha}_m^0 = \hat{\sigma}_m - 1$. The median of this price elasticity is 1.5, with an inter-quartile range of 0.8 to 2.4, implying a median and inter-quartile range for the elasticity of substitution of 2.5 and 1.8 to 3.4, respectively. The magnitude and distribution of these estimates are similar across the four models and are well-identified in most modules, with over 90 percent significant at the 95 percent level across all four models. The own-price elasticities found here are in the range of those estimated in Nevo (2000), Dube (2004), and Faber and Fally (2017).[47]

Columns [3] and [8] of Table 3 summarize the distribution of the estimated values for $\gamma_m$. All four models assume that all consumers agree on the relative quality of products, as described by the distribution of the $\beta_{mg}$ parameters for products $g \in \mathbf{G_m}$ within a module $m$. For positive values of $\gamma_m$, however, the utility weight that consumers place on this component of utility, relative to their idiosyncratic utility draw for each product or the quantity consumed, is increasing in their non-grocery expenditure $Z$. This implies that consumers with higher non-grocery expenditures have a higher willingness to pay for quality. In estimation, these parameters are identified by the fact that higher income consumers spend a relatively greater share of module expenditure on products with relatively high $\beta_{mg}$ estimates, that is, the products for which all consumers have a higher willingness to pay. Figure 5 shows that products with higher $\beta_{mg}$ estimates have higher expenditures at all income levels, but more so for the rich. Accordingly,

[46]The parameters were bounded as follows: $\alpha_m^0 \in (0.05, 30)$, $\alpha_m^1 \in (-5, 5)$, and $\gamma_m \in (-5, 5)$. See Appendix (D.2.1) for more detail on the steps taken to identify interior estimates.

[47]Nevo (2000), Dube (2004), and Faber and Fally (2017) also use the Hausman-IV approach and estimate the own-price elasticity of demand for similar categories of products to those considered here to be between in the range of -0.5 to -4. These estimates are lower than the own-price elasticities implied by the elasticity of substitution estimates in Broda and Weinstein (2010) and Hottman et al. (2016), who use the Feenstra (1994) methodology to identify the elasticity of substitution between products. In Section 6.4.1, I demonstrate the robustness of the estimated price indexes to a set of demand parameters estimated in a two-step procedure, first using the Feenstra (1994) method to estimate the average price elasticities, $\alpha_m^0$ , which are then held fixed when estimating the remaining parameters ($\alpha_m^1$, $\gamma_m$, and $\beta_{mg}$) using the moments described above.

Table 3: Summary Statistics for Parameter Estimates

| Model: | Homothetic | NH in Quality | | NH in Price | | NH in Quality and Price | | |
|---|---|---|---|---|---|---|---|---|
| Restrictions: | $\alpha_m^1 = 0$ & $\gamma_m = 0$ | $\alpha_m^1 = 0$ | | $\gamma_m = 0$ | | None | | |
| | [1] | [2] | [3] | [4] | [5] | [6] | [7] | [8] |
| Parameter: | $\alpha_m^0$ | $\alpha_m^0$ | $\gamma_m$ | $\alpha_m^0$ | $\alpha_m^1$ | $\alpha_m^0$ | $\alpha_m^1$ | $\gamma_m$ |
| Count | 404 | 470 | 470 | 421 | 421 | 494 | 494 | 494 |
| with t>1.96 | 375 | 463 | 406 | 421 | 231 | 433 | 90 | 420 |
| with t<-1.96 | 29 | 7 | 64 | 0 | 190 | 61 | 404 | 74 |
| Mean | 2.60 | 2.20 | 1.42 | 1.94 | 0.23 | 1.73 | -0.60 | 1.95 |
| p25 | 0.83 | 0.89 | 0.38 | 0.72 | -0.28 | 0.97 | -1.24 | 1.35 |
| p50 | 1.52 | 1.71 | 1.16 | 1.36 | 0.15 | 1.58 | -0.79 | 1.99 |
| p75 | 2.39 | 2.41 | 2.52 | 2.06 | 0.80 | 2.42 | -0.18 | 2.69 |

Notes: These tables report the summary statistics for the main module-level parameter estimates governing the elasticity of substitution and non-homotheticities in demand. Attention is limited to modules for which the estimation procedure converged at interior estimates for all relevant parameters. The second and third rows of the table show the number of modules in which the estimated t-statistic for the parameter was above or below 1.96. The mean and percentile statistics in the subsequent rows are weighted by module sales in the Nielsen store-level data.

Columns [3] and [8] of Table 3 show that the willingness to pay for quality (governed by $\gamma_m$) increases with income in over three-quarters of the modules represented in the data.[48] The demand for quality is therefore increasing with income in most grocery sectors.

Columns [5] and [7] of Table 3 summarize the distribution of the estimated values for $\alpha_m^1$ in each module. Recall this parameter governs how the elasticity of substitution varies across consumers with different non-grocery expenditures. For $\alpha_m^1 < 0$, high-income consumers will find other products to be less substitutable with their ideal variety and, therefore, substitute less across products in response to relative price changes. Comparing columns [5] and [7] of Table 3, we see that the majority of the $\alpha_m^1$ estimates are instead positive unless you control for non-homotheticity in the demand for quality. Column [5] shows that the majority of the $\alpha_m^1$ estimates, and even the majority of those that are statistically significant, are instead positive when $\gamma_m$ is constrained to be zero.[49] Column [7], on the other hand, shows that, in over 75

[48] The full distributions of the $\gamma_m$ and $\alpha_m^1$ estimates are depicted in Figures A.12 and A.13 in Appendix E.2.

[49] These estimates may be biased upwards by a correlation between unobserved income-specific product tastes and prices. Consider the model where $\gamma_m$ is restricted to equal zero for a degenerate store income distribution: $\ln s_{mgst} - \ln s_{m\bar{g}_m st} = (\alpha_m^0 + \alpha_m^1 y_{st})[(\beta_{mg} - \beta_{m\bar{g}_m}) - (\ln p_{mgst} - \ln p_{m\bar{g}_m st})] + \nu_{mg\bar{g}_m st}$. If the true $\gamma_m$ is positive, the error terms here will include any income-specific product tastes, $\gamma_m (\beta_{mg} - \beta_{m\bar{g}_m})$. If the chains at which high-income consumers shop set prices in accordance with these tastes such that $Corr(\gamma_m (\beta_{mg} - \beta_{m\bar{g}_m}), \ln p_{mgst} - \ln p_{m\bar{g}_m st}) \neq 0$, then the assumption that $\mathbb{E}[\mathbf{W}\xi] = 0$ will be violated. The fact that the $\alpha_m^1$ estimates are lower, and generally negative, in the model that allows for non-homotheticity in the demand for quality and the price sensitivity supports this theory, since this model directly controls for, $\gamma_m y_{st}(\beta_{mg} - \beta_{m\bar{g}_m})$. I do not, therefore, take the positive $\alpha_m^1$ estimates in the model that does not control for correlations in income-product specific tastes as evidence that high-income consumers are more price sensitive than low-income consumers. Instead, the positive $\alpha_m^1$ estimates highlight the difficulty in identifying the non-

Figure 5: Product Quality ($\beta_{mg}$) Estimates and High-vs.-Low Income Household Expenditures

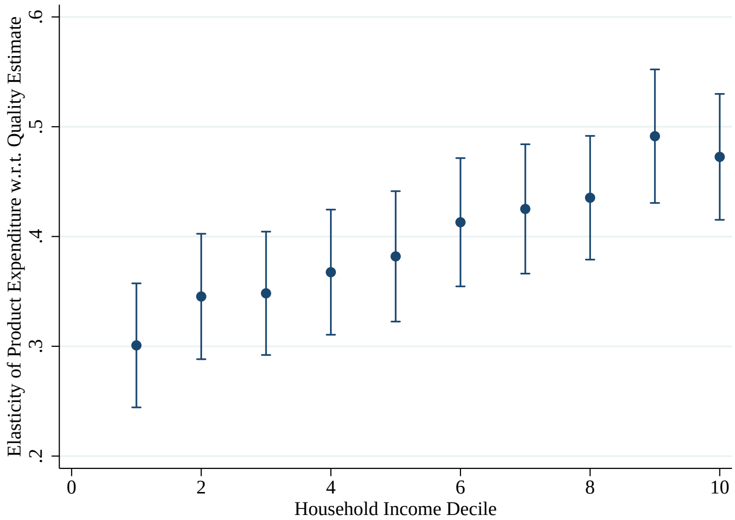


Note: Plots shows coefficient on log product-level expenditures by each income decile in the household-level (HMS) data regressed against the product quality ($\beta_{mg}$) estimates in the model that allows for non-homotheticity in quality but not price sensitivity (i.e., restricting $\alpha_m^1 = 0$ but allowing $\gamma_m \neq 0$). These regressions include product module fixed effects and observations are weighted by aggregate module sales. Attention is limited to estimates in the 470 modules where the estimation procedure converged at interior estimates.

percent of modules, high-income consumers are less price sensitive, or $\hat{\alpha}_m^1 < 0$, when you control for the fact that they also have a greater willingness to pay for quality.

The parameter estimates generally support that demand is non-homothetic within modules. In particular, high-income consumers have a greater willingness to pay for quality than low-income consumers and, when controlling for this non-homotheticity in the demand for quality, the results show that high-income consumers are also less price sensitive.[50]

The upper-level estimation yields between-module elasticity $\sigma$ estimates reported in Table 4. As expected, products in different modules are less substitutable than products in the same module, with between-module substitution elasticities close to one.

homotheticity related to price sensitivity in isolation from the non-homotheticity related to product quality.

[50]Appendix E.2.2 provides further evidence with moments demonstrating the out-of-sample fit of the model.

Table 4: Upper-Level Substitution Elasticity Estimates

| Model Name | $\sigma$ |
|---|---|
| Homothetic | 1.007 |
| | [0.010] |
| Non-Homothetic in Price | 1.007 |
| | [0.008] |
| Non-Homothetic in Quality | 1.027 |
| | [0.005] |
| Non-Homothetic in Quality and Price | 1.058 |
| | [0.009] |

Note: This table shows the estimates for the elasticity of substitution between modules. Standard errors (in brackets) have been adjusted for first-stage measurement error as described in Appendix D.2.2.

## 6.2 Model Selection

The model estimates above provide micro-evidence that high-income households have a stronger taste for high-quality products and, controlling for this, they are less price sensitive. Allowing for both forms of non-homotheticity introduces over 400 additional parameters to the model (one $\alpha_m^1$ or $\gamma_m$ for each module). These parameters will all be sources of error in the income-specific price indexes used to address the paper's main question in Section 6.3 below. Prior to undertaking this analysis, I therefore first attempt to determine whether this parametric flexibility is valuable enough to warrant these additional errors. To do this, I use the GMM-BIC model selection criterion that judges models using a trade-off between model fit and model complexity, measured using the number of parameters relative to the number of moments used in the estimation of those parameters.[51] Specifically, for each module, the GMM-BIC criterion selects the model and moment conditions that minimize the difference between the estimated $J$ statistic and the log of the number of observations multiplied by the number of over-identifying restrictions used in estimation. The results of this exercise are summarized in Table 5.

The model that permits non-homothetic demand for quality, but not for price, is the optimal model for 63 percent of modules, representing 70 percent of sample sales. Further, in bilateral model comparisons, the model that accounts for non-homothetic demand for quality had a lower GMM-BIC criterion in modules representing 82 and 79 percent of sales when compared to the models that account for non-homothetic demand for price or both price and quality, respectively.[52]

[51] This method was developed in Andrews (1999) as a moment selection criterion and is shown to be consistent for model selection in Andrews and Lu (2001).

[52] See Appendix Table A.4.

Table 5: Summary Statistics for BIC Model Selection Criteria

| | Non-Homotheticity in: | | |
|---|---|---|---|
| | Quality | Price | Price and Quality |
| Share of Modules where Model Dominates | 0.63 | 0.13 | 0.25 |
| Sales Share of Modules where Model Dominates | 0.70 | 0.11 | 0.19 |
| Average Statistic | 8,596 | 8,191 | 21,015 |
| Sales-Weighted Average Statistic | 622 | 1,889 | 789 |

Note: This table shows summary statistics describing the value of the Bayesian Information Criterion (BIC) statistic for the three non-homothetic models whose estimates are summarized in Table 3. Attention is limited to the set of modules that have interior estimates for all three models. See Appendix (D.3) for further details on the BIC calculation.

These results suggest that the salient form of non-homotheticity in grocery consumption is in the demand for quality. In the analysis below, I limit my attention to price indexes that account for this form of non-homotheticity alone when studying how grocery costs vary across local markets differently for consumers at different income levels. Any differences between the relative price indexes high- and low-income consumers face across cities and stores will reflect differences in the availability and prices of high- relative to low-quality products across these markets.[53]

## 6.3 Income-Specific Consumption Externalities

The analysis above has provided the inputs to market- and income-specific price indexes that represent how households at different income levels value the products and prices available to them in different U.S. cities and neighborhoods, as outlined in Section 5 above. I can now turn to answering the central question in this paper: do grocery costs vary differently across markets for consumers at different income levels?

To answer this question, I estimate the following regression:

$$\ln \hat{P}(\mathbb{P}_c, y_k) = \delta_k + \beta_1 y_c + \beta_2 (y_k - \bar{y}_k) y_c + \epsilon_{kc}, \tag{25}$$

where $\hat{P}(\mathbb{P}_c, y_k)$ is the grocery price index for a representative consumer with log income $y_k$ in each market $c$, obtained by plugging the market-specific price vector $\mathbb{P}_c$, income $y_k$, and model parameter estimates into equation (15); $\delta_k$ is an income-level fixed effect; $y_c$ is log per capita

[53]Conversely, these price indexes do not allow for non-homotheticity in consumer's price sensitivity (or idiosyncratic utility weight). So, while high-income consumers face relatively lower costs in markets with relatively more, and cheaper, high-quality products than low-quality products, all consumers get the same additional utility, and cost savings, in markets that offer more varieties and lower prices of both high- and low-quality products equally.

income in city $c$, and $\bar{y}_k$ is the mean log household income in the sample.

In the above specification, the coefficient on log city income ($\beta_1$) reflects the mean elasticity of grocery costs with respect to city income. The coefficient on the interaction of demeaned log consumer income and log city income ($\beta_2$) measures how the elasticity varies with household income. The grocery price index, $\hat{P}(\mathbb{P}_c, y_k)$, is calculated using a model that allows for non-homotheticity in the demand for quality, so the elasticity of grocery costs with respect to city income will vary with income, and $\beta_2$ will be non-zero, if the goods and prices available in each city are correlated with the tastes corresponding to the average income of the consumers living there. If wealthy cities offer more varieties of high-quality goods at lower prices than poorer cities, the price index faced by high-income consumers will decrease by more (or increase by less) than the price index faced by low-income consumers between poor and wealthy cities. This is because high-income consumers benefit more from the availability and lower prices of the goods that they prefer. Under this scenario, the elasticity of the price index faced by high-income consumers with respect to city income would be lower than the elasticity of the price index faced by low-income consumers with respect to city income yielding a negative $\beta_2$ estimate.[54]

Table 6 presents the results of the baseline regression estimated using income-specific price indexes calculated for price vectors reflecting the prices and products available at 100 random samples of 50 stores in each of the 125 CBSAs that have 50 or more stores.[55] The $\beta_1$ coefficient on CBSA per capita income is negative but not significant, reflecting the large degree of noise in the price indexes across CBSAs making it impossible to identify a systematic relationship between the mean price index that a household faces in a city and its per capita income. There is, on the other hand, strong evidence that the elasticity of the price index with respect to per capita income increases with household income: the $\beta_2$ coefficients on the interaction between CBSA per capita income and demeaned household income are negative and statistically significant. The magnitude of the $\beta_2$ estimate indicates that this variation is economically significant. A consumer who earns \$25,000 a year sees their per dollar grocery costs decrease by around 4 percent for each log unit increase in city per capita income, comparable to the gap between the wealthiest and poorest cities in the sample (Bridgeport-Stamford-Norwalk, CT with per capita income of \$49,688 and El Paso, TX with per capita income of \$18,684). On the other hand, the

[54]This regression characterizes an equilibrium relationship and should not be interpreted causally. The results presented here do not indicate whether, for example, grocery costs are lower for high-income consumers in wealthy cities because a high per capita income causes stores in a city to stock more high-quality products or because high-quality products attract more high-income inhabitants to a city, raising its per capita income.

[55]Formally, the regression estimated is:

$$\ln \hat{P}(\mathbb{P}_{cb}, y_k) = \delta_{kb} + \beta_1 y_c + \beta_2 (y_k - \tilde{y}_k) y_c + \epsilon_{kcb},$$

where $\mathbb{P}_{cb}$ denotes the set of prices available to consumers in the 50 stores in bootstrap sample $b$ for CBSA $c$ and $\delta_{kb}$ is a bootstrap sample-household income group fixed effect. Standard errors are clustered at the CBSA level.

per dollar grocery costs of a consumer with a yearly income of $200,000 decreases by over 50 percent for each log unit increase in city per capita income. A high-income household would experience approximately a 10 percent greater decrease in grocery costs than a low-income household when both move from a CBSA at the 25th percentile of the income distribution (e.g., Scranton, PA) to a CBSA at the 75th percentile of the income distribution (e.g., Pittsburgh, PA).

Table 6: City-Income Specific Price Index Regressions

| Dependent Variable: Ln(Price Index for Household in Income Group $k$ in CBSA $c$) | Local Prices | | National Prices | |
|---|---|---|---|---|
| | [1] | [2] | [3] | [4] |
| Ln(Per Capita Income$_c$) | -0.32 | -0.30 | -0.082 | -0.15 |
| | (0.21) | (0.20) | (0.19) | (0.21) |
| Ln(Per Capita Income$_c$)* Demeaned Ln(HH Income$_k$) | -0.25*** | -0.27*** | -0.22*** | -0.25*** |
| | (0.040) | (0.045) | (0.038) | (0.030) |
| Ln(Population$_c$) | | -0.0075 | | 0.020 |
| | | (0.047) | | (0.046) |
| Ln(Population$_c$)* Demeaned Ln(HH Income$_k$) | | 0.0050 | | 0.0093*** |
| | | (0.0031) | | (0.0018) |
| Income Group $k$*Bootstrap Sample FEs | Yes | Yes | Yes | Yes |
| Number of CBSAs ($c$) | 125 | 125 | 125 | 125 |
| Observations | 98,920 | 98,920 | 98,920 | 98,920 |
| adj. $R^2$ | 0.03 | 0.03 | 0.01 | 0.01 |

Notes: *** p< 0.01, ** p<0.05, * p<0.10; standard errors, clustered by bootstrap sample and CBSA, are in parentheses. This table presents results from regressions of household income- and CBSA-specific grocery price indexes against CBSA characteristics alone and interacted with demeaned log household income. The price indexes correspond to the baseline model that allows for non-homotheticity in the demand for quality but not in price sensitivity (i.e., restricting that $\alpha^1_m$=0) and measure how households at eight different income levels between $25,000 and $200,000 value the products and prices represented in each of 100 bootstrap samples of 50 stores in each of 125 CBSAs with 50 or more participating retailers. Bootstrap sample-CBSA-income group-level observations are weighted by CBSA population.

Market income is correlated with market size: wealthier sample cities are larger than poorer sample cities with a correlation coefficient of 0.35. Therefore, it is possible that a negative $\beta_2$ estimate in the baseline regression could result from grocery costs being lower for high-income households than for low-income households in larger, as opposed to wealthier, cities. In column [2] of Table 6, I therefore add controls for log population and log population interacted with log household income to the baseline regression. The $\beta_2$ coefficient is robust to these controls, whose coefficients are estimated as precise zeros. This evidence is consistent with the "within-group preference externalities" story in which higher income consumers receive relatively more consumption benefits from living in wealthier cities, as opposed to a story in which high-income consumers receive more consumption benefits from living in larger cities than low-income consumers.[56]

### Differentiating between Price and Variety Effects

The results above suggest that, relative to low-income households, high-income households receive higher consumption utility from the grocery bundles available in wealthier cities than from the grocery bundles available in poorer cities with the same population. The model allows for high-income households to have a stronger preference for high-quality goods than do low-income households. So, the fact that high-income households get relatively more utility from consuming grocery products in high-income cities must be either because there are more high-quality goods available in these locations or because the high-quality goods are sold at relatively lower prices in high-income cities, or for both reasons. I examine this issue by calculating income-specific price indexes for the set of products I observe in the 50-store sample for each city, as before, but setting the prices of each product equal to its national average price.

Columns [3] through [4] of Table 6 replicate columns [1] through [2] using these fixed-price indexes as the dependent variable. The coefficients on the interaction between per capita income and household income only decrease slightly in magnitude. High-income households would continue to find wealthy cities almost as cheap relative to poor cities, relative to low-income households, if products were sold in both locations at their national average price. This indicates that the difference in how high- and low-income households perceive the relative costs to vary across cities is almost all due to variety differences. The products that high-income consumers prefer to consume are sold at higher prices in wealthy cities than they are in poor cities, but high-income consumers are more than compensated for this price difference by the fact that more of these products are available to them in these locations.

### Semi-Parametric Estimates

The regression estimated above imposes that the elasticity of the income-specific price index with respect to city income is log-linear in income. There is no reason for this to be the case. To obtain non-parametric estimates of these elasticities at different income levels, I estimate the above regression specification but with a household income dummy interacted with per capita city income instead of the household income level interacted with per capita city income:

$$\ln \hat{P}(\mathbb{P}_c, y_k) = \delta_k + \beta_1 y_c + \beta_{2k} y_c + \epsilon_{kc}, \tag{26}$$

I estimate this regression separately for each set of 100 bootstrapped samples of 50 random stores from each CBSA. Figure 6 plots the mean of the resulting $\beta_{2k}$ elasticity parameter es-

[56]Table A.7 in Appendix E shows that these main results only change marginally when based on price indexes that account for non-homotheticities in both consumer's demand for quality and their price sensitivity.

timates against log household income, $y_k$. These results indicate that there is a non-linear relationship between this elasticity and household income, with the downward slope flattening out somewhat at the lower tail of the income distribution.[57]

Figure 6: Elasticity of Grocery Price Index with respect to CBSA Income for Households at Different Size-Adjusted Income Levels

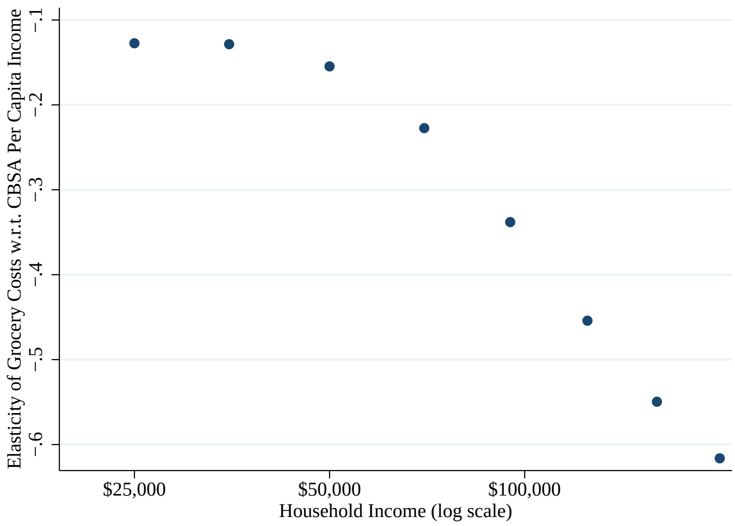


Notes: This figure plots the elasticity of the grocery price index for households at different income levels across different CBSAs with respect to CBSA per capita income. The price indexes correspond to the model that allows for non-homotheticity in the demand for quality but not in price sensitivity (i.e., restricting that $\alpha^1_m$=0). Each index measures how households at a given income level values the products and prices represented one of 100 bootstrap samples of 50 stores in each of 125 CBSAs with 50 or more participating retailers. The elasticities of these indexes with respect to CBSA per capita income are calculated by regressing the indexes against household income fixed effects and these fixed effects interacted with log CBSA per capita income separately for each bootstrap sample. Observations in these regressions are weighted by CBSA population. The figure plots the mean of these elasticities at each household-income level across bootstrap samples.

### Variation within CBSAs

We see similar variation in the per dollar grocery utility offered to high- and low-income households across stores in different neighborhoods as we did across CBSAs. Table 7 presents the elasticity estimates from equation (25) where market $c$ denotes a store.[58] Column [1] shows a similar qualitative pattern in the variation in the elasticity of price indexes with respect to household income across stores with different local per capita income as we saw across CBSAs with different per capita income. With these store-level indexes, we can consider whether sorting within CBSAs might enable households to mitigate some of the cross-CBSA variation in grocery availability. Column [2] shows that the elasticity of store-level indexes with respect

[57] Appendix Figure A.16 show that these elasticities vary much less across bootstrap samples than they do between incomes below and above $100,000, for example.

[58] For the store-level results, $\hat{P}(\mathbb{P}_c, y_k)$ reflects the grocery price index of a representative household earning $y_k$ faces in store $c$ and $y_c$ is the average size-adjusted income in the vicinity of store $c$, calculated using the non-parametric method described in Appendix A.2..

to household income is also increasing with CBSA income. Columns [3] and [4] show that this correlation is stronger when comparing the indexes for stores located in the high-income neighborhoods in different CBSAs. That is, the relationship between grocery costs and CBSA income is amplified for residents of high-income neighborhoods and mitigated for residents of low-income neighborhoods.

Table 7: Store Price Index Regressions

| | Dependent Variable: Ln(Price Index for Representative Consumer $k$ in Store $c$) | | | | | | | |
|---|---|---|---|---|---|---|---|---|
| | [1] | [2] | [3] | [4] | [5] | [6] | [7] | [8] |
| Ln(Per Capita Income$_c$) | -0.11*** | | | | -0.0040 | | | |
| | (0.0084) | | | | (0.0056) | | | |
| Ln(Per Capita Income$_c$)* Demeaned Ln(HH Income$_k$) | -0.033** | | | | -0.0014 | | | |
| | (0.0099) | | | | (0.0031) | | | |
| Ln(CBSA Per Capita Inc.) | | -0.13*** | -0.13* | -0.074 | | 0.0036 | 0.0023 | -0.0076 |
| | | (0.029) | (0.063) | (0.042) | | (0.016) | (0.022) | (0.017) |
| Ln(CBSA Per Capita Inc.)* Demeaned Ln(HH Income$_k$) | | -0.047** | -0.072** | -0.024*** | | 0.0072 | 0.0072 | 0.00034 |
| | | (0.014) | (0.021) | (0.0061) | | (0.0077) | (0.018) | (0.0041) |
| Income Group $k$ Fixed Effects | Yes | Yes | Yes | Yes | Yes | Yes | Yes | Yes |
| Chain x Income Group Fixed Effects | No | No | No | No | Yes | Yes | Yes | Yes |
| Store Set (Local Per Capita Income$_c$) | All | All | High-Inc. | Low-Inc. | All | All | High-Inc. | Low-Inc. |
| Number of Stores ($c$) | 9311 | 8859 | 4637 | 4222 | 9310 | 8858 | 4633 | 4212 |
| Number of CBSAs | - | 691 | 175 | 650 | - | 691 | 175 | 649 |
| Observations | 74,488 | 70,872 | 37,096 | 33,776 | 74,480 | 70,864 | 37,064 | 33,696 |
| adj. $R^2$ | 0.03 | 0.02 | 0.02 | 0.01 | 0.56 | 0.57 | 0.57 | 0.55 |

Notes: *** p< 0.01, ** p<0.05, * p<0.10; standard errors, clustered by store and household income in columns 1 and 5 and by CBSA and household income in columns 2 through 4 and 6 through 8, are in parentheses. This table presents results from regressions of household income- and store-specific grocery price indexes against measures of local store income alone and interacted with demeaned log household income. The price indexes correspond to the baseline model that allows for non-homotheticity in the demand for quality but not in price sensitivity (i.e., restricting that $\alpha^1_m$=0) and measure how households at eight different income levels between \$25,000 and \$200,000 value the products and prices represented in grocery stores in the Nielsen RMS sample. Store-by-income group observations are weighted by store sales.

The results in columns [5] through [8] show that the variation in columns [1] through [4] can be entirely explained by variation in the set of retail chains that locate in high- vs. low-income neighborhoods. Retail chains do not appear to significantly alter the mix of brands they offer across neighborhoods or CBSAs in a way that biases the attractiveness of their stores in higher-income locations to higher-income customers.

## 6.4 Robustness Checks

### 6.4.1 Robustness to Different Estimation Choices

Table 8 demonstrates the robustness of the demand parameter estimates used in the index calculation above to various decisions made in the course of estimation. Due to computation limitations, the main estimation procedure grouped any products with expenditure shares below the 60th percentile in a given store-month to an outside product for that store-month. The first row shows the distribution of parameter values under this base specification. The next two rows show the robustness of key parameter estimates – $\alpha_m^0$, $\gamma_m$, and the standard deviation of quality $\beta_{mg}$ across products – to allocating either fewer or more products (those below the 40th or 80th percentiles) to the outside product. The distributions of these $\alpha_m^0$ and $\gamma_m$ parameters are shifted marginally upward, while the variation in the quality estimates falls. Table 8 next shows that the parameters estimated using data at a weekly or quarterly, instead of monthly, frequency. The price elasticity estimates are higher in both cases, but more so when the frequency of the data is reduced, indicating that there is some attenuation bias in the parameter estimates. Finally, Table 8 shows the distribution of parameter estimates reached in a two-step procedure that uses the Feenstra (1994) method to identify the price coefficients and then – holding those price coefficients fixed – uses the moment conditions described above to identify the remaining model parameters.[59] As expected, the price coefficients estimated using this method are larger, with a median of 2.88. As in other robustness checks discussed above, the $\gamma_m$ parameters, that govern the non-homotheticity in demand for quality are also larger and there is less estimated variation in product quality.

Table 8: Robustness of Parameter Estimates

| | Price Elasticity ($\alpha_0$) | | | | Income-Quality Elasticity ($\gamma$) | | | | St. Dev. of Quality ($\sigma(\beta_{mg})$) | | | |
|---|---|---|---|---|---|---|---|---|---|---|---|---|
| | | Percentile | | | | Percentile | | | | Percentile | | |
| Specification | Mean | 25th | 50th | 75th | Mean | 25th | 50th | 75th | Mean | 25th | 50th | 75th |
| Base | 2.13 | 0.89 | 1.70 | 2.40 | 1.40 | 0.39 | 1.01 | 2.47 | 0.87 | 0.30 | 0.48 | 0.98 |
| Outside Good Cutoff = 40% | 3.55 | 1.40 | 2.36 | 4.26 | 1.72 | 0.49 | 1.72 | 2.65 | 0.67 | 0.25 | 0.46 | 0.73 |
| Outside Good Cutoff = 80% | 2.72 | 1.10 | 1.88 | 2.81 | 1.62 | 0.51 | 1.74 | 2.27 | 0.66 | 0.24 | 0.40 | 0.67 |
| Weekly Data | 2.79 | 1.19 | 1.90 | 3.24 | 1.71 | 0.64 | 1.69 | 2.57 | 0.68 | 0.27 | 0.43 | 0.57 |
| Quarterly Data | 3.72 | 1.33 | 2.29 | 4.39 | 1.64 | 0.72 | 1.39 | 2.51 | 0.75 | 0.27 | 0.37 | 0.56 |
| Feenstra | 3.30 | 1.35 | 2.88 | 4.61 | 1.41 | 0.51 | 1.38 | 2.17 | 0.67 | 0.26 | 0.43 | 0.71 |

Notes: This table presents the distribution of interior parameter estimates obtained for the baseline model that allows for non-homotheticity in the demand for quality but not in price sensitivity (i.e., restricting that $\alpha_m^1$=0) using different estimation samples and specifications discussed in the text.

The variation in parameter estimates across specifications is important insofar as it generates differences in how high- and low-income households perceive the relative product offerings across high- and low-income markets. To understand what drives these differences, consider a simple case in which there is a single module containing two types of varieties, indexed by $j \in \{H, L\}$. All products of each type have the same quality, denoted $\beta_j$ (where $\beta_H > \beta_L$) and, conditional on being available in a market, are sold at the same price, $p_j$. Markets, indexed by $t$, are only differentiated by the number of products of each type available, denoted $N_{Ht}$ and $N_{Lt}$.

[59]See Appendix D.4 for further details on how the Feenstra (1994) method is applied in this context.

Under these assumptions, the log price index for market $t$ for a household with income $y_k$ is

$$\ln P(\mathbb{P}_t, y_k) \quad = \quad -\frac{1}{\alpha^0} \ln \left[ N_{Ht} \left( \frac{p_H}{b_{Hk}} \right)^{-\alpha^0} + N_{Lt} \left( \frac{p_L}{b_{Lk}} \right)^{-\alpha^0} \right]$$

Consider how the price index changes differentially for households with different incomes, as the number of high-quality products available in the market increases. The cross-elasticity of the log price index with respect to the number of high-quality varieties $N_{Ht}$ and household income $y_k$ is:

$$(27) \qquad \frac{\partial \ln P(\mathbb{P}_t, y_k)}{\partial N_{Ht} \partial y_k} \quad = \quad -\gamma \left( \beta_H - \beta_L \right) \left\{ \frac{N_{Lt} \left( \frac{p_H}{b_{Hk}} \right)^{-\alpha^0} \left( \frac{p_L}{b_{Lk}} \right)^{-\alpha^0}}{\left[ N_{Ht} \left( \frac{p_H}{b_{Hk}} \right)^{-\alpha^0} + N_{Lt} \left( \frac{p_L}{b_{Lk}} \right)^{-\alpha^0} \right]^2} \right\}$$

This cross-elasticity is proportional to the quality-income gradient $\gamma$ and the difference in the quality of high- and low-quality products $(\beta_H - \beta_L)$. These two key estimates govern the differential gains that high-income consumers get from the higher quality product assortment in wealthier cities relative to low-income consumers are the income-quality elasticity ($\gamma_m$) and the degree of variation in the estimated quality across products, rather than the mean price elasticity ($\alpha_m^0$).[60]

Reassuringly, these estimates are relatively stable across specifications. If anything, the specifications that yield higher estimates for the mean price elasticity ($\alpha_m^0$) tend to also show higher estimates for the income-quality elasticity ($\gamma_m$) and less variation in the estimated quality across products.[61] These shifts balance out, making the relative gains to wealthy consumers from access to high-quality products quite robust across specifications. The variation in the relative grocery costs across different U.S. cities across households at different incomes is the same order of magnitude using price indexes based on estimates using the Feenstra method to identify price elasticities, the residualized instrument, quarterly frequency data, and a less restrictive cutoff for inclusion in the outside product. Across these specifications, high-income households are estimated to find wealthy cities between 20 and 50 percent less expensive than poor cities relative to low income households.[62]

[60]As is standard in CES setting, the welfare effect of having more variety – measured as the elasticity of the log price index with respect to $(N_{Ht} + N_{Lt})$ – is proportional to the inverse price elasticity, $1/\alpha^0 = 1/(\sigma - 1)$. This variety effect, however, cancels out when considering the differential effect of variety on the price index of households with different incomes. What matters for this differential welfare effect is the mix of varieties available.

[61]This makes sense. The quality parameters are scaled by the price (substitution) elasticity in the demand function, so when $\alpha_m^0$ is higher, less variation in $\beta_{mg}$ is required to rationalize the observed variation in sales across products. Conditional on there being less variation in quality across products, a higher income-quality elasticity gradient is required to match the bias of high-income purchases towards higher quality products.

[62]See Appendix Figure A.18.

### 6.4.2 Measurement Error in Quality

To estimate product quality, I have assumed that the quality of the outside good in each module is equal across markets. In practice, variation in the quality of the outside product across store-months will generate errors in the relative quality estimates ($\tilde{\beta}_{mg}$). One concern is that quality may be mis-measured in a way that biases the gradient of the quality elasticity with respect to income ($\gamma_m$). For example, suppose that high-income households tend to purchase products stocked in stores that also offer higher quality outside goods. $\tilde{\beta}_{mg}$ will then understate the relative quality of products that high-income households purchase, and overstate the relative quality of products that low-income households purchase. This could lead me to overstate the income-quality elasticity gradient ($\gamma_m$).[63]

I run two tests to gauge the degree of this error and its potential to bias the $\gamma_m$ estimates. The results, in Appendix E.1.3, show that these errors are typically small in magnitude. Importantly, I find that the errors are not correlated with the purchasing behavior of high- vs. low-income households in such a way that would bias the income-quality elasticity gradient. The robustness of the $\gamma_m$ and $\beta_{mg}$ parameter estimates to alternate definitions of the outside good in Table 8 is also reassuring.

### 6.4.3 Alternative Sources of Demand Heterogeneity

The price indexes calculated here account for how consumer tastes vary with income both across products in the same category and across categories of products. Income is a factor in determining a consumer's preferences over different types of breakfast cereal, for example, as well as in determining their willingness to pay for cereal relative to milk. In order to make this multi-sector analysis tractable, I have abstracted from a number of other ways in which demand and, therefore, aggregate costs could vary across heterogeneous households.

In particular, empirical micro-economists have shown that income is just one of a range of demographic characteristics that can be correlated with consumer demand for a variety of product characteristics, including brand quality. The model here is more stylized, allowing the willingness to pay for a single product characteristic, brand name, to vary with a single consumer characteristic, income. The benefit of such a simple framework is that it is generalizable: none of the variables are category-specific so it can be used to measure how demand varies systematically with consumer characteristics across products in many product categories. The

[63] Alternatively, if the bias is so large that the ordering of product quality is not maintained – such that products that high-income households favor are estimated to have lower relative quality than products low-income households favor when they are in fact higher quality (or vice versa) – I could estimate the wrong sign for the income-quality elasticity gradient ($\gamma_m$). In this case, the main result that high-income markets offer more of the products that high-income households favor and, therefore, provide high-income households with relatively lower grocery costs than low-income markets, would hold, but the interpretation that these products are higher quality (i.e., preferred on average by all households) would not.

drawback is that it imposes two types of strong assumptions on the consumer tastes.

The first is that households value units of products from the same brand and module equally, regardless of their flavor, texture, or the size and type of container they were packaged in. The cross-city price indexes I calculate account for the fact that high-quality brand name products are more available or sold at cheaper prices than low-quality brand name products in some cities than in others, but the prices of products in the same module and brand enter symmetrically, even if they have different sizes, container types, etc.. For violations of this assumption to bias the results of the paper, low-income tastes would need to be biased towards product characteristics that are disproportionately represented (or available at lower prices) in high-income cities. This is unlikely to be the case. I do not, for example, find any statistically significant correlations between either the price or availability of products with certain sizes and per capita income when controlling for product module and brand name.

The second simplification in the model above is that, controlling for size-adjusted household income, consumer demand does not vary systematically with other demographics, such as age, marital status, and education. The consumption patterns and parameter estimates above are consistent with non-homotheticities in demand but may instead pick up correlations between demand and these other demographics, to the extent that age, marital status, and education are also correlates of income. Similarly, the estimated patterns in product availability across high- and low-income markets are consistent with local firms catering to income-specific tastes, but could also be the result of preference externalities along other demographics or unrelated supply-side factors. It is important, therefore, to caution against interpreting these results causally. More work is needed to assess the role of preference externalities in grocery retail.

# 7 Conclusion

There is growing interest in the role of non-homothetic preferences and cross-market income differences in determining production patterns in macro, urban, and international economics. If preferences are income-specific and, further, if the products available in different markets are biased to the income-specific tastes in these markets, then consumers at different income levels will experience different changes in consumption utility across these markets. The results in this paper indicate that this is indeed the case: high-income households face greater grocery consumption gains from moving to high per capita income markets than do low-income households.

I show that high-income households face much lower grocery costs in wealthy cities than in poor cities, while low-income households face slightly higher grocery costs in these locations. Further work is required to extend the analysis presented here to other components of household expenditure in order to build income-specific aggregate spatial price indexes that can be used, for example, in real income measurement or in a Rosen-Roback framework to look at the role of these pecuniary consumption amenities, relative to skill-biased productivity spillovers, in explaining skill-biased agglomeration. Recent work by Atkin et al. (2019) suggests a promising path forward in this direction.

I do not expect that these grocery cost differentials are representative of the differentials that we would expect in other components of the typical consumer basket. For one, I expect that the availability of the food and fast-moving consumer goods represented in my sample varies less geographically than other parts of the consumption basket like non-tradable services and housing. If anything, I would expect the strength of consumption externalities to be higher in sectors that are less tradable. So, conditional on these other products having similar degrees of demand heterogeneity, I would consider my estimates to be a lower bound for the differentials we would expect to see in aggregate price indexes.[64]

[64]Even under the conservative assumption that preferences are homothetic within each of the households other consumption areas, the difference in relative grocery costs alone implies an economically-significant 1.4 percent gap between the aggregate living costs faced by high-income households in wealthy, relative to poor, cities and those faced by low-income households. High-income households, who spend around 2 percent of their annual income on groceries, would face 0.94 percent lower living costs in wealthy cities, whereas low-income households, who spend around 10 percent of their annual income on groceries, would face 0.4 percent higher living costs in these locations.

## References

Aguiar, M. and E. Hurst (2005). Consumption versus Expenditure. *Journal of Political Economy 113*(5), 919–948.

Albouy, D. (2009, August). The Unequal Geographic Burden of Federal Taxation. *Journal of Political Economy 117*(4), 635–667.

Albouy, D., G. Ehrlich, and Y. Liu (2016). Housing demand, cost-of-living inequality, and the affordability crisis. Technical report, National Bureau of Economic Research.

Almas, I. (2012). International Income Inequality: Measuring PPP Bias by Estimating Engel Curves for Food. *The American Economic Review 102*(2), 1093–1117.

Anderson, S. P., A. de Palma, and J.-F. Thisse (1987). The CES is a Discrete Choice Model? *Economics Letters 24*, 139–140.

Andrews, D. W. K. (1999, May). Consistent Moment Selection Procedures for Generalized Method of Moments Estimation. *Econometrica 67*(3), 543–564.

Andrews, D. W. K. and B. Lu (2001). Consistent Model and Moment Selection Procedures for GMM Estimation with Application to Dynamic Panel Data Models. *Journal of Econometrics 101*, 123–164.

Argente, D. and M. Lee (2017). Cost of living inequality during the great recession. *Kilts Center for Marketing at Chicago Booth–Nielsen Dataset Paper Series*, 1–032.

Atkin, D., B. Faber, T. Fally, and M. Gonzalez-Navarro (2019). A New Engel on Price Index and Welfare Estimation.

Auer, R. (2010, August). Consumer Heterogeneity and the Impact of Trade Liberalization: How Representative is the Representative Agent Framework? *Swiss National Bank Working Paper* (13).

Bai, J. and S. Ng (2010). Instrumental Variable Estimation in a Data Rich Environment. *Econometric Theory 26*, 1577–1606.

Berry, S., J. Levinsohn, and A. Pakes (1995). Automobile Prices in Market Equilibrium. *Econometrica 63*(4), 841–890.

Berry, S., J. A. Levinsohn, and A. Pakes (2004). Differentiated Products Demand Systems from a Combination of Micro and Macro Data: the New Car Market. *Journal of Pol 112*(1), 68–105.

Bils, M. and P. J. Klenow (2001). Quantifying Quality Growth. *The American Economic Review 91*(4), 1006–1030.

Black, D., N. Kolesnikova, and L. Taylor (2009). Earnings functions when wages and prices vary by location. *Journal of Labor Economics 27*(1), 21–47.

Broda, C., E. Leibtag, and D. E. Weinstein (2009). The role of prices in measuring the poor's living standards. *The Journal of Economic Perspectives 23*(2), 77–97.

Broda, C. and J. Romalis (2009). The Welfare Implications of Rising Price Disperison. NBER Working Paper No. 18314.

Broda, C. and D. E. Weinstein (2010). Product Creation and Destruction: Evidence and Price Implications. *American Economic Review 100*, 691.

Cage, R. A., T. Garner, and J. Ruiz-Castillo (2002). Constructing household specific consumer price indexes: An analysis of different techniques and methods.

Couture, V. and J. Handbury (2017). Urban revival in America, 2000 to 2010. Technical report, National Bureau of Economic Research.

Cravino, J. and A. A. Levchenko (2017). The distributional consequences of large devaluations. *American Economic Review 107*(11), 3477–3509.

de Palma, A. and K. Kilani (2007). Invariance of Conditional Maximum Utility. *Journal of Economic Theory 132*, 137–146.

Deaton, A. (2010). Price indexes, inequality, and the measurement of world poverty. *American Economic Review 100*(1), 5–34.

Deaton, A. and O. Dupriez (2011a, April). Purchasing Power Parity Exchange Rates for the Global Poor. *American Economic Journal: Applied Economics 3*(2), 137–166.

Deaton, A. and O. Dupriez (2011b). Spatial Price Differences within Large Countries. *Woodrow Wilson School of Public and International Affairs Working Paper* (1321).

Deaton, A. and J. Muellbauer (1980). *Economics and Consumer Behavior*. Cambridge University Press.

DellaVigna, S. and M. Gentzkow (2019). Uniform pricing in us retail chains. *The Quarterly Journal of Economics 134*(4), 2011–2084.

Diamond, R. (2016). The determinants and welfare implications of us workers' diverging location choices by skill: 1980-2000. *American Economic Review 106*(3), 479–524.

Dingel, J. I. (2016). The determinants of quality specialization. *The Review of Economic Studies 84*(4), 1551–1582.

Dube, J.-P. (2004). Multiple Discreteness and Product Differentiation: Demand for Carbonated Soft Drinks. *Marketing Science 23*(1), 66–81.

Dubé, J.-P. H., J. T. Fox, and C.-L. Su (2009). Improving the numerical performance of blp static and dynamic discrete choice random coefficients demand estimation. Technical report, National Bureau of Economic Research.

Faber, B. (2012, November). Trade Liberalization, the Price of Quality, and Inequality: Evidence from Mexican Store Prices. Mimeo London School of Economics.

Faber, B. (2014). Trade liberalization, the price of quality, and inequality: Evidence from mexican store prices. *UC-Berkeley Working Paper*.

Faber, B. and T. Fally (2017). Firm heterogeneity in consumption baskets: Evidence from home and store scanner data. Technical report, National Bureau of Economic Research.

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

Feenstra, R. C. (1994). New Product Varieties and the Measurement of International Prices. *The American Economic Review 84*(1), 157–177.

Feenstra, R. C. and J. Romalis (2012, August). International Prices and Endogenous Quality. Mimeo University of Chicago.

Feenstra, R. C. and J. Romalis (2014). International prices and endogenous quality. *The Quarterly Journal of Economics 129*(2), 477–527.

Gandhi, A., Z. Lu, and X. Shi (2013). Estimating demand for differentiated products with error in market shares. Technical report, cemmap working paper.

Garner, T. I., D. S. Johnson, and M. F. Kokoski (1996). An experimental consumer price index for the poor. *Monthly Lab. Rev. 119*, 32.

Glaeser, E. L., J. Kolko, and A. Saiz (2001). Consumer city. *Journal of Economic Geography 1*, 27–50.

Hallak, J. C. (2006). Product Quality and the Direction of Trade. *Journal of International Economics 68*(1), 238–265.

Handbury, J. and D. E. Weinstein (2014). Goods prices and availability in cities. *The Review of Economic Studies 82*(1), 258–296.

Hausman, J. A., G. Leonard, and J. D. Zona (1994). Competitive Analysis with Differentiated Products. *Annals D'Economie et Statistique 34*, 159–180.

Hendel, I. (1999). Estimating Multiple-Discrete Choice Models: An Application to Computerization Returns. *The Review of Economic Studies 66*(2), 423–446.

Hortaçsu, A. and J. Joo (2015). Semiparametric estimation of a CES demand system with observed and unobserved product characteristics. *arXiv preprint arXiv:1511.05600*.

Hottman, C. (2014). Retail markups, misallocation, and store variety in the US. *Mimeograph, Columbia University*.

Hottman, C. and R. Monarch (2018). Estimating unequal gains across US consumers with supplier trade data. *FRB International Finance Discussion Paper* (1220).

Hottman, C. J., S. J. Redding, and D. E. Weinstein (2016). Quantifying the sources of firm heterogeneity. *The Quarterly Journal of Economics 131*(3), 1291–1364.

Hummels, D. and V. Lugovskyy (2009). International Pricing in a Generalized Model of Ideal Variety. *Journal of Money, Credit and Banking 41*(1), 3–33.

Jaimovich, N., S. Rebelo, A. Wong, and M. B. Zhang (2019). Trading Up and the Skill Premium. Technical report, National Bureau of Economic Research.

Jaravel, X. (2018). The unequal gains from product innovations: Evidence from the us retail sector. *The Quarterly Journal of Economics 134*(2), 715–783.

Jorgenson, D., D. T. Slesnick, and D. Slottje (1989). Redistributional policy and the measurement of poverty. *Research on Economic Inequality*.

Kokoski, M. F. (1987). Consumer price indexes by demographic group. *Bureau of Labor Statistics (Washington DC) Working Paper* (167).

Li, N. (2012, January). An Engel Curve for Variety. Mimeo University of Toronto.

Manson, S., J. Schroeder, D. Van Riper, and S. Ruggles. *IPUMS National Historical Geographic Information System: Version 13.0 [Database](NHGIS). Minneapolis: University of Minnesota. 2018.*

Moretti, E. (2013). Real Wage Inequality. *American Economic Journal: Applied Economics 5*(1), 65–103.

Murphy, K. M. and R. H. Topel (1985). Estimation and Inference in Two-Step Econometric Models. *Journal of Business & Economic Statistics 3*(4), 370–379.

Nevo, A. (2000). Mergers with Differentiated Products: The Case of the Ready-to-Eat Cereal Industry. *The RAND Journal of Economics 31*(3), 395–421.

Nevo, A. (2001). Measuring Market Power in the Ready-to-Eat Cereal Industry. *Econometrica 69*(2), 307–342.

Newey, W. K. (1984). A Method of Moments Interpretation of Sequential Estimators. *Economics Letters 14*, 201–206.

Newey, W. K. and D. McFadden (1994). Large Sample Estimation and Hypothesis Testing. *Handbook of Econometrics IV*, 2111–2245.

Pinjari, A. R. and C. Bhat (2010). A multiple discrete-continuous nested extreme value (MDC-NEV) model: Formulatin and application to non-worker activity time-use and timing behavior on weekdays. *Transportation Research Part B: Methodological 44*(4), 562–583.

Simonovska, I. (2010). Income Differences and Prices of Tradables. *UC Davis Working Paper No. 10-15*.

Slesnick, D. T. (2002, May). Prices and Regional Variation in Welfare. *Journal of Urban Economics 51*(3), 446–468.

Snyder, E. M. (1956). Measuring Comparable Living Costs in Cities of Diverse Characteristics. *Monthly Labor Review*, 1187–1190.

Snyder, E. M. P. (1961). Cost of Living Indexes for Special Classes of Consumers. In *The Price Statistics of the Federal Goverment*, pp. 337–372. NBER.

Song, I. and P. K. Chintagunta (2007, November). A Discrete-Continuous Model for Multicategory Purchase Behavior of Households. *Journal of Marketing Research 54*, 595–612.

Verhoogen, E. (2008). Trade, Quality Upgrading, and Wage Inequality in the Mexican Manufacturing Sector. *Quarterly Journal of Economics 123*(2), 489–530.

Waldfogel, J. (2003). Preference Externalities: An Empirical Study of Who Benefits Whom in Differentiated-Product Markets. *The RAND Journal of Economics 34*(3), 557–568.

Winkelried, D. and R. J. Smith (2011). Principal components instrumental variable estimation.

Ziliak, J. P. (Ed.) (2011, May). *Cost of Living and the Supplemental Poverty Measure: A Research Forum*.

# Appendices for Online Publication

## A Data Appendix

### A.1 Data Cleaning

The estimation sample is cleaned in various ways. Throughout I define prices on a per unit basis, limiting my attention to products whose container size is expressed in the modal units for that module. I exclude any module whose modal container size is either not expressed in meaningful units (i.e., counts instead of weights or volume) or in the same units for at least 75% of UPCs. Approximately one quarter of modules do not satisfy these restrictions. I also exclude random weight items, whose quality can be variable over time.[65]

To control for data recording errors, I drop any market (store-month) in which I observe a UPC sold at a unit price greater than three times or less than a third of the median unit price paid per unit of any UPC within the same product or module categorization. The typical module loses 6.5 percent of product-store-month observations for this reason, and 5.6 percent of store-month markets.

For computational reasons, I group any products with small sales shares into a single outside product for each module. This implies that product quality is only identified for products that see non-negligible sales shares, on average, across markets. In the base specification, I allocate any product to the outside product if its average non-zero sales share across store-month markets falls above the 60th percentile of the products in its respective module. Using this cutoff, products grouped in the outside product account for 15 percent of the store sales observed in the data.[66] I drop any markets that sell less than two non-outside products.

Finally, for identification purposes, I limit my attention to products that are sold in 5 or more of the remaining markets where I observe HMS households shopping and to modules that have some overlap between the product-store-month RMS store sales data and the HMS household purchase data.

Of the 1,071 non-random weight modules in the RMS data, 807 have 75% or more of their UPC container sizes measured in useful units. 666 of the remaining modules appear in five or more markets where HMS panelists shop (for the purposes of identifying store neighborhood income used to calculate the macro moments), after removing markets with outlier prices or insufficient non-outside products. Of this 666 modules, 530 have some store-products represented in the household purchase data, i.e., observations with which I can calculate the micro moments to identify the non-homotheticity parameters.

The cleaned data contains approximately 200,000 UPCs categorized into 22,655 products across 530 product modules. Almost three quarters of these products are purchased by households in the HMS data. The median numbers of products and UPCs per module are 28 and 119, respectively.

[65]The quality of random weight items, such as fruit, vegetables, and deli meats, varies over time as the produce loses its freshness and it is likely that stores set prices to reflect this. This potential inter-temporal correlation between their unobserved quality of random weight products and their prices would introduce biases in the price elasticities estimated below, so they are excluded from this analysis.

[66]Gandhi et al. (2013) highlight a selection problem associated with this treatment of low and zero sales shares. To gauge the magnitude of this problem, I test the robustness of my estimates to higher and lower selection criteria for the main model in Section 6.4.1 of the paper.

Table A.1: Summary Statistics for the Nielsen Data Used in Estimation

Panel A: Full Sample

| Data: | RMS (Store) | | | | | | | HMS (HH) |
|---|---|---|---|---|---|---|---|---|
| | Total | Count Per Module | | | Count Per Product | | | Total |
| | Count | Min | Median | Max | Min | Median | Max | Count |
| Modules | 1,071 | - | - | - | - | - | - | 1,060 |
| Products | 188,549 | 1 | 122 | 4,844 | - | - | - | 107,455 |
| UPCs | 768,639 | 1 | 220 | 32,554 | 1 | 2 | 3,000 | 362,143 |

Panel B: Estimation Sample

| Data: | RMS (Store) | | | | | | | HMS (HH) |
|---|---|---|---|---|---|---|---|---|
| | Total | Count Per Module | | | Count Per Product | | | Total |
| | Count | Min | Median | Max | Min | Median | Max | Count |
| Modules | 530 | - | - | - | - | - | - | 530 |
| Products | 22,655 | 2 | 28 | 627 | - | - | - | 15,615 |
| UPCs | 203,049 | 2 | 119 | 6,595 | 1 | 6 | 1,347 | 106,118 |

Notes: This table shows the distribution of UPCs across product and module categories in the raw Nielsen RMS store sales and HMS household purchase data as well as the samples used for estimation. A product is defined as the set of UPCs within a module with the same brand. The estimation sample does not include the "outside" product (into which 60 percent of products are allocated, in the base specification). In the raw data, the products grouped into the "outside" product are reported as individual products.

## A.2 Estimating the Empirical Distribution of Store Customers

Nielsen provides the county in which each Nielsen sample store is located. To obtain a more precise measure of each store's clientele, I estimate the income distribution in the vicinity of the store by taking a distance-weighted average of the income distributions observed in the Census tracts within 30km of the centroid of the modal residential zip code of Nielsen panelists that report shopping there.

The income distribution of each sample store $F(Y|s)$ is a generalized beta distribution fitted to the average binned income distribution in nearby tracts.[67] The number of households in each income bin for each store is calculated combining tract-level income from the 2010-2014 5-year American Community Survey (ACS) 1% sample and household-store-level trip data from the Nielsen HMS sample for the same period. Let $N_t(k)$ denote the number of households that the ACS reports in each of 16 income brackets $k$ residing in a Census tract $t$ and $N_t$ denote the total number of households in the ACS sample for tract $t$. I estimate the share of store $s$ customers in income bracket $k$ as the weighted average of the density of households in each income bracket in each Census tract in the vicinity of store $s$:

$$d_s(k) = \frac{\sum_{\{t|d_{st})\leq 30\text{km}\}} w_s(d_{st})N_t(k)}{\sum_{\{t|d_{st})\leq 30\text{km}\}} w_s(d_{st})N_t}$$

Tract weights, $w_s(d_{st})$, are a store type-specific function of distance from the centroid of the tract to the centroid of the store zip (estimated to be the modal residential zip code of the store's customers observed in the Nielsen HMS data). Specifically, the weight for tract $t$ whose centroid is a distance $d_{st}$ from the centroid of the zip code for store $s$ is:

$$w_s(d_{st}) = \frac{pop_t\hat{s}_s(d)}{\sum_{\{t|d_{st})\leq 30\}} pop_t\hat{s}_s(d)}$$

where $pop_t$ is the total population in each tract $t$, also from the 2010-2014 5-year ACS and $\hat{s}_s(d)$ is the estimated density of sales for store $s$ as a function of customer distance. The sales density for stores of each type (grocery and mass merchandise in low, medium, and high population density zip codes) is interpolated using the observed densities of the shopping trips observed in the Nielsen HMS data for years 2010 to 2014. These curves are shown for each store type in Figure A.1.

[67]Income bins are as defined in the ACS data. To fit the binned income distributions for each store to a generalized beta distribution, I assume the income of the first 15 bins is the midpoint of the bracket and the income of the top bracket is the mean income estimated assuming a Pareto distribution.

Figure A.1: Sales Density by Store Type

a. Grocery Store

b. Mass Merchandise Stores

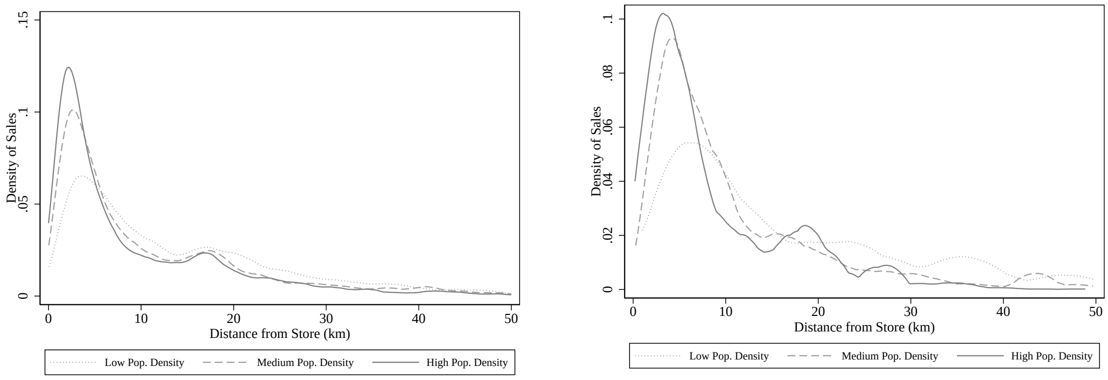


Notes: This table shows the density of sales at different distances from grocery stores (in a.) and mass merchandise stores (in b.) separately for stores in high, middle, and low population density zip codes.

## A.3 Representativeness of Nielsen Samples by Income

Figure A.2 compares the income distribution of the households in the Nielsen HMS sample to the national U.S. population. Figure A.3 compares the income distribution of the counties of stores in the Nielsen RMS sample to the counties of stores in the County Business Patterns dataset.

Figure A.2: Distribution of Household Income: Nielsen HMS (Homescan) versus ACS

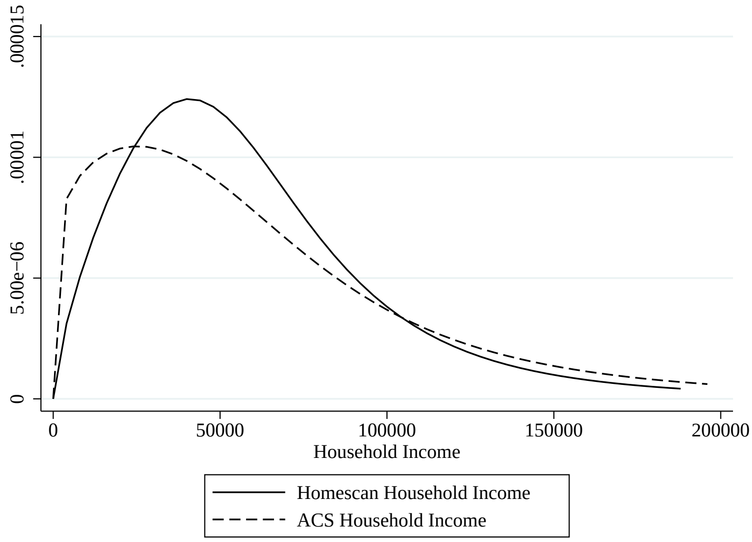


Notes: This figure compares the income distribution among Nielsen household panelists in 2012 with the national household income distribution in that year. The solid line depicts the fitted distribution of household income from the full 2012 Nielsen household (Homescan) sample; the dashed line depicts the fitted distribution of household income from the 2012 ACS single-year estimates.

Figure A.3: Distribution of Store Local Income: Nielsen RMS vs. County Business Patterns

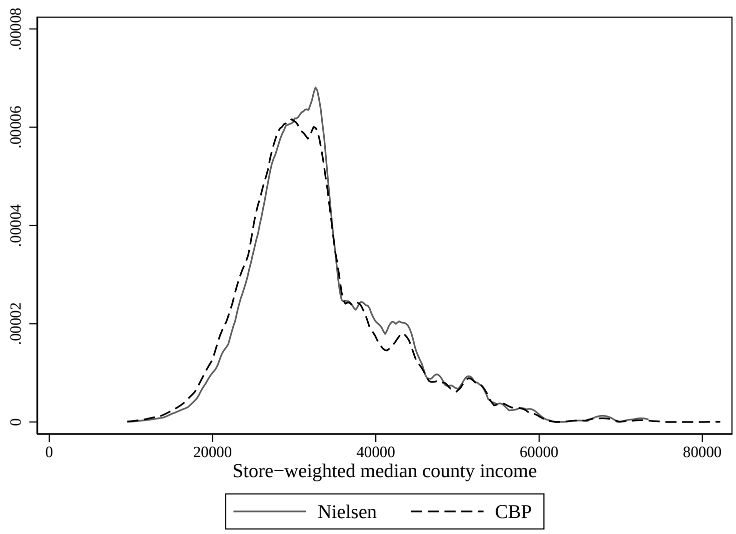


Notes: This figure compares the income distribution across the counties where Nielsen participating retailers are located with the income distribution across the counties where all grocery and non-durable stores are located. Each line depicts the distribution of median household income per county from the 2008-2012 ACS, weighted by the number of stores in the county. The solid line weights counties by the number of Nielsen RMS stores in the county, while the dashed line weights counties by the number of stores in the County Business Patterns, limiting attention to the following categories: 445110: Supermarket; 445120: Convenience stores; 446110: Pharmacies and Drug stores; 447110: Gasoline Stations with Convenience stores; 452910: Warehouse Clubs and Supercenters; and 452990: All Other General Merchandise Stores including Dollar stores.

## A.4 CBSA Statistics

This table shows the number of sample stores, population, and per capita income in each of the 125 CBSAs with 50 or more sample stores. The population and per capita income data five-year averages from the 2010-2014 ACS.

Table A.2: Sample Size, Population, and Income by CBSA

| CBSA Name | Store Count | Per Capita Income | Population |
|---|---|---|---|
| Akron, OH (AKR) | 76 | 27,823 | 703,017 |
| Albany-Schenectady-Troy, NY (ALB) | 127 | 32,069 | 875,567 |
| Albuquerque, NM (ABQ) | 102 | 26,144 | 899,137 |
| Allentown-Bethlehem-Easton, PA-NJ (ABE) | 94 | 29,397 | 826,260 |
| Asheville, NC (ASH) | 82 | 26,023 | 433,204 |
| Atlanta-Sandy Springs-Roswell, GA (ATL) | 620 | 28,880 | 5,455,053 |
| Augusta-Richmond County, GA-SC (AUG) | 97 | 23,905 | 575,669 |
| Austin-Round Rock, TX (AUS) | 136 | 32,035 | 1,835,016 |
| Bakersfield, CA (BAK) | 81 | 20,467 | 857,730 |
| Baltimore-Columbia-Towson, MD (BAL) | 305 | 35,613 | 2,753,396 |
| Baton Rouge, LA (BRI) | 96 | 26,639 | 814,805 |
| Birmingham-Hoover, AL (BIR) | 104 | 26,706 | 1,135,534 |
| Boise City, ID (BC) | 78 | 24,715 | 639,616 |
| Boston-Cambridge-Newton, MA-NH (BOS) | 562 | 39,572 | 4,650,876 |
| Bridgeport-Stamford-Norwalk, CT (BRI) | 87 | 49,688 | 934,215 |
| Buffalo-Cheektowaga-Niagara Falls, NY (BUF) | 163 | 28,171 | 1,135,667 |
| Canton-Massillon, OH (CAN) | 67 | 24,646 | 403,629 |
| Cape Coral-Fort Myers, FL (CC) | 70 | 27,578 | 647,554 |
| Charleston, WV (CHA) | 56 | 26,851 | 225,248 |
| Charleston-North Charleston, SC (CH) | 104 | 28,033 | 697,281 |
| Charlotte-Concord-Gastonia, NC-SC (CHA) | 449 | 28,403 | 2,298,915 |
| Chattanooga, TN-GA (CHA) | 99 | 25,315 | 537,397 |
| Chicago-Naperville-Elgin, IL-IN-WI (CHI) | 1082 | 31,488 | 9,516,448 |
| Cincinnati, OH-KY-IN (CIN) | 259 | 29,008 | 2,131,793 |
| Claremont-Lebanon, NH-VT (CLA) | 50 | 30,451 | 217,906 |
| Cleveland-Elyria, OH (CLE) | 245 | 28,499 | 2,067,490 |
| Colorado Springs, CO (CS) | 76 | 29,398 | 669,070 |
| Columbia, SC (COL) | 127 | 25,615 | 784,698 |
| Columbus, OH (CMH) | 218 | 29,145 | 1,948,188 |
| Dallas-Fort Worth-Arlington, TX (DAL) | 705 | 29,766 | 6,703,020 |
| Dayton, OH (DAY) | 102 | 26,345 | 801,259 |
| Deltona-Daytona Beach-Ormond Beach, FL (DAB) | 89 | 23,935 | 597,824 |
| Denver-Aurora-Lakewood, CO (DEN) | 310 | 34,173 | 2,651,392 |
| Des Moines-West Des Moines, IA (DM) | 123 | 31,342 | 590,741 |
| Detroit-Warren-Dearborn, MI (DET) | 507 | 28,182 | 4,292,647 |
| Durham-Chapel Hill, NC (DUR) | 77 | 30,945 | 525,050 |
| El Paso, TX (ELP) | 94 | 18,684 | 827,206 |
| Fayetteville, NC (PAY) | 62 | 22,647 | 374,036 |
| Fayetteville-Springdale-Rogers, AR-MO (FAY) | 62 | 25,291 | 483,396 |

*Continued on next page*

*Continued from previous page*

| CBSA Name | Store Count | Per Capita Income | Population |
|---|---|---|---|
| Flint, MI (FLI) | 82 | 22,536 | 418,654 |
| Fresno, CA (FRE) | 86 | 20,231 | 948,844 |
| Grand Rapids-Wyoming, MI (GRW) | 91 | 25,786 | 1,007,329 |
| Greensboro-High Point, NC (GHP) | 117 | 24,619 | 735,777 |
| Greenville-Anderson-Mauldin, SC (GRE) | 157 | 24,583 | 842,817 |
| Gulfport-Biloxi-Pascagoula, MS (GBP) | 52 | 23,006 | 378,972 |
| Harrisburg-Carlisle, PA (HAR) | 66 | 30,404 | 555,154 |
| Hartford-West Hartford-East Hartford, CT (HRT) | 116 | 35,991 | 1,215,159 |
| Hickory-Lenoir-Morganton, NC (HIC) | 62 | 21,385 | 363,936 |
| Houston-The Woodlands-Sugar Land, TX (HOU) | 690 | 29,594 | 6,204,141 |
| Huntington-Ashland, WV-KY-OH (HUN) | 51 | 23,326 | 364,514 |
| Indianapolis-Carmel-Anderson, IN (IND) | 195 | 27,778 | 1,931,182 |
| Jackson, MS (JAK) | 68 | 24,311 | 574,998 |
| Jacksonville, FL (JAC) | 228 | 27,950 | 1,380,995 |
| Kansas City, MO-KS (KC) | 152 | 30,101 | 2,040,869 |
| Kingsport-Bristol-Bristol, TN-VA (BRI) | 56 | 23,471 | 308,800 |
| Knoxville, TN (KNX) | 136 | 25,833 | 847,765 |
| Lafayette, LA (LAF) | 67 | 25,781 | 475,457 |
| Lakeland-Winter Haven, FL (LWH) | 66 | 21,157 | 617,323 |
| Lansing-East Lansing, MI (LAN) | 53 | 26,126 | 467,122 |
| Las Vegas-Henderson-Paradise, NV (LV) | 205 | 26,040 | 2,003,613 |
| Lexington-Fayette, KY (LEX) | 72 | 28,216 | 483,997 |
| Little Rock-North Little Rock-Conway, AR (LR) | 85 | 26,222 | 716,849 |
| Los Angeles-Long Beach-Anaheim, CA (LA) | 906 | 29,506 | 13,060,534 |
| Louisville/Jefferson County, KY-IN (LOU) | 182 | 27,488 | 1,253,305 |
| Madison, WI (MAD) | 82 | 32,778 | 620,368 |
| Manchester-Nashua, NH (MAN) | 77 | 34,767 | 402,776 |
| Memphis, TN-MS-AR (MEM) | 228 | 25,191 | 1,337,014 |
| Miami-Fort Lauderdale-West Palm Beach, FL (MIA) | 314 | 27,240 | 5,775,204 |
| Milwaukee-Waukesha-West Allis, WI (MIL) | 222 | 29,733 | 1,565,368 |
| Minneapolis-St. Paul-Bloomington, MN-WI (MIN) | 299 | 34,593 | 3,424,786 |
| Mobile, AL (MOB) | 69 | 23,009 | 414,045 |
| Myrtle Beach-Conway-North Myrtle Beach, SC-NC (MYR) | 86 | 24,709 | 396,187 |
| Nashville-Davidson–Murfreesboro–Franklin, TN (NAS) | 235 | 28,521 | 1,730,515 |
| New Haven-Milford, CT (NH) | 113 | 32,794 | 863,148 |
| New Orleans-Metairie, LA (NO) | 170 | 27,458 | 1,226,440 |
| New York-Newark-Jersey City, NY-NJ-PA (NYC) | 1697 | 36,078 | 19,865,045 |
| North Port-Sarasota-Bradenton, FL (NP) | 84 | 30,813 | 722,784 |
| Ogden-Clearfield, UT (OGD) | 56 | 24,890 | 614,521 |
| Oklahoma City, OK (OKC) | 94 | 26,994 | 1,297,998 |
| Omaha-Council Bluffs, NE-IA (OM) | 141 | 29,147 | 886,157 |
| Orlando-Kissimmee-Sanford, FL (ORL) | 271 | 24,876 | 2,226,835 |
| Oxnard-Thousand Oaks-Ventura, CA (OX) | 75 | 33,308 | 835,790 |
| Palm Bay-Melbourne-Titusville, FL (MEL) | 71 | 27,360 | 548,891 |
| Pensacola-Ferry Pass-Brent, FL (PEN) | 52 | 25,199 | 462,339 |
| Philadelphia-Camden-Wilmington, PA-NJ-DE-MD (PHL) | 802 | 32,850 | 6,015,336 |

*Continued on next page*

*Continued from previous page*

| CBSA Name | Store Count | Per Capita Income | Population |
|---|---|---|---|
| Phoenix-Mesa-Scottsdale, AZ (PHX) | 505 | 26,893 | 4,337,542 |
| Pittsburgh, PA (PIT) | 361 | 30,272 | 2,358,793 |
| Portland-South Portland, ME (POR) | 112 | 32,001 | 518,387 |
| Portland-Vancouver-Hillsboro, OR-WA (PVH) | 232 | 30,560 | 2,288,796 |
| Port St. Lucie, FL (PSL) | 54 | 27,481 | 433,646 |
| Providence-Warwick, RI-MA (PROV) | 257 | 30,218 | 1,604,317 |
| Raleigh, NC (RAL) | 210 | 31,468 | 1,189,579 |
| Richmond, VA (RIC) | 202 | 30,944 | 1,234,058 |
| Riverside-San Bernardino-Ontario, CA (RSB) | 338 | 22,571 | 4,345,485 |
| Roanoke, VA (ROA) | 52 | 27,505 | 310,934 |
| Rochester, NY (ROC) | 115 | 28,320 | 1,082,578 |
| Sacramento–Roseville–Arden-Arcade, CA (SAC) | 189 | 29,252 | 2,197,422 |
| St. Louis, MO-IL (STL) | 272 | 30,024 | 2,797,737 |
| Salisbury, MD-DE (SAL) | 90 | 27,353 | 381,868 |
| Salt Lake City, UT (SLC) | 93 | 26,516 | 1,123,643 |
| San Antonio-New Braunfels, TX (SA) | 233 | 25,298 | 2,239,222 |
| San Diego-Carlsbad, CA (SD) | 238 | 31,043 | 3,183,143 |
| San Francisco-Oakland-Hayward, CA (SF) | 365 | 42,540 | 4,466,251 |
| San Jose-Sunnyvale-Santa Clara, CA (SJ) | 139 | 42,176 | 1,898,457 |
| Savannah, GA (SAV) | 55 | 25,818 | 361,161 |
| Scranton–Wilkes-Barre–Hazleton, PA (SCR) | 70 | 25,304 | 562,644 |
| Seattle-Tacoma-Bellevue, WA (SEA) | 390 | 36,061 | 3,557,037 |
| Shreveport-Bossier City, LA (SHR) | 78 | 24,833 | 445,305 |
| Spartanburg, SC (SPA) | 59 | 22,055 | 317,057 |
| Spokane-Spokane Valley, WA (SPO) | 54 | 25,685 | 533,456 |
| Springfield, MA (SPR) | 81 | 27,179 | 626,775 |
| Springfield, MO (SGF) | 78 | 23,233 | 444,728 |
| Stockton-Lodi, CA (STL) | 52 | 22,642 | 701,050 |
| Syracuse, NY (SYR) | 98 | 27,741 | 662,236 |
| Tampa-St. Petersburg-Clearwater, FL (TAM) | 375 | 27,252 | 2,851,235 |
| Toledo, OH (TOL) | 97 | 25,312 | 608,847 |
| Tucson, AZ (TUC) | 150 | 25,524 | 993,144 |
| Tulsa, OK (TUL) | 109 | 26,635 | 954,055 |
| Virginia Beach-Norfolk-Newport News, VA-NC (VB) | 344 | 29,098 | 1,697,898 |
| Washington-Arlington-Alexandria, DC-VA-MD-WV (WAS) | 568 | 43,884 | 5,863,608 |
| Wichita, KS (WIC) | 66 | 25,848 | 636,095 |
| Wilmington, NC (WIL) | 54 | 28,435 | 263,804 |
| Winston-Salem, NC (WS) | 114 | 24,978 | 648,045 |
| Worcester, MA-CT (WOR) | 138 | 31,558 | 924,722 |
| Youngstown-Warren-Boardman, OH-PA (YOU) | 98 | 23,357 | 559,144 |

Notes: This table shows the number of Nielsen participating retailers, population, and per capita income in each of the 125 CBSAs with 50 or more participating retailers. The population and per capita income data five-year averages from the 2010-2014 ACS.

# B Stylized Facts Appendix

## B.1 Engel Curves in CBSAs with Different Per Capita Income

Figure A.4: Income-Specific Food Expenditure Shares Across Markets

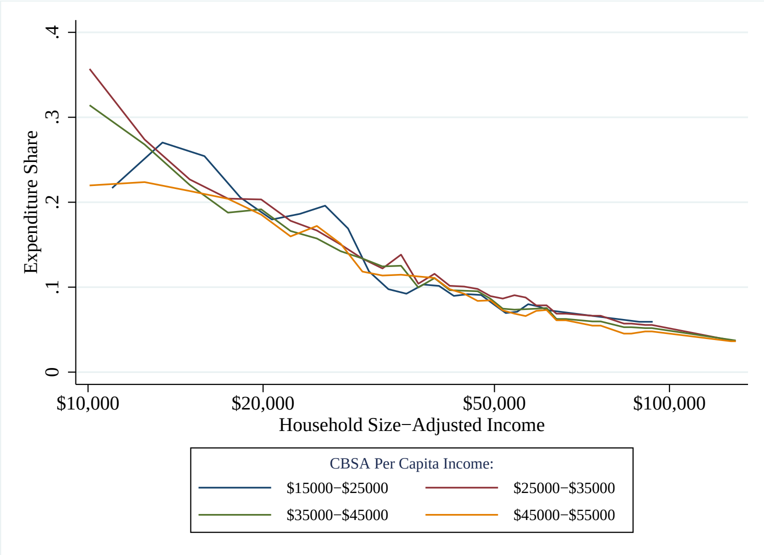


Note: This figure plots a kernel density of the Nielsen household expenditure share against size-adjusted income for panelists living in cities with different per capita incomes. The household expenditure share is calculated as the annual reported expenditures (for households reporting trips in all 12 months of the year) divided by their reported income.

## B.2 Store-Level Product Availability and Price Levels

Figure A.5: Difference in the Availability and Relative Price of High-Income and Low-Income Baskets Across Stores

a. Availability

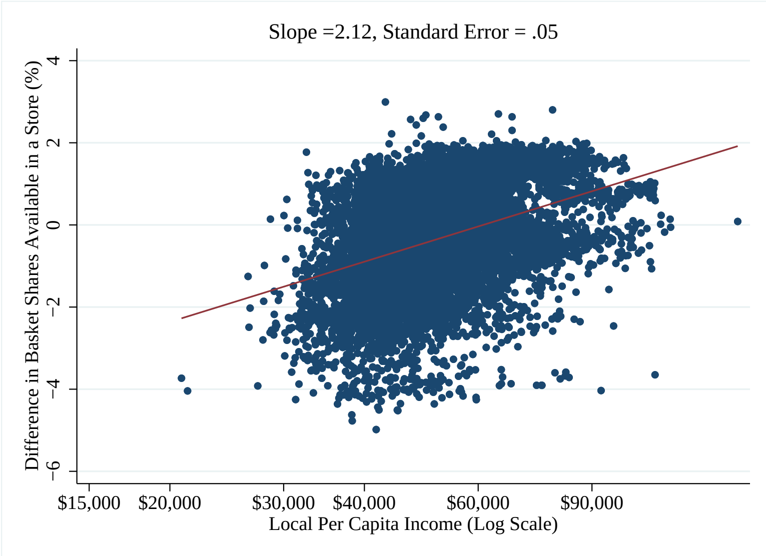


b. Relative Price

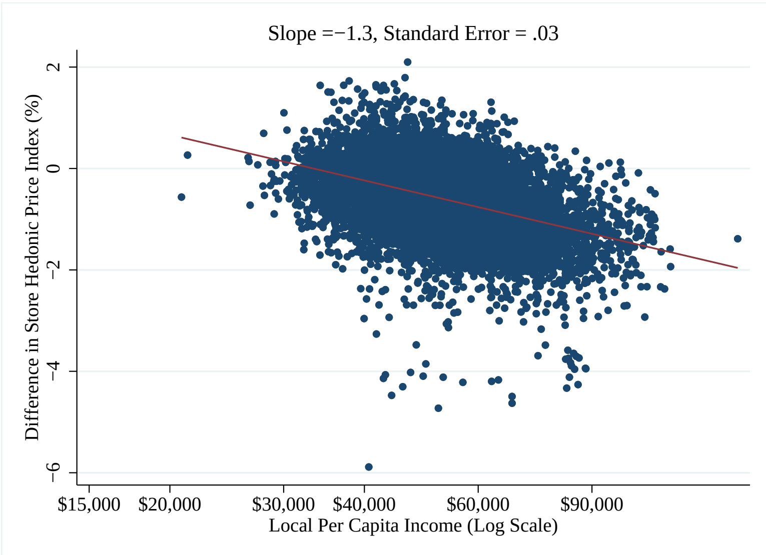


Notes: Figure a. plots store-level data for the difference between the expenditure shares of high-income Nielsen HMS panelists represented in the set of products sold by a store from the expenditure share of low-income Nielsen HMS panelists represented in that product set against local per capita income. The panelist expenditure shares are calculated for 2012. Figure b. plots store-level data for the difference between the average price level faced by consumers in the top income decile and the average price level faced by households in the bottom income decile against local per capita income. The price level in each store for a given income decile is calculated as the weighted average ratio between the price a product is sold for in a store relative to the price that product is sold at in the national sample where weights are defined as the value of the purchases of that product made by households in the respective income decile in the Nielsen household-level panel. Panelists are defined as high- (or low-) income if their size-adjusted income falls in the top (bottom) decile of panelist incomes. Local income is a distance-weighted average size-adjusted household income across tracts within 30km of the centroid of the modal residential zip of Nielsen panelist households that report shopping in the store. Household income adjusted for size using a square-root equivalence scale.

# C Model Appendix

## C.1 Non-Homotheticity Condition

Suppose that consumers select grocery consumption quantities, $\mathbb{Q} = \{\{q_{mg}\}_{g\in\mathbf{G_m}}\}_{m\in\mathbf{M}}$, and non-grocery expenditure, $Z$, by maximizing:

$$f(U_{iG}(\mathbb{Q},Z),Z) \quad \text{subject to} \quad \sum_{m\in\mathbf{M}}\sum_{g\in\mathbf{G_m}} p_{mg}q_{mg} + Z \leq Y_i,\ q_{mg} \geq 0\ \forall\ mg \in \mathbf{G} \tag{A.1}$$

I break this problem into two parts, first solving for the consumer's optimal grocery consumption quantities conditional on their non-grocery expenditure $Z$:

$$\begin{aligned} &\max_{\mathbb{Q},Z} U_{iG}(\mathbb{Q},Z) = \left\{\sum_{m\in M}\left[\sum_{g\in G_m} q_{mg}\exp(\gamma_m(Z)\beta_{mg} + \mu_m(Z)\varepsilon_{img})\right]^{\frac{\sigma-1}{\sigma}}\right\}^{\frac{\sigma}{\sigma-1}} \\ &\qquad\qquad \text{subject to} \quad \sum_{m\in\mathbf{M}}\sum_{g\in\mathbf{G_m}} p_{mg}q_{mg} \leq Y_i - Z,\ q_{mg} \geq 0\ \forall\ mg \in \mathbf{G} \end{aligned} \tag{A.2}$$

where $\gamma_m(Z) = (1 + \gamma_m \ln Z)$ and $\mu_m(Z) = \frac{1}{\alpha_m^0 + \alpha_m^1 \ln Z}$. Equations (8), and (9) define the optimal grocery bundle, $\mathbb{Q}^*(Z) = \left\{\left\{q^*_{mg}(Z)\right\}_{g\in\mathbf{G},}\right\}_{m\in\mathbf{M}}$ and can be summarized as follows:

$$q^*_{img}(Z) = \begin{cases} (Y_i - Z)\frac{[\tilde{p}_{img}]^{\sigma-1}}{P_i(Z)^{1-\sigma}}/p_{mg} & \text{if } g = \arg\max_{g\in\mathbf{G_m}} \tilde{p}_{img} \\ 0 & \text{otherwise} \end{cases}$$

where

$$P_i(Z) = \left[\left(\sum_{m\in\mathbf{M}}\max_{g\in\mathbf{G_m}} \tilde{p}_{img}\right)^{\sigma-1}\right]^{\frac{1}{1-\sigma}} \qquad \text{and} \qquad \tilde{p}_{img} = \frac{\exp(\gamma_m(Z)\beta_{mg} + \mu_m(Z)\varepsilon_{ig})}{p_{mg}}$$

Plugging this solution into $U_{iG}(\mathbb{Q},Z)$ yields the consumer's indirect utility from grocery consumption, conditional on their non-grocery expenditure:

$$\begin{aligned} \tilde{U}_{iG}(Z) &= U_{iG}(\mathbb{Q}^*(Z),Z) \\ &= \left\{\sum_{m\in M}\left[\left((Y_i - Z)\frac{(\tilde{p}_{img})^\sigma}{P_i(Z)^{1-\sigma}}\right)\mathbb{I}\left[g = \arg\max_{g\in\mathbf{G_m}} \tilde{p}_{img}\right]\right]^{\frac{\sigma-1}{\sigma}}\right\}^{\frac{\sigma}{\sigma-1}} \\ &= \frac{Y_i - Z}{P_i(Z)^{1-\sigma}}\left\{\sum_{m\in M}\left[\tilde{p}^\sigma_{img}\mathbb{I}\left[g = \arg\max_{g\in\mathbf{G_m}} \tilde{p}_{img}\right]\right]^{\frac{\sigma-1}{\sigma}}\right\}^{\frac{\sigma}{\sigma-1}} \\ &= \frac{Y_i - Z}{P_i(Z)^{1-\sigma}}\left\{\sum_{m\in M}\left(\max_{g\in\mathbf{G_m}} \tilde{p}_{img}\right)^{\sigma-1}\right\}^{\frac{\sigma}{\sigma-1}} \\ &= \frac{Y_i - Z}{P_i(Z)} \end{aligned} \tag{A.3}$$

We can now express problem (A.1) to be a choice over one variable, $Z$:

$$\max_{Z} f(\tilde{U}_{iG}(Z), Z) \tag{A.4}$$

The first order condition to the utility maximization problem defined in problem (A.4) with respect to $Z$ is:

$$f_1(\tilde{U}_{iG}(Z), Z)\frac{\partial \tilde{U}_{iG}(Z)}{\partial Z} + f_2(\tilde{U}_{iG}(Z), Z) = 0$$

Substituting the maximized grocery expenditure conditional on $Z$, $\tilde{U}_{iG}(Z)$, from equation (A.3) into this first order condition yields a function that implicitly defines the optimal non-grocery expenditure, $Z_i$, in terms of household income, $Y_i$, the consumer's idiosyncratic utility draws, $\varepsilon_i$, and model parameters:

$$Y_i \quad = \quad Z - \frac{P_i(Z)}{P_i'(Z)} + \frac{f_2(\tilde{U}_{iG}(Z), Z)}{f_1(\tilde{U}_{iG}(Z), Z)}\frac{P_i(Z)^2}{P_i'(Z)}$$

Taking the derivative of income with respect to non-grocery expenditure, $Z$, we can see that the non-grocery will be normal if the price vector and aggregate utility function are such that:

$$\frac{\partial}{\partial Z}\left[\frac{P_i(Z)}{P_i'(Z)} + \frac{f_2(\tilde{U}_{iG}(Z), Z)}{f_1(\tilde{U}_{iG}(Z), Z)}\frac{P_i(Z)^2}{P_i'(Z)}\right] < 1$$

It is computationally infeasible to show that this condition holds generally (there will be a different price index $P_i(Z)$ for each of universe of potential price vectors), but I can show that it holds in the data by simply demonstrating that non-grocery expenditures are increasing in household income. I annualize the observed grocery expenditure for each household and measure annual non-grocery expenditures as the difference between the mid-point of each household's reported income category and the household's annual grocery expenditures. After controlling for household demographics with dummies for household size, marital status, education and age of the male and female heads of household, race, and Hispanic origin, the elasticity of observed non-grocery expenditures, $Z_i$, with respect to household income, $Y_i$, is 1.19 with a standard error of 0.003.

Figure A.6 demonstrates that households earning higher incomes spend a smaller share of their income on grocery products. Within income groups, however, the average grocery expenditure share does not vary much across cities and, in particular, Table A.3 confirms that the average grocery share of an income group in a city does not vary systematically with city income.

Figure A.6: Income-Specific Grocery Expenditure Shares Across Markets

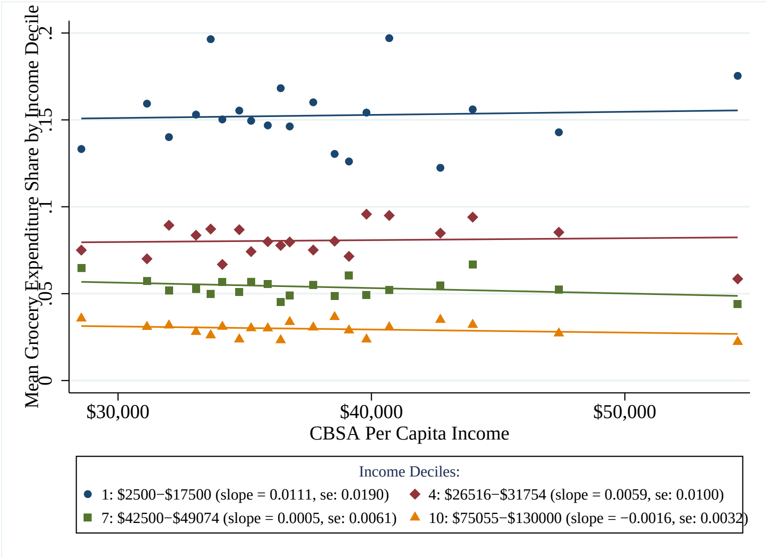


Note: Each point reflects the mean grocery expenditure share of households in each income decile that reside in households at each CBSAs at each vigntile of the CBSA per capita income distribution plotted against the mean CBSA per capital income of that vigntile. The household expenditure share is calculated as the annual reported expenditures on groceries (for households reporting trips in all 12 months of the year) divided by their reported income. For the purposes of visual clarity, only a representative sample of deciles are represented. The coefficient of variation of household grocery expenditure shares is 71 across all households in the sample, but drops to between 42 and 52 when you only consider households within each income decile. For the purposes of visual clarity, only a representative sample of deciles are represented.

Table A.3: Income-Specific Grocery Expenditure Shares Across Markets

| | Dependent Variable: Mean Grocery Expenditure Share of Households in Income Decile | | | | | | | | | |
|---|---|---|---|---|---|---|---|---|---|---|
| | [1] | [2] | [3] | [4] | [5] | [6] | [7] | [8] | [9] | [10] |
| Ln(CBSA PC Income) | 0.011 | -0.0034 | -0.0045 | 0.0059 | 0.0048 | -0.0051 | 0.00046 | 0.0075* | 0.0060 | -0.0016 |
| | (0.019) | (0.012) | (0.013) | (0.010) | (0.0095) | (0.0072) | (0.0061) | (0.0042) | (0.0056) | (0.0032) |
| Observations | 383 | 321 | 325 | 356 | 316 | 318 | 313 | 356 | 170 | 225 |

Notes: Standard errors in parentheses. * p<0.05, ** p<0.01, *** p<0.001. This table reports the correlation between the grocery expenditure share of Nielsen household panelists from each income decile and the per capita income of the CBSA where they reside. Observations are at the decile-by-CBSA level. The nth column reports regression for nth income decile.

## C.2 Derivations

### C.2.1 Within-Module Consumption Decision

Consumer $i$, spending $Z$ on the non-grocery items, chooses how to allocate expenditures between products within a module $m$ conditional on their expenditure in that module, $w_m$, to maximize

$$u_{im}(w_m, Z) = \sum_{g \in G_m} q_{mg} \exp(\gamma_m(Z)\beta_{mg} + \mu_m(Z)\varepsilon_{img})$$

subject to the module-level budget constraint, $\sum_{m \in \mathbf{M}} \sum_{g \in \mathbf{G_m}} p_{mg} q_{mg} \leq w_m$, and non-negativity constraints, $q_{mg} \geq 0 \;\; \forall mg \in \mathbf{G}$.

Recall that the additive log-logit functional form implies that consumers optimally purchase a positive quantity of only one product in a module. This product maximizes their marginal utility of expenditure in a module conditional on their non-grocery expenditure:[68]

$$g^*_{im}(Z) = \arg\max_{g \in \mathbf{G_m}} \frac{\exp(\gamma_m(Z)\beta_{mg} + \mu_m(Z)\varepsilon_{img})}{p_{mg}} \tag{A.5}$$

Since all of a consumer's module expenditure, $w_m$, is allocated to this optimal product, $g^*_{im}$, the consumer's optimal module bundle, $\mathbb{Q}^*_{im}(w_m, Z)$, can be written as:

$$\begin{aligned} \mathbb{Q}^*_{im}(w_m, Z) &= (q^*_{im1}(w_m, Z), \ldots, q^*_{imG_m}(w_m, Z)) \\ \text{where } q^*_{img}(w_m) &= \begin{cases} w_m/p_{mg} & \text{if } g = \arg\max_{g \in \mathbf{G_m}} \dfrac{\exp(\gamma_m(Z)\beta_{mg} + \mu_m(Z)\varepsilon_{img})}{p_{mg}} \\ 0 & \text{otherwise} \end{cases} \end{aligned} \tag{A.6}$$

That is, a consumer $i$ optimally consumes as much of their optimal product, $g^*_{im}(Z)$, as their module expenditure, $w_m$, will afford them and zero of any other product in the module.

### C.2.2 Across-Module Consumption Decision

Consumer $i$, spending $Z$ on non-grocery items, chooses how to allocate expenditures between modules by selecting $w_1, ..., w_M$ to maximize

$$U_i(w_1, \ldots, w_M) = \left\{ \sum_{m \in \mathbf{M}} [\tilde{u}_{im}(w_m, Z)]^{\rho_i} \right\}^{\frac{1}{\rho_i}} = \left\{ \sum_{m \in \mathbf{M}} \left[ w_m \max_{g \in \mathbf{G_m}} \frac{\exp(\gamma_m(Z)\beta_{mg} + \mu_m(Z)\varepsilon_{img})}{p_{mg}} \right]^{\frac{\sigma-1}{\sigma}} \right\}^{\frac{\sigma}{\sigma-1}}$$

subject to

$$\sum_{m \in \mathbf{M}} w_m \leq Y_i - Z$$

We simplify the expression for the target utility function by denoting consumer $i$'s marginal utility

[68]Note that the marginal utility of expenditure in a module and, therefore, the optimal product choice, $g^*_{im}$, depends on a consumer's non-grocery expenditure, $Z$, but is independent of their module expenditure, $w_m$.

from expenditure in module $m$ as the inverse of $A_{im}$:

$$\max_{g\in\mathbf{G_m}} \frac{\exp(\gamma_m(Z)\beta_{mg} + \mu_m(Z)\varepsilon_{img})}{p_{mg}} = \frac{1}{A_{im}} \tag{A.7}$$

The within-module allocation decision now simplifies to:

$$\mathbf{w_i^*}(Z) = (w_{i1}^*(Z), ..., w_{iM}^*(Z)) = \underset{\sum_{m\in\mathbf{M}} w_m \leq Y_i - Z}{\arg\max} \left\{ \sum_{m\in M} \left[\frac{w_m}{A_{im}}\right]^{\frac{\sigma-1}{\sigma}} \right\}^{\frac{\sigma}{\sigma-1}} \tag{A.8}$$

The utility function over module expenditures is concave in module expenditure for each module $m$. Therefore, there will be an interior solution to the maximization problem and it can be solved using the first order conditions with respect to expenditure in each module $m$. The first order condition for each module $m$ is:

$$\frac{\partial U_i(w_1, \ldots, w_M)}{\partial w_m} = \left\{ \sum_{m\in M} \left[\frac{w_m}{A_{im}}\right]^{\frac{\sigma-1}{\sigma}} \right\}^{\frac{1}{1-\sigma}} \frac{1}{A_{im}} \left[\frac{w_m}{A_{im}}\right]^{-\frac{1}{\sigma}} = \lambda$$

where $\lambda$ is the marginal utility of expenditure. This implies that the marginal utility of expenditure must be equal across modules. We use this equality across two modules, $m$ and $m'$, to solve for the optimal expenditure in module $m'$:

$$\begin{aligned}
\left\{ \sum_{m\in M} \left[\frac{w_m}{A_{im}}\right]^{\frac{\sigma-1}{\sigma}} \right\}^{\frac{1}{1-\sigma}} \frac{1}{A_{im'}} \left[\frac{w_{m'}}{A_{im'}}\right]^{-\frac{1}{\sigma}} &= \left\{ \sum_{m\in M} \left[\frac{w_m}{A_{im}}\right]^{\frac{\sigma-1}{\sigma}} \right\}^{\frac{1}{1-\sigma}} \frac{1}{A_{im}} \left[\frac{w_m}{A_{im}}\right]^{-\frac{1}{\sigma}} \\
\frac{1}{A_{im'}} \left[\frac{w_{m'}}{A_{im'}}\right]^{-\frac{1}{\sigma}} &= \frac{1}{A_{im}} \left[\frac{w_m}{A_{im}}\right]^{-\frac{1}{\sigma}} \\
w_{m'} &= w_m \left[\frac{A_{im'}}{A_{im}}\right]^{1-\sigma}
\end{aligned}$$

Imposing the budget constraint, $\sum_{m\in\mathbf{M}} w_{m'} = \sum_{m\in\mathbf{M}} w_m \leq Y_i - Z$, yields an expression for $w_m$ in terms of total expenditure, $Y_i - Z$, and an index of the $A_{im}$ terms:

$$\begin{aligned}
Y_i - Z &= \sum_{m'\in\mathbf{M}} w_{m'} \\
Y_i - Z &= \frac{w_m}{A_{im}^{1-\sigma}} \sum_{m'\in\mathbf{M}} [A_{im'}]^{1-\sigma} \\
w_m &= \frac{A_{im}^{1-\sigma}}{\sum_{m'\in\mathbf{M}} [A_{im'}]^{1-\sigma}} (Y_i - Z)
\end{aligned}$$

The solution to problem (A.8) is, therefore,

$$\mathbf{w_i^*}(Z) = (w_{i1}^*(Z), ..., w_{iM}^*(Z)) \quad \text{where} \quad w_{im}^* = \frac{A_{im}^{1-\sigma}}{P_i^{1-\sigma}}(Y_i - Z) \;\; \forall m \in \mathbf{M}$$

where $P_i(Z)$ is a CES price index over $A_{im}$ for all modules $m \in \mathbf{M}$ defined as:

$$P_i(Z) = \left[\sum_{m \in \mathbf{M}} A_{im}^{1-\sigma}\right]^{\frac{1}{1-\sigma}}$$

Substituting from equation (A.7) for $A_{img}$ yields consumer $i$'s optimal module expenditure vector, $\mathbf{w_i^*}(Z)$, as a function of total grocery expenditures, prices, and model parameters:

$$\mathbf{w_i^*}(\mathbf{Z}) = (w_{i1}^*(Z), ..., w_{iM}^*(Z)) \;\; \text{where} \;\; w_{im}^* = (Y_i - Z)\frac{\left[\max\limits_{g \in \mathbf{G_m}} \dfrac{\exp(\gamma_m(Z)\beta_{mg} + \mu_m(Z)\varepsilon_{img})}{p_{mg}}\right]^{\sigma-1}}{P_i(Z)^{1-\sigma}}$$

$$P_i(Z) = \left[\sum_{m \in \mathbf{M}} \left(\max_{g \in \mathbf{G_m}} \frac{\exp(\gamma_m(Z)\beta_{mg} + \mu_m(Z)\varepsilon_{img})}{p_{mg}}\right)^{\sigma-1}\right]^{\frac{1}{1-\sigma}}$$

## C.3 Connection to Nested CES Utility Function

In Section 4 of the paper, I model consumer demand assuming that a consumer $i$'s utility from grocery consumption, conditional on their non-grocery expenditure $Z$, is a CES aggregate over consumer-specific module-level utilities that are, in turn, additive in product-level log-logit utilities. This utility function is presented in equations (1), (2), and (3) and can be summarized as:

$$\begin{aligned} U_{iG}(\mathbb{Q}, Z) &= \left\{\sum_{m \in M} u_{im}\left(\mathbb{Q}_m, Z\right)^{\frac{\sigma-1}{\sigma}}\right\}^{\frac{\sigma}{\sigma-1}} \\ &= \left\{\sum_{m \in M}\left(\sum_{g \in \mathbf{G_m}} u_{img}(\mathbb{Q}_m, Z)\right)^{\frac{\sigma-1}{\sigma}}\right\}^{\frac{\sigma}{\sigma-1}} \\ &= \left\{\sum_{m \in M}\left(\sum_{g \in \mathbf{G_m}} q_{mg}\exp(\gamma_m(Z)\beta_{mg} + \mu_m(Z)\varepsilon_{img})\right)^{\frac{\sigma-1}{\sigma}}\right\}^{\frac{\sigma}{\sigma-1}} \\ &= \left\{\sum_{m \in M}\left(\sum_{g \in \mathbf{G_m}} q_{mg}\exp(\gamma_m(Z)\beta_{mg} + \frac{\varepsilon_{img}}{\sigma_m(Z)-1})\right)^{\frac{\sigma-1}{\sigma}}\right\}^{\frac{\sigma}{\sigma-1}} \end{aligned} \tag{A.9}$$

where $q_{mg}$ is the consumption quantity of each product $g$ in module $m$; $\beta_{mg}$ is the quality of product $g$ in module $m$; $\varepsilon_{img}$ is the idiosyncratic utility of consumer $i$ from product $g$ in module $m$; $\gamma_m(Z)$ and $\mu_m(Z) = \frac{1}{\sigma_m(Z)-1} > 0$ are weights that govern the extent to which consumers with non-grocery expenditure $Z$ care about product quality and their idiosyncratic utility draws; and $\sigma_m(Z)$ is the elasticity of substitution between products in the same module $m$ and $\sigma > 1$ is the elasticity of substitution between products in different modules for a consumer with non-grocery expenditure $Z$.

Consider the utility of the representative agent for consumers with non-grocery expenditure $Z$. This agent's utility function from grocery consumption is defined in equation (12) in Section 5.1 as follows:

$$U_G^{CES}(\mathbb{Q}, Z) = \left\{ \sum_{m \in M} \left[ \sum_{g \in \mathbf{G_m}} [q_{mg} \exp(\beta_{mg}\gamma_m(Z))]^{\frac{\sigma_m(Z)-1}{\sigma_m(Z)}} \right]^{\left(\frac{\sigma_m(Z)}{\sigma_m(Z)-1}\right)\left(\frac{\sigma-1}{\sigma}\right)} \right\}^{\frac{\sigma}{\sigma-1}}, \tag{A.10}$$

where $q_{mg}$, $\beta_{mg}$, $\gamma_m(Z)$, $\sigma_m(Z)$, and $\sigma > 1$ take the same definition as in equation (A.9) above. $\sigma$ is the elasticity of substitution between products in different modules for a consumer with non-grocery expenditure $Z$.

Suppose that this representative consumer with the nested-CES utility function $U_G^{CES}(\mathbb{Q}, Z)$ defined in equation (A.10) faces the same prices $\mathbb{P}$ and has the same non-grocery expenditure $Z$ as a group of "idiosyncratic" consumers with the CES-nested log-logit utility $U_{iG}(\mathbb{Q}, Z)$ defined in equation (A.9). A simple extension of Anderson et al. (1987) shows that the representative consumer and the group of "idiosyncratic" consumers will allocate expenditures across products within modules and across modules identically.

First consider the within-module expenditure allocations. Denote the share of module $m$ expenditures that the representative consumer allocates to product $g$ as $s_{mg|m}^{CES}(Z)$ and the share of total grocery expenditures the representative consumer allocates to module $m$ as $s_m^{CES}(Z)$. This share is equal to

$$s_{mg|m}^{CES}(Z) = \left[ \frac{\frac{p_{mg}}{\exp(\beta_{mg}\gamma_m(Z))}}{P_m^{CES}(Z, \mathbb{P}_m)} \right]^{1-\sigma_m(Z)}$$

where $P_m^{CES}(Z, \mathbb{P}_m)$ is a module-level CES price index. The relative log share that the representative consumer optimally allocates to product $g$ in module $m$ relative to some other product $\bar{g}$ in the same module is, therefore,

$$\ln s_{mg|m}^{CES}(Z) - \ln s_{m\bar{g}|m}^{CES}(Z) = (1 - \sigma_m(Z)) \left( (\ln p_{mg} - \ln p_{m\bar{g}}) - (\beta_{mg} - \beta_{m\bar{g}})\gamma_m(Z) \right) \tag{A.11}$$

The expected relative module expenditure share of a group of "idiosyncratic" consumers with non-grocery expenditure $Z$ facing the same prices $p_{mg}$ and $p_{m\bar{g}}$ is derived in Appendix (D.1.1) as:

$$\mathbb{E}_\varepsilon \left[ \ln(s_{img|m}(Z, \mathbb{P}_m)) - \ln(s_{im\bar{g}|m}(Z, \mathbb{P}_m)) \right] = (\sigma_m(Z) - 1) \left[ (\beta_{mg} - \beta_{m\bar{g}})\gamma_m(Z) - (\ln p_{mg} - \ln p_{m\bar{g}}) \right] \tag{A.12}$$

where I have substituted $\sigma_m(Z)$ and $\gamma_m(Z)$ for their log-linear parametric forms $\left(1 + \alpha_m^0 + \alpha_m^1 \ln Z\right)$ and $(1 + \gamma_m \ln Z)$, respectively. We can multiply both terms of the right-hand side of (A.12) to show that it is equivalent to the right-hand side of equation (A.11):

$$\begin{aligned}\mathbb{E}_\varepsilon\left[\ln(s_{img|m}(Z,\mathbb{P}_m))-\ln(s_{im\bar{g}|m}(Z,\mathbb{P}_m))\right] &= (\sigma_m(Z)-1)\left[(\beta_{mg}-\beta_{m\bar{g}})\gamma_m(Z)-(\ln p_{mg}-\ln p_{m\bar{g}})\right]\\ &= (1-\sigma_m(Z))\left((\ln p_{mg}-\ln p_{m\bar{g}})-(\beta_{mg}-\beta_{m\bar{g}})\gamma_m(Z)\right)\\ &= \ln s^{CES}_{mg|m}(Z)-\ln s^{CES}_{m\bar{g}|m}(Z)\end{aligned}$$

whereby showing that the representative consumer allocates expenditures across products in the same module identically to a group of the "idiosyncratic" consumers.

Now consider the between-module expenditure allocations. Denote the share of total grocery expenditures the representative consumer allocates to module $m$ as $s^{CES}_m(Z)$. The relative log share that the representative consumer optimally allocates to module $m$ relative to some other module $\bar{m}$ is

(A.13) $$\ln s^{CES}_m(Z)-\ln s^{CES}_{\overline{m}}(Z)=(1-\sigma)\left(\ln\left(P^{CES}_m(Z,\mathbb{P}_m)\right)-\ln\left(P^{CES}_{\bar{m}}(Z,\mathbb{P}_{\bar{m}})\right)\right)$$

where $P^{CES}_m(Z,\mathbb{P}_m)$ is a module-level CES price index defined as:

(A.14) $$P^{CES}_m(Z,\mathbb{P}_m)=\left[\sum_{g\in\mathbf{G_m}}\left(\frac{p_{mg}}{\exp(\beta_{mg}\gamma_m(Z))}\right)^{(1-\sigma_m(Z))}\right]^{\frac{1}{(1-\sigma_m(Z))}}$$

The expected relative module expenditure share of a group of "idiosyncratic" consumers with non-grocery expenditure $Z$ facing the same sets of prices $\mathbb{P}_m$ and $\mathbb{P}_{\bar{m}}$ faced by the representative consumer is derived in Appendix (D.1.2) as:

(A.15) $$\mathbb{E}_\varepsilon\left[\ln s_{im}(Z,\mathbb{P})-\ln s_{i\bar{m}}(Z,\mathbb{P})\right] = (\sigma-1)\left[\ln V_m(Z,\mathbb{P}_m)-\ln V_{\bar{m}}(Z,\mathbb{P}_{\bar{\mathbf{m}}})\right]$$

where $V_m(Z,\mathbb{P}_m)$ is a CES-style index over price-adjusted product qualities:

(A.16) $$V_m(Z,\mathbb{P}_m)=\left[\sum_{g\in\mathbf{G_m}}\left(\frac{\exp(\beta_{mg}\gamma_m(Z))}{p_{mg}}\right)^{(\sigma_m(Z)-1)}\right]^{\frac{1}{(\sigma_m(Z)-1)}}$$

To see that the right-hand sides of equations (A.13) and (A.15) are identical first note that we can re-write the equation (A.15) as

$$\begin{aligned}\mathbb{E}_\varepsilon\left[\ln s_{im}(Z,\mathbb{P})-\ln s_{i\bar{m}}(Z,\mathbb{P})\right] &= (1-\sigma)\left[-\ln V_m(Z,\mathbb{P}_m)+\ln V_{\bar{m}}(Z,\mathbb{P}_{\bar{\mathbf{m}}})\right]\\ &= (1-\sigma)\left[\ln\left([V_m(Z,\mathbb{P}_m)]^{-1}\right)-\ln\left([V_{\bar{m}}(Z,\mathbb{P}_{\bar{\mathbf{m}}})]^{-1}\right)\right]\end{aligned}$$

In fact, the right-hand sides of equations (A.13) and (A.15) will be identical as long as the quality-adjusted price levels defined in equation (A.14) are equal to the inverse of the price-adjusted quality levels defined in equation (A.16), *i.e.*, $P^{CES}_m(Z,\mathbb{P}_m)=[V_m(Z,\mathbb{P}_m)]^{-1}$. We can see this is the case below:

$$\begin{aligned}
P_m^{CES}(Z,\mathbb{P}_m) &= \left[\sum_{g\in\mathbf{G_m}}\left(\frac{p_{mg}}{\exp(\beta_{mg}\gamma_m(Z))}\right)^{(1-\sigma_m(Z))}\right]^{\frac{1}{(1-\sigma_m(Z))}} \\
&= \left[\sum_{g\in\mathbf{G_m}}\left(\frac{\exp(\beta_{mg}\gamma_m(Z))}{p_{mg}}\right)^{(\sigma_m(Z)-1)}\right]^{\frac{1}{(1-\sigma_m(Z))}} \\
&= \left\{\left[\sum_{g\in\mathbf{G_m}}\left(\frac{\exp(\beta_{mg}\gamma_m(Z))}{p_{mg}}\right)^{(\sigma_m(Z)-1)}\right]^{\frac{1}{(\sigma_m(Z)-1)}}\right\}^{-1} \\
&= [V_m(z,\mathbb{P}_m)]^{-1}
\end{aligned}$$

The representative consumer therefore allocates expenditures across modules in identical proportions to a group of the "idiosyncratic" consumers.

The algebra above has shown that the CES-nested log-logit utility function yields identical relative expenditure share equations, both across and within modules, to the nested-CES utility function assumed for the representative agent. In particular, note that the model parameters play identical roles in the nested-CES and CES-nested log-logit expenditure share equations, so the parameter estimates identified using moments based on these equations can be used as direct inputs into the nested-CES price indexes that form the basis for the main results presented above.

# D Empirical Strategy Appendix

## D.1 Derivations of Expenditure Share for Moment Equations

### D.1.1 Within-Module Market Expenditure Shares

Equation (A.6) states that:

$$\mathbb{Q}^*_{im}(w_m, Z) = (q^*_{im1}(w_m, Z), \ldots, q^*_{imG_m}(w_m, Z)) \text{ where } q^*_{img}(w_m, Z) = \begin{cases} w_m/p_{mg} & \text{if } g = \arg\max_{g \in \mathbf{G_m}} \tilde{p}_{img} \\ 0 & \text{otherwise} \end{cases}$$

where $\tilde{p}_{img} = \frac{\exp(\gamma_m(Z)\beta_{mg} + \mu_m(Z)\varepsilon_{ig})}{p_{mg}}$. If we rewrite consumer $i$'s optimal consumption quantity using an indicator function to identify which product is selected by the consumer, consumer $i$'s optimal consumption quantity of product $g$ in module $m$ is:

$$q^*_{img}(w_m, Z) = \frac{w_m}{p_{mg}} \mathbb{I}\left[g = \arg\max_{g \in \mathbf{G_m}} \tilde{p}_{img}\right]$$

We can use this definition to derive consumer $i$'s expenditure on product $g$ in module $m$:

$$w_{img}(w_m) = p_{mg} q^*_{img}(w_m, Z) = w_m \mathbb{I}\left[g = \arg\max_{g \in \mathbf{G_m}} \tilde{p}_{img}\right]$$

Dividing through by $w_m$ yields the consumer's expenditure share on product $g$ in module $m$, conditional on their non-grocery expenditure $Z$ and the vector of module prices they face, $\mathbb{P}_m$:

$$s_{img|m}(Z, \mathbb{P}_m) = \mathbb{I}\left[g = \arg\max_{g \in \mathbf{G_m}} \tilde{p}_{img}\right]$$

The expected value of this expenditure share is derived by integrating over the idiosyncratic utilities in module $m$, $\varepsilon_{\mathbf{im}}$:

$$\begin{aligned}
\mathbb{E}_\varepsilon[s_{img|m}(Z, \mathbb{P}_m)] &= \mathbb{E}_\varepsilon\left[\mathbb{I}\left[g = \arg\max_{g \in \mathbf{G_m}} \tilde{p}_{img}\right]\right] \\
&= Pr\left[\tilde{p}_{img} \geq \tilde{p}_{img'}, \ \ \forall g' \in \mathbf{G_m}\right] \\
&= Pr\left[\varepsilon_{img} - \varepsilon_{img'} \geq \frac{\gamma_m(Z)(\beta_{mg} - \beta_{mg'}) - (\ln p_{mg} - \ln p_{mg'})}{\mu_m(Z)}, \ \ \forall g' \in \mathbf{G_m}\right] \\
&= \frac{\tilde{p}_{img}}{\sum\limits_{g' \in \mathbf{G_m}} \tilde{p}_{img'}}
\end{aligned}$$

The final equality holds because the idiosyncratic utilities, $\varepsilon_{im}$, are iid draws from a type I extreme value distribution. Imposing the parametric forms for $\gamma_m(Z) = (1 + \gamma_m \ln Z)$ and $\mu_m(Z) = (\alpha^0_m + \alpha^1_m \ln Z)^{-1}$ from equations (4) and (5), respectively, ensures that the consumer's expected expenditure share is common with other consumers with the same income that face the same product prices:

$$\mathbb{E}_{\varepsilon}[s_{img|m}(Z,\mathbb{P}_m)] = \frac{\exp[(\alpha_m^0 + \alpha_m^1 \ln Z)((1+\gamma_m \ln Z)\beta_{mg} - \ln p_{mg})]}{\sum_{g' \in \mathbf{G_m}} \left(\exp[(\alpha_m^0 + \alpha_m^1 \ln Z)((1+\gamma_m \ln Z)\beta_{mg'} - \ln p_{mg'})]\right)}$$

I interpret the expected expenditure share function derived above as the expected share of expenditure that a group of households with the same non-grocery expenditure, $Z$, facing identical prices for products in module $m$ spend on product $g$. If the group of households is in the same market, then this expected expenditure share will be the income-specific market share of product $g$ in module $m$, which I denote by $s_{mg|m}(Z,\mathbb{P}_m)$. $s_{mg|m}(Z,\mathbb{P}_m)$ is the share of expenditure that a group of households with the non-grocery expenditure, $Z$, and facing a common vector of module prices, $\mathbb{P}_m$:

$$s_{mg|m}(Z,\mathbb{P}_m) = \mathbb{E}_{\varepsilon}[s_{img|m}(Z,\mathbb{P}_m)] = \frac{\exp[(\alpha_m^0 + \alpha_m^1 \ln Z)(\beta_{mg}(1+\gamma_m \ln Z) - \ln p_{mg})]}{\sum_{g' \in \mathbf{G_m}} \left(\exp[(\alpha_m^0 + \alpha_m^1 \ln Z)(\beta_{mg'}(1+\gamma_m \ln Z) - \ln p_{mg'})]\right)}$$

Dividing this market share for product $g$ in module $m$ by the market share for a fixed product $\bar{g}_m$ in the same module $m$ results in a relative market share that depends only on model parameters, consumer income, and the prices of product $g$ and $\bar{g}_m$:[69]

$$\frac{s_{mg|m}(Z,\mathbb{P}_m)}{s_{m\bar{g}|m}(Z,\mathbb{P}_m)} = \frac{\exp[(\alpha_m^0 + \alpha_m^1 \ln Z)(\beta_{mg}(1+\gamma_m \ln Z) - \ln p_{mg})]}{\exp[(\alpha_m^0 + \alpha_m^1 \ln Z)(\beta_{m\bar{g}}(1+\gamma_m \ln Z) - \ln p_{m\bar{g}})]}$$

I linearize the relative expenditure share equation by taking the log of both sides:

(A.17)

$$\ln(s_{mg|m}(Z,\mathbb{P}_m)) - \ln(s_{m\bar{g}|m}(Z,\mathbb{P}_m)) = (\alpha_m^0 + \alpha_m^1 \ln Z)\left[(\beta_{mg} - \beta_{m\bar{g}})(1+\gamma_m \ln Z) - (\ln p_{mg} - \ln p_{m\bar{g}})\right]$$

Equation (A.17) defines the expected within-module expenditure share of a set of households with non-grocery expenditure $Z$ facing prices $p_{mg}$ and $p_{m\bar{g}_m}$ on product $g$ in module $m$ relative to product $\bar{g}_m$ in the same module $m$ in terms of parameters $\alpha_m$, $\gamma_m$, and $(\beta_{mg} - \beta_{m\bar{g}_m})$. This equation is used to calculate moments for each product $g \neq \bar{g}_m$ in each module $m$, that are in turn used to estimate all of the $\alpha_m$ and $\gamma_m$ parameters, as well as each $\beta_{mg}$ parameter relative to $\beta_{m\bar{g}_m}$, *i.e.* $\{\beta_{mg} - \beta_{m\bar{g}_m}\}_{g \in \mathbf{G_m}}$.

### D.1.2 Between-Module Relative Market Expenditure Shares

I now want to generate a similar estimation equation that can be used to identify $\sigma$ and $\{\beta_{\bar{g}_m}\}_{g \in \mathbf{G_m}}$ using data on module-level income-specific market shares. Equations (8) and (9) together characterize the optimal cross-module expenditure allocation for consumer $i$ conditional on this consumer's idiosyncratic

[69]The utility function assumes weak separability between modules and the independence of irrelevant alternatives (IIA) property both across modules and across products with the same quality parameter. Although neither of these are realistic characteristics of consumer behavior, they are useful for the purposes of estimation as they imply that relative market expenditure shares can be derived as functions of observed variables, such as household income, expenditures, and transaction prices.

utility draws for each product in each module. These equations are:

$$\mathbf{w_i^*}(Z,\mathbb{P}) = (w^*_{i1}(Z,\mathbb{P}), ..., w^*_{iM}(Z,\mathbb{P})) \text{ where } w^*_{im} = (Y_i - Z)\frac{\left[\max\limits_{g\in\mathbf{G_m}} \tilde{p}_{img}\right]^{\sigma-1}}{P_i(Z)^{1-\sigma}}$$

$$P_i(Z,\mathbb{P}) = \left[\sum_{m\in\mathbf{M}} \left(\max_{g\in\mathbf{G_m}} \tilde{p}_{img}\right)^{\sigma-1}\right]^{\frac{1}{1-\sigma}}$$

where $\tilde{p}_{img} = \frac{\exp(\gamma_m(Z)\beta_{mg}+\mu_m(Z)\varepsilon_{img})}{p_{mg}}$. Dividing through by total grocery expenditure, $(Y_i - Z)$, I generate consumer $i$'s optimal module $m$ expenditure share, conditional on their non-grocery expenditure $Z$ and the vector of prices they face, $\mathbb{P}$:

$$s_{im}(Z,\mathbb{P}) = \frac{w^*_{im}(Z)}{Y_i - Z} = \frac{\left[\max\limits_{g\in\mathbf{G_m}} \tilde{p}_{img}\right]^{\sigma-1}}{P_i^{1-\sigma}}$$

When deriving the within-module relative market share, equation (A.17) above, I take the expectation of the consumer's expected product expenditure share over the idiosyncratic errors, $\mathbb{E}_\varepsilon[s_{img|m}(Z,\mathbb{P}_m)]$, to derive an expression for the market share of each product. I then divide these market shares by the market share of a module specific base product and take logs to linearize the equation. I change the order of this procedure when deriving the between-module relative market share equation, *i.e.* difference and take the log of the individual's expenditure shares before taking the expectation of these terms over the idiosyncratic errors. The reason for this reordering is that the consumer's module expenditure shares include a term, $P_i$, that depends non-linearly on all of the consumer's idiosyncratic utility draws. This term is common to all of the consumer's module shares, and thus drops out of the consumer's relative module expenditure shares, so that these relative shares are functions of the consumer's idiosyncratic utility draws in the two relevant modules. The log of this relative module expenditure share term is additive in terms that depend on the consumer's idiosyncratic utility draws in only one module at a time; that is, a term that depends on the consumer's idiosyncratic utility draws in module $m$ and a term that depends on the consumer's idiosyncratic utility draws in the base module $\bar{m}$. This makes the expectation of the consumer's log expenditure share in module $m$ relative to module $\bar{m}$ easier to derive than the expectation of the consumer's expenditure share for a single module $m$.[70]

[70]The order of the expectation, differencing, and log operations does not make a difference to the relative market share equation in the within-module case, that is:

$$\begin{aligned}\ln(s_{mg|m}(Z,\mathbb{P}_m)) - \ln(s_{m\bar{g}|m}(Z,\mathbb{P}_m)) &= \ln\left[\mathbb{E}_\varepsilon[s_{img|m}(Z,\mathbb{P}_m)]/\mathbb{E}_\varepsilon[s_{im\bar{g}|\bar{m}}(Z,\mathbb{P}_m)]\right)] \\ &= \mathbb{E}_\varepsilon\left[\ln(s_{img|m}(Z,\mathbb{P}_m)) - \ln(s_{im\bar{g}|m}(Z,\mathbb{P}_m))\right] \\ &= (\alpha^0_m + \alpha^1_m \ln Z)\left[(\beta_{mg} - \beta_{m\bar{g}})(1+\gamma_m \ln Z) - (\ln p_{mg} - \ln p_{m\bar{g}})\right]\end{aligned}$$

I derive the expression for the $Z$-specific market share of product $g$, $s_{mg|m}(Z,\mathbb{P}_m) = \mathbb{E}_\varepsilon[s_{img|m}(Z,\mathbb{P}_m)]$, before taking logs and differencing to generate the estimation equation (A.17), as it demonstrates the relationship between the term on the left-hand side of this equation, $\ln(s_{mg|m}(Z,\mathbb{P}_m)) - \ln(s_{m\bar{g}|m}(Z,\mathbb{P}_m))$, and its value in the data: the difference between the log of the expenditure consumers spending $Z$ on non-grocery items in a given market on product $g$ relative to the log of their expenditure on the base product $\bar{g}$ or, more succinctly, the log difference between the $Z$-specific market shares on products $g$ and $\bar{g}$.

I now generate the relative module market shares. As discussed above, I first divide consumer $i$'s module expenditure share, $s_{im}(Z,\mathbb{P})$, by his/her expenditure share in some fixed base module $\bar{m}$:

$$\frac{s_{im}(Z,\mathbb{P})}{s_{i\bar{m}}(Z,\mathbb{P})} = \frac{\left[\max_{g\in\mathbf{G_m}} \tilde{p}_{img}\right]^{\sigma-1}}{\left[\max_{g\in\mathbf{G_{\bar{m}}}} \tilde{p}_{i\bar{m}g}\right]^{\sigma-1}}$$

Since $P_i$ does not vary across modules for a given consumer $i$, it drops out of the relative module expenditure share expression. I take the log of this relative share expression to linearize the equation:

$$\ln s_{im}(Z,\mathbb{P}) - \ln s_{i\bar{m}}(Z,\mathbb{P}) = (\sigma-1)\ln\left(\max_{g\in\mathbf{G_m}} \tilde{p}_{img}\right) - (\sigma-1)\ln\left(\max_{g\in\mathbf{G_{\bar{m}}}} \tilde{p}_{i\bar{m}g}\right),$$

This equation is a linear function of two terms, the first of which depends on the consumer's idiosyncratic utility draws in only module $m$ and the second of which depends on the consumer's idiosyncratic utility draws in only module $\bar{m}$. The expectation of the log difference between the consumer's module expenditure shares can be split into the difference between two expected values:

(A.18)

$$\mathbb{E}_\varepsilon\left[\ln s_{im}(Z,\mathbb{P}) - \ln s_{i\bar{m}}(Z,\mathbb{P})\right] = (\sigma-1)\left\{\mathbb{E}_\varepsilon\left[\ln\left(\max_{g\in\mathbf{G_m}} \tilde{p}_{img}\right)\right] - \mathbb{E}_\varepsilon\left[\ln\left(\max_{g\in\mathbf{G_{\bar{m}}}} \tilde{p}_{i\bar{m}g}\right)\right]\right\}$$

Consider the two expectation terms in equation (A.18). Both take the same form, and thus I only solve for the first expectation:

$$\mathbb{E}_\varepsilon\left[\ln\left(\max_{g\in\mathbf{G_m}} \tilde{p}_{img}\right)\right] \tag{A.19}$$

The expectation term defined in equation (A.19) is the expected value of the log of a maximum. Since the log is a monotonically increasing function, we can switch the order of the log and maximum functions inside the expectation and linearize to yield:

$$\begin{aligned}
\mathbb{E}_\varepsilon\left[\ln\left(\max_{g\in\mathbf{G_m}} \tilde{p}_{img}\right)\right] &= \mathbb{E}_\varepsilon\left[\ln\left(\max_{g\in\mathbf{G_m}} \frac{\exp(\gamma_m(Z)\beta_{mg} + \mu_m(Z)\varepsilon_{img})}{p_{mg}}\right)\right] \\
&= \mathbb{E}_\varepsilon\left[\max_{g\in\mathbf{G_m}} \ln\left(\frac{\exp(\gamma_m(Z)\beta_{mg} + \mu_m(Z)\varepsilon_{img})}{p_{mg}}\right)\right] \\
&= \mathbb{E}_\varepsilon\left[\max_{g\in\mathbf{G_m}} \gamma_m(Z)\beta_{mg} - \ln p_{mg} + \mu_m(Z)\varepsilon_{img}\right] \\
&= \mu_m(Z)\mathbb{E}_\varepsilon\left[\max_{g\in\mathbf{G_m}} (\gamma_m(Z)\beta_{mg} - \ln p_{mg})/\mu_m(Z) + \varepsilon_{img}\right] \qquad \text{(A.20)}
\end{aligned}$$

De Palma and Kilani (2007) show that, for an additive random utility model with $u_i = \nu_i + \varepsilon_i$, $i =$

$1, \ldots, n$ and $\varepsilon_i \overset{\text{iid}}{\sim} F(x)$ a continuous CDF with finite expectation, the expected maximum utility is:

$$\mathbb{E}_\varepsilon[\max_i \nu_i + \varepsilon_i] = \int_{-\infty}^{\infty} z d\phi(z) \text{ where } \phi(z) = Pr[\max_k \nu_k \leq z] = \prod_{k=1}^{n} F(z - \nu_k)$$

Since the expectation in equation (A.20) takes the form $\mathbb{E}_\varepsilon[\max_g \nu_{img} + \varepsilon_{img}]$, with $\nu_{img} = (\gamma_m(Z)\beta_{mg} - \ln p_{mg})/\mu_m(Z)$, and since I have assumed that $\varepsilon_{img} \overset{\text{iid}}{\sim} F(x)$ for $F(x) = \exp(-\exp(-x))$, I can use the de Palma and Kilani (2007) result to solve for the expectation as follows, dropping the $i$ and $m$ subscripts for the notational convenience:

$$\begin{aligned}
\mathbb{E}_\varepsilon\left[\max_{g \in \mathbf{G_m}} v_g + \varepsilon_g\right] &= \int_{-\infty}^{\infty} z d\phi(z) \\
&= \int_{-\infty}^{\infty} z d\left[\prod_{g=1}^{G_m} \exp(-\exp(v_g - z))\right] \\
&= \int_{-\infty}^{\infty} z d\left[\exp\left(\sum_{g=1}^{G_m} -\exp(v_g - z)\right)\right] \\
&= \int_{-\infty}^{\infty} z \left(\sum_{g=1}^{G_m} \exp(v_g - z)\right) \exp\left(\sum_{g=1}^{G_m} -\exp(v_g - z)\right) dz
\end{aligned}$$

Let $V = \ln\left[\sum_{g=1}^{G_m} \exp(v_g)\right]$ and $x = \sum_{g=1}^{G_m} \exp(v_g - z) = \left[\sum_{g=1}^{G_m} \exp(v_g)\right] \exp(-z) = V\exp(-z)$. I solve the above integral by substituting for $z = V - \ln x$, where $dz = -(1/x)dx$ :

$$\begin{aligned}
\mathbb{E}_\varepsilon\left[\max_{g \in \mathbf{G_m}} v_g + \varepsilon_g\right] &= \int_{-\infty}^{\infty} z \left(\sum_{g=1}^{G_m} \exp(v_g - z)\right) \exp\left(\sum_{g=1}^{G_m} -\exp(v_g - z)\right) dz \\
&= \int_{-\infty}^{\infty} z \exp\left(\sum_{g=1}^{G_m} -\exp(v_g - z)\right)\left(\sum_{g=1}^{G_m} \exp(v_g - z)\right) dz \\
&= \int_{\infty}^{0} (V - \ln x)\exp(-x)\, x(-1/x)dx \\
&= \int_{0}^{\infty} (V - \ln x)\exp(-x)\, dx \\
&= V
\end{aligned}$$

Since we have defined $\nu_{img} = (\gamma_m(Z)\beta_{mg} - \ln p_{mg})/\mu_m(Z)$ and $V = \ln\left[\sum_{g=1}^{G_m} \exp(v_g)\right]$, we can use the above result to solve for the expectation in equation (A.19):

$$\mathbb{E}_\varepsilon\left[\ln\left(\max_{g\in\mathbf{G_m}}\tilde{p}_{img}\right)\right] = \mu_m(Z)\ln\left[\sum_{g\in\mathbf{G_m}}\exp((\gamma_m(Z)\beta_{mg}-\ln p_{mg})/\mu_m(Z))\right]$$

$$= \mu_m(Z)\ln\left[\sum_{g\in\mathbf{G_m}}\left(\frac{\exp(\gamma_m(Z)\beta_{mg})}{p_{mg}}\right)^{\frac{1}{\mu_m(Z)}}\right]$$

$$= \ln\left[\sum_{g\in\mathbf{G_m}}\left(\frac{\exp(\gamma_m(Z)\beta_{mg})}{p_{mg}}\right)^{\frac{1}{\mu_m(Z)}}\right]^{\mu_m(Z)} \tag{A.21}$$

Plugging this result back into equation (A.18) yields the expected relative module expenditure share for consumer $i$ in terms of product prices and model parameters:

$$\begin{aligned}
&\mathbb{E}_\varepsilon\left[\ln s_{im}(Z,\mathbb{P})-\ln s_{i\bar{m}}(Z,\mathbb{P})\right]\\
&\quad=(\sigma-1)\mathbb{E}_\varepsilon\left[\ln\left(\max_{g\in\mathbf{G_m}}\frac{\exp(\gamma_m(Z)\beta_{mg}+\mu_m(Z)\varepsilon_{img})}{p_{mg}}\right)\right]\\
&\quad\quad-(\sigma-1)\mathbb{E}_\varepsilon\left[\ln\left(\max_{g\in\mathbf{G_{\bar{m}}}}\frac{\exp(\gamma_{\bar{m}}(Z)\beta_{\bar{m}g}+\mu_{\bar{m}}(Z)\varepsilon_{i\bar{m}g})}{p_{\bar{m}g}}\right)\right]\\
&\quad=(\sigma-1)\ln\left[\sum_{g\in\mathbf{G_m}}\left(\frac{\exp(\gamma_m(Z)\beta_{mg})}{p_{mg}}\right)^{\frac{1}{\mu_m(Z)}}\right]^{\mu_m(Z)}\\
&\quad\quad-(\sigma-1)\ln\left[\sum_{g\in\mathbf{G_{\bar{m}}}}\left(\frac{\exp(\gamma_{\bar{m}}(Z)\beta_{\bar{m}g})}{p_{\bar{m}g}}\right)^{\frac{1}{\mu_{\bar{m}}(Z)}}\right]^{\mu_{\bar{m}}(Z)}
\end{aligned}$$

This function only varies by consumer through their non-grocery expenditure. All consumers with the same non-grocery expenditure and facing the same prices, $\mathbb{P}$, will have the same expected relative module expenditure share:

$$\mathbb{E}_\varepsilon\left[\ln s_{im}(Z,\mathbb{P})-\ln s_{i\bar{m}}(Z,\mathbb{P})\right] = -(\sigma-1)\left[\ln V_m(Z,\mathbb{P}_m)-\ln V_{\bar{m}}(Z,\mathbb{P}_{\bar{\mathbf{m}}})\right] \tag{A.22}$$

where $V_m(Z,\mathbb{P}_m)$ is a CES-style index over price-adjusted product qualities:

$$V_m(Z,\mathbb{P}_m) = \left[\sum_{g\in\mathbf{G_m}}\left(\frac{\exp(\beta_{mg}(1+\gamma_m\ln Z))}{p_{mg}}\right)^{(1-\sigma)}\right]^{\frac{1}{1-\sigma}} \tag{A.23}$$

where I have substituted in the parametrizations for $\gamma_m(Z)=(1+\gamma_m\ln Z)$ and $\mu_m(Z)=1/\left(\alpha_m^0+\alpha_m^1\ln Z\right)$. Equations (A.22) and (A.23) together define the expected relative module expenditure share of a set of households with income $Y_i$ that face prices $\mathbb{P}_m$ and $\mathbb{P}_{\bar{m}}$ in terms of parameters $\alpha^0$, $\alpha^1$, as well as $\alpha_m$, $\gamma_m$, $\beta_{mg}$ for all $g\in G_m$, and $\alpha_{\bar{m}}$, $\gamma_{\bar{m}}$, $\beta_{\bar{m}g}$ for all $g\in G_{\bar{m}}$.

### D.1.3 Extracting Second Stage Estimates $\theta_2$ From the Inclusive Value Function

The expected log expenditure share in module $m$ relative to $\bar{m}$ for a group of households with the same non-grocery expenditure, $Z_i$, facing a common vector of grocery prices, $\mathbb{P}$, is defined above in Equations (A.22) and (A.23). Adjusting these expressions to reflect time-varying store-specific pricing and promotion activity yields:

$$\mathbb{E}_\varepsilon \left[\ln s_{imst} - \ln s_{i\bar{m}st}\right] = (\sigma - 1) \ln \tilde{V}_m(Z_i, \mathbb{P}_{mst}, \mathbb{P}_{\bar{m}st}) \tag{A.24}$$

where $\tilde{V}_m(Z_i, \mathbb{P}_{mst}, \mathbb{P}_{\bar{m}st}) = V_m(Z_i, \mathbb{P}_{mst})/V_{\bar{m}}(Z_i, \mathbb{P}_{\bar{m}st})$. $V_m(Z_i, \mathbb{P}_{mst})$ is a CES-style index over price-adjusted product qualities:

$$V_m(Z_i, \mathbb{P}_{mst}) = \left[\sum_{g \in \mathbf{G_m}} \left(\frac{\exp(\gamma_{im}\beta_{mgst})}{p_{mgst}}\right)^{-\alpha_{im}}\right]^{\frac{1}{-\alpha_{im}}} \tag{A.25}$$

for $\alpha_{im} = \left(\alpha_m^0 + \alpha_m^1 \ln Z_i\right)$ and $\gamma_{im} = (1 + \gamma_m \ln Z_i)$. Note that the inclusive value is a function of the parameters estimated in both the first and second stage, i.e., $\theta_1$ and $\theta_2$. Specifically, each store-specific product quality shock, $\beta_{mgst}$, is the the sum of $(\beta_{mgst} - \beta_{m\bar{g}_m st})$, estimated in stage 1, and an unknown base product quality shock, $\beta_{m\bar{g}_m st}$. We can express the inclusive value function as the product of the base product quality parameter, $\beta_{m\bar{g}_m st}$, to be estimated in the second stage and an inclusive value function calculated using only elements of $\theta_{1m}$ estimated in the first stage:

$$V_m(Z_i, \mathbb{P}_{mst}) = \exp(\gamma_{im}\beta_{m\bar{g}_m st}) V_{1m}(Z_i, \mathbb{P}_{mst})$$

where

$$V_{1m}(Z_i, \mathbb{P}_{mst}) = \left[\sum_{g \in \mathbf{G_m}} \left(\frac{\exp(\gamma_{im}\tilde{\beta}_{mgst})}{p_{mgst}}\right)^{-\alpha_{im}}\right]^{\frac{1}{-\alpha_{im}}} \tag{A.26}$$

and $\tilde{\beta}_{mgst} = \beta_{mgst} - \beta_{m\bar{g}_m st}$. Under the normalization that $\beta_{\bar{m}\bar{g}_{\bar{m}} st} = 0$ for all $s, t$, and using the decomposition of the inclusive value function above, we can now rewrite equation (A.24) as:

$$\mathbb{E}_\varepsilon \left[\ln s_{imst} - \ln s_{i\bar{m}st}\right] = (\sigma - 1) \left(\gamma_{im}\beta_{m\bar{g}_m st} + \ln \tilde{V}_{1mst}(Z_i, \mathbb{P}_{mst}, \mathbb{P}_{\bar{m}st})\right) \tag{A.27}$$

where $\ln \tilde{V}_{1mst}(Z_i, \mathbb{P}_{mst}, \mathbb{P}_{\bar{m}st}) = \ln V_{1m}(Z_i, \mathbb{P}_{mst}) - \ln V_{1\bar{m}}(Z_i, \mathbb{P}_{\bar{m}st})$.

The predicted log expenditure share of module $m$ relative to module $\bar{m}$ at store $s$ at time $t$ is obtained by aggregating $i$-specific expected relative shares over the units purchased by customers at each non-grocery expenditure level:

(A.28)

$$\mathbb{E}_z \left[\mathbb{E}_\varepsilon \left[\ln s_{imst} - \ln s_{i\bar{m}st}\right]\right] = \int (\sigma - 1) \left(\gamma_{im}\beta_{m\bar{g}_m st} + \ln \tilde{V}_{1mst}(Z_i, \mathbb{P}_{mst}, \mathbb{P}_{\bar{m}st})\right) dF(Z|s, t)$$

where $F(Z|s, t)$ is the distribution of non-grocery expenditures over the households shopping in store $s$ at time $t$.

Notice that this function is linear in the unobserved base product quality for module $m$, $\beta_{m\bar{g}_m st}$, and the relative inclusive value function, so we can derive the following linear estimating equation:

$$\mathbb{E}_z\left[\mathbb{E}_\varepsilon\left[\ln s_{imst} - \ln s_{i\bar{m}st}\right]\right] = \beta_{m\bar{g}_m st}\left(\sigma - 1\right)\overline{\gamma}_{mst} + \left(\sigma - 1\right)\overline{\tilde{v}}_{mst} \tag{A.29}$$

where $\overline{\gamma}_{mst} = \int \gamma_{im} dF(Z|s,t)$ and $\overline{\tilde{v}}_{mst} = \int \ln \tilde{V}_{1mst}(Z_i, \mathbb{P}_{mst}, \mathbb{P}_{\bar{m}st}) dF(Z|s,t)$ can be calculated using price data, estimates of the store-time-level income distributions, and stage 1 parameter estimates.

## D.2 Estimation Procedure

In this appendix I describe the details involved in the estimation and statistical inference of the parameter vector $\theta$. I estimate the parameters of the model sequentially. Recall that the full set of demand parameters, $\theta$, are partitioned into $M$ sets of lower-level module-specific parameters, $\theta_{1m}$ for each module $m$, that are identified using module-specific sub-samples of the data, and a single set of parameters, $\theta_2$, whose identification requires data from all modules.

### D.2.1 Step 1: Parallel Estimation of $\theta_1 = \{\alpha_m^0, \alpha_m^1, \gamma_m, \{\beta_{mg} - \beta_{m\bar{g}_m}\}\}_{m=1,...,M}$

The first step in my estimation is to obtain estimates for $\theta_1 = \left\{\alpha_m^0, \alpha_m^1, \gamma_m, \{\tilde{\beta}_{mg}\}_{g \in \mathbf{G_m}, g \neq \bar{g}_m}\right\}_{m=1,...,M}$, where $\tilde{\beta}_{mg}$ denotes $\beta_{mg} - \beta_{m\bar{g}_m}$. I obtain $\hat{\theta}_1$ using a two-stage GMM procedure based on the following exogeneity restriction:

$$\mathbb{E}[g(\mathbf{X}; \theta_1)] = 0 \tag{A.30}$$

where $g(\mathbf{X}; \theta_1)$=$[g^1(\mathbf{X}; \theta), g^2(\mathbf{X}; \theta), g^3(\mathbf{X}; \theta)]$ consists of three vectors of module-specific moments, $g^k(\mathbf{X}; \theta) = [g^k(\mathbf{X}_1; \theta_1), ..., g^k(\mathbf{X}_M; \theta_M)]$.

The first vector of moments is calculated using store-level data. They are defined as:

$$\bar{g}^1(\mathbf{X}_m; \theta_{1m}) = \frac{1}{n} \sum_{mg,s,t} g^1_{mgst}(\mathbf{X}_m; \theta_{1m}) = \frac{1}{n} \sum_{mg,s,t} \tilde{\xi}_{mgst}(\mathbf{X}_m; \theta_{1m}) \tilde{\mathbf{Z}}^1_{mgst}$$

where $n$ is the number of store-product-time observations; $\xi_{mgst}(\mathbf{X}_m; \theta^{NL}_{1m})$ are transient store-time-specific product taste shocks defined below; and $\mathbf{Z}^1_{mgst}$ is a vector of $L^1_m$ pre-determined variables including product fixed effects and price instruments. The tilde denotes that a variable has been differenced from the respective value for the base product in each module, $\bar{g}_m$, e.g., $\tilde{\xi}_{mgst}(\mathbf{X}_m; \theta_{1m}) = \xi_{mgst}(\mathbf{X}_m; \theta_{1m}) - \xi_{m\bar{g}_m st}(\mathbf{X}_m; \theta_{1m})$.

The second and third vectors of moments are designed to employ the Nielsen data on household-level product choices. The second set of moments equalizes the predicted uncentered covariance between product quality and household non-grocery expenditure for Nielsen HMS sample households. The sample analog of this covariance is:

$$\bar{g}^2(\mathbf{X_m}; \theta_{1m}) = \frac{1}{N_m} \sum_{mg} g^2_{mg}(\mathbf{X_m}; \theta_{1m}) = \frac{1}{N_m} \sum_{mg} N_{mg} \beta_{mg} \left\{ \frac{1}{N_{mg}} \sum_{i_{mg}=1}^{n_{mg}} Z_{i_{mg}} - E[Z|y = mg, \theta] \right\}$$

where $i_{mg}$ denotes one of the $N_{mg}$ units of product $g$ in module $m$ that is purchased in the Nielsen HMS sample; $i$ denotes one of the $N_m$ units of any product in module $m$ that is purchased in the Nielsen HMS sample; and $Z_i$ denotes the non-grocery expenditure of the Nielsen HMS panelist purchasing unit $i$. Similarly, the third set of moments equalizes the predicted uncentered covariance between unit price paid and household non-grocery expenditure. The sample analog of this covariance is:

$$\hat{\hat{g}}^3(\mathbf{X}_m;\theta_{1m}) = \frac{1}{N_m}\sum_i \left(Z_i - \bar{Z}\right)\sum_{s,t}\left((\tilde{p}_{imst} - E[\tilde{p}_{imst}|\theta_{1m}]) - \frac{1}{N_m}\sum_i\sum_{s,t}(\tilde{p}_{imst} - E[\tilde{p}_{imst}|\theta_{1m}])\right)$$

The sample analogs of the three moment conditions defined above are:

$$\begin{aligned}
\hat{\hat{g}}^1(\mathbf{X}_m;\theta_{1m}) &= \frac{1}{n}\sum_{mg,s,t}\hat{\hat{\xi}}_{mgst}(\mathbf{X}_m;\theta_{1m})\tilde{\mathbf{Z}}_{mgst} \\
\hat{\hat{g}}^2(\mathbf{X}_m;\theta_{1m}) &= \frac{1}{N_m}\sum_{mg} N_{mg}\left\{\beta_{mg}\left[\frac{1}{N_{mg}}\sum_{i_{mg}=1}^{N_{mg}} Z_{i_{mg}} - \frac{1}{N_m}\sum_{i=1}^{N_m} Z_i P_{mg}(Z_i,\mathbb{P}_{st},\theta_{1m},\hat{\beta}_{st})\right]\right\}^2 \\
\hat{\hat{g}}^3(\mathbf{X}_m;\theta_{1m}) &= \frac{1}{N_m}\sum_i \left(Z_i - \bar{Z}\right)\sum_{s,t}\left((\tilde{p}_{imst} - E[\tilde{p}_{imst}|\theta_{1m}]) - \frac{1}{N_m}\sum_i\sum_{s,t}(\tilde{p}_{imst} - E[\tilde{p}_{imst}|\theta_{1m}])\right)
\end{aligned}$$

where $\bar{Z} = \frac{1}{N_m}\sum_i \bar{Z}_i$ is the unit-weighted mean non-grocery expenditure of sample households; $\tilde{p}_{imst} = (p_{imgst} - \bar{p}_{mst})$ is the relative unit value paid by a household $i$ in module $m$ in store $s$ at time $t$, where $\bar{p}_{mst} = \sum_{g\in\mathbf{G}_{mst}} w_{mgst}p_{mgst}$ and $w_{mgst} = s_{mg}/\sum_{g\in\mathbf{G}_{mst}} s_{mg}$, and $E[\tilde{p}_{imst}|\theta_{1m}]$ is the predicted relative unit value paid by household $i$ in module $m$ in store $s$ and time $t$ defined as:[71]

$$E[\tilde{p}_{imst}|\theta_{1m}] = \sum_{g\in\mathbf{G}_{mst}}\tilde{p}_{mgst}P_{mg}(Z_i,\mathbb{P}_{st},\theta_{1m},\hat{\beta}_{st})$$

To obtain estimates for the quality parameters $\tilde{\beta}_{mg}(\theta_{1m}^{NL})$ that enter the micro moments, I first follow Berry et al. (1995) inverting simulated market shares to obtain the vector product- and market-specific taste parameters $\tilde{\beta}_{mgst}(\theta_{1m}^{NL})$ that rationalizes the observed product shares in each store and time period conditional on a given set of non-linear parameter vector $\theta_{1m}^{NL} = \left\{\alpha_m^0, \alpha_m^1, \gamma_m\right\}$. Details on the simulation and inversion procedure are provided below.[72] I project the estimated taste parameters, $\hat{\hat{\xi}}_{mgst}(\theta_{1m}^{NL})$, on brand as well as market dummies to control for market-level variation in the quality of the products included in the base good. The coefficients on the brand dummies are used as estimates for the product-specific quality parameters, $\tilde{\beta}_{mg}(\theta_{1m}^{NL})$, employed in the quality micro moment. The residuals from these regressions provide estimates for the transitory shocks, $\xi_{mgst}(\theta_{1m}^{NL})$, which are in turn used to calculate the macro (store-level) moment conditions.

The fact that all three sets of moments depend only on module-specific data, $\mathbf{X}_m$, and parameters,

[71] I can only calculate the probability of purchase, $P_{mg}(Z_i,\mathbb{P}_{st},\theta_{1m},\hat{\beta}_{st})$, employed in the calculation of the micro moments ($\hat{\hat{g}}^2(\mathbf{X}_m;\theta_{1m})$ and $\hat{\hat{g}}^3(\mathbf{X}_m;\theta_{1m})$), when I observe the full choice set available to the Nielsen household panelist $i$; that is, the set of products and prices available to the customer in the store and time period that they are observed to make their purchase ($\mathbb{P}_{st}$). I observe these choice sets for the stores and time periods in the Nielsen RMS data, so calculate the micro moments using household transactions in these stores and time periods alone.

[72] I also attempted estimating these taste shocks using a fourth set of moments equalizing the predicted expenditure shares of a simulated set of customers at each store in each time period with the observed sales shares for the respective stores and time periods following Dubé et al. (2009)'s implementation of Berry et al. (1995). I ran into difficulties getting this model to converge across many modules, however, given the non-linearity of the problem.

$\theta_{1m}$, enables me to partition A.30 into module-specific auxiliary moments:

$$\mathbb{E}[g(\mathbf{X}_m;\theta_{1m})] = 0$$

This partition allows me to estimate the $K_{1m}$ parameters, $\theta_{1m} = \left\{\alpha_m^0, \alpha_m^1, \gamma_m, \{\tilde{\beta}_{mg}\}_{g\in\mathbf{G_m}, g\neq\bar{g}_m}\right\}$, for each module $m$ in separate but parallel minimization procedures. Consistent estimates of the elasticity parameters, $\theta_{1m}^{NL} = \left\{\alpha_m^0, \alpha_m^1, \gamma_m\right\}$, are obtained by minimizing module-specific GMM objective functions as follows:

$$\hat{\theta}_{1m}^{NL} = \arg\min_{\theta_{1m}^{NL}} \hat{\bar{g}}(\mathbf{X}_m;\theta_{1m})'\hat{\mathbf{W}}_{1m}\hat{\bar{g}}(\mathbf{X}_m;\theta_{1m})$$

where $\hat{\bar{g}}(\mathbf{X}_m;\theta_{1m})$ is the sample analog of the $L_m^1 + 1 \geq K_{1m}$ moments, $\bar{g}(\mathbf{X}_m;\theta_{1m})$ and $\hat{\mathbf{W}}_{1m}$ is the efficient weighting matrix.

The weighting matrix, $\hat{\mathbf{W}}_{1m}^1$, is block-diagonal since the three moments are calculated using different datasets:

$$\hat{\mathbf{W}}_{1m}^1 = \begin{bmatrix} \hat{W}_{1m}^{11}(\mathbf{X}_m;\tilde{\theta}_{1m}) & 0 & 0 \\ 0 & \hat{W}_{1m}^{12}(\mathbf{X}_m;\tilde{\theta}_{1m}) & 0 \\ 0 & 0 & \hat{W}_{1m}^{13}(\mathbf{X}_m;\tilde{\theta}_{1m}) \end{bmatrix}^{-1}$$

for

$$\begin{aligned} \hat{W}_{1m}^{11}(\mathbf{X}_m;\tilde{\theta}_{1m}) &= \frac{1}{n}\sum_{mg,s,t} \hat{g}_{mgst}^1(\mathbf{X}_m;\tilde{\theta}_{1m})\hat{g}_{mgst}^1(\mathbf{X}_m;\tilde{\theta}_{1m})' \\ \hat{W}_{1m}^{12}(\mathbf{X}_m;\tilde{\theta}_{1m}) &= \frac{1}{N_m}\sum_{mg} \hat{g}_{mg}^2(\mathbf{X}_m;\tilde{\theta}_{1m})\hat{g}_{mg}^2(\mathbf{X}_m;\tilde{\theta}_{1m})' \\ \hat{W}_{1m}^{13}(\mathbf{X}_m;\tilde{\theta}_{1m}) &= \frac{1}{N_m}\sum_{mg} \hat{g}_{mg}^3(\mathbf{X}_m;\tilde{\theta}_{1m})\hat{g}_{mg}^3(\mathbf{X}_m;\tilde{\theta}_{1m})' \end{aligned}$$

Each of these components is calculated using consistent first-stage estimates of $\theta_{1m}^{NL}$:

$$\tilde{\theta}_{1m}^{NL} = \arg\min_{\theta_{1m}^{NL}} \hat{\bar{g}}(\mathbf{X}_m;\theta_{1m})'\mathbf{W}_{1m}\hat{\bar{g}}(\mathbf{X}_m;\theta_{1m})$$

for

$$\mathbf{W}_{1m} = \begin{bmatrix} \left[\frac{1}{n}\sum_{mg,s,t}\sum_{g\in\mathbf{G_{m}}_{st}} \tilde{Z}_{mgst}^1\left(\tilde{Z}_{mgst}^1\right)'\right]^{-1} & 0 & 0 \\ 0 & 1 & 0 \\ 0 & 0 & 1 \end{bmatrix}.$$

After estimating the non-linear parameters, $\hat{\theta}_{1m}^{NL}$, I project the product-store-time specific taste shocks implied by these parameters, $\tilde{\beta}_{mgst}(\hat{\theta}_{1m}^{NL})$, onto brand dummies in order to extract estimates of the product quality parameters, $\{\tilde{\beta}_{mg}\}_{g \in \mathbf{G_m}, g \neq \bar{g}_m}$.

Assuming that the random components of the $M$ module-specific auxiliary models are independent, the variance-covariance matrix of $\hat{\theta}_1$, $\Omega_1$, can be written as:

$$\Omega_{\theta_1} = \begin{bmatrix} \Omega_{\theta_{11}} & & & & 0 \\ & \ddots & & & \\ & & \Omega_{\theta_{1m}} & & \\ & & & \ddots & \\ 0 & & & & \Omega_{\theta_{1M}} \end{bmatrix}$$

where $\Omega_{\theta_{1m}}$ is the variance-covariance matrix of $\theta_{1m}$ for each $m = 1, ..., M$. The consistent estimator for each of these sub-matrices is:

$$\hat{\Omega}_{\theta_{1m}} = \left( \hat{F}_{\theta_{1m}} \hat{V}_{ff}^{-1} \hat{F}'_{\theta_{1m}} \right)^{-1}$$

where $\hat{F}_{\theta_{1m}} = \begin{bmatrix} \hat{F}^1_{\theta_{1m}} & \hat{F}^2_{\theta_{1m}} \end{bmatrix}'$ for

$$\hat{F}^1_{\theta_{1m}} = \frac{1}{n} \sum_{mg,s,t} \nabla_{\theta_{1m}} \hat{g}^1_{mgst}(\mathbf{X}_m; \hat{\theta}_{1m})$$

and

$$\hat{F}^2_{\theta_{1m}} = \frac{1}{N_m} \sum_{mg} \nabla_{\theta_{1m}} \hat{g}^2_{mg}(\mathbf{X}_m; \hat{\theta}_{1m})$$

and

$$\hat{V}_{ff} = \begin{bmatrix} \frac{1}{n} \sum_{mg,s,t} \hat{g}^1_{mgst}(\mathbf{X}_m; \hat{\theta}_{1m}) \hat{g}^1_{mgst}(\mathbf{X}_m; \hat{\theta}_{1m})' & 0 \\ 0 & \frac{1}{N_m} \sum_{mg} \hat{g}^2_{mg}(\mathbf{X}_m; \hat{\theta}_{1m}) \hat{g}^2_{mg}(\mathbf{X}_m; \hat{\theta}_{1m})' \end{bmatrix}$$

**Inversion Algorithm** In order to evaluate the objective function at a given parameter vector $\theta^{NL}_{1m}$, it is necessary to invert the following system of non-linear equations:

(A.31) $$\beta_{mgst}(\theta_{1m}) \to \ln s_{mgst}(\beta_{st}; \theta^{NL}_{1m}) = \ln \hat{s}_{mgst}$$

where $s_{mgst}(\beta_{st}; \theta^{NL}_{1m})$ is the model predicted market share of product $g$ in store $s$ at time $t$, $\theta^{NL}_{1m} = \left\{ \alpha^0_m, \alpha^1_m, \gamma_m \right\}$ is the subset of elasticity parameters that must be estimated using non-linear moments, and $\hat{s}_{mgst}$ is the observed share. For each guess of $\theta^{NL}_{1m}$, I calculate the model predicted market share as the average probability of purchase predicted for a quadrature of $K$ points from the store-specific income distribution (recall that income is used to proxy for non-grocery expenditure $Z_i$) each with income $Y_k$ and weight $w_k$:

(A.32) $$s_{mgst}(\beta_{st}; \theta^{NL}_{1m}) = \sum_{k=1}^{K} w_k P_{mg}(Y_k, \mathbb{P}, \theta_m)$$

It is well known that this inversion does not work for products with small sales shares (see, e.g., Gandhi et al. (2013)). I therefore group all of the products that fall into the left tail of the average sales distribution an outside product. This grouping could impact my estimates in three ways. First, Gandhi et al. (2013) have demonstrated that ignoring the low end of the sales distribution in this manner yields a downward bias on price elasticity estimates. Second, variation in the quality of the outside goods sold in different stores could bias my average product quality estimates as discussed under identification in Section 5.3.1. Finally, I will not estimate product quality parameters for products that always appear in the low end of the sales distribution and, therefore, am unable to include them in the market price indexes. To test the impact of these biases on my results, I study how the estimated price elasticities and product quality gradients vary depending on the share of products that are grouped into this outside product, varying this set between 40, 60, and 80 percent of products in each store-week (reflecting 6, 15, and 33 percent of aggregate product sales, respectively) in the robustness exercises presented in Section E.3.5.

**Starting Values** I estimate a linear approximation of the store-level market share equation to obtain starting values for the non-linear parameters, $\theta_{1m}^{NL} = \{\alpha_m^0, \alpha_m^1, \gamma_m\}$. When the optimization routine returns estimates within 0.03 log units of the bounds for these non-linear estimates –$\alpha_m^0 \in (0.05, 30)$, $\alpha_m^1 \in (-5, 5)$, and $\gamma_m \in (-5, 5)$ – or otherwise fails, I instead conduct a grid search. Specifically, I run the optimization routine using a range of starting values for the mean price elasticity, $\alpha_m^{0,start}$ between 1 and 4, keeping the starting values for the non-homotheticity parameters of $\gamma_m^{start} = 1.5$ and $\alpha_m^{1,start} = 2$ (or zero, in the constrained model). If this yields multiple sets of interior estimates, I select the estimates minimize the objective function. If this routine does not yield interior estimates, I try estimating demand using a less restrictive retention rule for the instrument PCA, using $\delta$=-1.4 instead of $\delta$=-2, where the eigenvalue cutoff is sum(latent)*(size(latent,1)^$\delta$ ). The initial estimation with fixed starting values yields interior estimates for between 300 and 400 modules out of 530, depending on the model being estimated. The grid search yields interior estimates for between 50 and 100 additional modules. Loosening the instrument retention rule yields the remaining interior estimates reported in Table 3.

### D.2.2 Step 2: Sequential Estimation of $\theta_2 = \{\sigma, \{\beta_{m\bar{g}_m}\}_{m=1,\dots,M,m\neq\bar{m}}\}$

In the second step of the sequential estimation procedure, I estimate $\theta_2 = \{\sigma, \{\beta_{m\bar{g}_m}\}_{m=1,\dots,M,m\neq\bar{m}}\}$. These $K_2 = 1 + M$ parameters are identified by the following exogeneity restriction:

$$G = \mathbb{E}[h(\mathbf{X}; \theta_1, \theta_2)] = 0 \tag{A.33}$$

where $h(\mathbf{X}; \theta_1, \theta_2) = \mathbf{Z}_2(\mathbf{X}) \cdot u(\mathbf{X}; \theta_1, \theta_2)$. $\mathbf{Z}_2(\mathbf{X})$ is a set of $L_2$ instruments ($L_2 \geq K_2$) and $u(\mathbf{X}; \theta_1, \theta_2)$ is the error in the relative across-module expenditure share equation derived above.

Specifically, for module $m$ and store $s$ in time $t$ this error is derived above in equation (A.29) as:

$$u_{mst}(\mathbf{X}; \theta_1, \theta_2) \quad = \quad \ln\left(s_{mst}/s_{\bar{m}st}\right) - \beta_{m\bar{g}_m}\left(\sigma - 1\right)\overline{\gamma}_{mst}(\hat{\theta}_1) - \left(\sigma - 1\right)\overline{\overline{v}}_{mst}(\hat{\theta}_1)$$

where $s_{mst}$ and $s_{\bar{m}st}$ are data on the respective sales shares of module $m$ and $\bar{m}$ in store $s$ in time

$t$; each $\bar{x}_{mst} = \int x_{imst} dF(Z|s,t)$ is calculated by integrating $x_{imst}$ over the same local income distribution employed in the first-stage of estimation described above, for $\gamma_{im} = (1 + \gamma_m \ln Z_i)$ and $\tilde{v}_{mst} = \ln V_{1m}(Z_i, \mathbb{P}_{mst}, \theta_{1m}) - \ln V_{1\bar{m}}(Z_i, \mathbb{P}_{\bar{m}st}, \theta_{1m})$ where

$$V_{1m}(Z_i, \mathbb{P}_{mst}, \theta_{1m}) = \left[ \sum_{g \in \mathbf{G_m}} \left( \frac{\exp(\gamma_{im} \tilde{\beta}_{mgst})}{p_{mgst}} \right)^{-(\alpha_m^0 + \alpha_m^1 \ln Z_i)} \right]^{\frac{1}{-(\alpha_m^0 + \alpha_m^1 \ln Z_i)}}$$

is the inclusive value for a household with non-grocery expenditure $Z_i$ in module $m$ in market $st$ calculated using first-stage parameter estimates, $\hat{\theta}_1$ .

$\mathbf{Z}_2(\mathbf{X})$ is a vector of pre-determined variables including module fixed effects interacted with the market average quality weight, $\overline{\gamma}_{mst}$, and an instrument for the average relative inclusive value, $\overline{\tilde{v}}_{mst}(\hat{\theta}_1)$, faced by the store's customers. This instrument is identical to the data analog of $\overline{\tilde{v}}_{mst}(\hat{\theta}_1)$ but calculated using the same contemporaneous chain-specific national cost shock instruments that are used in the module-level estimation in place of store-time-specific price data.

The upper-level parameters are estimated using two-step GMM:

$$\hat{\theta}_2 = \arg\min_{\theta_2} \hat{h}(\mathbf{X}; \hat{\theta}_1, \theta_2)' \hat{\mathbf{W}}_\mathbf{2} \hat{h}(\mathbf{X}; \hat{\theta}_1, \theta_2)$$

where $\hat{\mathbf{W}}_2 = \left[ \frac{1}{\sum_{s,t} N_{st}} \sum_{s,t} \sum_{m \in \mathbf{M}_{st}} h_{mst}(\mathbf{X}; \hat{\theta}_1, \tilde{\theta}_2) h_{mst}(\mathbf{X}; \hat{\theta}_1, \tilde{\theta}_2)' \right]^{-1}$ is the optimal weighting matrix, for $\tilde{\theta}_2$ the consistent first-stage estimates of $\theta_2$ that minimize a GMM objective function as follows:

$$\tilde{\theta}_2 = \arg\min_{\theta_2} \hat{h}(\mathbf{X}; \hat{\theta}_1, \theta_2)' \tilde{\mathbf{W}}_\mathbf{2} \hat{h}(\mathbf{X}; \hat{\theta}_1, \theta_2)$$

where $\tilde{\mathbf{W}}_2 = \left[ \frac{1}{\sum_{k,t} N_{st}} \sum_{s,t} \sum_{m \in \mathbf{M}_{st}} \mathbf{Z}_{2mst} \mathbf{Z}'_{2mst} \right]^{-1}$.

Newey and McFadden (1994) show how to obtain a consistent covariance matrix for estimates that are obtained sequentially and Murphy and Topel (1985) describe the assumptions under which this method can be extended to the case in which the first-step estimates are obtained from different models estimated using sub-samples of the data. The naive variance-covariance matrix of the $\hat{\theta}_2$ estimates that does not account for the measurement error from the use of the first stage estimates, treating $\theta_1$ as known, is defined as:

$$\tilde{\Omega}_{\theta_2} = \left( \hat{H}_{\theta_2} \hat{V}_{hh}^{-1} \hat{H}'_{\theta_2} \right)^{-1}$$

where

$$\hat{H}_{\theta_2} = \frac{1}{\sum_{s,t} N_{st}} \sum_{s,t} \sum_{m \in \mathbf{M}_{st}} \nabla_{\theta_2} h_{mst}(\mathbf{X}; \hat{\theta}_1, \hat{\theta}_2) \; (K_2 \times L_2)$$

and

$$\hat{V}_{hh} = \hat{\mathbf{W}}_2 = \frac{1}{\sum_{s,t} N_{st}} \sum_{s,t} \sum_{m \in \mathbf{M}_{st}} h_{mst}(\mathbf{X}; \hat{\theta}_1, \hat{\theta}_2) h_{mst}(\mathbf{X}; \hat{\theta}_1, \hat{\theta}_2)' \; (L_2 \times L_2).$$

In order to account for the measurement error from the use of first stage estimates, we need to treat $\theta_1$ as unknown, calculating the variance-covariance of the full vector of $\hat{\theta}$ estimates:

$$\Omega_\theta = \begin{bmatrix} \Omega_{\theta_1} & \Omega_{\theta_2\theta_1} \\ \Omega_{\theta_1\theta_2} & \Omega_{\theta_2} \end{bmatrix} = \left( C_\theta V_{cc}^{-1} C_\theta' \right)^{-1}$$

where:

$$C_\theta = \begin{bmatrix} F_{\theta_1} & 0 \\ H_{\theta_1} & H_{\theta_2} \end{bmatrix} \text{ and } V_{cc} = \begin{bmatrix} V_{ff} & V_{hf} \\ V_{fh} & V_{hh} \end{bmatrix}$$

The correct covariance matrix for the second stage estimates is the lower right-hand block of this full covariance matrix, $\Omega_{\theta_2}$.[73] I obtain it by estimating the full covariance matrix, $\hat{\Omega}_{\theta_2}$, where $\hat{\Omega}_{\theta_1}$, $\tilde{\Omega}_{\theta_2}$, $\hat{H}_{\theta_2}$, and $\hat{F}_{\theta_1}$ are as defined above;

$$\hat{H}_{\theta_1} = \frac{1}{\sum_{s,t} N_{st}} \sum_{s,t} \sum_{m \in \mathbf{M}_{st}} \nabla_{\theta_1} h_{mst}(\mathbf{X}; \hat{\theta}_1, \hat{\theta}_2) \; (K_1 \times L_2);$$

and

$$\hat{V}_{fg} = \hat{V}_{gf}' = \frac{1}{\sum_{s,t} N_{st}} \sum_{s,t} \sum_{m \in \mathbf{M}_{st}} h_{mst}(\mathbf{X}; \hat{\theta}_1, \hat{\theta}_2) [\frac{1}{N_{mst}} \sum_{g \in \mathbf{G_m}_{st}} f_{gst}(\mathbf{X}; \hat{\theta}_{1m})]' \; (L_2 \times L_1).$$

## D.3 Model Selection Criterion

In the Section 6.1, I present estimates of the parameters that govern the within-module product choice for each module $m$, $\hat{\theta}_{1m}$, in a separate GMM estimation procedure under different sets of parameter restrictions. For the most flexible "full" version of the model, all elements of $\theta_{1m}$ are estimated. These include $\alpha_m^0$, $\alpha_m^1$, $\gamma_m$, and a relative quality parameter $(\beta_{mg} - \beta_{m\bar{g}})$ for each brand represented in the module except for the brand of the base product $\bar{g}$. The full model allows for non-homotheticity in both the price sensitivity and the demand for quality by letting both $\alpha_m^1$ and $\gamma_m$ be non-zero.

[73]Newey (1984) shows that, when $L_1 = K_1$ and $L_2 = K_2$, the asymptotic covariance matrix $\Omega_{\theta_2}$ of the second step estimator $\hat{\theta}_2$ is given by:

$$\hat{\Omega}_{\theta_2} = \tilde{\Omega}_{\theta_2} + \hat{H}_{\theta_2}^{-1} \hat{H}_{\theta_1} \hat{\Omega}_{\theta_{1m}} \hat{H}_{\theta_1}' (\hat{H}_{\theta_2}^{-1})' - \hat{H}_{\theta_2}^{-1} \left( \hat{H}_{\theta_1} \hat{F}_{\theta_1}^{-1} \hat{V}_{fh} + \hat{V}_{hf} (\hat{F}_{\theta_1}^{-1})' (\hat{H}_{\theta_1}^{-1}) \right)$$

where $\hat{\Omega}_{\theta_1}$, $\tilde{\Omega}_{\theta_2}$, $\hat{H}_{\theta_2}$, and $\hat{F}_{\theta_1}$ are as defined above and $\hat{H}_{\theta_1} = \frac{1}{\sum_{s,t} N_{st}} \sum_{s,t} \sum_{m \in \mathbf{M}_{st}} \nabla_{\theta_1} h_{mst}(\mathbf{X}; \hat{\theta}_1, \theta_2)$. This equation cannot be applied directly to estimate $\Omega_{\theta_2}$ here since both models estimated here are over-identified, such that $L_1 > K_1$ and $L_2 > K_2$ (and neither $\hat{F}_{\theta_1}$ or $\hat{H}_{\theta_2}$ are invertible).

In Section 6.2, I compare the GMM-BIC criterion for this model with the other models that allow for only one form of non-homotheticity by restricting either $\alpha_m^1$ or $\gamma_m$ to be zero. The selection criterion minimizes the following GMM-BIC function:

$$\text{GMM-BIC}_m^M(\hat{\theta}_{1m}^M) = n_m G_m(\hat{\theta}_{1m}^M, \bar{\theta}_{1m}^M)' W_m^* G_m(\hat{\theta}_{1m}^M, \bar{\theta}_{1m}^M) - \ln(n_m)(L_m^* - K_m^M) \tag{A.34}$$

where $G_m(\hat{\theta}_{1m}^M, \bar{\theta}_{1m}^M)$ are the moments for model $M$ evaluated at the estimated values for free parameters $\hat{\theta}_{M1m}^M$ and zero for the restricted parameters, $\bar{\theta}_{1m}^M$; $K_m^M$ is the number of free parameters in model $M$ for module $m$; and $n_m$ and $L_m^*$ are the number of observations and instruments, respectively, used to estimate all models for module $m$. The same set of instruments is used to calculate each moment condition, and thus the number of moments is also common between models for each module. $W_m^*$ is the optimal weighting matrix for the full model.

I evaluate models by calculating the unweighted and sales-weighted share of modules for which a given model minimizes the GMM-BIC criterion. The results of this model selection test are presented in Table 5 and Appendix Table A.4 below.

Table A.4: Bilateral Model Comparisons

| | | Model A | | |
|---|---|---|---|---|
| | | NHQ | NHP | Both |
| Model B | NHQ | - | 0.14 | 0.26 |
| | NHP | 0.82 | - | 0.52 |
| | Both | 0.79 | 0.59 | - |

Note: This table shows the share of modules in which Model 1 (the column model) has a lower Bayesian Information Criterion (BIC) statistic to Model 2 (the row model). The numbers above the diagonal are weighted by 2012 module sales in the RMS data. Those below the diagonal are unweighted. Attention is limited to the set of modules that have interior estimates for all three non-homothetic models.

## D.4 Integrating the Feenstra (1994) Methodology

For robustness of the main model estimates (which allow the demand for quality but not price sensitivity to vary with income), I employ the Feenstra method to calculate the within-module substitution elasticities and, re-estimate the other model parameters, holding these mean elasticities fixed. I first apply the Feenstra method to estimating elasticities with store-level data. The moments for this estimation method are built from the following system of demand and supply equations:

$$\Delta^{\bar{g}_m} \ln s_{mgst} = -\left(\sigma_m - 1\right) \Delta^{\bar{g}_m} \ln p_{mgst} + \varepsilon_{mgst}^{\bar{g}_m}$$

$$\Delta^{\bar{g}_m} \ln p_{mgst} = \frac{\omega_m}{1+\omega_m} \Delta^{\bar{g}_m} \ln s_{mgst} + \delta_{mgst}^{\bar{g}_m}$$

where $s_{mgst}$ and $p_{mgst}$ are the sales share and unit price of product $g$ in module $m$ in store $s$ in month $t$ and each equation is twice-differenced over time and relative to a module-specific base product $\bar{g}_m$, selected from the set that is observed most frequently across store-month markets ($\Delta^{\bar{g}_m} x_{mgst} = \Delta x_{mgst} - \Delta x_{m\bar{g}_m st}$ and $\Delta x_{mgst} = x_{mgst} - x_{mgst-1}$). Demand shocks, $\varepsilon_{mgt}^{\bar{g}_m}$, include the impact of $\Delta\beta_{mgst}$ (since the level is differenced out) as well as a second component, related to the fact that the

Feenstra method does not account for any variation in tastes across incomes. The GMM objective function is:

$$\hat{\kappa}_m = \arg\min_{\kappa_m \in K} G^*(\kappa_m)' \, W G^*(\kappa_m)$$

where $G^*(\kappa_m) = \sum_{s,t} \varepsilon_{mgst}^{\bar{g}_m} \delta_{mgst}^{\bar{g}_m}$, $\kappa_m = \begin{pmatrix} \sigma_m \\ \omega_m \end{pmatrix}$, and $W$ weights moments using product-level sales.

This procedure yields price elasticity estimates, $\alpha_m^0 = \sigma_m - 1$, that are then held fixed in a second estimation step where the remaining parameters, $\gamma_m$, $\beta_{mg}$, and $\alpha_o$ are estimated using same within- and across-module estimation procedure described above. This procedure results in higher estimates for the price elasticity, but also higher estimates for the income-quality gradient ($\gamma$). On net, indexes calculated using these "Feenstra-method" estimates demonstrate the same qualitative patterns across households and cities with different incomes as we observe in the baseline indexes.

# E Results Appendix

## E.1 Identification

### E.1.1 Summary Statistics for Estimation Data Across Modules

Figure A.7: Raw Data Summary Statistics

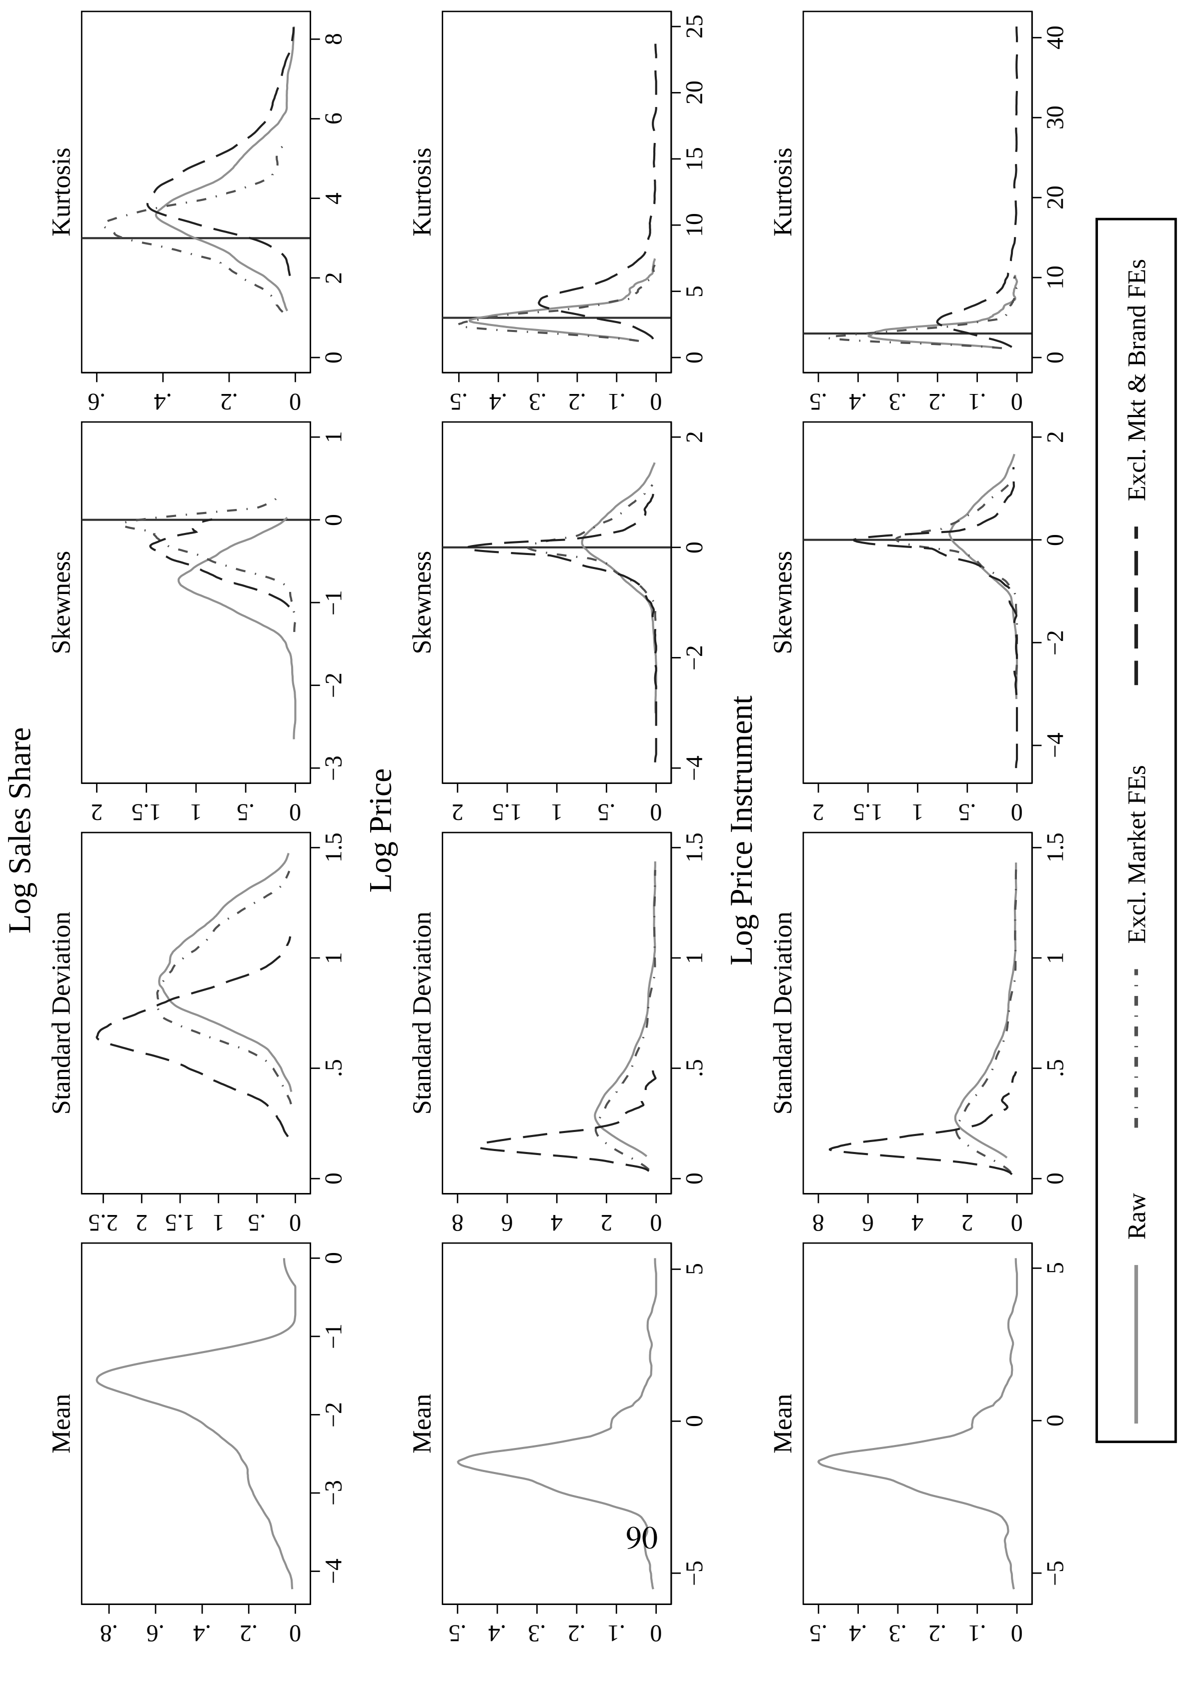


Notes: The plots above depict the distribution of moments of the module-level summary statistics of the data used in the within-module analysis. The vertical lines in the third and fourth columns of plots show the skew and kurtosis of a normal distribution (0 and 3, respectively).

### E.1.2 Distribution of First-Stage Regression Results Across Modules

Figure A.8: Summary Statistics for First Stage Results

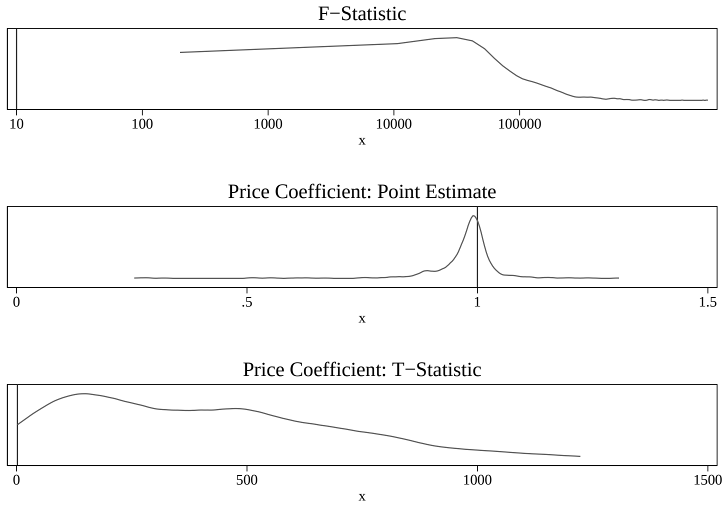


Notes: The above plots depict the distribution of the price instrument coefficients and F-statistics in the module-level first-stage regression of log relative price paid against price instruments, brand dummies, and all of the above interacted with the log median income of the county in which a store is located.

### E.1.3 Measurement Error in Product Quality Estimates

In practice, the quality of each product, relative to the outside good, $\tilde{\beta}_{mg} = \beta_{mg} - \beta_{m\bar{g}_m}$, is calculated as the mean of store-month-specific quality shocks, $\tilde{\beta}_{mgst}(\hat{\theta}^{NL}_{1m}) = \beta_{mgst}(\hat{\theta}^{NL}_{1m}) - \beta_{m\bar{g}_m st}(\hat{\theta}^{NL}_{1m})$, that rationalize the relative sales shares on that product relative to the outside product given the non-linear parameter estimates, across the store-months in which the product is sold; i.e., $\hat{\tilde{\beta}}_{mg} = \frac{1}{N_g}\sum_{st}\tilde{\beta}_{mgst}(\hat{\theta}^{NL}_{1m})$. Variation in the quality of the outside product across store-months will generate measurement error in the quality estimates. $\tilde{\beta}_{mg}$, for example, may understate the relative quality of products that tend to be sold in stores that sell higher quality outside products. If this measurement error is correlated with the relative purchase probability of high- vs. low-income households, it might yield biases in the income-quality elasticity gradient ($\gamma_m$).

To gauge the degree of this error and associated bias, I calculate the relative qualities of "inside" products in two ways. First, I difference the base quality estimate for each product $g$ from the quality estimate for a common product in each module, $\bar{g}^1_m$, i.e., $\tilde{\beta}_{mg} - \tilde{\beta}_{m\bar{g}^1_m}$. This relative quality estimate will be subject to the measurement error noted above (i.e., if $g$ is sold in stores with higher quality outside products than the stores in which $\bar{g}^1_m$ is sold, $\tilde{\beta}_{mg} - \tilde{\beta}_{m\bar{g}^1_m}$ will be biased downwards than the true relative quality of product $g$ relative to product $\bar{g}^1_m$).

I then calculate an alternative measure of the quality of $g$ relative to $\bar{g}^1_m$ that is not subject to this measurement error. Specifically, I difference the market-level quality estimates for product $g$ relative to that for product $\bar{g}^1_m$ within each market and then take the average of this mean across the $N_{g\bar{g}^1_m}$ stores that sell both $g$ and and the common product $\bar{g}^1_m$, i.e., $\frac{1}{N_{g\bar{g}^1_m}}\sum_{st}\left(\tilde{\beta}_{mgst}(\hat{\theta}^{NL}_{1m}) - \tilde{\beta}_{m\bar{g}^1_m st}(\hat{\theta}^{NL}_{1m})\right)$. This procedure purges the relative quality estimate from any variation in the outside product quality level across markets, which appears in both the $\tilde{\beta}_{mgst}(\theta_{1m}^{NL})$ and $\tilde{\beta}_{m\bar{g}_m^1 st}(\theta_{1m}^{NL})$ so is differenced out before averaging.[74]

Comparing these two quality measures assuages concerns that measurement error induced by the variable quality of the outside good across markets generates biases in the estimates. Figure A.9 shows that the two quality measures are highly correlated: the median correlation coefficient across products within modules is 0.8 and over 0.5 in more than 90 percent of modules). More importantly, Figure A.10 shows that there is no systematic variation in the implicit errors in the base quality estimates (i.e., the difference between the base and alternative relative quality measures) across the consumption baskets of high- and low-income households that might generate a bias in the $\gamma_m$ estimates.

Figure A.9: Correlation between Base and Alternative Relative Product Quality Estimates

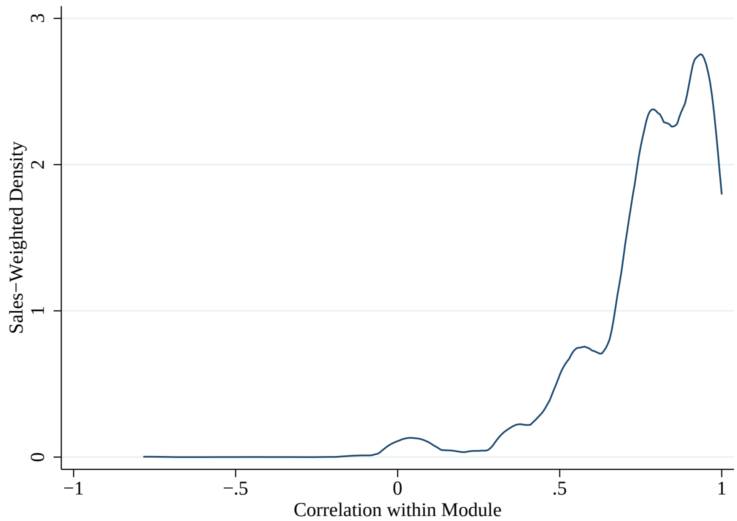


Notes: The above plots depict the distribution of the module-level correlations between two relative quality measures. The first is equal to the mean quality for each product across the stores in which it is sold differenced from the mean quality for a common product across the stores in which it is sold. The second is the difference of the quality of each product in the module in a store from the quality of the common product in that store, averaged over all of the stores in which both products are sold. Module-level correlations are weighted by sales.

[74]I do not obtain my base quality estimates via this procedure because it limits the sample of markets I can use for estimation to those that have a common product. To maximize the number of store-month markets included in the calculations described above and, in turn, the number of products for which this alternate quality measure is feasible, I select as the common product, $\bar{g}_m^1$, the product in each module that appears in the highest number of sample markets. Still, over twenty percent of products are dropped from the analysis entirely because they are not sold in any of the subset of the 5000 randomly-sampled markets that sell the most commonly-sold product for that module. In over a quarter of modules, less than half of the subset of the 5000 randomly-sampled markets that sell the most commonly-sold product for that module. Limiting the sample in this respect might result in the sample becoming biased towards one or two chains that carry similar products.

Figure A.10: Correlation between Base and Alternative Relative Product Quality Estimates

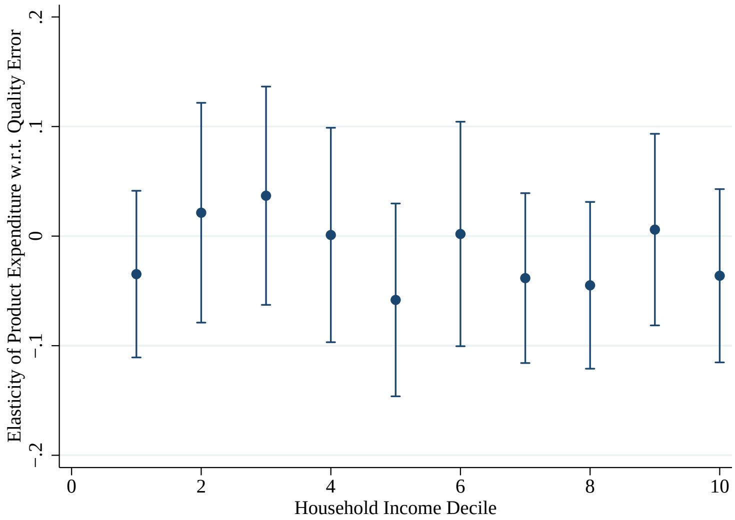


Notes: The above plot shows the elasticity of the expenditure of RMS panelists in different deciles of size-adjusted income with respect to the errors in relative product quality estimated using the method outlined in Appendix Section E.1.3.

### E.1.4 Distribution of Price Elasticities Estimated in OLS and IV

Figure A.11: Distribution Price Coefficients Across Modules with Different Price Instruments

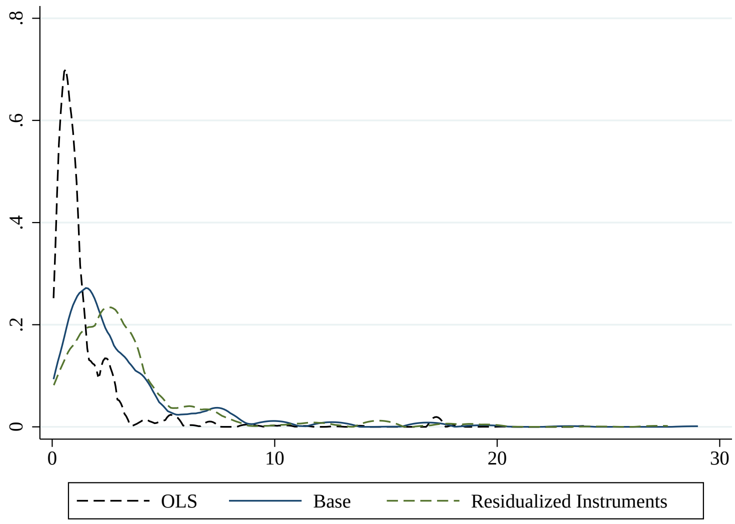


Notes: The above plot depicts the distribution of estimates of the module-level $\alpha_m^0$ parameters in the baseline model that allows for non-homotheticity in the demand for quality but not in price sensitivity (i.e., restricting that $\alpha_m^1$=0). The three kernel densities show the distribution of estimates obtained in OLS specification as well as instead using the two different price instruments described under Identification in Section 5.3.1 of the paper.

## E.2 Parameter Estimates

### E.2.1 Distribution of Non-Homotheticity Parameters Across Modules

Table A.5: Summary Statistics for Parameter Estimates with $abs(t - statistic) > 1.96$

| Model: | Homothetic | NH in Quality | | NH in Price | | NH in Quality and Price | | |
|---|---|---|---|---|---|---|---|---|
| Restrictions: | $\alpha_m^1 = 0$ & $\gamma_m = 0$ | $\alpha_m^1 = 0$ | | $\gamma_m = 0$ | | None | | |
| | [1] | [2] | [3] | [4] | [5] | [6] | [7] | [8] |
| Parameter: | $\alpha_m^0$ | $\alpha_m^0$ | $\gamma_m$ | $\alpha_m^0$ | $\alpha_m^1$ | $\alpha_m^0$ | $\alpha_m^1$ | $\gamma_m$ |
| Count | 375 | 463 | 466 | 421 | 417 | 433 | 405 | 438 |
| p25 | 1.09 | 0.89 | 0.38 | 0.72 | -0.28 | 1.04 | -1.44 | 1.35 |
| p50 | 1.73 | 1.71 | 1.16 | 1.36 | 0.15 | 1.69 | -0.82 | 1.99 |
| p75 | 2.45 | 2.41 | 2.52 | 2.06 | 0.80 | 2.50 | -0.23 | 2.74 |

Notes: These tables report the summary statistics for the main module-level parameter estimates governing the elasticity of substitution and non-homotheticities in demand. Attention is limited to modules for which the estimation procedure converged at interior estimates for all relevant parameters. The mean and percentile statistics are weighted by module sales in the Nielsen store-level data.

## Figure A.12: Distribution of $\gamma_m$ Parameter Estimates Across Modules

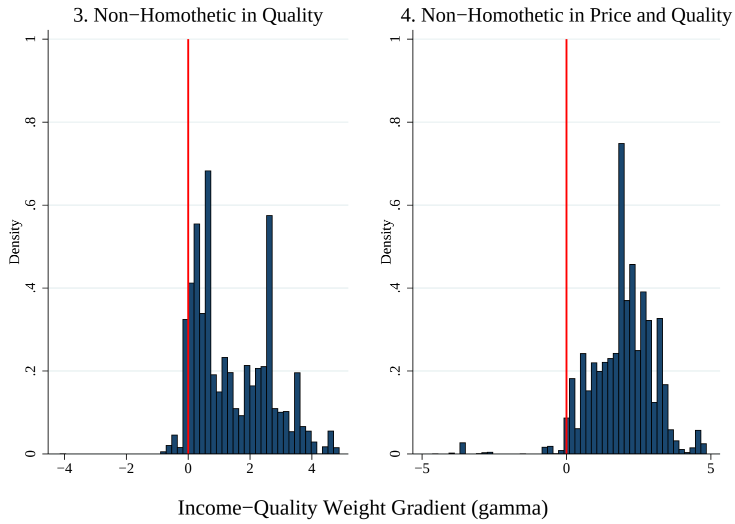


Notes: The plots above depict the distribution of the $\gamma_m$ estimates, for the model allowing for non-homotheticity in the demand for quality alone (i.e., restricting that $\alpha_m^1$=0) on the left and for the model allowing non-homotheticity in both the demand for quality and price sensitivity (i.e., allowing both $\gamma_m$ and $\alpha_m^1$ to be non-zero) on the right. Attention is limited to modules for which the estimation procedure converged at interior estimates for all parameters.

## Figure A.13: Distribution of $\alpha_m^1$ Parameter Estimates Across Modules

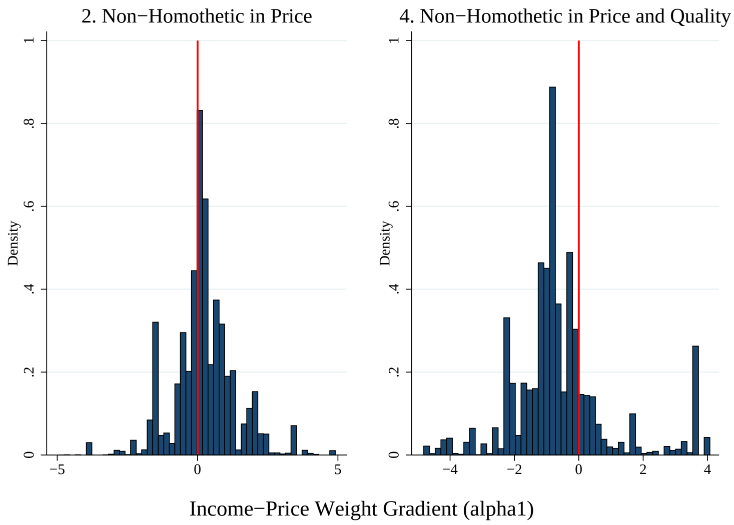


Notes: The plots above depict the distribution of the $\alpha_m^1$ estimates, for the model allowing for non-homotheticity in price sensitivity alone (i.e., restricting that $\gamma_m$ =0) on the left and for the model allowing non-homotheticity in both the demand for quality and price sensitivity (i.e., allowing both $\gamma_m$ and $\alpha_m^1$ to be non-zero) on the right. Attention is limited to modules for which the estimation procedure converged at interior estimates for all relevant parameters.

### E.2.2 Out-of-Sample Fit

The model is currently estimated using data describing sales in a sample of 5000 store-month markets for each product module. This leaves plenty of data to study the out-of-sample fit. The analysis below studies the out-of-sample fit for the baseline model used for the price index analysis (i.e., the model that allows non-homotheticity in the demand for quality, but not price sensitivity).

Figure A.14 compares the distribution of the unexplained component of store-month sales, which take the structural interpretation of transient taste shocks, in the estimation sample with that in a secondary sample of 5000 store-month markets for each product module. The two distributions–truncated at the 1st and 99th percentiles–are very similar to one another.

This fit is summarized in the J-statistics of the macro moments.[75] Figure A.15 compares the J-statistics calculated using the model estimates for $\alpha_m^0$ and $\gamma_m$ in the secondary sample to the J-statistics for the estimation sample. The average fit is, as expected, worse out-of-sample, but, barring some outliers, the fit of the macro moments is highly correlated across modules between the estimation and secondary samples.

[75] The store-month sampling procedure prioritizes store-months where HMS households are observed to make product purchases, so there is not a secondary sample of household purchases with which I can calculate out-of-sample micro moments.

Figure A.14: Transient Taste Shocks ($\xi_{mgst} - \beta_{mg}$) Predicted In-Sample and Out-of-Sample

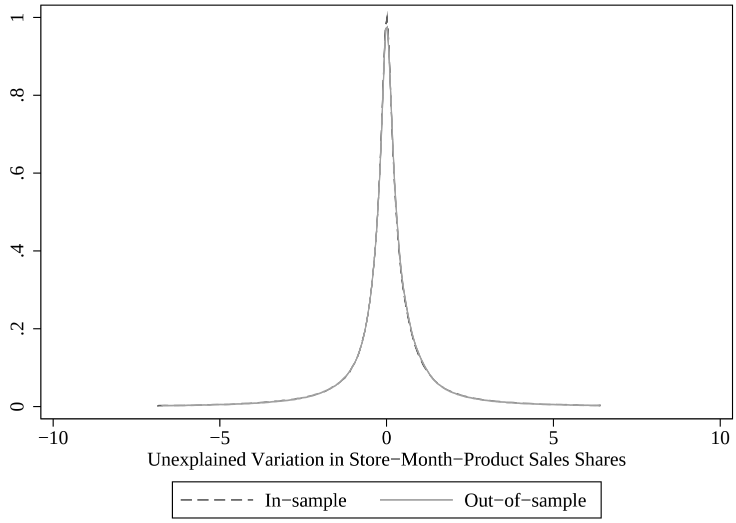


Notes: This plot shows the distribution of transient store-month tastes for products, estimated using sales in the base sample of 5000 store-month markets (in-sample) and then calculated using the same non-linear parameter estimates in a hold-back sample of 5000 different store-month markets (out-of-sample). This out-of-sample check is for the baseline model that allows for non-homotheticity in the demand for quality but not in price sensitivity (i.e., restricting that $\alpha_m^1$=0).

Figure A.15: J-Statistics for Store-Level Moments In-Sample and Out-of-Sample

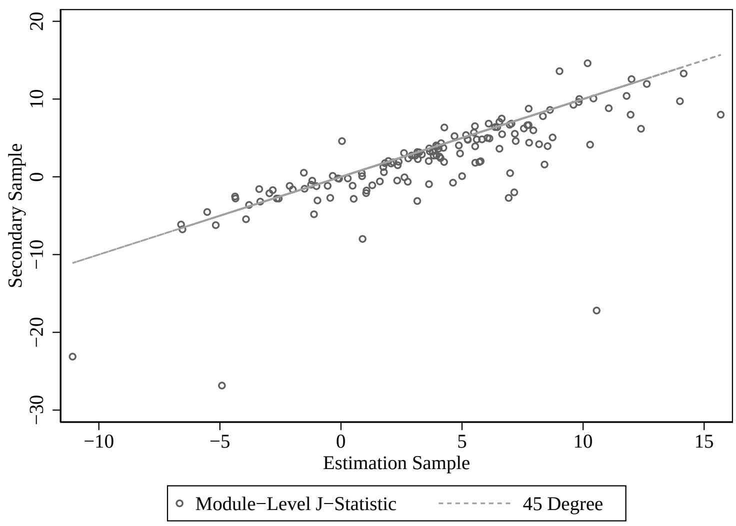


Notes: This plot compares the fit of the store-level moments estimated using sales in the base sample of 5000 store-month markets (the "estimation" sample) and then calculated using the same non-linear parameter values but for a hold back sample of 5000 different store-month markets (the "secondary" sample) across different modules. The fit of these moments in each sample is summarized with a module-level J statistic calculated with the weighting matrix and store-level moment conditions described above in Appendix Section D.2.1. This out-of-sample check is for the baseline model that allows for non-homotheticity in the demand for quality but not in price sensitivity (i.e., restricting that $\alpha_m^1$=0).

## E.3 Price Indexes

### E.3.1 Variation in Non-Parametric Price Index Results Across Bootstrap Samples

Figure A.16: Variation Across Bootstrap Samples in the Elasticity of Grocery Price Index with respect to CBSA Income for Households at Different Size-Adjusted Income Levels

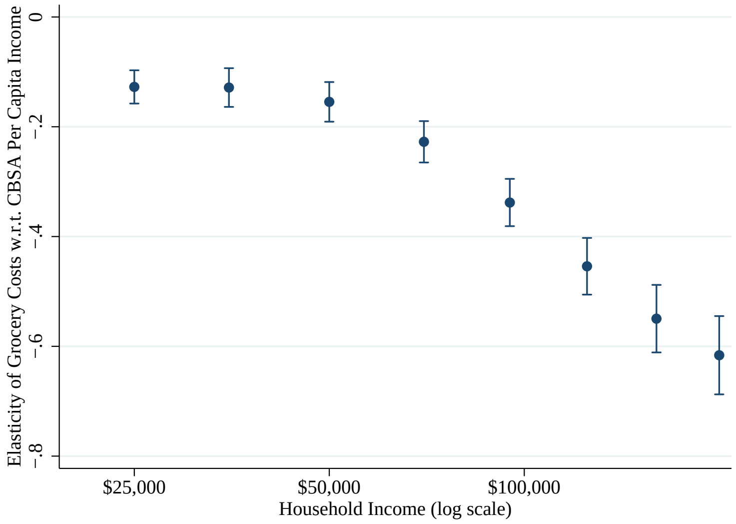


Notes: This plot shows the elasticity of income- and CBSA-specific price indexes with respect to CBSA per capita income for households at compares the different income levels. The point shows the mean elasticity estimated across 100 bootstrap iterations of price index calculations (each drawing a random sample of 50 stores in each CBSA) and the bands show the 95 percent confidence intervals around this mean. The price indexes are calculated using the baseline model that allows for non-homotheticity in the demand for quality but not in price sensitivity (i.e., restricting that $\alpha_m^1$=0).

Figure A.17: Variation Across Bootstrap Samples in the Elasticity of Grocery Price Index with respect to Household Income for CBSAs with Different Per Capita Income

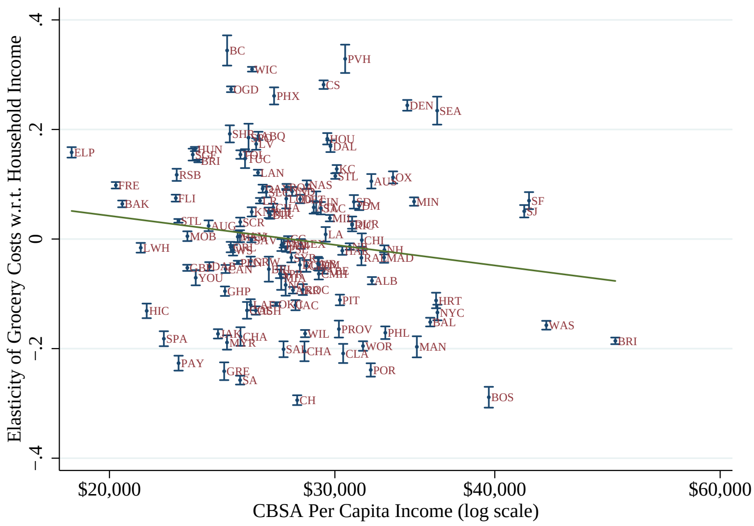


Notes: This plot shows the elasticity of household income- and CBSA-specific price indexes with respect to household income in CBSAs with different per capita incomes. The point shows the mean elasticity estimated across 100 bootstrap iterations of price index calculations (each drawing a random sample of 50 stores in each CBSA) and the bands show the 95 percent confidence intervals around this mean. The price indexes are calculated using the baseline model that allows for non-homotheticity in the demand for quality but not in price sensitivity (i.e., restricting that $\alpha_m^1$=0). The marker labels for each CBSA are acronyms linked to the full CBSA name in Appendix A.4.

### E.3.2 Price Indexes in High- vs. Low-Coverage CBSAs

Table A.6: City-Income Specific Price Index Regressions in High- and Low-Coverage CBSAs

| Dependent Variable: Ln(Price Index for Household in Income Group $k$ in CBSA $c$) | | | | | | |
|---|---|---|---|---|---|---|
| Sample: | All CBSAs | | High Coverage | | Low Coverage | |
| | [1] | [2] | [3] | [4] | [5] | [6] |
| Ln(Per Capita Income$_c$) | -0.32 | -0.30 | -0.81*** | -0.75** | -0.16 | -0.20 |
| | (0.21) | (0.20) | (0.20) | (0.23) | (0.35) | (0.40) |
| Ln(Per Capita Income$_c$)* Demeaned Ln(HH Income$_k$) | -0.25*** | -0.27*** | -0.29 | -0.30 | -0.31*** | -0.38*** |
| | (0.040) | (0.045) | (0.23) | (0.23) | (0.047) | (0.053) |
| Ln(Population$_c$) | | -0.0075 | | -0.037 | | 0.0088 |
| | | (0.047) | | (0.041) | | (0.076) |
| Ln(Population$_c$)* Demeaned Ln(HH Income$_k$) | | 0.0050 | | 0.0091 | | 0.019*** |
| | | (0.0031) | | (0.023) | | (0.0010) |
| Income Group $k$*Bootstrap Sample FEs | Yes | Yes | Yes | Yes | Yes | Yes |
| Number of CBSAs ($c$) | 125 | 125 | 28 | 28 | 44 | 44 |
| Observations | 98,920 | 98,920 | 22,400 | 22,400 | 34,656 | 34,656 |
| adj. $R^2$ | 0.03 | 0.03 | 0.14 | 0.14 | 0.03 | 0.03 |

Notes: *** p< 0.01, ** p<0.05, * p<0.10; standard errors, clustered by CBSA and bootstrap sample, are in parentheses. The table replicates columns [1] and [2] of Table A.7 from the main text using data for different samples of CBSAs. Columns [1] and [2] present results of the regression estimated in a sample containing 125 with 50 or more participating retailers. Columns [3] and [4] show results estimated in a sub-sample of these CBSAs that are identified as being located in DMAs where the Nielsen sample has high-coverage (accounts for over 50 percent of sales, on average across grocery, drug, and mass-merchandisers). Columns [5] and [6] show the results estimated on the sub-sample of CBSAs that are located in DMAs where the Nielsen has low-coverage (accounts for less than 50 percent of sales). Observations are weighted by CBSA population.

### E.3.3 Analysis of Price Indexes Based on Full Model with Non-Homotheticity in Price and Quality

Table A.7: City-Income Specific Price Index Regressions Accounting for Non-Homotheticity in Price and Quality

| Dependent Variable: Ln(Price Index for Household in Income Group $k$ in CBSA $c$) | | | | |
|---|---|---|---|---|
| | Local Prices | | National Prices | |
| | [1] | [2] | [3] | [4] |
| Ln(Per Capita Income$_c$) | -0.47 | -0.66* | -0.33 | -0.55 |
| | (0.29) | (0.29) | (0.29) | (0.29) |
| Ln(Per Capita Income$_c$)* Demeaned Ln(HH Income$_k$) | -0.45*** | -0.71*** | -0.44*** | -0.70*** |
| | (0.078) | (0.14) | (0.076) | (0.14) |
| Ln(Population$_c$) | | 0.058 | | 0.065 |
| | | (0.067) | | (0.067) |
| Ln(Population$_c$)* Demeaned Ln(HH Income$_k$) | | 0.079 | | 0.079 |
| | | (0.045) | | (0.045) |
| Income Group $k$*Bootstrap Sample FEs | Yes | Yes | Yes | Yes |
| Number of CBSAs ($c$) | 125 | 125 | 125 | 125 |
| Observations | 97,856 | 97,856 | 97,856 | 97,856 |
| adj. $R^2$ | 0.03 | 0.05 | 0.02 | 0.04 |

Notes: *** p< 0.01, ** p<0.05, * p<0.10; standard errors, clustered by CBSA, are in parentheses. The table replicates Table A.7 from the main text using price indexes that allow for non-homotheticities in both price and quality (i.e., allowing both $\alpha_m^1$ and $\gamma_m$ to take non-zero values).

### E.3.4 Price Indexes in Large CBSAs

## Table A.8: City-Specific Price Indexes for Cities with Population Over 1 Million

| CBSA Name | Per Capita Income | Population | Homothetic Price Index | Non-Homothetic Price Index at Household Income: $25,000 | $35,000 | $50,000 | $70,000 | $95,000 | $125,000 | $160,000 | $200,000 | $(P(150K) - P(25K))$ /stdev(P) |
|---|---|---|---|---|---|---|---|---|---|---|---|---|
| Atlanta-Sandy Springs-Roswell, GA | $28,880 | 5,455,053 | 1.43 | 1.34 | 1.44 | 1.51 | 1.54 | 1.57 | 1.58 | 1.58 | 1.57 | 81.35% |
| Austin-Round Rock, TX | $32,035 | 1,835,016 | 0.96 | 0.89 | 0.87 | 0.87 | 0.90 | 0.96 | 1.02 | 1.08 | 1.12 | 85.26% |
| Baltimore-Columbia-Towson, MD | $35,613 | 2,753,396 | 1.11 | 1.10 | 1.16 | 1.18 | 1.13 | 1.04 | 0.96 | 0.89 | 0.85 | -90.33% |
| Birmingham-Hoover, AL | $26,706 | 1,135,534 | 0.95 | 0.93 | 0.91 | 0.89 | 0.90 | 0.92 | 0.95 | 0.99 | 1.03 | 35.43% |
| Boston-Cambridge-Newton, MA-NH | $39,572 | 4,650,876 | 0.96 | 1.09 | 1.09 | 1.04 | 0.94 | 0.84 | 0.76 | 0.70 | 0.67 | -150.93% |
| Buffalo-Cheektowaga-Niagara Falls, NY | $28,171 | 1,135,667 | 0.53 | 0.65 | 0.59 | 0.55 | 0.54 | 0.54 | 0.56 | 0.58 | 0.60 | -20.20% |
| Charlotte-Concord-Gastonia, NC-SC | $28,403 | 2,298,915 | 0.99 | 1.14 | 1.15 | 1.11 | 1.02 | 0.94 | 0.88 | 0.84 | 0.82 | -115.15% |
| Chicago-Naperville-Elgin, IL-IN-WI | $31,488 | 9,516,448 | 1.15 | 0.97 | 1.00 | 1.02 | 1.02 | 1.01 | 1.00 | 1.00 | 1.00 | 13.45% |
| Cincinnati, OH-KY-IN | $29,008 | 2,131,793 | 1.01 | 0.99 | 1.00 | 1.02 | 1.07 | 1.12 | 1.16 | 1.17 | 1.17 | 66.67% |
| Cleveland-Elyria, OH | $28,499 | 2,067,490 | 0.82 | 0.85 | 0.85 | 0.85 | 0.84 | 0.83 | 0.81 | 0.80 | 0.78 | -24.54% |
| Columbus, OH | $29,145 | 1,948,188 | 1.07 | 1.12 | 1.11 | 1.09 | 1.08 | 1.06 | 1.04 | 1.02 | 1.00 | -44.02% |
| Dallas-Fort Worth-Arlington, TX | $29,766 | 6,703,020 | 1.26 | 1.07 | 1.12 | 1.19 | 1.27 | 1.36 | 1.43 | 1.49 | 1.52 | 163.80% |
| Denver-Aurora-Lakewood, CO | $34,173 | 2,651,392 | 1.48 | 1.09 | 1.19 | 1.31 | 1.44 | 1.58 | 1.69 | 1.77 | 1.81 | 260.97% |
| Detroit-Warren-Dearborn, MI | $28,182 | 4,292,647 | 0.95 | 0.95 | 0.95 | 0.97 | 1.00 | 1.03 | 1.06 | 1.09 | 1.10 | 54.50% |
| Grand Rapids-Wyoming, MI | $25,786 | 1,007,329 | 0.69 | 0.83 | 0.75 | 0.71 | 0.70 | 0.71 | 0.72 | 0.73 | 0.75 | -28.10% |
| Hartford-West Hartford-East Hartford, CT | $35,991 | 1,215,159 | 0.50 | 0.63 | 0.55 | 0.51 | 0.49 | 0.48 | 0.49 | 0.50 | 0.51 | -43.49% |
| Houston-The Woodlands-Sugar Land, TX | $29,594 | 6,204,141 | 1.20 | 1.02 | 1.06 | 1.12 | 1.22 | 1.32 | 1.40 | 1.46 | 1.48 | 167.91% |
| Indianapolis-Carmel-Anderson, IN | $27,778 | 1,931,182 | 1.15 | 1.16 | 1.19 | 1.21 | 1.24 | 1.29 | 1.34 | 1.38 | 1.40 | 85.05% |
| Jacksonville, FL | $27,950 | 1,380,995 | 1.30 | 1.32 | 1.30 | 1.24 | 1.16 | 1.11 | 1.08 | 1.08 | 1.08 | -87.80% |
| Kansas City, MO-KS | $30,101 | 2,040,869 | 1.31 | 1.29 | 1.35 | 1.41 | 1.47 | 1.55 | 1.62 | 1.67 | 1.70 | 147.24% |
| Las Vegas-Henderson-Paradise, NV | $26,040 | 2,003,613 | 1.57 | 1.13 | 1.25 | 1.35 | 1.44 | 1.52 | 1.58 | 1.63 | 1.65 | 191.93% |
| Los Angeles-Long Beach-Anaheim, CA | $29,506 | 13,060,534 | 1.13 | 0.95 | 1.01 | 1.03 | 1.03 | 1.02 | 1.02 | 1.03 | 1.03 | 29.27% |
| Louisville/Jefferson County, KY-IN | $27,488 | 1,253,305 | 0.95 | 0.97 | 0.98 | 1.00 | 1.04 | 1.09 | 1.14 | 1.17 | 1.19 | 77.94% |
| Memphis, TN-MS-AR | $25,191 | 1,337,014 | 0.94 | 0.92 | 0.92 | 0.93 | 0.95 | 0.97 | 0.97 | 0.95 | 0.92 | 2.43% |
| Miami-Fort Lauderdale-West Palm Beach, FL | $27,240 | 5,775,204 | 0.78 | 0.78 | 0.73 | 0.69 | 0.68 | 0.69 | 0.70 | 0.72 | 0.74 | -16.65% |
| Milwaukee-Waukesha-West Allis, WI | $29,733 | 1,565,368 | 1.06 | 1.20 | 1.21 | 1.19 | 1.18 | 1.20 | 1.24 | 1.28 | 1.31 | 41.63% |
| Minneapolis-St. Paul-Bloomington, MN-WI | $34,593 | 3,424,786 | 1.26 | 1.24 | 1.28 | 1.29 | 1.29 | 1.31 | 1.36 | 1.42 | 1.47 | 83.26% |
| Nashville-Davidson–Murfreesboro–Franklin, TN | $28,521 | 1,730,515 | 1.19 | 1.19 | 1.23 | 1.26 | 1.30 | 1.34 | 1.40 | 1.44 | 1.47 | 104.53% |
| New Orleans-Metairie, LA | $27,458 | 1,226,440 | 0.74 | 0.79 | 0.74 | 0.71 | 0.69 | 0.69 | 0.69 | 0.70 | 0.71 | -30.09% |
| New York-Newark-Jersey City, NY-NJ-PA | $36,078 | 19,865,045 | 0.78 | 0.84 | 0.84 | 0.83 | 0.80 | 0.76 | 0.72 | 0.69 | 0.68 | -60.16% |
| Oklahoma City, OK | $26,994 | 1,297,998 | 0.54 | 0.69 | 0.62 | 0.56 | 0.52 | 0.51 | 0.51 | 0.53 | 0.54 | -55.86% |
| Orlando-Kissimmee-Sanford, FL | $24,876 | 2,226,835 | 0.93 | 0.97 | 0.93 | 0.89 | 0.88 | 0.88 | 0.90 | 0.93 | 0.96 | -6.19% |
| Philadelphia-Camden-Wilmington, PA-NJ-DE-MD | $32,850 | 6,015,336 | 1.22 | 1.14 | 1.23 | 1.26 | 1.20 | 1.09 | 0.99 | 0.92 | 0.87 | -99.47% |
| Phoenix-Mesa-Scottsdale, AZ | $26,893 | 4,337,542 | 1.67 | 1.09 | 1.20 | 1.33 | 1.47 | 1.62 | 1.76 | 1.88 | 1.96 | 315.24% |
| Pittsburgh, PA | $30,272 | 2,358,793 | 0.86 | 0.92 | 0.90 | 0.87 | 0.84 | 0.80 | 0.77 | 0.76 | 0.75 | -62.37% |
| Portland-Vancouver-Hillsboro, OR-WA | $30,560 | 2,288,796 | 1.37 | 1.05 | 1.13 | 1.26 | 1.45 | 1.68 | 1.91 | 2.11 | 2.25 | 435.68% |
| Providence-Warwick, RI-MA | $30,218 | 1,604,317 | 0.85 | 0.92 | 0.89 | 0.85 | 0.81 | 0.76 | 0.72 | 0.69 | 0.68 | -86.44% |
| Raleigh, NC | $31,468 | 1,189,579 | 1.14 | 1.16 | 1.21 | 1.23 | 1.23 | 1.21 | 1.18 | 1.16 | 1.15 | -3.55% |
| Richmond, VA | $30,944 | 1,234,058 | 1.15 | 1.12 | 1.15 | 1.19 | 1.22 | 1.23 | 1.22 | 1.21 | 1.19 | 24.56% |
| Riverside-San Bernardino-Ontario, CA | $22,571 | 4,345,485 | 1.41 | 1.14 | 1.24 | 1.28 | 1.28 | 1.32 | 1.39 | 1.46 | 1.53 | 142.41% |
| Rochester, NY | $28,320 | 1,082,578 | 0.65 | 0.81 | 0.72 | 0.66 | 0.64 | 0.64 | 0.64 | 0.65 | 0.66 | -51.86% |
| Sacramento–Roseville–Arden-Arcade, CA | $29,252 | 2,197,422 | 0.91 | 0.89 | 0.87 | 0.86 | 0.88 | 0.91 | 0.95 | 0.98 | 1.02 | 46.15% |
| Salt Lake City, UT | $26,516 | 1,123,643 | 1.26 | 1.26 | 1.38 | 1.44 | 1.44 | 1.46 | 1.50 | 1.54 | 1.57 | 113.73% |
| San Antonio-New Braunfels, TX | $25,298 | 2,239,222 | 0.52 | 0.68 | 0.58 | 0.50 | 0.44 | 0.41 | 0.40 | 0.40 | 0.41 | -99.49% |
| San Diego-Carlsbad, CA | $31,043 | 3,183,143 | 1.14 | 0.93 | 0.99 | 1.03 | 1.04 | 1.06 | 1.09 | 1.11 | 1.13 | 71.53% |
| San Francisco-Oakland-Hayward, CA | $42,540 | 4,466,251 | 0.79 | 0.79 | 0.75 | 0.74 | 0.75 | 0.78 | 0.84 | 0.89 | 0.94 | 57.28% |
| San Jose-Sunnyvale-Santa Clara, CA | $42,176 | 1,898,457 | 0.68 | 0.76 | 0.71 | 0.69 | 0.69 | 0.72 | 0.76 | 0.81 | 0.85 | 36.16% |
| Seattle-Tacoma-Bellevue, WA | $36,061 | 3,557,037 | 1.27 | 1.01 | 1.06 | 1.14 | 1.27 | 1.42 | 1.57 | 1.70 | 1.79 | 283.27% |
| St. Louis, MO-IL | $30,024 | 2,797,737 | 0.89 | 0.85 | 0.83 | 0.85 | 0.89 | 0.94 | 0.99 | 1.03 | 1.05 | 73.56% |
| Tampa-St. Petersburg-Clearwater, FL | $27,252 | 2,851,235 | 0.97 | 1.00 | 0.98 | 0.97 | 0.97 | 0.97 | 0.98 | 0.99 | 1.00 | 0.99% |
| Virginia Beach-Norfolk-Newport News, VA-NC | $29,098 | 1,697,898 | 1.34 | 1.13 | 1.20 | 1.26 | 1.28 | 1.25 | 1.19 | 1.14 | 1.09 | -15.75% |
| Washington-Arlington-Alexandria, DC-VA-MD-WV | $43,884 | 5,863,608 | 1.05 | 1.16 | 1.21 | 1.20 | 1.14 | 1.06 | 0.98 | 0.92 | 0.89 | -99.39% |

Note: This table shows the price indexes calculated for the sample cities with more than 1 million residents reported in the 2010-2014 ACS data. The homothetic index is calculated using the model parameters estimated when restricting $\alpha_m^1$ and $\gamma_m$ to equal zero. The non-homothetic indexes are calculated using the model parameters estimated for the preferred model that allows non-homotheticity in the demand for quality but not in price (i.e., allowing $\gamma_m \neq 0$ but restricting $\alpha_m^1 = 0$).

### E.3.5 Variation in Price Indexes based on Alternative Specifications

Figure A.18: Robustness: Elasticity of Grocery Price Index with respect to CBSA Income for Households at Different Size-Adjusted Income Levels

(a) Feenstra Estimates

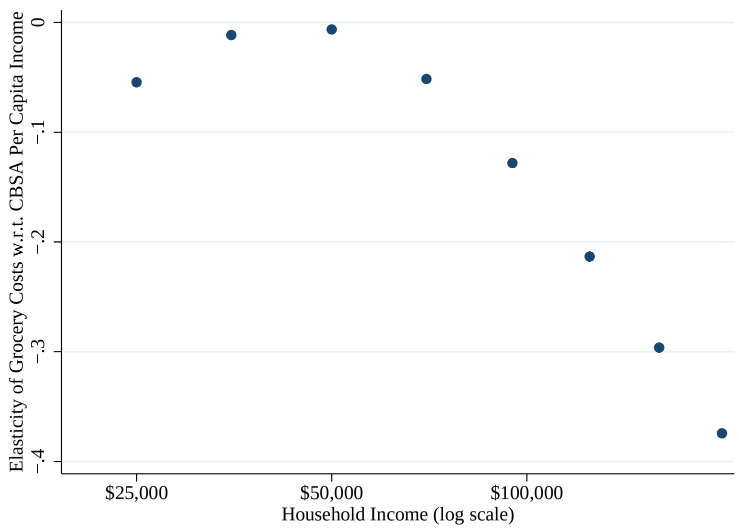


(b) Residualized Instruments

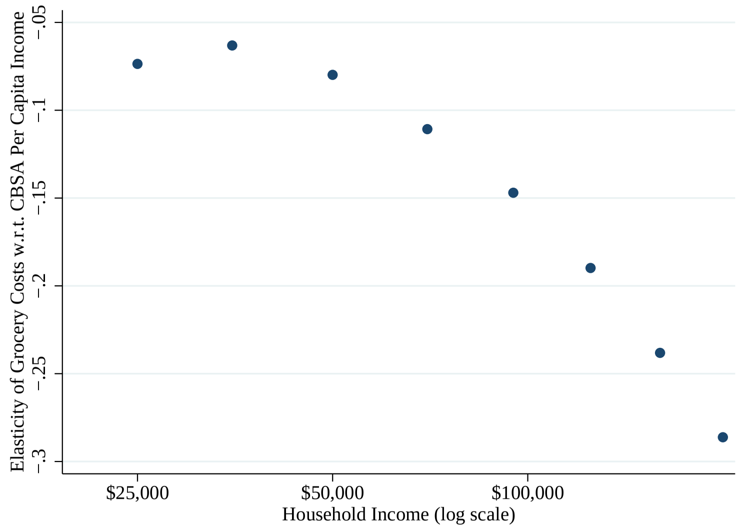


(c) Quarterly Frequency Data

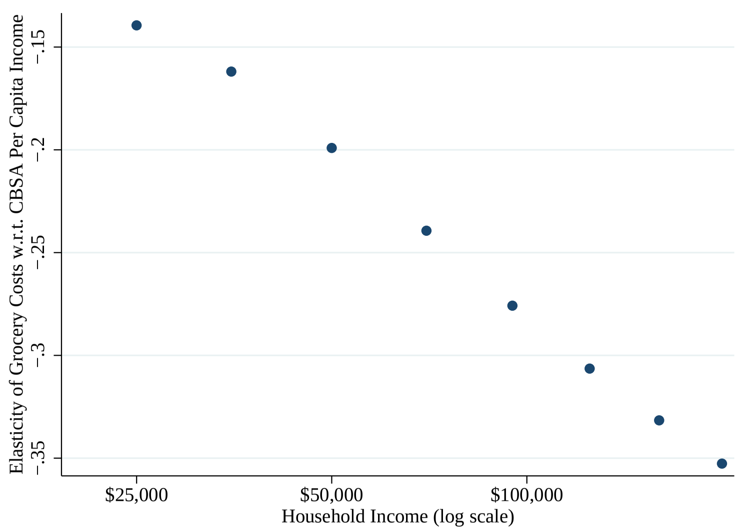


(d) 40th Percentile Outside Product Cutoff

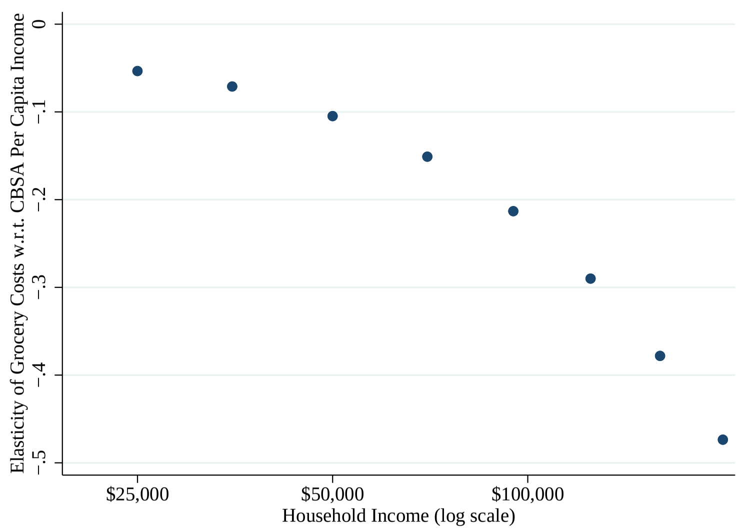


Notes: This figure replicates Figure 6 from the main text but for price indexes calculated using different parameter estimates. Each panel plots the elasticity of a grocery price index for households at different income levels across different CBSAs with respect to CBSA per capita income. The grocery price indexes are calculated using the structural price index formula discussed in the text, using parameter estimates obtained under the two-step routine outlined in Appendix D.4 in plot (a), and the main estimation procedure but with the residualized price instrument, quarterly instead of monthly data, and a 40th percentile instead of 60th percentile outside product cutoff in plots (b), (c), and (d) respectively. The elasticities plotted are calculated by regressing the resulting household income- and CBSA-specific log grocery price indexes against household income fixed effects and these fixed effects interacted with log CBSA per capita income. Observations in these regressions are weighted by CBSA population.