Are Poor Cities Cheap for Everyone?
Non-Homotheticity and the Cost of Living Across U.S. Cities*

Jessie Handbury†

December 17, 2012

Abstract

Standard cost-of-living indexes assume that preferences are homothetic, ignoring the well-established fact that tastes vary with income. This paper considers how assuming homotheticity biases our estimates of spatial price indexes for consumers at different income levels. I use Nielsen household-level purchase data in over 500 categories of food products to calculate micro-founded income- and city-specific price indexes that account for non-homotheticity, as well as city-specific price indexes that do not. I find that the income-specific cross-city price indexes vary widely across income groups. Grocery costs are 20 percent lower in a poor city relative to a wealthy city for a low-income household, but they are 20 percent higher in the poor city for a high-income household. The homothetic price indexes perform well in predicting the cross-city variation in prices for low- and middle-income households, but poorly for high-income households. These results suggest that using homothetic cost-of-living indexes understate the relative price level in poor locations for rich households.

*I am grateful to my advisors David Weinstein and Kate Ho for their guidance and encouragement throughout. Further thanks to David Bieri, Chris Conlon, Don Davis, Jonathan Dingel, Gilles Duranton, Galina Hale, Amit Khandelwal, Joan Monras, Eleonora Patacchini, Mike Riordan, Molly Schnell, Ina Simonovska, Todd Sinai, Daniel Sturm, Eric Verhoogen, and Jonathan Vogel, as well as numerous seminar and conference participants, for valuable discussions and comments. I gratefully acknowledge financial support from the Program of Economic Research at Columbia. All errors are my own.

†The Wharton School, University of Pennsylvania 3620 Locust Walk, 1463 Steinberg-Dietrich Hall, Philadelphia PA 19094-6302, handbury@wharton.upenn.edu.

1 Introduction

Cost-of-living indexes are central to measuring real incomes and expenditures of households living in different locations. The price indexes commonly used in this analysis provide only one number to represent the relative cost of living across locations for all consumers and, thus implicitly assume that tastes do not vary systematically across consumers.[1] Numerous household-level studies in multiple countries have shown that this assumption is violated in one important respect: tastes vary with consumer income.[2]

This paper measures the extent to which assuming homotheticity biases our estimates of spatial price indexes for consumers at different income levels. To measure these biases, I first develop a utility framework that characterizes non-homothetic preferences across many sectors of differentiated products. I use data with household-level purchases in over 500 grocery categories to structurally estimate to the parameters of models that allow for correlations between a household's income and their demand for quality, their price sensitivity, neither, or both, demonstrating the salience of non-homothetic demand for quality in this context. I use the estimates of the model that allows for this form of non-homotheticity to calculate price indexes characterizing how the grocery component of the cost of living varies across cities in the U.S. differently for consumers at different income levels. This analysis yields three main results.

First, I find that there are large differences in how high- and low-income households perceive the price levels across U.S. cities. For example, a low-income household earning \$15,000 a year faces approximately 20 percent higher grocery costs in cities with relatively high per capita income like San Francisco relative to cities with half that per capita income, such as New Orleans. But the exact opposite is true for high-income households earning \$100,000 a year. Their grocery costs are 20 percent lower in the city with the higher per capita income.

Second, I show that these differences are related to cross-city variation in product variety, rather than prices. High-income households are better off in wealthier cities because more varieties of the high-quality products that high-income consumers prefer to consume are available in these locations.

Finally, I find that a standard homothetic price index does a better job of predicting the distribution of costs across locations for low- and middle-income households than it does for high-income households. The homothetic price index is highly correlated with the grocery costs I have measured for households with incomes below \$70,000, but negatively correlated with the grocery costs for households with incomes above \$100,000. The homoethic index systematically underestimates the costs faced by high-income consumers in wealthy relative to poor cities.

[1] Notable exceptions include Deaton and Dupriez (2011) 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.

[2] The prevalence of this result was noted in Deaton and Muellbauer (1980). More recent direct evidence includes Bils and Klenow (2001) and Broda, Liebtag, and Weinstein (2009) for the U.S. and Li (2012) for India.

My first two results contribute to a large urban economics literature that seeks to understand why large cities are more skilled.[3] A number of papers explain that high-skill workers co-locate in large cities because they enjoy greater productivity spillovers than low-skill workers (see, *e.g.*, Glaeser and Mare (2001); Wheeler (2001); Davis and Dingel (2012)). A complementary explanation focuses on the role of these workers as consumers: high-skill, high-income consumers enjoy more utility from urban consumption amenities than low-skill, low-income consumers. Since consumers with similar income levels have similar tastes, they sort geographically by income in order to maximize the local production of the products that suit their tastes. Papers highlighting the role of non-homotheticities in skill-biased agglomeration have previously relied on evidence showing that the city-size wage premium is lower for high-skill workers than for low-skill workers (Adamson, Clark, and Partridge, 2004; Lee, 2010; Black, Kolesnikova, and Taylor, 2009). This evidence indicates that productivity spillovers are higher for low-skill workers; placing the consumption- and production-side explanations for skill-biased agglomeration at odds with one another. My first result instead provides direct evidence of the consumption-side story that is consistent with the widely-held and empirically-founded belief that skill-biased productivity spillovers play a role here too.

The idea that both the nominal wage and consumption benefits of large, skilled cities accrue disproportionately to high-skilled workers is consistent with a spatial equilibrium framework as long as some urban disamenity or congestion cost is disproportionately borne by high-income earners. Diamond (2012) develops such a model, providing complementary, indirect evidence of skill-biased consumption spillovers. Her framework allows for locations to have an endogenous amenity that is increasing in the local college employment ratio and she shows empirically that high-skill workers are more willing to pay for this amenity, in terms of wages and housing costs, than low-skill workers. My work suggests that these endogenous skill-biased amenities are, at least in part, due to preference externalities. Other research supporting this idea includes Waldfogel (2003) and Fajgelbaum, Grossman, and Helpman (2011), who show that, in markets with increasing returns and demand heterogeneity, differentiated product firms cater to the prevalent tastes in a market such that the composition of demand in a location impacts the value of consuming there.[4] The results here provide the first structural estimates of these pecuniary consumption externalities.

My final result is particularly relevant to economists studying the welfare implications of spatial price variation. Recent work, for example, has highlighted biases that arise from ignoring spatial price variation in the measurement of key statistics on poverty, income inequality, and tax incidence (Almås, 2012;

[3]This pattern has been well-documented by, *e.g.*, Combes, Duranton, and Gobillon (2008); Bacolod, Blum, and Strange (2009); Glaeser and Resseger (2010).

[4]This distribution of product prices and availability is also consistent with a comparative advantage story, where skilled, high-income workers are more productive in the production of high-quality goods, independent of the fact that they like to consume these goods. I do not differentiate between these two stories in my analysis. The utility effects I estimate should, therefore, be interpreted as the combined effect of the consumption spillovers and productivity advantages that person brings to a city.

Moretti, forthcoming; Deaton, 2010; Albouy, 2009). These authors use standard homothetic price indexes to account for cost-of-living differences across locations, implicitly ignoring that households with different incomes have different tastes and, therefore, may perceive these relative costs differently. Although I find these differeces to be large, I also show that a homothetic price index represents the spatial variation in prices quite closely for households at all points but the upper tail of the income distribution. This result is consistent with Deaton and Dupriez (2011) who find that reweighting the International Comparison Project (ICP)'s purchasing power parity (PPP) indexes to reflect the consumption pattens of the world's poor does not change the indexes or, therefore, poverty counts dramatically. Taken together, this evidence suggests that it might be reasonable 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)). That said, my results indicate that it is necessary to account for income-specific tastes when measuring the real incomes and expenditures of high-income households living in different locations. Recent work by Albouy (2009) and Moretti (forthcoming) demonstrates how ignoring intranational price variation biases measures of real income inequality and the geographic distribution of real tax expenditures in the U.S..[5] My results suggest that income-specific consumption externalities mitigate these biases. If consumers with high nominal incomes find the wealthy cities to be less expensive than the poor cities that the remainder of the population finds to be cheap, then they will have higher real income and face a lower real tax burden in these locations.[6]

The main methological challenge I overcome in this paper is in empirically summarizing the costs that consumers face across many product categories in a way that parsimoniously accounts for the non-homothetic tastes observed in the micro-level household behavior. To do this, I draw from two literatures. I develop a framework that nests a log-logit utility function, similar to those used in empirical industrial organization (IO) to model how within-sector differentiated product demand varies with a range of consumer characteristics, including income, in a constant elasticity of substitution (CES) demand system, commonly used by international trade and macroeconomists to model representative agent demand across many categories of differentiated products. The log-logit utility function governs how consumers allocate expenditures between products within product categories, while a CES superstructure governs the substitutability across products in different categories. These functional form assumptions impose strong restrictions on consumer behavior that I will discuss in more detail below, but they also make the analysis below more tractable.[7]

[5] In the same spirit but an international context, Almås (2012) accounts for international price variation in the measurement of worldwide real income inequality.

[6] Complementary work studies how inter-temporal price indexes, or inflation, varies with income and have also found large variation across income levels. Broda and Romalis (2009) use the same Nielsen dataset, but a different methodology, to calculate income-specific U.S. inflation indexes and find that half of the increase in conventional measures of U.S. income inequality was due to a bias caused by ignoring the variation in consumption behavior across income groups.

[7] In particular, 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, de Palma, and Thisse

To add non-homotheticity to this framework, I borrow from a recent international trade literature that seeks to understand how representative agent demand varies with income across many categories of differentiated products. The model is non-homothetic because the amount that consumers care about the price of the product they are purchasing and the quality of its brand depends on their expenditure on a composite of outside, non-food products that I assume to be normal.[8] 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 grocery products generally and, in addition, spend more on those products that are ranked as high quality by all consumers. As in Hummels and Lugovskyy (2009), high-income households are less price elastic be they have a stronger attachment to their ideal variety and are, therefore, more willing to pay for it. As in Fajgelbaum, Grossman, and Helpman (2011), there are complementarities between product quality and outside good expenditure so, while all households agree on the quality ranking of products, high-income households get relatively more utility from and are, therefore, more willing to pay for higher quality products. Income and product quality enter into the consumer demand function in the same way as in Hallak (2006) and Feenstra and Romalis (2012), who calculate cross-country price indexes similar to those estimated here.[9]

This paper proceeds as follows. In Section 2, I introduce the dataset. In Section 3, I outline the model I use to estimate preferences. In Section 4, I outline the procedures used to estimate model parameters and demonstrate how I use the parameters to measure relative welfare across markets. I estimate four different four different models, the first allowing for non-homotheticity in both quality and price sensitivity, the second and third allowing for just one form of non-homotheticity, or the final homothetic version of the model. In Section 5, I present the parameter estimates for each of these models, the model selection criteria I use to determine which of these models explains the observed consumption behavior in the most parsimonious way, and finally my analysis of how the cross-city price indexes implied by this model vary with income.

2 Data

The results in this paper are based on analysis of detailed household consumption data from the Nielsen HomeScan database. This data includes all food product purchases in grocery, drug, mass merchandise,

(1987), whose proof was extended to models that account for product quality in Verhoogen (2008) and again to this context in Appendix A below.

[8]Evidence shows that 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.

[9]The Feenstra and Romalis (2012) indexes non-homothetic in the sense that they are based on parameter estimates from a non-homothetic utility function. They are not, however, income-specific and are instead calculated to reflect the indirect utility of a single representative agent, whose taste for any available product is equal to the taste of a consumer with the mean income of the countries that consume the product.

and other stores for a demographically representative, but unbalanced, panel of over 40,000 households in 52 markets across the United States between 2003 and 2005. The households in the sample were provided with barcode scanners and instructed to collect information such as the Universal Product Code (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, their income category, the number of members, the ages of all members, and the occupation and education levels of the female and male head of household.

There are around 400,000 UPCs represented in the sample. Nielsen categorizes UPCs into 640 modules and 65 groups. The dataset includes these categorizations as well as detailed data on the brand, size (including units), container, flavor, form, formula, variety, style, organic seal, and salt content of the UPC. Of the 640 modules included in the grocery database, 46 are for random weight items and are excluded from the analysis.[10] I determine the manufacturer of each UPC by matching the first 7 digits of the UPC code with a list of manufacturers downloaded from www.upcdatabase.com.

The data includes the exact price paid by each household for each of the UPCs they purchase. Combined with the prices paid for UPCs by other households in the same market, this data is useful both for constructing market-level expenditure shares for estimation and for defining the choice set used to measure city-level grocery costs. This dataset is uniquely suited for estimating a non-homothetic utility function because it links the characteristics of the UPCs a household purchases with the demographics of the household. I discuss how I use the Nielsen data on product characteristics and household demographics below.

For the structural demand estimation, I aggregate UPCs into a broader level of classification that I call a "product." One product identifier is assigned to each set of UPCs within a product module with the same brand, manufacturer, container size, salt content of a product, diet and organic categorization, and number of containers sold in a pack (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 multipack). For example, in the product module "SOFT DRINKS - CARBONATED", there are 15 UPCs that refer to non-diet, non-organic, regular salt, and single-pack 12 ounce containers sold under the brand "COCA-COLA CLASSIC R" that are produced at "COCA-COLA USA OPERATIONS." Products are defined such that this set of UPCs belong to the same product.

The utility function presented below assumes that consumers do not differentiate between UPCs in the same product. Since firms use UPCs to monitor their distribution and sales, otherwise identical products might have different UPCs because, for example, they are distributed through different channels. It is appropriate to assume that, conditional on price, consumers do not differentiate between these products. The assumption is stronger in cases where different UPCs that I have defined to be the same product

[10]One source of bias in the parameter estimates is unobserved correlation between the component of product quality that varies across markets or time and the prices at which these products are sold. 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.

are differentiated by their label or flavor.

Table 1 shows summary statistics for the sample used for estimation. This sample has been cleaned to control for data recording errors[11] and contains 538 product modules. There are between 6 and 5,284 products in each module, and the median number of products per module is 190. The median number of UPCs per module at 289. Although there are much fewer products in each module than there are UPCs, the typical product only contains one UPC.[12]

For the purposes of this paper, the most important demographic information contained in the dataset relates to household income. Nielsen classifies households into 16 categories based on annual income. I drop the households in the three categories with annual income under $10,000 and adjust the income of the remaining households for the number of household members. To adjust for household size, I first assign a numerical value to the income of each household. This income variable is equal to the mid-point of the bounds of a household's income category when both bounds exist and is equal to $150,000 for those households in the "above $100,000" income category. To adjust this income variable for the number of members in each household, I regress log household income against dummies for the ages, years of education, marital status, and race of the female and male heads-of-household, as well as fixed effects for every level of household size. I subtract the estimated household size fixed effects from log household income and add back the one-member household fixed effect to all observations to get a projection of the log household income of each household if it were to have only one member. The distribution of the mean-size household income projections are shown in Figure 1. The bulk of the distribution is between $10,000 and $80,000, which seems reasonable given that the per capita incomes of the MSAs represented in the sample range from $21,446 in New Orleans to $54,191 in San Francisco.

The Nielsen data contains information on grocery purchases of over 40,000 households at various times in 2003, 2004, and 2005. I only use a subsample of these households in the analysis below. This sample excludes any households not included from the income adjustment regression discussed above, either because they have low reported incomes or are missing demographic data. To estimate demand, I consider the quarterly expenditure shares of households grouped by size-adjusted income and market. To ensure that the purchases included in these expenditure shares are representative of households' expenditures over the entire quarter, I also exclude households in any quarter if they do not report over a period of over two weeks in the quarter.

[11] I drop any purchase observation for which the price paid for a UPC was greater than three times or less than a third of the median price paid per unit of any UPC within the same product categorization. I also drop any purchase observations for which the price paid for a UPC was greater than three times or less than a third of the median price paid per unit of any UPC within the same product module. I also limit the sample to products that are purchased by 20 or more households.

[12] To check the extent to which consumers differentiate between UPCs within product categories, I compare the coefficient of variation for the unit value paid for each UPC with the coefficient of variation for the unit value paid for the set of UPCs with the same product categorization. The median UPC-level coefficient of variation is 0.14, which is only slightly lower than the median product-level coefficient of variation at 0.15. This indicates that there is little variation in the prices charged for UPCs within the same product.

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

Total Count

Quarters	12
Metropolitan Statistical Areas (MSAs)	49
Modules	538
Brands	12,194
Products	181,072
Universal Product Codes (UPCs)	318,825

Count of Unique UPCs Per Category

Category	Minimum	Median	Maximum
Module	6	289	9,464
Brand	1	2	431
Product	1	1	153

Count of Unique Products Per Category

Category	Minimum	Median	Maximum
Module	6	190	5,284
Brand	1	2	135

Coefficient of Variation of Unit Price Paid for a UPC

	Minimum	Median	Maximum
Within the same module	0.15	0.39	0.72
Within the same brand	0.00	0.20	1.35
Within the same product	0.00	0.15	1.34
Within the same UPC	0.00	0.14	0.98

Figure 1: Distribution of Size-Adjusted Household Income

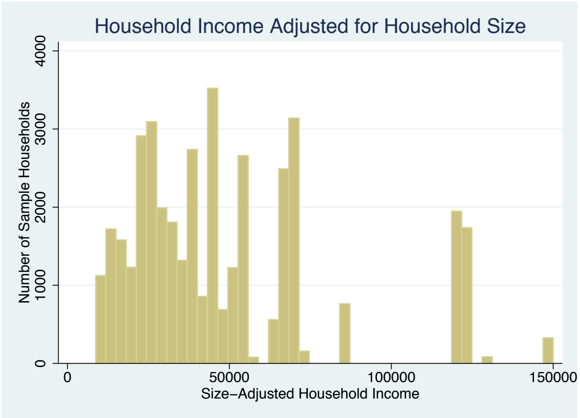

Table 2 shows how the resulting sample of 32,553 households are distributed across U.S. cities, as well as the population and per capita income for each of these locations from the 2000 U.S. census.[13] Although per capita income is correlated with population across the sample MSAs, three of the ten largest cities in the sample (Los Angeles, Detroit, and Chicago) and three of the ten smallest (Des Moines, Omaha, and Richmond) have similar per capita incomes, between \$35,000 and \$40,000. This variation will help to separately identify whether the observed variation in household income- and city-specific grocery costs is related to a story in which high-income households benefit more from city size than low-income consumers, or one in which all consumers benefit more from living in locations with per capita incomes closer to their own.

I only work with purchase data for between 116 and 1,477 households in each city and, therefore, do not observe the full set of products available in each city. The reason for this, as I will discuss in more detail below, is that the CES and log-logit functional forms assumed to govern demand imply that I can identify the parameters that govern how consumers value the prices and varieties that I observe nationally without observing all of the prices and varieties available in each location. The sample size issue is more problematic when I use these demand parameters to calculate price indexes for each city. I am limited to calculating price indexes over the set of varieties that I observe in each market; that is, the set of varieties purchased by the households sampled in each market. The concern here is that the set of households I observe varies systematically across cities with different per capita incomes, potentially biasing my estimates of the relative price levels across these locations.[14] In my main analysis, I control

[13]Nielsen groups households into 52 markets. I instead classify cities at the level of Consolidated Metropolitan Statistical Area (CMSA) where available, and the Metropolitan Statistical Area (MSA) otherwise. For example, where Nielsen classifies urban, suburban, and ex-urban New York separately, I group them all as New York-Northern New Jersey-Long Island CMSA. In the two cases in which Nielsen groups two MSAs into one market, I count the two MSAs as one city, using the sum of the population and the population-weighted per capita income.

[14]The correlations of the sample household count with market population and market income are 0.57 and 0.43, respectively.

Table 2: Summary Statistics on Market Size and Income

Market ID	Market Name	Household Count				Population	Per Capita Income
		Total	Low Income	Middle Income	High Income		
1	Des Moines	143	45	58	40	456,022	37,650
2	Little Rock	348	185	104	59	583,845	33,289
3	Omaha	116	41	50	25	716,998	37,869
4	Syracuse	164	73	57	34	732,117	31,445
5	Albany	135	47	48	40	875,583	24,811
6	Birmingham	565	179	230	156	921,106	35,448
7	Richmond	192	51	77	64	996,512	37,082
8	Louisville	449	150	192	107	1,025,598	34,162
9	Grand Rapids	198	84	66	48	1,088,514	31,966
10	Jacksonville	128	39	45	44	1,100,491	35,439
11	Memphis	328	117	125	86	1,135,614	34,052
12	Raleigh-Durham	238	79	86	73	1,187,941	35,585
13	Nashville	397	154	156	87	1,231,311	36,056
14	Salt Lake City	245	92	103	50	1,333,914	33,469
15	Charlotte	1,100	392	403	305	1,499,293	36,580
16	Columbus	1,095	402	400	293	1,540,157	34,777
17	San Antonio	911	301	345	265	1,592,383	31,189
18	Indianapolis	243	86	92	65	1,607,486	36,160
19	Orlando	246	108	78	60	1,644,561	31,828
20	Milwaukee	169	57	64	48	1,689,572	37,361
21	Hartford-New Haven	193	62	69	62	1,725,259	41,705
22	Kansas City	219	83	80	56	1,776,062	35,839
23	Sacramento	1,044	297	352	395	1,796,857	35,318
24	New Orleans-Mobile	406	168	151	87	1,877,984	21,446
25	Oklahoma City-Tulsa	520	226	177	117	1,886,581	34,068
26	Cincinnati	316	97	112	107	1,979,202	35,326
27	Portland, Or	259	115	87	57	2,265,223	34,921
28	Buffalo-Rochester	1,121	476	384	261	2,268,312	33,020
29	Pittsburgh	300	150	100	50	2,358,695	36,159
30	Tampa	1,162	468	420	274	2,395,997	33,678
31	Denver	1,118	313	391	414	2,581,506	42,476
32	St. Louis	1,174	394	449	331	2,603,607	35,991
33	San Diego	138	47	38	53	2,813,833	40,383
34	Cleveland	360	137	130	93	2,945,831	35,555
35	Minneapolis	1,229	367	425	437	2,968,806	42,457
36	Phoenix	1,204	478	420	306	3,251,876	32,660
37	Seattle	1,139	302	431	406	3,554,760	42,356
38	Miami	1,069	315	395	359	3,876,380	38,342
39	Atlanta	1,029	253	393	383	4,112,198	35,262
40	Houston	1,065	295	354	416	4,669,571	40,734
41	Dallas	891	237	303	351	5,221,801	38,089
42	Detroit	991	261	372	358	5,456,428	37,204
43	Boston	1,271	348	439	484	5,819,100	47,491
44	Philadelphia	1,169	327	422	420	6,188,463	40,948
45	San Francisco	991	175	309	507	7,039,362	54,191
46	Washington, DC-Baltimore	1,106	238	348	520	7,608,070	46,725
47	Chicago	1,056	279	358	419	9,157,540	39,454
48	Los Angeles	1,143	289	392	462	16,373,645	37,441
49	New York	1,460	340	451	669	21,199,865	46,221

for any biases resulting from the fact that I observe systematically fewer households in larger, and often wealthier, cities by calculating price indexes using the prices and varieties recorded in the purchases of a random sample of 850 households in each of the 23 cities with 850 or more households in the Nielsen sample.[15]

An additional concern is that Nielsen samples a demographically-representative set of households in each city, so I observe the purchases of more high-income households in the samples for high-income cities than for low-income cities. Since tastes are identified using the same sample of purchases, the sampling pattern might lead me to conclude that wealthy cities have more varieties that favor wealthy tastes, simply because I have identified these tastes using the purchases of the wealthy households who are disproportionately sampled in wealthy cities. I deal with this potential sample bias by considering only the purchases of a stratified sample of households in each city. These stratified samples include a randomly-selected set of 570 households, 190 from each tercile of the full-sample income distribution, for each of the 22 cities for which I observe 190 or more households in each tercile.[16]

In addition to the Nielsen and 2000 U.S. census data, I also use data from the 2007 economic census on the number of grocery stores in each county where purchases are made in order to construct price instruments that help identify the parameters of the model.

3 Model

In this section, I introduce the model by first outlining some notation and then presenting the utility function that governs consumer choices. I then discuss how the model yields non-homotheticities in price sensitivity and demand for quality, and how they relate to the sources of non-homotheticities that have been modeled in previous literature. Finally, I solve the consumer's utility maximization problem to find the demand and indirect utility functions. Indirect utility is equal to the ratio of grocery expenditures and a price index over the products and prices available to a consumer in a market, their income, and model parameters. I use the demand function to derive income-specific market shares that will be used to empirically identify the parameters of the model. The indirect utility function provides an expression for the measure the marginal utility from grocery expenditure which forms the basis of the income- and city-specific price indexes that I calculate below.

[15] Handbury and Weinstein (2011) use the same Nielsen household consumption data to calculate homothetic cross-city price indexes based the Feenstra (1994) price index methodology. These indexes account for cross-city variety differences with a semi-parametric adjustment term based on expenditure shares which, the authors show, can be adjusted to account for potential sample size biases in observed variety. Unfortunately the parametric price index methodology used here does not allow for similar adjustments.

[16] This includes every city in which I observe 850 of more households in total except for San Francisco, where I only observe 145 low-income households.

3.1 Notation

Figure 2 shows how consumers choose to allocate expenditures. At the upper-most level, a consumer i allocates their income Y_i between expenditure W on a set of grocery products, denoted $\mathbf{G}$, and expenditure Z on a set of other goods, denoted $\mathbf{Z}$. I do not explicitly model the upper-level expenditure allocation decision, but it is crucial in one respect: preferences over grocery products are non-homothetic because they depend on non-grocery expenditures. This is generically the case as long as optimal non-grocery expenditures are normal.

In Appendix B , I solve for an implicit restriction on utility and prices under which the optimal non-grocery expenditure, Z_i^*, will be increasing in income. Although I cannot show that this restriction holds generally, I am able to show that it holds in the data. To do so, 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. The elasticity of observed non-grocery expenditures, Z_i, with respect to household income, Y_i, is 1.05 with a standard error of 0.0003.[17]

Figure 2: Consumer Choices

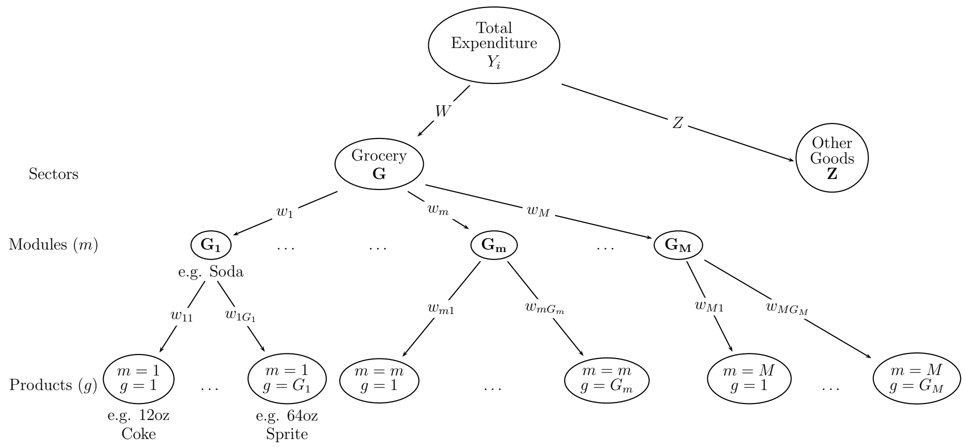

This paper focuses on the choices that consumers make within the grocery sector; that is, how consumers allocate their grocery expenditure W between product modules, $m = 1, ..., M$, and their module expenditure w_m between the varieties of grocery products, $g_m = 1, ..., G_m$, for each module m.

I refer to the set of product modules as $\mathbf{M}$ and index modules with the subscript $m = 1, ..., M$. Consumers allocate some expenditure w_m to products in module m, under the constraint that their module expenditures sum to their total grocery expenditure allocation W; that is, $\sum_{m \in \mathbf{M}} w_m = W$.

[17]There is also an Engel curve relationship between grocery expenditures and income. The median ratio of grocery expenditures to household income decreasing from 0.15 in the lowest income category of households to 0.05 in the highest income category.

I denote the set of all products in a module m as $\mathbf{G_m}$, where $\mathbf{G}$ is the union of these sets over all modules. I identify products in a module by the module index, m, and a product index, $g = 1, \ldots, G_m$. The product index uniquely identifies products within, but not across, modules. A consumer chooses to spend some w_{mg} on each product g in module m. The consumer purchases $q_{mg} = w_{mg}/p_{mg}$ units of product g in module m, where units are module-specific and p_{mg} is the price of product g in module m. I denote the set of observed grocery prices and consumption 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 consumption quantity sets over all modules.

The consumer's across-module and within-module expenditure allocation decisions are linked by the fact that the consumer cannot allocate more than their total module expenditure, w_m, between products $g \in \mathbf{G_m}$; that is, $\sum\limits_{g \in \mathbf{G_m}} w_{mg} = w_m$. The consumer's grocery expenditure allocation decision is, therefore, to allocate grocery expenditure W between products in the set $\mathbf{G}$ such that $\sum\limits_{mg \in \mathbf{G}} w_{mg} = \sum\limits_{m \in \mathbf{M}} \sum\limits_{g \in \mathbf{G_m}} w_{mg} = W = Y_i - Z$, where the final equality is due to the consumer's budget constraint in their grocery/non-grocery expenditure allocation decision.

3.2 Consumption Utility

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

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

where $\sigma(Z) > 1$ is the elasticity of substitution between modules for a consumer with outside good expenditure Z.

The assumption that the cross-module substitution patterns are governed by a CES utility function implies that consumers will optimally consume a positive amount in each module. In the data, the typical household buys a positive amount of a product in only 190 of 538 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, typically one year. Under this scenario, households will purchase a positive quantity of products in these modules in expectation, but not in every time period. That said, it is unlikely that all households consume products in every module, even in expectation. Unfortunately, models that reflects these more realistic cross-module consumption patterns, either by accounting for dynamic purchase behavior or explicitly modeling consumer's discrete-continuous preferences over modules, would be difficult to estimate given the dimensions of the problem that this paper addresses.[18] The moments used to estimate the model parameters are based on expected expenditure shares and calculated using the purchases of a market of multiple households, so there will be fewer zeros at this market level than at the household level. In estimation, the fact that some households do not purchase products in certain modules during a given quarter will be reflected in the fact that the modules have low market 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.

Consumer i's utility from consumption in module m, conditional on their outside good expenditure Z, is additive in their consumer-specific product-level utilities:

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

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

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

where β_{mg} is the quality of product g in module m; ε_{img} is the idiosyncratic utility of consumer i from product g in module m; $\gamma_m(Z)$ and $\mu_m(Z) > 0$ are weights that govern the extent to which consumers with outside good expenditure Z care about product quality and their idiosyncratic utility draws.[19]

Assuming that module utility is additive in product utilities implies that product-level utilities are perfectly substitutable with the utility from each of the other products within the same module. Consumers will, therefore, purchase positive quantities of only the product(s) that maximize their marginal utility from product expenditure, $\exp(\gamma_m(Z)\beta_{mg} + \mu_m(Z)\varepsilon_{img})/p_{mg}$. I assume that each consumer's product-specific idiosyncratic utility draws, ε_{img}, are drawn from a continuous type I extreme value distribution, with scale 0 and shape 1, so there will be a unique product that maximizes the marginal utility of expenditure for each household within each module; each household will allocate all of their module expenditure to only one product. This matches the discrete-continuous behavior of the typical household in the data, who purchases one or more units of exactly one product per module in each quarter.[20]

[18] A less-restrictive model of multiple-discrete purchasing behavior would allow consumers to purchase products in some, but not all, modules. For example, one could assume nested-logit preferences across all grocery products, $\mathbf{G}$, with module-level nesting and allow for consumers to purchase multiple units of a variety of products on each purchase occasion in order to provide for multiple consumption occasions, as in Hendel (1999) and Dube (2004). This multiple-discrete consumption behavior could also result from a static model by assuming that consumers are endowed with a certain level of utility in each module. This method, used in Song and Chintagunta (2007) and Pinjari and Bhat (2010), implies that a consumer's maximum marginal utility from expenditure in a module must be greater than a certain cut-off value for a consumer to want to purchase any products in that module. The maximum likelihood methods these authors use to estimate these models are not computationally feasible given the size of the choice set and the number of households in the dataset.

[19] 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.

[20] That is, the median number of products purchased in a module by a household in a quarter is equal to one.

3.3 Discussion of 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, β_{mg}, and their idiosyncratic utility draw for the product, ε_{img}. Preferences will be non-homothetic when at least one of the weights on these scalars, $\gamma_m(Z)$ or $\mu_m(Z)$, or the elasticity of substitution between modules, $\sigma(Z)$, varies with outside good expenditure and, as discussed above, this expenditure is normal. In order to interpret how these weights vary with income empirically, I make further functional form assumptions.

I interpret $\gamma_m(Z)$ to be the valuation for product quality, β_{mg}, for product g in module m shared by consumers who spend Z on the outside good. I assume that $\gamma_m(Z)$ is log-linear in outside good expenditure, Z, with a module specific slope, γ_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 $\gamma_m > 0$.

For estimation purposes, these quality parameters are assumed to be common across products with the same brand name in a module and estimated as the average willingness to pay for products with this brand across consumers. These parameters are identified using a revealed preference approach. The idea here is that product g in module m is estimated as having high quality, β_{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 products with the same brand as g than on products with the same brand as $\tilde{g}$. All consumers agree on this ranking but, for $\gamma_m > 0$, consumers who spend more on the outside good 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 γ_m would imply 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, Grossman, and Helpman (2011) show the theoretical implications of non-homothetic demand with a model that allows for complementarities between product quality and outside good expenditure. These complementarities imply that the elasticity of demand for quality is increasing with income, as in Hallak (2006) and Feenstra and Romalis (2012).

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, ε_{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 elasticity of substitution between products in a module m, equal to the inverse of the idiosyncratic utility draw weight, is log linear in non-grocery expenditures:

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

For $\alpha_m^1 < 0$, $\mu_m(z)$ increases 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 tendancy to switch between products in response to relative price changes is decreasing in consumer income. I impose the same functional form assumption on the elasticity of substitution between product modules, $\sigma(Z) \equiv 1 + \alpha^0 + \alpha^1$, so, for $\alpha^1 < 0$, high-Z consumers will also be less sensitive to changes in the aggregate quality-adjusted price across modules. In the expenditure equations derived below, it is easy to see that both the idiosyncratic utility weight, $\mu_m(Z)$, and the elasticity of substitution between modules, $\sigma(Z)$, will both govern the elasticity of consumer demand with respect to price.

The non-homothetic price sensitivity is also similar to recent international trade models. Hummels and Lugovskyy (2009), for example, develop a Lancaster ideal variety utility function where the disutility from distance between a product and a consumer's ideal type is an increasing function of their consumption quantity q_ω^γ 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. Income-specific price elasticities are also generated by the translated additive-log utility function used in Simonovska (2010) and translated CES utility functions more generally, but under different modeling assumptions that I will address below.

The model described above accounts 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 ways in which demand and, therefore, aggregate costs could vary across heterogeneous households. Below, I will consider whether ignoring forms of demand heterogeneity that other authors have found to be empirically relevant may impact the aggregate costs this paper estimates.

First, it is worth noting two forms of non-homotheticity that are prevalent in the recent literature, but do not appear to drive the differences in the purchase behavior of high- and low-income households observed in the Nielsen data and, therefore, are unlikely to have a large impact on how aggregate costs vary across cities differently for these different households.

One of these forms of non-homotheticity relates to shopping behavior and price sensitivity. Using the same Nielsen dataset, Broda, Liebtag, and Weinstein (2009) find that low-income households pay less for identical products than high-income households. The results in Aguiar and Hurst (2005) indicate that this is due to differences in shopping behavior and search costs. These differences could, in theory, enable low-income households to mitigate the high prices in wealthy cities at a lower cost than high-income households. Broda, Liebtag, and Weinstein (2009), however, additionally find that these identical product price differences are a small component of the overall differences in the prices paid for non-identical products in the same category, with the bulk of this difference being attributable to differences in how low- and high-income households allocate expenditures between products. Here, I abstract from differences in shopping behavior and focus instead on non-homotheticities that could drive differences in households' inter-product, rather than inter-temporal, expenditure allocations.

A second form of non-homotheticity that has found empirical support in household-level purchase (Li, 2012) and helps explain international pricing-to-market (Simonovska, 2010) is based on utility functions that yield hierarchic demand. Both hierarchic demand and the model presented below imply that consumers' price sensitivity is decreasing with income, but for different reasons. Under hierarchic demand, this result arises because the range of products consumed increases with income. Although I find that, consistent with Li (2012), households who spend more on food products tend to purchase more varieties, I do not observe that households who earn higher incomes or spend more on non-grocery products purchase more varieties of food products. In the model presented here, therefore, all households consume the same number of products. Instead, high-income households are less price sensitive and have a greater love of variety in the sense that they are more attached to their ideal varieties.

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. Here, I am using a more stylized model that allows the willingness to pay for a single product characteristic, brand name, to vary with a single consumer characteristic, income. This simplicity means that the framework 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 drawback of such a simple model is are its strong assumptions on consumer tastes.

First, the utility function presented above does not account for the components of product quality that are not correlated with brand name, such as product size, container type, or module-specific characteristics like flavor or texture. The cross-city price indexes I calculate will 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; however, the prices of products in the same module and brand will enter the price index symmetrically, even if they have different sizes, container types, etc. The assumption here 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.

Second, correlations between consumer demand and demographics, such as age, marital status, and household size, will be picked up by the model only to the extent that these demographics are correlated with size-adjusted household income, in which case these patterns will be attributed to non-homotheticities, biasing the estimates of the model parameters that govern them. These omitted variable biases will only carry through to bias the cross-city price indexes if demographics vary across wealthy and poor cities differently for households earning different incomes.[21] I do not expect these biases to be

[21] Suppose, for example, that high-income households are more educated and that demand for quality is correlated with education, and not income. The model and estimates presented below will attribute this correlation to non-homotheticities. The

large as the data do not show these correlations. Table 3 looks at how demographics vary across sample households that earn the same size-adjusted income but live in different cities. This table indicates that households with the same size-adjusted income do not in general have systematically different demographic compositions in wealthy, relative to poor, cities. In particular, any patterns that do exist are no stronger for high-income households than they are for low-income households.

Table 3: Cross-City Variation in Household Demographics within Income Groups

Panel A: OLS Regression

	Ln(Household Size)		Ln(Age (yrs))		Ln(Education (yrs))	
Ln(Per Capita Income)	-0.003	0.025	-0.026	0.233	0.012	0.144
	[0.004]	[0.083]	[0.032]	[0.199]	[0.014]	[0.133]
Ln(HH Income)*Ln(PC Income)		-0.003		-0.024		-0.012
		[0.008]		[0.018]		[0.012]
Constant	1.132***	1.089***	4.182***	3.784***	2.418***	2.217***
	[0.046]	[0.124]	[0.333]	[0.451]	[0.145]	[0.302]
Observations	39,767	39,767	39,767	39,767	39,767	39,767
(Psuedo) R-squared	0.968	0.968	0.187	0.187	0.202	0.202
Income Decile FE	Yes	Yes	Yes	Yes	Yes	Yes

Panel B: Logit Regression

	I(Married)		I(White)		I(Hispanic)	
Ln(Per Capita Income)	-1.175***	0.227	-0.996	2.825	0.686	2.627
	[0.265]	[1.964]	[0.799]	[4.674]	[1.127]	[5.276]
Ln(HH Income)*Ln(PC Income)		-0.133		-0.358		-0.181
		[0.172]		[0.414]		[0.430]
Constant	12.141***	10.091**	11.896	5.904	-10.775	-13.867
	[2.762]	[4.690]	[8.387]	[12.000]	[11.837]	[16.772]
Observations	39,767	39,767	39,767	39,767	39,744	39,744
(Psuedo) R-squared	0.458	0.458	0.0229	0.0230	0.0454	0.0454
Income Decile FE	Yes	Yes	Yes	Yes	Yes	Yes

Robust standard errors, clustered by city, in brackets.
*** p<0.01, ** p<0.05, * p<0.1
Notes:
[1] Age and education refer to the mean age and education of one or two heads of household, respectively.
[2] This sample includes all Nielsen households sampled from 2003 to 2005 reporting an aggregate annual income above $10,000.

3.4 Individual Utility Maximization Problem

The utility function defined in equation (1) is specific to the individual through consumer income and the consumer's idiosyncratic utility draws. In this section, I solve for consumer i's optimal demand for each product g and their indirect utility as a function of his/her outside good expenditure, Z, and his/her

price index estimates will still be a valid representation of the relative price levels faced by households at different income levels across cities insofar as education levels do not vary systematically within income groups across cities. If the education gap is wider in wealthier cities because high-income households are even more educated in high-income cities, and more high-quality products are available at cheaper prices in these locations, then the price index estimates will tend to overestimate the costs high-income households face in wealthy, relative to poor, cities.

idiosyncratic utility draws. In the next section, I aggregate these demand and indirect utility functions across these unobserved idiosyncratic utility draws for consumers with the same income in order to obtain expressions for demand and indirect utility in terms of (estimable) model parameters and data.

I assume that consumers draw an idiosyncratic utility ε_{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, defined in equations (1)-(3), subject to budget and non-negativity constraints:

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

I can solve this optimization problem in two steps. To see this, first note that

The consumer's optimal module bundle, $\mathbb{Q}^*_{im}(w_m, Z)$, is a function of their expenditure in that module, w_m, and their non-grocery expenditure, Z:

$$\mathbb{Q}^*_{im}(w_m, Z) = \underset{\substack{\mathbb{Q}_m \geq 0, \text{ s.t.} \\ \sum_{g \in \mathbf{G_m}} p_{mg} q_{mg} \leq w_m}}{\arg\max} \sum_{g \in G_m} q_{mg} d_{img}(Z) \tag{7}$$

where I use $d_{img}(Z) = \exp(\gamma_m(Z)\beta_{mg} + \mu_m(Z)\varepsilon_{img})$ to denote the marginal utility that consumer i receives from consuming one unit of product g in module m when spending Z on non-groceries. 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 outside good expenditure:[22]

$$g^*_{im}(Z) = \underset{g \in \mathbf{G_m}}{\arg\max} \frac{d_{img}(Z)}{p_{mg}} \tag{8}$$

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

$$\begin{aligned} \mathbb{Q}^*_{im}(w_m, Z) &= (q^*_{im1}(w_m, Z), \ldots, q^*_{imG_m}(w_m, Z)) \\ \text{where } \; q^*_{img}(w_m) &= \begin{cases} w_m / p_{mg} & \text{if } g = \underset{g \in \mathbf{G_m}}{\arg\max} \dfrac{d_{img}(Z)}{p_{mg}} \\ 0 & \text{otherwise} \end{cases} \end{aligned} \tag{9}$$

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.

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

I can now solve for the consumer's optimal module expenditures, $\mathbf{w_i^*}(Z) = (w_{i1}^*(Z), ..., w_{iM}^*(Z))$, as a function of their income, Y_i, and their non-grocery expenditure, Z:

$$\mathbf{w_i^*}(Z) = (w_{i1}^*(Z), ..., w_{iM}^*(Z)) = \underset{\sum_{m \in \mathbf{M}} w_m \leq Y_i - Z}{\arg\max} \left\{ \sum_{m \in M} [\tilde{u}_{im}(w_m)]^{\frac{\sigma(Z)-1}{\sigma(Z)}} \right\}^{\frac{\sigma(Z)}{\sigma(Z)-1}} \tag{10}$$

where the indirect utility for consumer i in module m is the value of the within-module utility, defined in equation (2), calculated at the optimal module bundle, defined in equation (9), is:

$$\tilde{u}_{im}(w_m, Z) = u_{im}\left(\mathbb{Q}_{im}^*(w_m, Z)\right) \quad = \quad w_m \left(\max_{g \in \mathbf{G_m}} \frac{d_{img}(Z)}{p_{mg}} \right) \tag{11}$$

The solution to this problem is derived in Appendix D.2:

$$\begin{aligned} \mathbf{w_i^*}(Z) \quad &= \quad (w_{i1}^*(Z), ..., w_{iM}^*(Z)) \\ \text{where } \; w_{im}^*(Z) \quad &= \quad (Y_i - Z) \frac{\left(\max_{g \in \mathbf{G_m}} \frac{d_{img}(Z)}{p_{mg}} \right)^{\sigma(Z)-1}}{P(\mathbb{P}, Z, \varepsilon_{\mathbf{i}})^{1-\sigma(Z)}} \end{aligned} \tag{12}$$

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:

$$P(\mathbb{P}, Z, \varepsilon_{\mathbf{i}}) = \left[\sum_{m \in \mathbf{M}} \left(\max_{g \in \mathbf{G_m}} \frac{d_{img}(Z)}{p_{mg}} \right)^{\sigma(Z)-1} \right]^{\frac{1}{1-\sigma(Z)}} \tag{13}$$

Note that the price index faced by a consumer i is a function of the prices and product variety in their consumption opportunity set, prices and product variety available to them, $\mathbb{P} = \{p_{mg}\}_{g \in \mathbf{G}}$, their outside good expenditures, Z, and idiosycratic utility draws, $\varepsilon_{\mathbf{i}} = \{\varepsilon_{img}\}_{g \in \mathbf{G}}$.

Plugging the optimal product choices and module expenditures derived above into the direct utility function defined in equations (1)-(3) yields the indirect utility of consumer i from grocery consumption:

$$V(\mathbb{P}, Y_i, Z, \varepsilon_{\mathbf{i}}) = \frac{(Y_i - Z)}{P(\mathbb{P}, Z, \varepsilon_{\mathbf{i}})} \tag{14}$$

4 Empirical Strategy

In this section, I provide details on how I measure the city-specific price indexes for consumers at different income levels. These price indexes represent the grocery component of the cost of living for an agent with the preferences described in Section 3 as a function of their income. I first outline the framework I use to calculate the indexes. These indexes require two key components: vectors of the prices that provide a comparable representation of the prices and product variety available across U.S. cities, and estimates for model parameters that govern consumer's perceptions of these price vectors. In the remainder of this section, I describe how I use the Nielsen data to obtain each of these components.

4.1 Measuring Relative Utility Across Markets

A main goal of this paper is to use the model estimates to measure how variation in the prices and product availability across U.S. cities differentially impacts the utility of consumers at different income levels. In Section (3.4) above, I solved for the indirect utility of a consumer i earning income Y_i from grocery consumption in a market offering a vector of prices $\mathbb{P}$, conditional on their non-grocery expenditures Z and their idiosyncratic utility draws $\varepsilon_{\mathbf{i}}$. Denoting the set of prices and products offered in a market t as $\mathbb{P}_t = \{p_{mgt}\}_{g \in \mathbf{G}_t}$ and the optimal non-grocery expenditures of household i in this market as Z_{it}, the indirect utility of a consumer i can be re-written as:

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

This indirect utility function is consumer-specific via three channels: it depends on a consumer's income, Y_i, on their optimal non-grocery expenditures, Z_{it}, and on their idiosyncratic utility draws, $\varepsilon_{\mathbf{i}}$. In order to compare grocery utility (or costs) across markets for consumers at different income levels, I need to summarize this indirect utility function across households so that it varies with i only through income, Y_i.

I first use the somewhat surprising empirical regularity in Engel curves across cities, with slopes ranging from 0.069 and -0.037, intercepts from 0.41 to 0.76.[23] Further, Figure 3 shows that, while households earning higher incomes spend a smaller share of their income on grocery products, the average grocery expenditure share does not vary greatly within income groups across cities.[24] I, therefore, approximate household non-grocery expenditures by assuming that non-grocery expenditures, Z_{it}, vary only with household income, Y, such that $Z_{it} = Z(Y)$.[25]

[23]I obtain these estimates by taking each sample household's annualized observed grocery expenditures as a share of reported income and regressing the resulting household food expenditure share against household log income and demographic controls over the households in each sample city.

[24]The coefficient of variation of household grocery expenditure shares is 78 across all households in the sample, but drops to between 32 and 52 when you only consider households within each income decile.

[25]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. In AppendixC, I show that, if an Almost Ideal Demand System governs grocery and non-grocery aggregate expenditure allocations, the invariance of the non-grocery expenditure share across markets, implies a particular relationship between the non-grocery and grocery price indexes faced by consumers within income groups across cities.

Figure 3: Income-Specific Food Expenditure Shares Across Markets

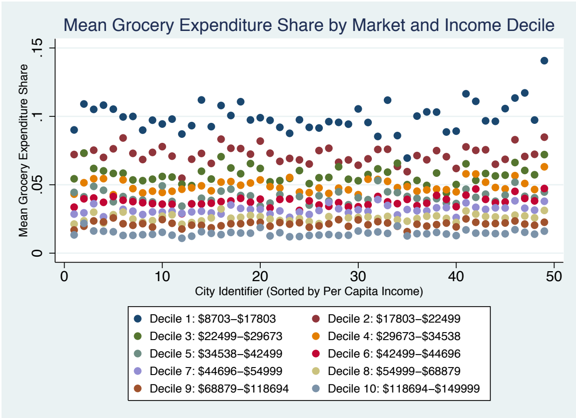

This assumption implies that a consumer's indirect utility is a function of market prices, $\mathbb{P}_t$, consumer income, Y_i, and idiosyncratic utility draws, $\varepsilon_{\mathbf{i}}$:

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

In particular, this assumption imples that 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:

$$\frac{V(\mathbb{P}_t, Y_i, \varepsilon_{\mathbf{i}})}{V(\mathbb{P}_{t'}, Y_i, \varepsilon_{\mathbf{i}})} = \frac{P(\mathbb{P}_{t'}, Z(Y_i), \varepsilon_{\mathbf{i}})}{P(\mathbb{P}_t, Z(Y_i), \varepsilon_{\mathbf{i}})}$$

The magnitude of the relative price index in market t relative to market t' above (or below) one indicates how much lower (or higher) the consumer's grocery utility is in market t relative to market t'.

I cannot measure these price indexes directly, since they depend on unobserved idiosyncratic utility draws. These draws imply that no two consumers, even with the same income, Y_i, will get the same utility from a set of consumption opportunities, $\mathbb{P}_t$. Since I wish to study the systematic variation in utility across consumers earning different incomes, I abstract from this random variation in consumer utility, summarizing over the idiosyncratic utility draws to approximate price indexes and, therefore, indirect utility as functions of market prices and consumer income alone.

The most direct way to summarize the indirect utilities of consumers with the same income level would be to take the expectation of the indirect utility over the idiosyncratic draws. There is no analytic solution to this problem, and numerical solutions are computationally intensive, [26] so I instead measure

[26]The log-logit module-level utility function is linear in these idiosyncratic draws, so it is possible to derive an analytic function for the expected module-level utility for a consumer conditional on his/her income. These module-level utilities, however, are nested within a non-linear CES aggregator, making it difficult to derive an analytic function for the expected income-specific utility from consumption in many modules. Numerically integrating the price index in each market over $\varepsilon_{\mathbf{i}}$ is computationally intensive so I reserved it for a robustness exercise, examining only the relative price index between the highest and lowest income cities in the full Nielsen sample: San Francisco and New Orleans. Using the numerical integration method,

the relative utility of households at various income levels across different markets by measuring the utility of an income-specific representative consumer. I assume that the representative consumer for households with income Y facing a set of prices $\mathbb{P}$ has utility

$$U(\mathbb{P},Y)=\left\{\sum_{m\in M}\left[\sum_{g\in\mathbf{G_m}}\left[q_{mg}\exp(\beta_{mg}\gamma_m(Z(Y)))\right]^{\rho_m(Z(Y))}\right]^{\frac{\rho(Z(Y))}{\rho_m(Z(Y))}}\right\}^{\frac{1}{\rho(Z(Y))}},$$

where $Z(Y)$ is the approximation of the optimal outside good expenditure of a household with income Y; and the remaining parameters and variables are as defined in the utility function above. That is, q_{mg} is the quantity consumed of product g in module m, β_{mg} is a parameter representing its quality; $\gamma_m(Z)=(1+\gamma_m\ln Z)$ governs the extent to which a consumer cares about product quality; and$\rho_m(Z)=\frac{1-\sigma_m(Z)}{\sigma_m(Z)}$ for $\sigma_m(Z)=1+\alpha_m^0+\alpha_m^1\ln Z$ and $\rho(Z)=\frac{1-\sigma(Z)}{\sigma(Z)}$ for $\sigma(Z)=1+\alpha^0+\alpha^1\ln Z$ govern the extent to which consumers differentiate horizontally between products in the same and different modules, respectively. In Appendix A, 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.[27]

Since the utility of the representative agent for households with income Y_i is CES, their indirect utility is the product of their grocery expenditure and a market-specific price index that summarizes the set of prices for available products in the market. I denote this indirect utility as $V^{CES}(\mathbb{P}_t,Y_i)$, and it is expressed as

$$V^{CES}(\mathbb{P}_t,Y_i)=\frac{(Y_i-Z(Y_i))}{P^{CES}(\mathbb{P}_t,Y_i)}, \tag{15}$$

where

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

It is worth taking a moment here to note that this approach to measuring income-specific cross-city price indexes is different from the approach that Broda and Romalis (2009) use to calculate income-specific inflation with the same Nielsen Homescan data. Broda and Romalis (2009) use the Feenstra (1994) methodology to calculate price indexes that are exact to nested-CES utility functions that are income-specific in that they allow households at different income levels to have wholly different perceptions of the relative quality of products. In the framework above, the Broda and Romalis (2009)

I found a gap of approximately 60 percent between the relative grocery costs faced by a high-income consumer between these markets and those face by a low-income consumer. Using the CES method, I found similar results; high-income households face 75 percent lower grocery costs in San Francisco relative to New Orleans relative to low-income households.

[27]The link between the log-logit and CES utility functions was first shown in Anderson, de Palma, and Thisse (1987). This proof was first extended to models that account for product quality in Verhoogen (2008). Appendix A further extends this result to show that the non-homothetic log-logit preferences presented have a Dixit-Stiglitz CES counterpart in that both models yield identical aggregate demand functions within groups of consumers at the same income level.

approach is equivalent to allowing β_{mg} to vary directly with household income, making γ_m irrelevant, and imposing that the substitution elasticities, σ_m and σ, do not vary with income. These price indexes are semi-parametric in that they are calculated using estimates for the substitution elasticities alone. The impact of the income-specific quality parameters on inflation is accounted for non-parametrically using aggregate market share information. In more recent work, the authors above take a more parametric approach to calculate cross-country price indexes. The Feenstra and Romalis (2012) 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. Where Broda and Romalis (2009) implicitly allow for the distribution of products' revealed qualities to vary systematically with income, Feenstra and Romalis (2012) and the model presented above both impose that households agree on the distribution of quality across products, and instead only allow them to value quality differentially.

4.2 Parameter Estimation

The indirect utility function derived in equation (15) uses model parameters to 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 θ defined as:

$$\theta = \{(\beta_1, \ldots, \beta_M), (\gamma_1, \ldots, \gamma_M), (\alpha_1, \ldots, \alpha_M), \alpha^0, \alpha^1\}$$

where $\alpha_m = \{\alpha_m^0, \alpha_m^1\}$ and $\beta_m = \{\beta_{m1}, \ldots, \beta_{mG_m}\}$. The estimation procedure outlined below identifies the parameter values that minimize the distance between the model's predictions for the expenditure share of each product and module in the data, conditional on consumer income, and the values of these expenditure shares that are observed in the data for market-specific groups of households with similar income levels. The moments used in this analysis are based on two types of relative expenditure share equations. The first defines the market share of a product within a module and the second defines the market share of a module in total grocery expenditures. I estimate a subset of module-specific parameters using the first of these relative expenditure share equations in parallel module-specific estimation procedures. The output from this first step of estimation is then to calculate the moment based on the module market share equation that is used in a second estimation step to identify the remaining module-specific parameters and the inter-module substitution parameters, α^0 and α^1.

4.2.1 Within-Module Estimating Equation

The within-module expenditure share on product g in module m for a group of households with the same outside good expenditure, Z_i, facing a common vector of module prices, $\mathbb{P}_m$, derived in the Appendix

D.3, is:

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

Note that the denominator of this market share will not vary across products within a module m and, therefore, drops out when I take the log of this expenditure share for any product g in module m and difference from the log expenditure share for a fixed product $\bar{g}_m$ in the same module:[28]

$$\ln\left(\frac{s_{mg|m}(Z_i, \mathbb{P}_m)}{s_{m\bar{g}_m|m}(Z_i, \mathbb{P}_m)}\right) = (\alpha_m^0 + \alpha_m^1 \ln(Z_i)) \left[(\beta_{mg} - \beta_{m\bar{g}_m})(1 + \gamma_m \ln(Z_i)) - (\ln p_{mg} - \ln p_{m\bar{g}_m})\right] \tag{16}$$

Equation (16) defines the expected within-module expenditure share of a set of households with outside good expenditure Z_i facing prices p_{mg} and $p_{m\bar{g}_m}$ on product g relative to product $\bar{g}_m$ in terms of module-specific price sensitivity parameters (α_m^0, α_m^1) and quality taste-income gradients (γ_m), as well as, product-specific relative quality parameters ($\beta_{mg} - \beta_{m\bar{g}_m}$) for each product $g \in \mathbf{G}$.[29] , *i.e.* $\{\beta_{mg} - \beta_{m\bar{g}_m}\}_{g \in \mathbf{G_m}}$. I denote this set of parameters by θ_1:

$$\theta_1 = \left\{\alpha_m^0, \alpha_m^1, \gamma_m, \{\beta_{mg} - \beta_{m\bar{g}_m}\}_{g \in \mathbf{G}_m}\right\}_{m \in \mathbf{M}}$$

To estimate these parameters, I calculate moments based on the relative expenditure shares I observe for markets in the data. For the purposes of estimation, I proxy outside good expenditure, Z_i, with household income, Y_i. Furthermore, I split the sample of households for each quarter-MSA market, t, into income quintiles, denoted by k. I calculate the relative expenditure shares implied by the model using the median log income of households in each income quintile k in each market t, y_{kt}, and the unit values paid by households in income quintile k in market t, p_{kgt}, for each product g purchased by a household in an income quintile k in market t. I also calculate the observed expenditure shares of income quintile k in market t, s_{kgt}, for each product g purchased by a household in that quintile and market. It is worth noting here that the estimation is based on a relative share equation, so I do not need to observe every product available to these households in order to calculate the moment condition.

The following equation, based on equation (16), defines the relationship between the sample relative within-module product shares and the model's prediction for these shares:[30]

$$\ln\left(\frac{s_{kmgt}}{s_{km\bar{g}_mt}}\right) = \left(\alpha_m^0 + \alpha_m^1 y_{kt}\right)\left[\left(\beta_{mg} - \beta_{m\bar{g}_m}\right)\left(1 + \gamma_m y_{kt}\right) - \ln\left(\frac{p_{kmgt}}{p_{km\bar{g}_mt}}\right)\right] + \nu_{kg\bar{g}_mt} \tag{17}$$

[28] 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.

[29] I restrict β_{mg} to be equal across all products g with the same brand name.

[30] In a slight abuse of notation, I denote the coefficients on log income using the same notation used for the coefficients on log outside good expenditure in equation 16 above. These new coefficients are in fact approximations of the original coefficient multiplied by the elasticity of outside good expenditure with respect to household income.

Here, $\bar{g}_m \in \mathbf{G_m}$ is a fixed base product for each module, which I define to be the product that is sold in the largest number of markets in its module. $\nu_{kg\bar{g}_m t}$ is the error in the predicted value of the relative product shares. This error includes differences between the mean prices paid and median incomes of households in quintile k in market t (p_{kmgt} and y_k) and the actual prices observed and incomes earned by these households. This error also includes measurement error in the collection of the raw data.

The θ_1 parameters are estimated in separate non-linear GMM procedures for each module using moments based on the relative expenditure share equation defined above and optimal weighting matrices calculated using the conventional two-step procedure. The relative brand quality parameters, $(\beta_{mg} - \beta_{m\bar{g}_m})$, are identified by variation in the average market shares of products within a brand, relative to the market share of the brand of the base product, conditional on price. The idea here is that, if products with two different brands sell at the same price, but products under brand A have higher average relative market shares across all income quintiles and MSA-quarter markets than products under brand B, then brand A will be assigned a higher quality parameter relative to the brand of the base product in the module. The quality-income gradient γ_m parameters that govern how demand for quality varies with price are identified by the extent to which the relative market shares of high-income quintile households are even more biased towards products under "high-quality" brands, *i.e.* those that have higher market shares across all income quintiles, than low-income quintile market shares. Conditional on brand quality, the base price sensitivity parameter, α_m^0 is identified by the extent to which relative market shares covary with the components of relative price variation captured by the price instruments, and the α_m^1 parameter is identified by the extent to which the relative market shares of high-income quintiles covary more (or less) with these relative price changes.

For each module, I construct moments by assuming that $\mathbb{E}[\mathbf{Z}_1 \cdot \nu] = 0$ for a set of instruments $\mathbf{Z}_1$. These instruments are constructed using a set of brand dummies, a set of price instruments, and interaction terms between these sets of variables and quintile income, y_k. The set of brand dummies includes one dummy for each brand except the base product in each module $\bar{g}_m$. I do not use prices as instruments because they might be correlated with the error term, $\nu_{kg\bar{g}_m t}$, across markets through market-specific product tastes, or across products within brands through product-specific tastes, neither of which are accounted for in the model.

For a given market, I instrument for price using four measures based on the prices observed in surrounding geographic or temporal markets. First, I use average price paid by consumers in the same income quintile in all other geographic markets in the same time period. I expect that this instrument will be correlated with price through national cost shocks, such as the increase in the price of wheat. Second, I use the average price paid by consumers in the same income quintile in all geographic markets in the same region in the same time period. This proxy captures any regional cost shocks and relies on the assumption that market-specific taste shocks are not shared across regional MSAs.[31] These instruments

[31] The region categorizations are provided in Table 15 in Appendix E.

are similar to those used in Hausman, Leonard, and Zona (1994) and Nevo (2000). Finally, I use the price paid for a product by consumers in the same income quintile in the same market in both the lead and lag time quarter. These two instruments are similar to those used in Asker (2004) and are intended to capture persistent local cost shocks, such as increases in sales taxes or wages. The strength of this instrument relies on MSA-specific cost shocks that persist over more than one quarter and its validity relies on the assumption that MSA-specific demand shocks do not persist for more than one quarter.

The instruments above are intended to capture temporal cost shocks that are correlated either across geography or over time. To capture spatial shocks in the level of market competition, I use county-level data on the number of grocery stores per capita. In each quarter-MSA market t, consumers in a given income quintile will purchase products in stores in a range of counties within their MSA. For each product, market, and income quintile, I take the weighted average of the county-level per capita store count, weighting by the purchases of a given product that consumers in a given income quintile make in each of the counties in that MSA-quarter market. The final price instrument is the number of products sold in a market for each module. This is intended to capture the level of competition in each module-market.

To test the strength of these instruments, I also estimated the lower-level of the demand system using subsets of the instruments described above, as well as using non-linear least squares (NLLS), ignoring the potential price endogeneity. This analysis yields two main results. First, the price coefficients are sensitive to the use of price instruments: ignoring price endogeneity in the NLLS specification, attenuates the base, α_m^0, price coefficient for the typical module by between 25 and 26 percent, depending on the parameter restrictions imposed. Second, the price coefficients are primarily identified by variation in the first two national and regional cost shock instruments. Excluding all of the remaining instruments changes the typical base, α_m^0, price coefficient estimates by between 0.75 and 0.85 percent, depending on the instruments excluded and parameter restrictions imposed. Instead of excluding the remaining weaker instruments, I deal with this issue by constructing $\mathbf{Z}_1$ using principal component analysis, as suggested in Bai and Ng (2010).[32]

4.2.2 Across-Module Estimating Equation

The remaining model parameters, α^0, α^1, and $\{\beta_{\bar{g}m}\}_{g\in\mathbf{G_m}}$, are identified using moments based on the model's prediction for module-level market shares. Specifically, the expected log expenditure share in module m relative to $\bar{m}$ for a group of households with the same outside good expenditure, Z_i, facing a common vector of grocery prices, $\mathbb{P}$, is derived in Appendix D.4 to be equal to:

$$\mathbb{E}_\varepsilon\left[\ln s_{im} - \ln s_{i\bar{m}}\right] = -(\alpha^0 + \alpha^1 \ln(Z_i))\left[\ln V_m(Z_i, \mathbb{P}_m) - \ln V_{\bar{m}}(Z_i, \mathbb{P}_{\bar{\mathbf{m}}})\right] \tag{18}$$

[32]Specifically, I construct the set of instrumental variables, $\mathbf{Z}_1$, used in estimation as the set of factors that explain the 99.5 percent of the variation in the full set of price instruments, brand dummies, and the interactions of both with quintile income.

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

$$V_m(Z_i, \mathbb{P}_m) = \left[\sum_{g \in \mathbf{G_m}} \left(\frac{d_{img}(Z_i)}{p_{mg}} \right)^{-(\alpha_m^0 + \alpha_m^1 \ln(Z_i))} \right]^{\frac{1}{-(\alpha_m^0 + \alpha_m^1 \ln(Z_i))}} \tag{19}$$

where $d_{img}(Z) = \exp(\gamma_m(Z)\beta_{mg} + \mu_m(Z)\varepsilon_{img})$. Together equations (18) and (19) define the expected relative module expenditure share of a set of households with outside good expenditure Z_i that face prices $\mathbb{P}_m$ and $\mathbb{P}_{\bar{m}}$ in terms of parameters α^0 and α^1, as well as α_m, γ_m, β_{mg} for all $g \in G_m$, and $\alpha_{\bar{m}}$, $\gamma_{\bar{m}}$, $\beta_{\bar{m}g}$ for all $g \in G_{\bar{m}}$. I use these expressions as the basis for calculating moments for each module $m \neq \bar{m}$, which, in turn, will be used to estimate the cross-module substitution parameters, α^0 and α^1, as well as 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 $\bar{m}$.[33] I denote this set of parameters by θ_2:

$$\theta_2 = \left\{ \alpha^0, \alpha^1, \{\beta_{m\bar{g}_m}\}_{m \in \mathbf{M}, m \neq \bar{m}} \right\}$$

To estimate these parameters, I again proxy for outside good expenditures with household income and aggregate the household-level purchase data to the income-quintile, MSA-quarter level. I define the difference between the relative module expenditure shares in the sample, where the module m expenditure share of households in income quintile k in market t is denoted as s_{kmt}, and the model's prediction for these shares as:

$$\ln\left(\frac{s_{kmt}}{s_{k\bar{m}t}}\right) = (\alpha^0 + \alpha^1 y_{kt}) \left[\ln V_m(y_{kt}, \mathbb{P}_{kmt}) - \ln V_{\bar{m}}(y_{kt}, \mathbb{P}_{k\bar{m}t})\right] + u_{km\bar{m}t} \tag{20}$$

where:

$$V_m(y_{kt}, \mathbb{P}_{kmt}) = \left[\sum_{g \in \mathbf{G_m}} \left(\frac{d_{kmgt}}{p_{kmgt}} \right)^{-(\alpha_m^0 + \alpha_m^1 y_{kt})} \right]^{\frac{1}{-(\alpha_m^0 + \alpha_m^1 y_{kt})}} \tag{21}$$

and $d_{kmgt} = \exp(\beta_{mg}(1 + \gamma_m y_{kt}))$. The parameters estimated in the first stage of estimation, θ_1, enter the equation above through the price-adjusted quality index, or inclusive value, for each module $V_m(y_k, \mathbb{P}_{kmt})$. I construct sample moments based on the above equation using data and estimates from the first stage, $\hat{\theta}_1$, to estimate the remaining parameters, θ_2.[34] These parameters include the parameters governing the elasticity of substitution between module, α^0 and α^1, and the quality parameters for the base product in each module, $\{\beta_{m\bar{g}_m}\}_{m \in \mathbf{M}}$. All but one of these quality parameters are identified in the equations above. I therefore normalize the quality of the base product in the base module $\bar{m}$ to zero.

Note that the inclusive value is a function of the parameters in both θ_1 and θ_2. Specifically, each product quality, β_{mg}, parameter is the the sum of $(\beta_{mg} - \beta_{m\bar{g}_m})$, a component of θ_1, estimated using

[33] I normalize the quality of the base product in the base module to equal zero.

[34] The point estimates for θ_1, which I will denote $\hat{\theta}_1$, are used to generate right-hand side variables used in the second stage of estimation. Since the first step of estimation yields consistent point estimates for θ_1, the second step of estimation will also yield consistent point estimates for the θ_2 parameters. However, the covariance matrix for $\hat{\theta}_2$ will need to be adjusted to account for the error in the calculation of $\hat{\theta}_1$ in order to get consistent standard errors for $\hat{\theta}_2$.

equation (17), and $\beta_{m\bar{g}_m}$, component of θ_2. We can rewrite the inclusive value function so that it is log linear in the $\beta_{m\bar{g}_m}$ parameters to be estimated:

$$\ln V_m(y_{kt}, \mathbb{P}_{kmt}) = \ln \left[\sum_{g \in \mathbf{G_m}} \left(\frac{d_{kmg\bar{g}t}}{p_{kmgt}} \right)^{-(\alpha_m^0 + \alpha_m^1 y_{kt})} \right]^{\frac{1}{-(\alpha_m^0 + \alpha_m^1 y_{kt})}} + \beta_{m\bar{g}_m}(1 + \gamma_m y_{kt})$$

where $d_{kmg\bar{g}t} = d_{kmgt}/\exp(\beta_{m\bar{g}_m}(1 + \gamma_m y_{kt})) = \exp((\beta_{mg} - \beta_{m\bar{g}_m})(1 + \gamma_m y_{kt}))$. Under the normalization that $\beta_{\bar{m}\bar{g}_{\bar{m}}} = 0$, and using the decomposition of the inclusive value function above, we can rewrite equation (20) as:

$$\ln \left(\frac{s_{kmt}}{s_{k\bar{m}t}} \right) = (\alpha^0 + \alpha^1 y_{kt}) \left[\Delta V_{1m\bar{m}}(y_{kt}, \mathbb{P}_{kmt}, \mathbb{P}_{k\bar{m}t}, \theta_1) + \beta_{m\bar{g}_m}(1 + \gamma_m y_{kt}) \right] + u_{km\bar{m}t},$$

where $\Delta V_{1m\bar{m}}(y_{kt}, \mathbb{P}_{kmt}, \mathbb{P}_{k\bar{m}t}, \theta_1) = \ln V_{1m}(y_{kt}, \mathbb{P}_{kmt}, \theta_1) - \ln V_{1\bar{m}}(y_{kt}, \mathbb{P}_{k\bar{m}t}, \theta_1)$ and

$$V_{1m}(y_{kt}, \mathbb{P}_{kmt}, \theta_1) = \left[\sum_{g \in \mathbf{G_m}} \left(\frac{d_{kmg\bar{g}t}}{p_{kmgt}} \right)^{-(\alpha_m^0 + \alpha_m^1 y_{kt})} \right]^{\frac{1}{-(\alpha_m^0 + \alpha_m^1 y_{kt})}}$$

The $u_{km\bar{m}t}$ errors are equal to the difference between the observed and predicted values of relative module shares in each market when the predicted values are calculated using the true values for the θ_1 parameters. In practice, the predicted values of the relative module shares in each market will be calculated using first-stage estimates for θ_1. There will be additional errors in the sample predicted values due to the fact that $\Delta V_{1m\bar{m}}(y_{kt}, \mathbb{P}_{kmt}, \mathbb{P}_{k\bar{m}t}, \theta_1) \neq \Delta V_{1m\bar{m}}(y_{kt}, \mathbb{P}_{kmt}, \mathbb{P}_{k\bar{m}t}, \hat{\theta}_1)$ and $\gamma_m \neq \hat{\gamma}_m$. I will denote these errors by $\nu_{km\bar{m}t}$. Taking these additional errors into account, the estimating equation defined in equations (20) and (21) becomes:

$$\ln \left(\frac{s_{kmt}}{s_{k\bar{m}t}} \right) = (\alpha^0 + \alpha^1 y_{kt}) \left[\Delta V_{1m\bar{m}}(y_{kt}, \mathbb{P}_{kmt}, \mathbb{P}_{k\bar{m}t}, \hat{\theta}_1) + \beta_{m\bar{g}_m}(1 + \hat{\gamma}_m y_{kt}) \right] + \nu_{km\bar{m}t} + u_{km\bar{m}t} \quad (22)$$

The θ_2 parameters are estimated using a single two-step non-linear GMM procedure. I construct moments by assuming that $\mathbb{E}[\mathbf{Z}_2 \cdot (\nu + u)] = 0$ for a set of instruments $\mathbf{Z}_2$. $\mathbf{Z}_2$ includes a set of dummies for all modules except the base module, two instruments for $\Delta V_{1m\bar{m}}(y_{kt}, \mathbb{P}_{kmt}, \mathbb{P}_{k\bar{m}t}, \hat{\theta}_1)$ calculated using the same national and regional cost shock instruments for prices that are used in the module-level estimation,[35] and each of these dummies and instruments interacted with income, y_{kt}.

The base α^0 substitution elasticity parameter is identified by the extent to which relative module shares react to components of the module inclusive value that are correlated with variation in national and

[35]These two instruments are $\Delta V_{1m\bar{m}}(y_{kt}, \tilde{\mathbb{P}}^N{}_{kmt}, \tilde{\mathbb{P}}^N_{k\bar{m}t}, \hat{\theta}_1)$ and $\Delta V_{1m\bar{m}}(y_{kt}, \tilde{\mathbb{P}}^R{}_{kmt}, \tilde{\mathbb{P}}^R_{k\bar{m}t}, \hat{\theta}_1)$. These differenced quality-adjusted price indexes are constructed in the same way as $\Delta V_{1m\bar{m}}(y_{kt}, \mathbb{P}_{kmt}, \mathbb{P}_{k\bar{m}t}, \hat{\theta}_1)$, using different price vectors. In place of the mean prices paid by income quintile k households in geographic-temporal market t, $\tilde{\mathbb{P}}^N{}_{kmt}$ and $\tilde{\mathbb{P}}^N_{k\bar{m}t}$ include mean prices paid by income quintile k households in all other geographic markets the same time period for the same set of module m and $\bar{m}$ products purchased by income quintile k households in market t. $\tilde{\mathbb{P}}^R{}_{kmt}$ and $\tilde{\mathbb{P}}^R_{k\bar{m}t}$ instead replace the mean prices paid by income quintile k households in market t with the mean prices paid for the same set of products by income quintile k households in all other cities in the same region in the same time period.

regional prices. The extent to which high- relative to low-income quintile module shares are more or less correlated with relative module price changes identifies the α^1 parameter that governs how much more or less substitutable high-income consumers find products in different modules. The relative inclusive value, $\Delta V_{1m\bar{m}}$, function is scaled up or down by the quality of the brand of the base product, $\bar{g}_m$, in a module m relative to the quality of the brand of the base product, $\bar{g}_{\bar{m}}$, in the base module $\bar{m}$, milk, which is normalized to equal zero. Any difference between the relative expenditure share of module m relative to milk and what would be expected given the relative inclusive value of the two modules and the estimates of the α^0 and α^1 parameters identified by the inter-market correlations discussed above will identify the quality of the brand of the base product in the module, $\beta_{m\bar{g}_m}$. Together with the relative brand quality estimates from the first stage of estimation, $\beta_{mg} - \beta_{m\bar{g}_m}$, these base brand quality estimates define the quality of each brand in the dataset relative to the quality of the base brand in the milk module.

This estimation procedure yields consistent estimates for θ_2, but the variance-covariance matrix of these parameters will be biased due to the presence of the first-step estimates for θ_1 in the ν 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).[36] The adjusted variance-covariance matrix yields consistent standard errors for the θ_2 estimates.

4.3 Inferring Prices and Product Availability

The second requirement to calculate the indirect utility functions implied by equation (15) above is the market-specific price vector, $\mathbb{P}_t$, representing the set of prices and products available to consumers in a city t. The Nielsen Homescan dataset includes the purchases of a between 100 and 1,500 households in each of the 49 MSAs included in the analysis. I proxy for the set of prices and products available to consumers in each city in 2005 using the set of products and unit values represented in the purchases of a sub-sample of the households surveyed by Nielsen. In order to deal with the potential sampling biases, discussed above in Section 2, I restrict the sub-samples of households used to build the $\mathbb{P}_t$ price vectors for each city in two ways.

In the main analysis, I construct the price vector for a city using the purchase records of 850 randomly-selected households from the city. Fixing the number of households whose purchases are included in the sample across cities mitigates any biases that could be generated by more households being surveyed in wealthy cities. This sampling method limits the utility analysis to the 23 cities whose Nielsen samples contain 850 or more households.

As a robustness check, I instead construct the price vector of a city by randomly selecting 150 sample households from each tercile of the national sample income distribution for each of the 22 cities where Nielsen has surveyed 150 or more households from each tercile. Fixing the income distribution of house-

[36] Appendix F details how these adjustments are calculated.

holds whose purchases are included in the sample across cities helps to mitigate any biases that could be generated by the fact that Nielsen's sample of households is demographically representative, such that more high-income households, whose purchases will naturally be biased towards products high-income households prefer to consume, will be sampled in high-income cities.

5 Results

5.1 Parameter Estimates

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

Table 4 summarizes the estimates for the module-level parameters in each of these four models across over 500 modules. Table 5 summarizes parameter estimates that are statistically significant at the 95 percent level.[37] Columns [1] through [3] of each table summarize the parameter estimates for the unrestricted version of the model. The corresponding utility function that governs consumer i's preferences between products $g \in \mathbf{G_m}$ within each module m is obtained by subbing the parametrizations for $\gamma_m(Z)$ and $\mu_m(Z)$, provided in equations 4 and 5, respectively, into equation 2:

$$u_{im}\left(\mathbb{Q}_m, Z\right) = \sum_{g \in \mathbf{G_m}} q_{mg} \exp\left(\beta_{mg}(1 + \gamma_m \ln Z) + \frac{\varepsilon_{img}}{\alpha_m^0 + \alpha_m^1 \ln Z}\right), \tag{23}$$

where β_{mg} characterizes the quality of product g relative to other products and ε_{img} is consumer i's idiosyncratic utility from product g. In this model, the weights that consumers place on quality and their idiosyncratic utility when determining their product ranking are functions of their non-grocery expenditure, Z, which is proxied in estimation by consumer income, $Y = \exp(y)$. The estimate for the elasticity of substitution between products is a function of the estimated weight placed on the idiosyncratic utility, $\hat{\sigma}_m(y) = 1 + \hat{\alpha}_m^0 + \hat{\alpha}_m^1 y$. Therefore, when the estimated value for this weight varies with income, or $\alpha_m^1 \neq 0$, the elasticity of substitution will also vary with income. The first column of Table 4 reports the elasticity of substitution of a consumer with the mean log income level in the sample for each module, or $\hat{\bar{\sigma}}_m = 1 + \hat{\alpha}_m^0 + \hat{\alpha}_m^1 \bar{y}_m$. The median of this elasticity is 2.09, with an inter-quartile range of 1.58 to 2.59. The magnitude and distribution of these estimates is similar across all four models. These estimates imply a median price elasticity of -1.09 and an inter-quartile range of -0.58 to -1.59. These parameter estimates are well-identified in most modules, with over 380 out of 504 significant at the 95 percent level in all four models.

[37] Statistical significance implies that the lower bound of the 95 percent confidence interval for the σ_m estimates is greater than one or that the lower (upper) bound of the 95 percent confidence intervals for α_m^1 and γ_m estimates are greater (less than) zero.

We can get some sense of the reasonableness of these point estimates by comparing them to those found in other studies. The own-price elasticities found here are slightly smaller than those estimated in Nevo (2000), who finds the own-price elasticity across cereal products to be between -2.2 and -4.2 whereas I estimate the own-price elasticity for cereals to be -1.13. On the other hand, the own-price elasticities here are slightly larger than those estimated in Dube (2004), who finds the own-price elasticity of demand for carbonated beverages to be between -0.42 and -0.85, whereas I find the own-price elasticity for carbonated beverages to be -1.96. All of these estimates are, however, much lower in magnitude than the own-price elasticities implied by the elasticity of substitution estimates in Broda and Weinstein (2008).[38]

Table 4: Summary Statistics for Parameter Estimates

Model:	NH in Quality and Price			NH in Quality		NH in Price		Homothetic
Restrictions:	None			$\alpha_m^1 = 0$		$\gamma_m = 0$		$\gamma_m = 0$ & $\alpha_m^1 = 0$
	[1]	[2]	[3]	[4]	[5]	[6]	[7]	[8]
Parameter:	$\bar{\sigma}_m$	γ_m	α_m^1	$\bar{\sigma}_m$	γ_m	$\bar{\sigma}_m$	α_m^1	$\bar{\sigma}_m$
Count	504	504	504	504	504	504	504	504
p25	1.58	-0.04	-0.19	1.62	-0.08	1.66	-0.03	1.64
p50	2.09	0.15	-0.05	2.13	0.09	2.14	0.04	2.11
p75	2.59	0.30	0.05	2.60	0.19	2.62	0.15	2.61

Table 5: Summary Statistics for Statistically Significant Parameter Estimates

Model:	NH in Quality and Price			NH in Quality		NH in Price		Homothetic
Restrictions:	None			$\alpha_m^1 = 0$		$\gamma_m = 0$		$\gamma_m = 0$ & $\alpha_m^1 = 0$
	[1]	[2]	[3]	[4]	[5]	[6]	[7]	[8]
Parameter:	$\bar{\sigma}_m$	γ_m	α_m^1	$\bar{\sigma}_m$	γ_m	$\bar{\sigma}_m$	α_m^1	$\bar{\sigma}_m$
Count	383	254	171	380	301	393	263	391
p25	1.87	0.15	-0.31	1.93	0.05	1.88	-0.05	1.88
p50	2.27	0.23	-0.19	2.28	0.15	2.26	0.10	2.26
p75	2.72	0.35	-0.10	2.73	0.23	2.73	0.19	2.69

[38] Broda and Weinstein (2008) use the Feenstra (1994) methodology to identify the elasticity of substitution between products. This method uses the assumption that all products in the same module share the same price elasticity of supply to help identify the price elasticity of demand. When I estimate the above model with this method, the median, as implied by the elasticity of substitution estimates for the largest 100 modules by sales value, rises from 2.4 to 7.7. I do not use this method more broadly because it does not allow me to identify the brand quality β_{mg} parameters.

Columns [2] and [5] of Table 4 summarize the distribution of the estimated values for γ_m across all modules. All four models assume that all consumers agree on the relative quality of products, as described by the distribution of the β_{mg} parameters for products $g \in \mathbf{G_m}$ within a module m. For positive values of γ_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 outside good expenditure z_i. This implies that consumers with higher expenditures on the outside good 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 β_{mg} estimates, that is, the products for which all consumers have a higher willingness to pay. In over half of the modules represented in the data, the willingness to pay for quality increases with income.

Columns [2] and [5] of Table 5 show that over half of the estimates for γ_m, or 254 and 301 estimates out of 504, are statistically significant in the models that allow for non-homothetic demand for quality and price sensitivity and for non-homothetic demand for quality but not price sensitivity, respectively. Over 75 percent of these statistically significant γ_m estimates are positive in the model that allows for non-homotheticity in demand for quality alone. Figure 5 shows that almost all of the statistically significant γ_m estimates are positive in the model that allows for non-homothetic demand for both quality and price sensitivity. These results indicate that non-homotheticity is important in many sectors. Since the demand for quality is increasing with income in most grocery sectors, richer households appreciate product quality more than poorer households.

Figure 4: Distribution of γ_m Parameter Estimates Across Modules

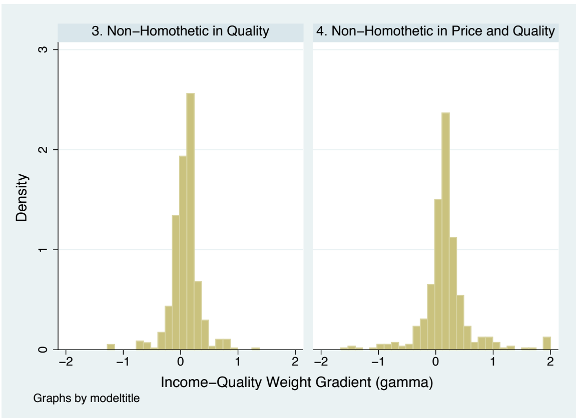

Columns [3] and [7] of Table 4 summarize the distribution of the estimated values for α_m^1 in each module. In equation 23, one can see that the weight that consumers place on their idiosyncratic utility is increasing in their outside good expenditure, z_i, for negative values of α_m^1. Suppose that two consumers draw very high values for ε_{img^*}, such that both consumers select to consume product g^* at the current

Figure 5: Distribution of Statistically Significant γ_m Parameter Estimates Across Modules

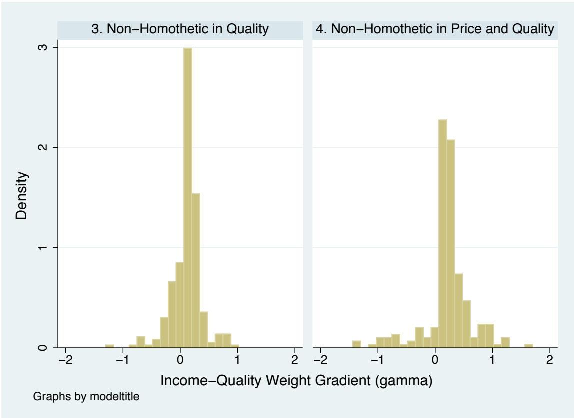

market price. If the price of product g^* increases, then the consumer who places a higher weight on his/her idiosyncratic utility draw will be less likely to switch to another product for which he/she drew a lower value for ε_{img} relative to the product's quality, β_{mg}. For $\alpha_m^1 < 0$, high-income consumers will place higher weights on their idiosyncratic utility draws and their expenditure shares will, therefore, be less sensitive to price changes. Column [3] of Table 4 shows that, for the majority 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. While Column [3] of Table 5 indicates that 171, or fewer than half, of the 504 α_m^1 estimates are statistically significant, the right hand panel of Figure 7 shows that the vast majority of these statistically significant α_m^1 estimates are less than zero. While the evidence that price sensitivity varies with income is less prevalent across modules, the price sensitivity is in the expected direction in modules where there is statistically significant variation in the price sensitivity by income.

If we focus instead on the model that allows for non-homothetic price sensitivity but not non-homothetic demand for quality, Column [7] of Tables 4 and 5 show that the majority of the α_m^1 estimates, and even the majority of those that are statistically significant, are positive when γ_m is constrained to be zero. These estimates may be biased upwards by a correlation between unobserved income-specific product tastes and prices. Consider the model: $\ln s_{kgt} - \ln s_{k\bar{g}_m t} = (\alpha_m^0 + \alpha_m^1 y_k)[(\beta_{mg} - \beta_{m\bar{g}_m}) - (\ln p_{kgt} - \ln p_{k\bar{g}_m t})] + \nu_{kg\bar{g}_m t}$. Here, the error terms include any income-specific product tastes, $\beta_{kmg} - \beta_{km\bar{g}_m}$. If the stores at which high-income consumers shop set prices in accordance with these tastes, such that $Corr(\beta_{kmg} - \beta_{km\bar{g}_m}, \ln p_{kgt} - \ln p_{k\bar{g}_m t}) \neq 0$, then the assumption that $\mathbb{E}[\mathbf{Z}_1 \cdot \nu] = 0$ will be violated. The fact that the α_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 includes a term that varies by product and income, $(\beta_{mg} - \beta_{m\bar{g}_m})\gamma_m y_k$, and therefore does not include the full value of $\beta_{kmg} - \beta_{km\bar{g}_m}$ in the errors. I do not, therefore, take the positive α_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 α_m^1 estimates highlight the difficulty in identifying the non-homotheticity related to price sensitivity in isolation from the non-homotheticity related to product quality.

Figure 6: Distribution of α_m^1 Parameter Estimates Across Modules

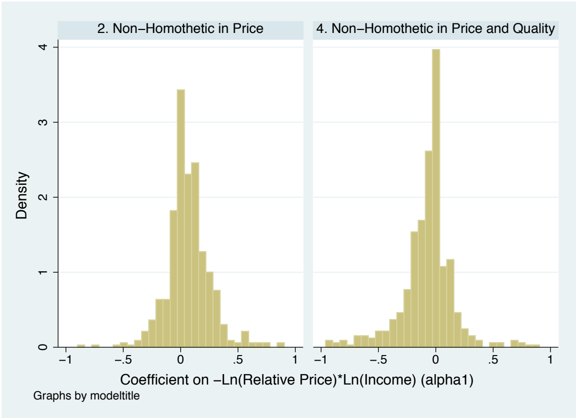

Figure 7: Distribution of Statistically Significant α_m^1 Parameter Estimates Across Modules

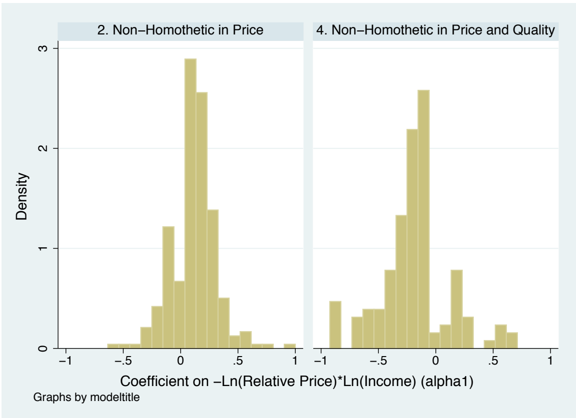

Despite the results from this one model, the parameter estimates generally show convincing evidence of non-homothetic demand. Specifically, high-income consumers have a greater willingness to pay for quality than low-income consumers and, when controlling for this non-homotheticity in the demand for quality, the results show that high-income consumers are also less price sensitive.

The upper-level between-module estimation equation, 22, yields the elasticity estimates reported in Table 6.[39] $\bar{\sigma} = 1 + \alpha^0 + \alpha^1 \bar{y}$ reflects the elasticity of substitution for the household with the mean income in the markets used in analysis. As expected, products in different modules are less substitutable

[39]Standard errors are reported in parentheses underneath the point estimates. These standard errors have been adjusted for measurement error in the first stage estimates as described in Appendix F.

than products in the same module, with between module substitution elasticities of between 1.03 and 1.09. The model that allows for non-homotheticity in price sensitivity alone indicates that high-income consumers are more price sensitive than low-income households but, as with the within-module estimates reported above, this pattern reverses to indicate that high-income consumers are less sensitive to cross-module price changes once you control for the fact that they are have a stronger demand for higher quality products and, therefore, the modules containing these brands.

Table 6: Upper-Level Substitution Elasticity Estimates

		Parameter	
	Model Name	$\bar{\sigma}$	α^1
1.	Homothetic	1.036	-
		[0.00201]	
2.	Non-Homothetic in Price	1.033	0.00043
		[0.00003]	[0.00202]
3.	Non-Homothetic in Quality	1.087	-
		[0.00945]	
4.	Non-Homothetic in Quality and Price	1.047	-0.0018
		[0.00044]	[0.00127]

5.2 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 an additional 534 parameters to the model (one α_m^1 for each module, as well as the aggregate cross-module α^1). These parameters will all be sources of error in the income-specific price indexes used to address the paper's main question in Section 5.3 below. Prior to undertaking this anlaysis, 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.[40]

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. In Section 5.1 above, I presented estimates of the parameters that govern the within-module product choice for each module m, denoted θ_m^1, in a separate GMM

[40]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).

estimation procedure under the four sets of parameter restrictions corresponding to the four models. For the most flexible version of the model, all elements of θ^1_m are estimated. These include α^0_m, α^1_m, γ_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 $\bar{g}$. For the brand of the base product $\bar{g}$, $\beta_{mg} - \beta_{m\bar{g}}$ equals zero. Each of the models with parameter restrictions are nested in the full model, which allows for non-homotheticity in both the demand for quality and price sensitivity. All four models have been estimated using the optimal weighting matrix for the full model, which I denote by W^*. 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. I denote this set of instruments by L^*.

The selection criterion minimizes the following GMM-BIC function:

$$\text{GMM-BIC}_{Mm}(\hat{\theta}^1_{Mm}) = n_m G_m(\hat{\theta}^1_{Mm}, \bar{\theta}^1_M)' W^*_m G_m(\hat{\theta}^1_{Mm}, \bar{\theta}^1_M) - \ln(n_m)(L^*_m - K_{Mm}) \tag{24}$$

Here, $G_M(\hat{\theta}^1_m, \bar{\theta}^1_M)$ are the moments for model M evaluated at the estimated values for free parameters $\hat{\theta}^1_{Mm}$ and zero for the restricted parameters, $\bar{\theta}^1_M$; K_{Mm} is the number of free parameters in model M for module m; and n is the number of observations. I evaluate models by calculating the unweighted and sales-weighted share of modules for which that model minimizes the GMM-BIC criterion. The results of this model selection test are presented in Table 7.

Table 7: Share of Modules in which GMM-BIC Criterion Selects Each Model

M	Model Name	Parameter Restrictions	Unweighted Share	Sales Weighted Share
1.	Homothetic	$\gamma_m = 0, \alpha^1_m = 0$	0.15	0.15
2.	Non-Homothetic in Price	$\gamma_m = 0$	0.15	0.10
3.	Non-Homothetic in Quality	$\alpha^1_m = 0$	0.47	0.39
4.	Non-Homothetic in Quality and Price	None	0.26	0.36

The model that permits non-homothetic demand for quality, but not for price, is the optimal model for almost half of the modules. These modules represent 39 percent of sample sales. The most flexible model has the next-highest share of "winning" modules, representing 36 percent of sample sales. This indicates that the most flexible model performs better in the larger modules. In fact, the most flexible model performs the best in 61 per cent of the largest 50 modules by sales, while the least restrictive, or homothetic model, performs the best in 61 percent of the smallest 200 modules by sales. This is most likely related to the number of observations used in the estimation of the largest, relative to the smallest, modules.

In the trade literature, non-homothetic utility functions are either non-homothetic in the demand for quality or non-homothetic in price sensitivity. The J statistics for these two models can be compared directly in Figure 8.[41] The J statistic is lower for the model that permits non-homotheticity in the demand for quality in the majority of modules. This indicates that the quality model (Model 2) has a lower GMM-BIC criterion than the price sensitivity model (Model 3) in the majority of those modules for which the most flexible model has the lowest GMM-BIC criterion overall. That is, the model allowing for non-homotheticity in the demand for quality alone explains more of the cross-income variation in consumer behavior than the model that allows for non-homotheticity in price sensitivity alone.

Figure 8: Distribution of Log Relative J Statistic for Model 2 Relative to Model 3

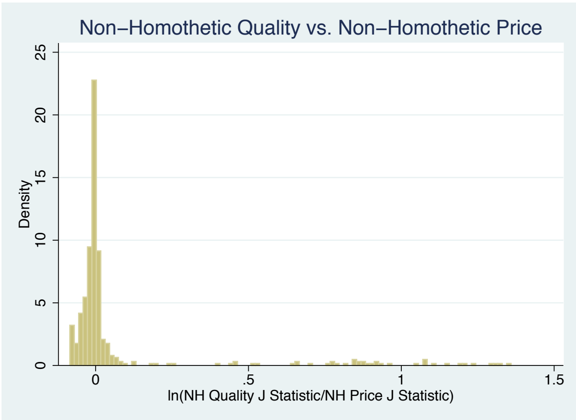

Table 8 shows the results of these bilateral model comparisons across all four models. We see that the model that accounts for non-homothetic demand for quality has a lower GMM-BIC criterion in modules representing approximately two-thirds of sales when compared to all three of the other models.

Table 8: Bilateral Model Comparisons: Sales Share of Modules where GMM-BIC(M)<GMM-BIC(N)

		Model N			
Model M		1.	2.	3.	4.
1.	Homothetic	-	0.18	0.28	0.32
2.	Non-Homothetic in Price	0.82	-	0.38	0.40
3.	Non-Homothetic in Quality	0.72	0.62	-	0.68
4.	Non-Homothetic in Quality and Price	0.68	0.60	0.32	-

The analysis above suggests that the salient form of non-homotheticity is in the demand for quality. In the analysis below, I study how price indexes that account for this form of non-homotheticity alone vary across cities differently for consumers at different income levels. Any differences between the

[41]The two elements of the GMM-BIC criterion defined in equation 24 that vary by model are the J statistic, $nG_m(\hat{\theta}^1_{Mm}, \bar{\theta}^1_M)'W^*_m G_m(\hat{\theta}^1_{Mm}, \bar{\theta}^1_M)$, and the number of free parameters, K_{Mm}. The models with non-homotheticity in quality or price alone each have one parameter restriction and, therefore, the same number of free parameters.

cross-city price indexes for high- and low-income consumers will reflect cross-city differences in the availability and prices of high- relative to low-quality products. 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. Appendix G shows that the main results presented below only change marginally when based on price indexes that account for non-homotheticities in both consumer's demand for quality and their price sensitivity.

5.3 Income-Specific Consumption Externalities

As outlined in Section 4.1, I use market- and income-specific price indexes as a measure of representative consumer utility from grocery consumption. I compare how the estimated price index for a representative consumer with log income y_k in city c, $\hat{P}(\mathbb{P}_c, y_k)$, varies city-to-city using the following baseline regression model:

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

where δ_k is an income-level fixed effect and y_c is log per capita income in city c.

The goal of this analysis is to determine how grocery costs differentially vary across cities for consumers at different income levels. Specifically, equation (25) measures how the elasticity of grocery costs in a city with respect to its per capita income varies with household income. The grocery cost price index is calculated using a model that allows for non-homotheticity in the demand for quality, and thus this elasticity will vary with income if the goods available in each city are correlated with the tastes of the incomes of the consumers living there. Suppose, for example, that wealthy cities sell more varieties of high-quality goods at lower prices than poorer cities. If this is the case, the price index faced by high-income consumers will decrease at a faster rate (or increase at a slower rate) than the price index faced by low-income consumers who move from poor to wealthy cities. This is because high-income consumers benefit from the availability and lower prices of the goods that they prefer. This implies that the elasticity of the price index faced by high-income consumers with respect to city income will be lower than the elasticity of the price index faced by low-income consumers with respect to city income. In the above specification, the elasticity of grocery costs with respect to city income is equal to the coefficient on log city income added to the product of log consumer income and by the coefficient on the income interaction term: $\varepsilon_{P_{ck}, y_c} = \beta_1 + \beta_2 y_k$. The first column of Table 9 presents the results of this regression. The estimates for β_2 are negative and statistically significant at the 95 percent level confirming that the elasticity of the price level with respect to per capita income varies across income levels. The magnitude of the β_2 estimate indicates that this variation is economically significant. A consumer who earns \$15,000 a year sees his/her price index rise by around 30 percent for each log unit increase in city per capita income, approximately equivalent to the log difference between San Francisco (per capita income of \$54,191) and New Orleans (per capita income of \$21,446). On the other hand, the price index of a consumer with a yearly income of \$100,000 decreases by around 9 percent for each log unit increase in city per capita income. Therefore, a high-income household experiences a 40 percent greater increase in grocery consumption utility than a low-income household when both move from a poor city to a city with double the per capita income.

Table 9: City-Income Specific Price Index Regressions

Dependent Variable: Ln(Price Index for Representative Consumer k in City c)				
Ln(Per Capita Income$_c$)	β_1	2.412***	-	2.290*
		[0.996]	-	[1.22]
Ln(Per Capita Income$_c$)	β_2	-0.217**	-	-0.201*
*Ln(Household Income$_k$)		[0.0915]	-	[0.112]
Ln(Population$_c$)	β_3	-	0.28	0.040
		-	[0.190]	[0.232]
Ln(Population$_c$)	β_4	-	-0.027	-0.005
*Ln(Household Income$_k$)		-	[0.018]	[0.021]
Implied Elasticity of Price with respect to City Income ($\beta_1 + \beta_2 y_k$):				
- $y_k = \ln(\$15,000)$		0.325	-	0.357
- $y_k = \ln(\$50,000)$		0.064	-	0.115
- $y_k = \ln(\$100,000)$		-0.086	-	-0.069
Household Income Fixed Effects		Yes	Yes	Yes
Observations		230	230	230
R-Squared		0.03	0.012	0.036

*** p<0.01, ** p<0.05, * p<0.1

Standard errors in brackets.

The regression above imposes that the elasticity of the income-specific price index with respect to city income varies linearly with income. There is no reason that this needs 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:

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

Figure 9 plots the estimates of the β_{2k} elasticity parameters against 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 at the lower and upper tails of the income distribution. The price index of a consumer who earns \$15,000 per year increases by almost 20 percent with each log unit increase in city income, whereas the price index for a consumer who earns \$100,000 per year decreases by around 20 percent with each log unit increase in city income.

Figure 9: Variation in Elasticity of Grocery Costs with respect to City Income Across Household Income Levels

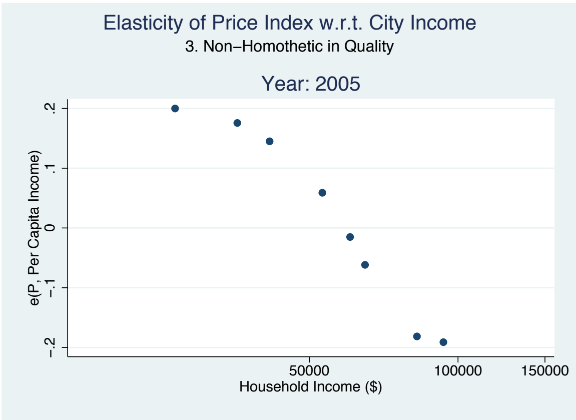

The data in Table 1 indicate that market income is correlated with market size, that is, wealthier U.S. cities are larger than poorer U.S. cities (the correlation coefficient is 0.47). Therefore, it is possible that the negative coefficient on the market income and household income interaction term in the base line regression (equation 25) is due to the fact that grocery costs are lower for high-income households than for low-income households in larger, as opposed to wealthier, cities. If this were the case, the results above would support a story in which high-income consumers receive more consumption benefits from living in larger cities than low-income consumers, as opposed to the "preference externalities" story in which high-income consumers receive more consumption benefits from living in wealthier cities and low-income consumers receive more consumption benefits from living in poorer cities. I test between these two theories by including log population and log population interacted with log household income in the regressions. The results from these regressions are presented in the second and third columns of Table 9. When the log price indexes are regressed against these population variables and household income dummies, the coefficients on log population and log population interacted with log household income are not statistically significant. When I include these extra variables in the baseline model that controls for city-size effects, the coefficients on the controls remain insignificant. More importantly, the coefficients on log per capita income and log per capita income interacted with log household income are similar in magnitude to the estimates in the baseline model, although their statistical significance has been reduced from the 5 to the 10 percent level.

There is reason to be concerned that sampling bias drives the results above. The prices and products included in the indexes for each city are based on the purchases of the households Nielsen sampled in each city for 2005. As discussed in the Data section above, the number of households Nielsen samples varies systematically with city size and, therefore, city income. I control for this source of sample bias by using only the purchases of a random sample of a fixed number of households from each city. An additional concern is that Nielsen samples demographically representative sets of households in each city. This implies that, even holding the number of households constant across cities, we will observe more purchases of products that high-income households prefer to consume in high-income cities. This bias could mechanically generate the results above. To check whether this is the case, I re-calculate the price indexes using a stratified sample of households for each city, including 190 randomly-sampled households from each tercile of the income distribution. This limits the number of cities included in the analysis to 22, since San Francisco has fewer than 190 low-income households. Figure 10 indicates that the magnitude and direction of the results presented above is robust to stratified sampling. Low-income households find the cost of the basket of prices and products observed in the stratified sample for high-income cities to be approximately 25 percent higher than it is for low-income cities, while high-income households find these costs to be approximately 10 percent lower. This indicates that both high- and low-income households find the stratified sample bundles in high-income cities to be more expensive relative to that in low-income cities than they did when comparing the non-stratified sample bundles. On the whole, however, the gap between how high- and low-income households perceive the relative costs between high- and low-income cities is similar at 35 percent with the stratified sample, relative to 40 percent with the base sample.

Figure 10: Variation in Elasticity of Grocery Costs with respect to City Income Across Household Income Levels Using a Stratified Sample of Households

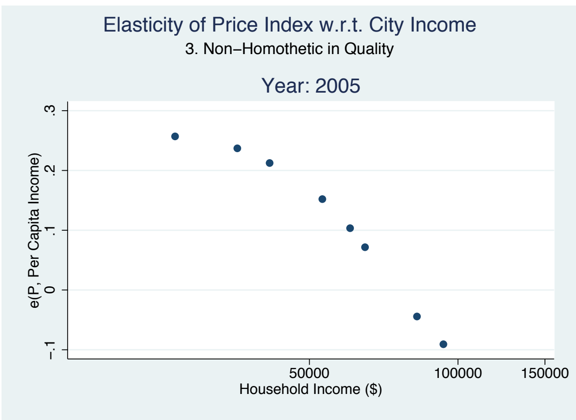

Taken together, the results above suggest that, relative to low-income households, high-income households receive higher consumption utility from the grocery bundles available in wealthier cities than from the grocery bundles available in poorer cities with the same population size. This pattern is consistent with theories that predict that the composition of demand affects the value of being in a location. Waldfogel (2003) and Fajgelbaum, Grossman, and Helpman (2011) posit that the gathering of consumers with similar tastes will yield greater consumption benefits to consumers with these tastes than to those with different tastes. The authors refer to this phenomenon as "within-group preference externalities" or "home market effects," respectively.

The model allows for high-income households to have a stronger preference for high-quality goods than do low-income households. 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 850-household sample for each city, as before, but setting the prices of these products equal to its national average price. The final two columns of Table 11 report the estimates of the baseline regression model run using these fixed-price indexes as the dependent variable. The first two columns replicate the results from the baseline regression. We observe that the coefficient on the interaction between per capita income and household income is more negative and more statistically significant at a higher level when the fixed-price indexes are used as the dependent variable. High-income households would find wealthy cities to be even less expensive than poor cities, relative to low-income households, if products were sold in both locations at their national average price. This suggests that the entire difference between how high- and low-income households perceive the relative costs to vary across cities is 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.

5.4 Comparison with Homothetic Index

I now turn to addressing the extent to which homotheticity biases the estimates of cross-city price indexes for consumers at different income levels. If we assume that preferences are homothetic such that all households get the same utility from the consumption baskets available in one market relative to another, we only need one homothetic price index to compare the utility that households get in one city relative to another. By allowing preferences to be non-homothetic, I allow households at different income levels to get different relative utilities from the consumption baskets available in different locations and, therefore, calculate a different price index to measure these relative utilities for each income-level. The analysis above has shown that there is economically significant variation in how these non-homothetic price

Table 11: City-Income Specific Price Index Regressions

Dependent Variable: Ln(Price Index for Representative Consumer k in City c)					
		City-Specific Prices		National Average Prices	
Ln(Per Capita Income$_c$)	β_1	2.412***	2.290*	3.136***	3.421***
		[0.996]	[1.22]	[0.879]	[1.10]
Ln(Per Capita Income$_c$)	β_2	-0.217**	-0.201*	-0.294***	-0.311***
*Ln(Household Income$_k$)		[0.0915]	[0.112]	[0.081]	[0.101]
Ln(Population$_c$)	β_3	-	0.040	-	-0.087
		-	[0.232]	-	[0.197]
Ln(Population$_c$)	β_4	-	-0.005	-	0.005
*Ln(Household Income$_k$)		-	[0.021]	-	[0.018]
Household Income Fixed Effects		Yes	Yes	Yes	Yes
Observations		230	230	230	230
R-Squared		0.03	0.036	0.098	0.099

*** p<0.01, ** p<0.05, * p<0.1

Standard errors in brackets.

indexes vary across cities for consumers at different income levels. A homothetic price index captures none of this variation, but it may match the cross-city variation in prices for consumers at some income levels better than others. To consider this question, I first calculate a homothetic price index for each city using the parameter estimates for the model that does not permit either the demand for quality or the price sensitivity of a household to vary with income. In Table 13 I compare these cross-city homothetic price indexes to the income-specific cross-city price indexes calculated using the parameter estimates for the selected model, which permits the demand for quality to vary with income. The homothetic price index is highly correlated with the non-homothetic price indexes calculated for households earning below $70,000 per year. The correlation between the homothetic price index and the non-homothetic indexes is highest with a coefficient of 0.97 for households earning around $50,000 per year. This indicates that the homothetic price index does a good job at predicting the cities in which low- and middle-income households will gain the most, and the least, from the grocery consumption bundles available there. The correlation coefficient drops to 0.21 for households earning around $80,000 per year and is negative for households earning around $150,000 per year. The homothetic price index, therefore, does a poor job of predicting which cities high-income households find the most and the least expensive.

Table 13: Correlation of City-Specific Price Indexes Calculated with Homothetic and Non-Homothetic Models

Household Income	Correlation Between Homothetic Index and Non-Homothetic Index
$16,896	0.83
$26,715	0.88
$35,715	0.94
$41,526	0.96
$53,103	0.97
$60,442	0.92
$64,805	0.83
$82,576	0.21
$93,411	0.03
$146,566	-0.11

Table14 further illustrates these facts. While the homothetic model does not perfectly predict the rankings of cities for households at any income level, it performs very poorly in predicting the most and least expensive cities for high-income households. The homothetic model predicts that Chicago, San Antonio, Sacramento, and San Francisco, are among the six most expensive cities for purchasing groceries. The non-homothetic model, however, predicts that these four cities are among the six cheapest for a household earning \$150,000 per year. Conversely, the homothetic model predicts that Atlanta, Detroit, and Columbus are among the five cheapest cities for purchasing groceries, while the non-homothetic model predicts that these cities are among the five most expensive cities for households earning either \$93,000 or \$150,000 per year.[42]

The evidence above shows that non-homothetic preferences yield economically significant variation in living costs: wealthy consumers benefit more from the consumption baskets available in wealthy cities than poor consumers. This variation is particularly relevant for economists who currently use homothetic price indexes to measure real income inequality. Moretti (forthcoming), for example, argues that one should adjust nominal income by location-specific prices when measuring national real income inequality. Since my results show that high-income consumers face vastly different location-specific prices than low-income consumers, they suggest that economists should adjust nominal income by prices that are both location- and income-specific when measuring national real income inequality. Moretti (forthcoming) finds that the U.S. college wage premium is lower in real terms than in nominal terms because college graduates are concentrated in metropolitan areas where homothetic price indexes are high. The results above indicate that the non-homothetic price indexes faced by high-income college graduates will vary across cities differently than those faced by high school graduates, who earn lower incomes. If homothetic price indexes are negatively correlated with the non-homothetic price index for high-income consumers, as the results in Table 13 suggest is the case for groceries, then Moretti (forthcoming) will tend to underestimate the real income of college graduates relative to high school graduates. Measuring the size of this bias will require cost-of-living indexes that account for non-homotheticity in demand for all products, services and housing and is left to future research.

6 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 their utilities across these markets. The results in this paper indicate that this is indeed the case: high-income households face greater consumption gains from moving to high per capita income markets than do low-income households.

I show that high-income households face 20 percent lower grocery costs in wealthy cities than in poor cities, while low-income households face 20 percent 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 spatial price indexes that can be used, for example, in real income mea-

[42] See Tables 18 and19 in Appendix Section E the the levels and ranks of the homothetic and non-homothetic grocery price indexes in all 23 markets with samples of 850 or more households.

Table 14: Cities Ranked According to Grocery Costs Calculated Using Homothetic and Non-Homothetic Models

City Ranks (Least Expensive to Most Expensive)

Six Least Expensive Cities According to Homothetic Model

		Non-Homothetic									
Market	Homothetic	$16,896	$26,715	$35,715	$41,526	$53,103	$60,442	$64,805	$82,576	$93,411	$146,566
Philadelphia	1	1	1	1	1	1	1	1	1	3	3
Washington, DC-Baltimore	2	4	5	2	2	2	2	2	5	11	14
Atlanta	3	7	4	3	3	3	3	3	14	18	21
Detroit	4	11	10	9	7	10	14	19	22	22	22
Columbus	5	2	2	4	4	4	4	6	17	20	20
Tampa	6	3	3	5	5	6	12	13	19	17	9

Six Most Expensive Cities According to Homothetic Model

		Non-Homothetic									
Market	Homothetic	$16,896	$26,715	$35,715	$41,526	$53,103	$60,442	$64,805	$82,576	$93,411	$146,566
Chicago	18	17	18	18	18	17	18	17	12	7	4
San Antonio	19	12	13	16	17	18	17	15	8	6	6
Sacramento	20	19	19	20	21	21	19	16	3	1	1
Minneapolis	21	20	21	21	20	20	21	21	23	23	23
Los Angeles	22	22	22	22	22	22	23	23	20	19	12
San Francisco	23	23	23	23	23	23	22	22	6	4	5

surement 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. The gap in relative grocery costs is large enough to suggest that this analysis is worthwhile. Even if we assume that preferences are homothetic within each of the households other consumption areas, the difference in relative grocery costs alone implies an economically-significant 2.4 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.[43]

The main goal of this paper is to measure how living costs vary across cities differently for consumers earning different incomes. In doing so, however, it also provides a methodological framework that could be applied more generally, for example, in analyzing how consumption costs vary differently across countries and over time for consumers earning different incomes. Though the detailed household-level data used in this paper might not be available in these contexts, it is conceivable that simulation techniques pioneered in the IO field could be used to identify the parameters of the model presented here using aggregate market- or country- level data. Where I have provided some sense of the large differences in cross-city price indexes across income levels, these extensions could shed light on how other key economic statistics, such as purchasing power parity, inflation, and the consumption gains from trade, vary with income.

[43] High-income households, who spend around 2 percent of their annual income on groceries, would face 0.4 percent lower living costs in wealthy cities, whereas low-income households, who spend around 10 percent of their annual income on groceries, would face 2 percent higher living costs in these locations.

Appendix

A Connection to Nested CES Utility Function

Consider the related utility function for some consumer i who is representative of all consumers with outside good expenditure Z:

$$U_i = \left\{ \sum_{m \in M} \left[\sum_{g \in G_m} [q_{mg} \exp(\gamma_m(Z)\beta_{mg})]^{\rho_m(Z)} \right]^{\frac{\rho(Z)}{\rho_m(Z)}} \right\}^{\frac{1}{\rho(Z)}}$$

where $\gamma_m(Z) = (1 + \gamma_m \ln Z)$, $\rho_m(Z) = \frac{1-\sigma_m(Z)}{\sigma_m(Z)}$ for $\sigma_m(Z) = 1 + \alpha_m^0 + \alpha_m^1 \ln Z$, $\rho(Z) = \frac{1-\sigma(Z)}{\sigma(Z)}$ for $\sigma(Z) = 1 + \alpha^0 + \alpha^1 \ln Z$. Suppose that this representative consumer faces the same prices $\mathbb{P}$ and has the same outside good expenditure Z as a group of consumers with the CES-nested log-logit utility defined in equation (1). A simple extension of the Anderson, de Palma, and Thisse (1987) result implies that the representative consumer's within-module expenditure shares will be identical to the within-module market shares of a group of consumers with the same outside good expenditure. Below, I show that the same is true for the representative consumer's between-module expenditure shares. For the sake of comparison, consider the representative consumer's log expenditure in module m relative to module $\bar{m}$:

$$\ln s_{im} - \ln s_{i\bar{m}} = -(\alpha^0 + \alpha^1 \ln Z) \left[\ln P_m(Z, \mathbb{P}_m) - \ln P_{\bar{m}}(Z, \mathbb{P}_{\bar{\mathbf{m}}}) \right] \tag{27}$$

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

$$P_m(z, \mathbb{P}_m) = \left[\sum_{g \in \mathbf{G_m}} \left(\frac{p_{mg}}{\exp(\beta_{mg}(1 + \gamma_m z))} \right)^{-(\alpha_m^0 + \alpha_m^1 z)} \right]^{\frac{1}{-(\alpha_m^0 + \alpha_m^1 z)}} \tag{28}$$

The representative consumer's relative log expenditure share is inversely proportional to the difference in the quality-adjusted price levels in the modules. However, the idiosyncratic consumer's relative log expenditure share is proportional to the difference in price-adjusted quality levels for each module. These relative log expenditure shares are equivalent because the quality-adjusted price levels defined in equation (28) are equal to the inverse of the price-adjusted quality levels defined in equation (48). The income-specific quality-adjusted price of a product g in module m is equal to the inverse of the income-specific price-adjusted quality of the same product:

$$\begin{aligned}
P_m(Z, \mathbb{P}_m) &= \left[\sum_{g \in \mathbf{G_m}} \left(\frac{p_{mg}}{\exp(\beta_{mg}(1 + \gamma_m \ln Z))} \right)^{-(\alpha_m^0 + \alpha_m^1 \ln Z)} \right]^{\frac{1}{-(\alpha_m^0 + \alpha_m^1 \ln Z)}} \\
&= \left[\sum_{g \in \mathbf{G_m}} \left(\left(\frac{\exp(\beta_{mg}(1 + \gamma_m \ln Z))}{p_{mg}} \right)^{-1} \right)^{-(\alpha_m^0 + \alpha_m^1 \ln Z)} \right]^{\frac{1}{-(\alpha_m^0 + \alpha_m^1 \ln Z)}} \\
&= \left[\sum_{g \in \mathbf{G_m}} \left(\frac{\exp(\beta_{mg}(1 + \gamma_m \ln Z))}{p_{mg}} \right)^{(\alpha_m^0 + \alpha_m^1 \ln Z)} \right]^{\frac{1}{-(\alpha_m^0 + \alpha_m^1 \ln Z)}} \\
&= \left\{ \left[\sum_{g \in \mathbf{G_m}} \left(\frac{\exp(\beta_{mg}(1 + \gamma_m \ln Z))}{p_{mg}} \right)^{(\alpha_m^0 + \alpha_m^1 \ln Z)} \right]^{\frac{1}{(\alpha_m^0 + \alpha_m^1 \ln Z)}} \right\}^{-1} \\
&= [V_m(z, \mathbb{P}_m)]^{-1}
\end{aligned}$$

B 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:

$$f(U_{iG}(\mathbb{Q},Z),Z) \quad \text{subject to} \quad \sum_{m\in\mathbf{M}}\sum_{g\in\mathbf{G_m}} p_{mg}q_{mg} + Z \leq Y_i,\ q_{mg} \geq 0\ \forall\ mg \in \mathbf{G} \tag{29}$$

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

$$\max_{\mathbb{Q},Z} U_{iG}(\mathbb{Q},Z) = \left\{\sum_{m\in M}\left[\sum_{g\in G_m} q_{mg}\exp(\gamma_m(Z)\beta_{mg} + \mu_m(Z)\varepsilon_{img})\right]^{\frac{\sigma(Z)-1}{\sigma(Z)}}\right\}^{\frac{\sigma(Z)}{\sigma(Z)-1}}$$

$$\text{subject to} \quad \sum_{m\in\mathbf{M}}\sum_{g\in\mathbf{G_m}} p_{mg}q_{mg} \leq Y_i - Z,\ q_{mg} \geq 0\ \forall\ mg \in \mathbf{G} \tag{30}$$

where $\gamma_m(Z) = (1+\gamma_m \ln Z)$, $\mu_m(Z) = \frac{1}{\alpha_m^0 + \alpha_1^1 \ln Z}$, and $\sigma(Z) = 1 + \alpha^0 + \alpha^1 \ln Z$. The solution to this problem for a consumer i with idiosyncratic utility draws ε_i is solved in the paper. Equations (9), (12), and (13) define the optimal grocery bundle, $\mathbb{Q}^*(Z) = \left\{\left\{q^*_{mg}(Z)\right\}_{g\in\mathbf{G},}\right\}_{m\in\mathbf{M}}$:

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

where

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

and

$$\tilde{p}_{img} = \frac{\exp(\gamma_m(Z)\beta_{mg} + \mu_m(Z)\varepsilon_{ig})}{p_{mg}}$$

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

$$\begin{aligned}
\tilde{U}_{iG}(Z) &= U_{iG}(\mathbb{Q}^*(Z),Z) \\
&= \left\{\sum_{m\in M}\left[\left((Y_i - Z)\frac{(\tilde{p}_{img})^{\sigma(Z)}}{P_i(Z)^{1-\sigma(Z)}}\right)\mathbb{I}\left[g = \arg\max_{g\in\mathbf{G_m}}\tilde{p}_{img}\right]\right]^{\frac{\sigma(Z)-1}{\sigma(Z)}}\right\}^{\frac{\sigma(Z)}{\sigma(Z)-1}} \\
&= \frac{Y_i - Z}{P_i(Z)^{1-\sigma(Z)}}\left\{\sum_{m\in M}\left[\tilde{p}^{\sigma(Z)}_{img}\mathbb{I}\left[g = \arg\max_{g\in\mathbf{G_m}}\tilde{p}_{img}\right]\right]^{\frac{\sigma(Z)-1}{\sigma(Z)}}\right\}^{\frac{\sigma(Z)}{\sigma(Z)-1}} \\
&= \frac{Y_i - Z}{P_i(Z)^{1-\sigma(Z)}}\left\{\sum_{m\in M}\left(\max_{g\in\mathbf{G_m}}\tilde{p}_{img}\right)^{\sigma(Z)-1}\right\}^{\frac{\sigma(Z)}{\sigma(Z)-1}} \\
&= \frac{Y_i - Z}{P_i(Z)^{1-\sigma(Z)}}P_i(Z)^{\sigma(Z)} \\
&= \frac{Y_i - Z}{P_i(Z)}
\end{aligned} \tag{31}$$

We can now express problem (29) to be a choice over one variable, Z:

$$\max_Z f(\tilde{U}_{iG}(Z),Z) \tag{32}$$

The first order condition to the utility maximization problem defined in problem (32) with respect to Z is:

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

Substituting the maximized grocery expenditure conditional on Z, $\tilde{U}_{iG}(Z)$, from equation (31) 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, ε_i, and model parameters:

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

If we assume that $f(U_{iG}(\mathbb{Q},Z), Z)$ is additive in grocery utility and non-grocery utility, i.e. $f(U_{iG}(\mathbb{Q},Z), Z) = U_{iG}(\mathbb{Q},Z) + Z$, the formula above simplifies to:

$$Y_i \quad = \quad Z + \frac{P_i(Z)^2 - P_i(Z)}{P_i'(Z)}$$

Taking the derivative of income with respect to outside good expenditure, Z, we can see that the outside good will be normal if the price vector is such that:

$$\frac{\partial}{\partial Z} \frac{P_i(Z)^2 - P_i(Z)}{P_i'(Z)} > -1$$

C Condition for Consistency of Grocery Expenditure Share Across Cities Under AIDS

Suppose that consumers' expenditure allocation between grocery and non-grocery products is governed by an Almost Ideal Demand System (AIDS). The Engel curve, or expenditure share on grocery products, implied by this system is given by

$$w_{i,c}^G = \alpha_1 + \beta \left(\ln y_i - \ln P_c(y_i)\right) + \gamma \left(\ln P_c^G(y_i) - \ln P_c^Z(y_i)\right) \tag{33}$$

where $w_{i,c}^G$ is the budget share for grocery products by household i in city c; y_i is the nominal size-adjusted household income of household i; $P_c^G(y_i)$ and $P_c^Z(y_i)$ are the price indexes faced by households with income y_i in city c for grocery and non-grocery goods, respectively; and $P_c(y_i)$ is the composite price of consumption in city c for a household with income y_i.

For households, different cities with the same size-adjusted income to have, on average, the same food expenditure shares, we must have that:

$$\ln\left(\frac{P_c(y_i)}{P_{c'}(y_i)}\right) = -\frac{\gamma}{\beta}\left[\ln\left(\frac{P_c^G(y_i)}{P_c^Z(y_i)}\right) - \ln\left(\frac{P_{c'}^G(y_i)}{P_{c'}^Z(y_i)}\right)\right] \tag{34}$$

for all y_i across each city pair c and c'. This expression implies that the difference between the log ratios of grocery-to-non-grocery prices in each city must be linear in the difference between the log composite price indexes. In the AIDS, the composite price index for each city c and income level y_i is a function of the income-specific grocery and non-grocery price indexes in that city:

$$\ln P_c(y_i) = \alpha_0 + \ln P_c^G(y_i) + (1-\alpha_1)\left(\ln P_c^Z(y_i) - \ln P_c^G(y_i)\right) + \frac{\gamma}{2}\left(\ln P_c^G(y_i) - \ln P_c^Z(y_i)\right)^2 \tag{35}$$

The left-hand side of condition 34 can, therefore, be expressed in terms of the relative grocery and relative non-grocery price indexes:

$$\ln\left(\frac{P_c(y_i)}{P_{c'}(y_i)}\right) = \ln\left(\frac{P_c^G(y_i)}{P_{c'}^G(y_i)}\right) - (1-\alpha_1)\left[\ln\left(\frac{P_c^G(y_i)}{P_c^Z(y_i)}\right) - \ln\left(\frac{P_{c'}^G(y_i)}{P_{c'}^Z(y_i)}\right)\right] + \frac{\gamma}{2}\left[\left(\ln\left(\frac{P_c^G(y_i)}{P_c^Z(y_i)}\right)\right)^2 - \left(\ln\left(\frac{P_{c'}^G(y_i)}{P_{c'}^Z(y_i)}\right)\right)^2\right] \tag{36}$$

Substituting into condition 34 yields

$$\ln\left(\frac{P_c^G(y_i)}{P_{c'}^G(y_i)}\right) = \left((1-\alpha_1) - \frac{\gamma}{\beta}\right)\left[\ln\left(\frac{P_c^G(y_i)}{P_c^Z(y_i)}\right) - \ln\left(\frac{P_{c'}^G(y_i)}{P_{c'}^Z(y_i)}\right)\right] - \frac{\gamma}{2}\left[\left(\ln\left(\frac{P_c^G(y_i)}{P_c^Z(y_i)}\right)\right)^2 - \left(\ln\left(\frac{P_{c'}^G(y_i)}{P_{c'}^Z(y_i)}\right)\right)^2\right] \tag{37}$$

Note that this condition indicates that the ratio between grocery and non-grocery prices need not, however, be identical across cities in order for consumers with the same incomes to make the same upper-level expenditure allocation decision.

In the special case where this occurs, such that the ratio between grocery and non-grocery prices was identical across two cities, condition 37 implies that grocery expenditure shares will be equal only if the grocery (and, therefore, also non-grocery and composite) price index is identical across the cities. This need not be the case more generally, however, in order for the observed empirical regularity in cross-city grocery expenditure shares to hold.

D Derivations

D.1 Within-Module Consumption Decision

Consumer i, spending Z on the outside good, 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 module-level budget and non-negativity constraints:

$$\sum_{m \in \mathbf{M}} \sum_{g \in \mathbf{G_m}} p_{mg} q_{mg} \leq w_m, q_{mg} \geq 0$$

Recall that the additive log-logit functional form implies that consumers optimally purchase a positive quantity only one product in a module. This product maximizes their marginal utility of expenditure in a module conditional on their outside good expenditure:[44]

$$g^*_{im}(Z) = \arg\max_{g \in \mathbf{G_m}} \frac{\exp(\gamma_m(Z)\beta_{mg} + \mu_m(Z)\varepsilon_{img})}{p_{mg}} \tag{38}$$

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:

$$\begin{aligned} \mathbb{Q}^*_{im}(w_m, Z) &= (q^*_{im1}(w_m, Z), \ldots, q^*_{imG_m}(w_m, Z)) \\ \text{where} \quad q^*_{img}(w_m) &= \begin{cases} w_m / p_{mg} & \text{if } g = \arg\max\limits_{g \in \mathbf{G_m}} \dfrac{\exp(\gamma_m(Z)\beta_{mg} + \mu_m(Z)\varepsilon_{img})}{p_{mg}} \\ 0 & \text{otherwise} \end{cases} \end{aligned} \tag{39}$$

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.

D.2 Across-Module Consumption Decision

Consumer i, spending Z on the outside good, chooses how to allocate expenditures between modules by selecting $w_1, ..., w_M$ to maximize

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

subject to

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

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

We simplify the expression for the target utility function by denoting consumer i's marginal utility from expenditure in module m as the inverse of A_{im}:

$$\max_{g\in\mathbf{G_m}} \frac{\exp(\gamma_m(Z)\beta_{mg}+\mu_m(Z)\varepsilon_{img})}{p_{mg}} = \frac{1}{A_{im}} \tag{40}$$

The within-module allocation decision now simplifies to:

$$\mathbf{w_i^*}(Z) = (w_{i1}^*(Z), ..., w_{iM}^*(Z)) = \underset{\sum_{m\in\mathbf{M}} w_m \le Y_i - Z}{\arg\max} \left\{ \sum_{m\in M} \left[\frac{w_m}{A_{im}}\right]^{\frac{\sigma(Z)-1}{\sigma(Z)}} \right\}^{\frac{\sigma(Z)}{\sigma(Z)-1}} \tag{41}$$

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:

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

where λ 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':

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

$$\frac{1}{A_{im'}} \left[\frac{w_{m'}}{A_{im'}}\right]^{-\frac{1}{\sigma(Z)}} = \frac{1}{A_{im}} \left[\frac{w_m}{A_{im}}\right]^{-\frac{1}{\sigma(Z)}}$$

$$w_{m'} = w_m \left[\frac{A_{im'}}{A_{im}}\right]^{1-\sigma(Z)}$$

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

$$Y_i - Z = \sum_{m'\in\mathbf{M}} w_{m'}$$

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

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

The solution to problem (41) is, therefore,

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

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

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

Substituting from equation (40) 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:

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

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

D.3 Within-Module Market Expenditure Shares

Equation (9) states that:

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

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 outside good 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, $\varepsilon_{\mathbf{im}}$:

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

The final equality holds because the idiosyncratic utilities, ε_{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:

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

I interpret the expected expenditure share function derived above as the expected share of expenditure that a group of households with the same outside good 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 outside good expenditure, Z, and facing a common vector of module prices, $\mathbb{P}_m$:

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

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 $\bar{g}_m$[45]:

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

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

$$\ln(s_{mg|m}(Z,\mathbb{P}_m)) - \ln(s_{m\bar{g}|m}(Z,\mathbb{P}_m)) = (\alpha_m^0+\alpha_m^1\ln Z)\left[(\beta_{mg}-\beta_{m\bar{g}})(1+\gamma_m\ln Z)-(\ln p_{mg}-\ln p_{m\bar{g}})\right] \quad (42)$$

Equation (42) defines the expected within-module expenditure share of a set of households with outside good 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 α_m, γ_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 α_m and γ_m parameters, as well as each β_{mg} parameter relative to $\beta_{m\bar{g}_m}$, *i.e.* $\{\beta_{mg}-\beta_{m\bar{g}_m}\}_{g\in\mathbf{G_m}}$.

D.4 Between-Module Relative Market Expenditure Shares

I now want to generate a similar estimation equation that can be used to identify α^0, α^1, and $\{\beta_{\bar{g}_m}\}_{g\in\mathbf{G_m}}$ using data on module-level income-specific market shares. Equations (12) and (13) together characterize the optimal cross-module expenditure allocation for consumer i conditional on this consumer's idiosyncratic utility draws for each product in each module. These equations are:

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

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

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

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

When deriving the within-module relative market share, equation (42) 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 taking 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

[45]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 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.[46]

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 $\bar{m}$:

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

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:

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

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:

$$\mathbb{E}_\varepsilon\left[\ln s_{im}(Z,\mathbb{P}) - \ln s_{i\bar{m}}(Z,\mathbb{P})\right] = (\sigma(Z)-1)\left\{\mathbb{E}_\varepsilon\left[\ln\left(\max_{g\in\mathbf{G_m}} \tilde{p}_{img}\right)\right] - \mathbb{E}_\varepsilon\left[\ln\left(\max_{g\in\mathbf{G_{\bar{m}}}} \tilde{p}_{i\bar{m}g}\right)\right]\right\} \tag{43}$$

Consider the two expectation terms in equation (43). Both take the same form, and thus I only solve for the first expectation:

$$\mathbb{E}_\varepsilon\left[\ln\left(\max_{g\in\mathbf{G_m}} \tilde{p}_{img}\right)\right] \tag{44}$$

The expectation term defined in equation (44) 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:

$$\begin{aligned}
\mathbb{E}_\varepsilon\left[\ln\left(\max_{g\in\mathbf{G_m}} \tilde{p}_{img}\right)\right] &= \mathbb{E}_\varepsilon\left[\ln\left(\max_{g\in\mathbf{G_m}} \frac{\exp(\gamma_m(Z)\beta_{mg} + \mu_m(Z)\varepsilon_{img})}{p_{mg}}\right)\right] \\
&= \mathbb{E}_\varepsilon\left[\max_{g\in\mathbf{G_m}} \ln\left(\frac{\exp(\gamma_m(Z)\beta_{mg} + \mu_m(Z)\varepsilon_{img})}{p_{mg}}\right)\right] \\
&= \mathbb{E}_\varepsilon\left[\max_{g\in\mathbf{G_m}} \gamma_m(Z)\beta_{mg} - \ln p_{mg} + \mu_m(Z)\varepsilon_{img}\right] \\
&= \mu_m(Z)\mathbb{E}_\varepsilon\left[\max_{g\in\mathbf{G_m}} (\gamma_m(Z)\beta_{mg} - \ln p_{mg})/\mu_m(Z) + \varepsilon_{img}\right]
\end{aligned} \tag{45}$$

[46]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:

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

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 (42), 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 the outside good 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 $\bar{g}$.

De Palma and Kilani (2007) show that, for an additive random utility model with $u_i = \nu_i + \varepsilon_i$, $i = 1, \dots, n$ 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 (45) 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:

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

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 $dz = -(1/x)dx$:

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

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 (44):

$$\begin{aligned}
\mathbb{E}_\varepsilon\left[\ln\left(\max_{g \in \mathbf{G_m}} \tilde{p}_{img}\right)\right] &= \mu_m(Z)\ln\left[\sum_{g \in \mathbf{G_m}} \exp((\gamma_m(Z)\beta_{mg} - \ln p_{mg})/\mu_m(Z))\right] \\
&= \mu_m(Z)\ln\left[\sum_{g \in \mathbf{G_m}} \left(\frac{\exp(\gamma_m(Z)\beta_{mg})}{p_{mg}}\right)^{\frac{1}{\mu_m(Z)}}\right] \\
&= \ln\left[\sum_{g \in \mathbf{G_m}} \left(\frac{\exp(\gamma_m(Z)\beta_{mg})}{p_{mg}}\right)^{\frac{1}{\mu_m(Z)}}\right]^{\mu_m(Z)} \qquad (46)
\end{aligned}$$

Plugging this result back into equation (43) yields the expected relative module expenditure share for consumer i in terms of product prices and model parameters:

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

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

$$\mathbb{E}_{\varepsilon}\left[\ln s_{im}(Z,\mathbb{P}) - \ln s_{i\bar{m}}(Z,\mathbb{P})\right] \quad = \quad -(\alpha^0 + \alpha^1 \ln Z)\left[\ln V_m(Z,\mathbb{P}_m) - \ln V_{\bar{m}}(Z,\mathbb{P}_{\bar{\mathbf{m}}})\right] \tag{47}$$

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

$$V_m(Z,\mathbb{P}_m) = \left[\sum_{g\in\mathbf{G_m}} \left(\frac{\exp(\beta_{mg}(1+\gamma_m \ln Z))}{p_{mg}}\right)^{\alpha_m^0+\alpha_m^1 \ln Z}\right]^{\frac{1}{\alpha_m^0+\alpha_m^1 \ln Z}} \tag{48}$$

Equations (47) and (48) 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 α^0, α^1, as well as α_m, γ_m, β_{mg} for all $g \in G_m$, and $\alpha_{\bar{m}}$, $\gamma_{\bar{m}}$, $\beta_{\bar{m}g}$ for all $g \in G_{\bar{m}}$.

E Market Regions

Table 15: Regional Categorizations for Sample Markets

Market Code	Market	Region	Neighboring Region
1	Des Moines	MW	SC
2	Little Rock	SE	SC
3	Omaha	NW	MW
4	Syracuse	NE	MW
5	Albany	NE	MW
6	Birmingham	SE	SC
7	Richmond	NE	SE
8	Louisville	MW	SE
9	Grand Rapids	MW	NE
10	Jacksonville	SE	SC
11	Memphis	SE	SC
12	Raleigh-Durham	SE	SC
13	Nashville	SE	SC
14	Salt Lake City	SW	SC
15	Charlotte	SE	SC
16	Columbus	MW	NE
17	San Antonio	SW	SC
18	Indianapolis	SE	MW
19	Orlando	SE	SC
20	Milwaukee	MW	NW
21	Hartford-New Haven	NE	MW
22	Kansas City	MW	SE
23	Sacramento	SW	NW
24	New Orleans-Mobile	SE	SC
25	Oklahoma City-Tulsa	SC	MW
26	Cincinnati	MW	NE
27	Portland, Or	NW	SW
28	Buffalo-Rochester	NE	MW
29	Pittsburgh	MW	NE
30	Tampa	SE	SC
31	Denver	SW	NW
32	St. Louis	MW	NE
33	San Diego	SW	NW
34	Cleveland	MW	NE
35	Minneapolis	MW	NW
36	Phoenix	SW	SC
37	Seattle	NW	SW
38	Miami	SE	SC
39	Atlanta	SE	SC
40	Houston	SC	SC
41	Dallas	SC	SW
42	Detroit	MW	NE
43	Boston	NE	MW
44	Philadelphia	NE	MW
45	San Francisco	NW	SW
46	Washington, DC-Baltimore	NE	SE
47	Chicago	MW	NE
48	Los Angeles	SW	SC
49	New York	NE	MW

F Procedure for Obtaining Standard Errors of Upper-Level Demand Parameters

I estimate the parameters of the model sequentially. Recall that the full set of demand parameters, θ, are partitioned into M sets of lower-level module-specific parameters, θ_{1m} for each module m, that are identified using module-specific sub-samples of the data and a single set of parameters, θ_2, whose identification requires data from all modules. 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 subsamples of the data. In this Appendix, I outline how I apply these methods to calculate the covariance matrix of the upper-level demand parameters.

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

The first step in my estimation is to obtain estimates for $\theta_1 = \left\{\alpha_m^0, \alpha_m^1, \gamma_m, \{\tilde{\beta}_{mg}\}_{g\in\mathbf{G_m}, g\neq\bar{g}_m}\right\}_{m=1,\dots,M}$, 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:

$$\mathbb{E}[f(\mathbf{X};\theta_1)] = 0 \tag{49}$$

where $f(\mathbf{X};\theta_1) = \mathbf{Z}_1(\mathbf{X})'\nu(\mathbf{X};\theta_1)$, $\mathbf{Z}_1(\mathbf{X})$ is a stacked vector of L_{1m} module-specific instruments, $\mathbf{Z}_{1m}(\mathbf{X})$, for each module m. $\nu(\mathbf{X};\theta_1)$ is the error in the relative within-module expenditure share equation. For income group, k, market t, and product g in module m, this error is defined as:

$$\nu_{gkt}(\mathbf{X_m};\theta_{1m}) = \ln\left(\frac{s_{gkt}}{s_{\bar{g}kt}}\right) - \left(\alpha_m^0 + \alpha_m^1 y_{kt}\right)\left[\tilde{\beta}_{mg}\left(1+\gamma_m y_{kt}\right) - \ln\left(\frac{p_{gkt}}{p_{\bar{g}_m kt}}\right)\right] \tag{50}$$

The fact that these errors depend only on module-specific data, $\mathbf{X_m}$, and parameters, θ_{1m}, enables me to partition 49 into module-specific auxiliary moments:

$$\mathbb{E}[f(\mathbf{X_m};\theta_{1m})] = 0$$

for $f(\mathbf{X_m};\theta_{1m}) = \mathbf{Z}_{1m}(\mathbf{X_m})'\nu(\mathbf{X_m};\theta_{1m})$.

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 are obtained by minimizing module-specific GMM objective functions as follows:

$$\hat{\theta}_{1m} = \arg\min_{\theta_{1m}} \hat{f}(\mathbf{X_m};\theta_{1m})'\hat{\mathbf{W}}_{\mathbf{1m}}\hat{f}(\mathbf{X_m};\theta_{1m})$$

where $\hat{f}(\mathbf{X_m};\theta_{1m}) = \frac{1}{\sum_{k,t} N_{mkt}} \sum_{k,t}\sum_{g\in\mathbf{G_m}_{kt}} f_{gkt}(\mathbf{X_m};\theta_{1m})$ is the sample analog of $\mathbb{E}[f(\mathbf{X_m};\theta_{1m})]$; $f_{gkt}(\mathbf{X_m};\theta_{1m}) = \mathbf{Z}_{1gkt}(\mathbf{X_m})'\nu_{gkt}(\mathbf{X_m};\theta_{1m})$; $\mathbf{G_m}_{kt}$ is the set of N_{mkt} module m non-base (*i.e.*, $g \neq \bar{g}_m$) products purchased by income-group k households in market t; and $\mathbf{Z}_{1gkt}$ is the $1 \times L_{1m}$ ($L_{1m} \leq K_{1m}$) vector of instruments for a product g-income group k-market t observation. $\hat{\mathbf{W}}_{1m} = \left[\frac{1}{\sum_{k,t} N_{mkt}} \sum_{k,t}\sum_{g\in\mathbf{G_m}_{kt}} f_{gkt}(\mathbf{X_m};\tilde{\theta}_{1m}) f_{gkt}(\mathbf{X_m};\tilde{\theta}_{1m})'\right]^{-1}$ is the efficient weighting matrix, calculated using consistent first-stage estimates of θ_{1m}:

$$\tilde{\theta}_{1m} = \arg\min_{\theta_{1m}} \hat{f}(\mathbf{X_m};\theta_{1m})'\tilde{\mathbf{W}}_{\mathbf{1m}}\hat{f}(\mathbf{X_m};\theta_{1m})$$

for $\tilde{\mathbf{W}}_{1m} = \left[\frac{1}{\sum_{k,t} N_{mkt}} \sum_{k,t}\sum_{g\in\mathbf{G_m}_{kt}} \mathbf{Z}_{1gkt}\mathbf{Z}_{1gkt}'\right]^{-1}$.

Assuming that the random components of the M module-specific auxiliary models are independent, the variance-covariance matrix of $\hat{\theta}_1$, Ω_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 θ_{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

$$\hat{F}_{\theta_{1m}} = \frac{1}{\sum_{k,t} N_{mkt}} \sum_{k,t} \sum_{g \in \mathbf{G_m}_{kt}} \nabla_{\theta_{1m}} f_{gkt}(\mathbf{X_m}; \hat{\theta}_{1m}) \; (K_1 \times L_1)$$

and

$$\hat{V}_{ff} = \hat{\mathbf{W}}_1 = \frac{1}{\sum_{k,t} N_{mkt}} \sum_{k,t} \sum_{g \in \mathbf{G_m}_{kt}} f_{gkt}(\mathbf{X_m}; \hat{\theta}_{1m}) f_{gkt}(\mathbf{X_m}; \hat{\theta}_{1m})' \; (L_1 \times L_1).$$

F.2 Step 2: Sequential Estimation of $\theta_2 = \{\alpha^0, \alpha^1, \{\beta_{m\bar{g}_m}\}_{m=1,...,M, m \neq \bar{m}}\}$

In the second step of the sequential estimation procedure, I estimate $\theta_2 = \{\alpha^0, \alpha^1, \{\beta_{m\bar{g}_m}\}_{m=1,...,M, m\neq\bar{m}}\}$. These $K_2 = 1 + M$ parameters are identified by the following exogeneity restriction:

$$G = \mathbb{E}[h(\mathbf{X}; \theta_1, \theta_2)] = 0 \tag{51}$$

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. For income group, k, market t, and module m, this error is defined as:

$$u_{kmt}(\mathbf{X}; \theta_1, \theta_2) = \ln\left(\frac{s_{mkt}}{s_{\bar{m}kt}}\right) - (\alpha^0 + \alpha^1 y_{kt}) \left[\Delta V_{1m\bar{m}}(y_{kt}, \mathbb{P}_{mkt}, \mathbb{P}_{\bar{m}kt}, \theta_1) + \beta_{m\bar{g}_m}(1 + \gamma_m y_{kt})\right],$$

where

$$\Delta V_{1m\bar{m}}(y_{kt}, \mathbb{P}_{mkt}, \mathbb{P}_{\bar{m}kt}, \theta_1) = \ln V_{1m}(y_{kt}, \mathbb{P}_{mkt}, \theta_1) - \ln V_{1\bar{m}}(y_{kt}, \mathbb{P}_{\bar{m}kt}, \theta_1),$$

and

$$V_{1m}(y_{kt}, \mathbb{P}_{mkt}, \theta_1) = \left[\sum_{g \in \mathbf{G_m}} \left(\frac{\exp(\tilde{\beta}_{mg}(1 + \gamma_m y_{kt}))}{p_{gkt}}\right)^{-(\alpha_m^0 + \alpha_m^1 y_{kt})}\right]^{\frac{1}{-(\alpha_m^0 + \alpha_m^1 y_{kt})}}$$

The first stage $\hat{\theta}_1$ estimates are inputs into the sample moment condition used to estimate the $1 \times K_2$ vector of θ_2 parameters, denoted $\hat{\theta}_2$. These upper-level parameters are estimated using two-step GMM:

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

wheree $\hat{h}(\mathbf{X}; \hat{\theta}_1, \theta_2) = \frac{1}{\sum_{k,t} N_{kt}} \sum_{k,t} \sum_{m \in \mathbf{M}_{kt}} h_{mkt}(\mathbf{X}; \hat{\theta}_1, \theta_2)$ is the sample analog of $\mathbb{E}[h(\mathbf{X}; \hat{\theta}_1, \theta_2)]$; $h_{mkt}(\mathbf{X}; \hat{\theta}_1, \theta_2) = \mathbf{Z}'_{2mkt} u_{mkt}(\mathbf{X}; \hat{\theta}_1, \theta_2)$; $\mathbf{M}_{kt}$ is the set of N_{kt} non-base modules (*i.e.*, $m \neq \bar{m}$) purchased by income-group k households in market t; and $\mathbf{Z}_{2mkt}$ is the $1 \times L_2$ vector of instruments for a module m-income group k-market t observation. $\hat{\mathbf{W}}_2 = \left[\frac{1}{\sum_{k,t} N_{kt}} \sum_{k,t} \sum_{m \in \mathbf{M}_{kt}} h_{mkt}(\mathbf{X}; \hat{\theta}_1, \tilde{\theta}_2) h_{mkt}(\mathbf{X}; \hat{\theta}_1, \tilde{\theta}_2)'\right]^{-1}$ is the optimal weighting matrix, where $\tilde{\theta}_2$ are consistent first-stage estimates of θ_2 that minimize a GMM objective function as follows:

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

for $\tilde{\mathbf{W}}_2 = \left[\frac{1}{\sum\limits_{k,t} N_{kt}} \sum_{k,t} \sum_{m \in \mathbf{M}_{kt}} \mathbf{Z}_{2mkt}\mathbf{Z}'_{2mkt}\right]^{-1}$.

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 θ_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

$$\hat{H}_{\theta_2} = \frac{1}{\sum\limits_{k,t} N_{kt}} \sum_{k,t} \sum_{m \in \mathbf{M}_{kt}} \nabla_{\theta_2} h_{mkt}(\mathbf{X}; \hat{\theta}_1, \hat{\theta}_2 \; (K_2 \times L_2)$$

and

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

In order to account for the measurement error from the use of first state estimates, we need to treat θ_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, Ω_{θ_2}.[47] 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;

$$\hat{H}_{\theta_1} = \frac{1}{\sum\limits_{k,t} N_{kt}} \sum_{k,t} \sum_{m \in \mathbf{M}_{kt}} \nabla_{\theta_1} h_{mkt}(\mathbf{X}; \hat{\theta}_1, \hat{\theta}_2) \; (K_1 \times L_2);$$

and

$$\hat{V}_{fg} = \hat{V}'_{gf} = \frac{1}{\sum\limits_{k,t} N_{kt}} \sum_{k,t} \sum_{m \in \mathbf{M}_{kt}} h_{mkt}(\mathbf{X}; \hat{\theta}_1, \hat{\theta}_2)[\frac{1}{N_{mkt}} \sum_{g \in \mathbf{G_m}_{kt}} f_{gkt}(\mathbf{X}; \hat{\theta}_{1m})]' \; (L_2 \times L_1).$$

[47] Newey (1984) shows that, when $L_1 = K_1$ and $L_2 = K_2$, the asymptotic covariance matrix Ω_{θ_2} of the second step estimator $\hat{\theta}_2$ is given by:

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

where $\hat{\Omega}_{\theta_1}$, $\tilde{\Omega}_{\theta_2}$, $\hat{H}_{\theta_2}$, and $\hat{F}_{\theta_1}$ are as defined above and $\hat{H}_{\theta_1} = \frac{1}{\sum\limits_{k,t} N_{kt}} \sum_{k,t} \sum_{m \in \mathbf{M}_{kt}} \nabla_{\theta_1} h_{mkt}(\mathbf{X}; \hat{\theta}_1, \theta_2)$. This equation cannot be applied directly to estimate Ω_{θ_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).

G Cross-city Price Comparisons Accounting for Non-Homothetic Demand for Quality and Price Sensitivity

Table 16: City-Income Specific Price Index Regressions

Dependent Variable: Ln(Price Index for Representative Consumer k in City c)					
		Model allowing for NH in:			
		Quality		Price and Quality	
Ln(Per Capita Income$_c$)	β_1	2.412***	2.290*	2.620**	2.523*
		[0.996]	[1.22]	[1.235]	[1.426]
Ln(Per Capita Income$_c$)	β_2	-0.217**	-0.201*	-0.249***	-0.226*
*Ln(Household Income$_k$)		[0.0915]	[0.112]	[0.105]	[0.131]
Ln(Population$_c$)	β_3	-	0.040	-	0.056
		-	[0.232]	-	[0.256]
Ln(Population$_c$)	β_4	-	-0.005	-	-0.007
*Ln(Household Income$_k$)		-	[0.021]	-	[0.024]
Household Income Fixed Effects		Yes	Yes	Yes	Yes
Observations		230	230	230	230
R-Squared		0.03	0.036	0.04	0.035

*** p<0.01, ** p<0.05, * p<0.1

Standard errors in brackets.

Figure 11: Variation in Elasticity of Grocery Costs with respect to City Income Across Household Income Levels Allowing for Non-Homotheticity in Price and Quality

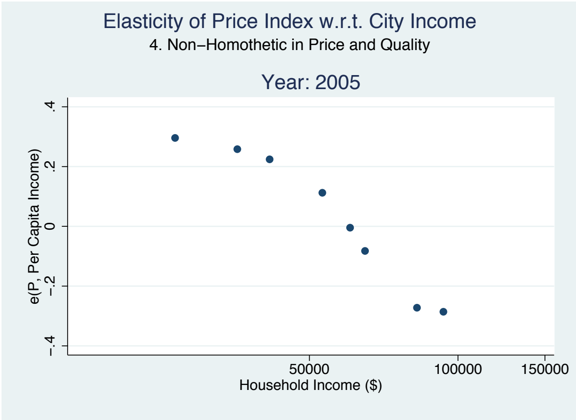

H Income- and City-Specific Price Indexes

See subsequent pages.

Table 18: City-Specific Price Indexes Calculated Using Homothetic and Non-Homothetic Models

	Price Index (Relative to the Homothetic Index for New York)										
		Non-Homothetic									
Market	Homothetic	$16,896	$26,715	$35,715	$41,526	$53,103	$60,442	$64,805	$82,576	$93,411	$146,566
Philadelphia	0.83	0.77	0.81	0.83	0.84	0.84	0.84	0.84	0.84	0.84	0.82
Washington, DC-Baltimore	0.83	0.86	0.87	0.86	0.85	0.85	0.87	0.88	0.96	0.99	1.03
Atlanta	0.87	0.88	0.85	0.86	0.86	0.88	0.89	0.91	1.00	1.05	1.16
Detroit	0.93	0.98	0.98	0.95	0.95	0.98	1.02	1.06	1.18	1.21	1.18
Columbus	0.94	0.77	0.83	0.88	0.90	0.92	0.93	0.94	1.03	1.07	1.15
Tampa	0.95	0.79	0.83	0.88	0.91	0.96	1.00	1.02	1.07	1.05	0.97
Houston	0.97	0.89	0.91	0.94	0.95	0.97	0.98	0.98	0.98	0.97	0.96
St. Louis	0.98	1.12	1.05	1.00	0.99	0.97	0.97	0.98	1.02	1.04	1.07
Phoenix	0.98	0.94	0.95	0.96	0.97	0.96	0.95	0.94	0.93	0.94	0.98
Dallas	0.98	0.88	0.90	0.93	0.94	0.95	0.96	0.96	1.00	1.03	1.10
New York	1.00	0.88	0.99	1.04	1.04	1.00	0.95	0.93	0.84	0.82	0.81
Miami	1.01	0.94	0.94	0.95	0.96	0.98	0.98	0.98	0.97	0.98	1.03
Boston	1.02	1.04	1.06	1.05	1.04	1.02	1.01	1.00	1.02	1.03	1.05
Charlotte	1.02	1.06	1.01	0.99	0.99	1.01	1.03	1.05	1.05	1.02	0.98
Denver	1.02	1.04	1.02	1.02	1.03	1.04	1.03	1.03	1.00	0.98	0.96
Buffalo-Rochester	1.03	1.27	1.19	1.11	1.08	1.07	1.08	1.09	1.13	1.13	1.08
Seattle	1.03	1.05	1.03	1.03	1.02	1.01	1.00	0.99	0.99	1.01	1.05
Chicago	1.05	1.09	1.10	1.08	1.07	1.05	1.04	1.04	1.00	0.97	0.90
San Antonio	1.08	1.00	1.02	1.05	1.06	1.06	1.04	1.03	0.98	0.96	0.95
Sacramento	1.10	1.21	1.15	1.13	1.13	1.11	1.07	1.04	0.86	0.80	0.74
Minneapolis	1.13	1.23	1.20	1.15	1.12	1.09	1.09	1.10	1.19	1.24	1.26
Los Angeles	1.15	1.29	1.23	1.18	1.17	1.17	1.17	1.16	1.10	1.06	1.01
San Francisco	1.22	1.33	1.29	1.26	1.24	1.20	1.14	1.10	0.96	0.93	0.93
Correlation with Homothetic Index		0.83	0.88	0.94	0.96	0.97	0.92	0.83	0.21	0.03	-0.11

Table 19: Cities Ranked According to Grocery Costs Calculated Using Homothetic and Non-Homothetic Models

	City Ranks (Least Expensive to Most Expensive)										
		Non-Homothetic									
Market	Homothetic	\$16,896	\$26,715	\$35,715	\$41,526	\$53,103	\$60,442	\$64,805	\$82,576	\$93,411	\$146,566
Philadelphia	1	1	1	1	1	1	1	1	1	3	3
Washington, DC-Baltimore	2	4	5	2	2	2	2	2	5	11	14
Atlanta	3	7	4	3	3	3	3	3	14	18	21
Detroit	4	11	10	9	7	10	14	19	22	22	22
Columbus	5	2	2	4	4	4	4	6	17	20	20
Tampa	6	3	3	5	5	6	12	13	19	17	9
Houston	7	8	7	7	8	8	10	10	9	8	8
St. Louis	8	18	16	12	12	9	8	9	16	16	17
Phoenix	9	9	9	10	10	7	5	5	4	5	11
Dallas	10	5	6	6	6	5	7	7	13	14	19
New York	11	6	11	15	15	12	6	4	2	2	2
Miami	12	10	8	8	9	11	9	8	7	10	13
Boston	13	14	17	17	16	15	13	12	15	15	15
Charlotte	14	16	12	11	11	14	16	18	18	13	10
Denver	15	13	14	13	14	16	15	14	11	9	7
Buffalo-Rochester	16	21	20	19	19	19	20	20	21	21	18
Seattle	17	15	15	14	13	13	11	11	10	12	16
Chicago	18	17	18	18	18	17	18	17	12	7	4
San Antonio	19	12	13	16	17	18	17	15	8	6	6
Sacramento	20	19	19	20	21	21	19	16	3	1	1
Minneapolis	21	20	21	21	20	20	21	21	23	23	23
Los Angeles	22	22	22	22	22	22	23	23	20	19	12
San Francisco	23	23	23	23	23	23	22	22	6	4	5

References

ADAMSON, D. W., D. E. CLARK, AND M. D. PARTRIDGE (2004): "Do Urban Agglomeration Effects and Household Amenities Have a Skill Bias?," *Journal of Regional Science*, 44, 201–224.

AGUIAR, M., AND E. HURST (2005): "Consumption versus Expenditure," *Journal of Political Economy*, 113(5), 919–948.

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

ALMÅS, 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): "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.

ASKER, J. (2004): "Measuring Cost Advantages from Exclusive Dealing: An Empirical Study of Beer Distribution," *Harvard mimeo*.

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

BACOLOD, M., B. S. BLUM, AND W. C. STRANGE (2009): "Skills in the city," *Journal of Urban Economics*, 65(2), 136–153.

BAI, J., AND S. NG (2010): "Instrumental Variable Estimation in a Data Rich Environment," *Econometric Theory*, 26, 1577–1606.

BILS, M., AND P. J. KLENOW (2001): "Quantifying Quality Growth," *The American Economic Review*, 91(4), 1006–1030.

BLACK, D., N. KOLESNIKOVA, AND L. J. TAYLOR (2009): "Earnings Functions When Wages and Prices Vary by Location," *Journal of Labor Economics*, 27, 21.

BRODA, C., E. LIEBTAG, AND D. E. WEINSTEIN (2009): "The Role of Prices in Measuring the Poor‘s Living Standards," *Journal of Economic Perspectives*, 23(2), 77–97.

BRODA, C., AND J. ROMALIS (2009): "The Welfare Implications of Rising Price Disperison," .

BRODA, C., AND D. E. WEINSTEIN (2008): "Prices, poverty, and inequality: Why Americans are better off than you think," AEI Press.

COMBES, P.-P., G. DURANTON, AND L. GOBILLON (2008): "Spatial Wage Disparities: Sorting Matters!," *Journal of Urban Economics*, 63, 723–742.

DAVIS, D. R., AND J. I. DINGEL (2012): "A Spatial Knowledge Economy," *NBER Working Paper*, (18188).

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 (2011): "Purchasing Power Parity Exchange Rates for the Global Poor," *American Economic Journal: Applied Economics*, 3(2), 137–166.

DEATON, A., AND J. MUELLBAUER (1980): *Economics and Consumer Behavior*. Cambridge University Press.

DIAMOND, R. (2012): "The determinants and welfare implications of US workers' diverging location choices by skill: 1980-2000," *Harvard mimeo*.

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

FAJGELBAUM, P., G. M. GROSSMAN, AND E. HELPMAN (2011): "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): "International Prices and Endogenous Quality," .

GLAESER, E. L., AND D. C. MARE (2001): "Cities and Skills," *Journal of Labor Economics*, 19(2).

GLAESER, E. L., AND M. G. RESSEGER (2010): "The complementarity between cities and skills," *Journal of Regional Science*, 50(1), 221–244.

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 (2011): "Is New Economic Geography Right? Evidence from Price Data," *NBER Working Paper 17067*.

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.

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.

LEE, S. (2010): "Ability Sorting and Consumer City," *Journal of Urban Economics*, 68, 20–33.

LI, N. (2012): "An Engel Curve for Variety," .

MORETTI, E. (forthcoming): "Real Wage Inequality," *American Economic Journal: Applied Economics*.

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.

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 (MD-CNEV) 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): "Prices and Regional Variation in Welfare," *Journal of Urban Economics*, 51(3), 446–468.

SONG, I., AND P. K. CHINTAGUNTA (2007): "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.

WHEELER, C. H. (2001): "Search, sorting, and urban agglomeration," *Journal of Labor Economics*, 19(4), 879–899.

ZILIAK, J. P. (ed.) (2011): *Cost of Living and the Supplemental Poverty Measure: A Research Forum*University of Kentucky Center for Poverty Research.