HOUSEHOLD FRUIT AND VEGETABLE DEMAND ESTIMATION AND FORECASTING: A REVEALED PREFERENCE APPROACH

by

Joseph Blumberg

A Thesis Submitted to the Faculty of the

DEPARTMENT OF AGRICULTURAL AND RESOURCE ECONOMICS

In Partial Fulfillment of the Requirements

For the Degree of

MASTER OF SCIENCE

In the Graduate College

THE UNIVERSITY OF ARIZONA

2019

THE UNIVERSITY OF ARIZONA
GRADUATE COLLEGE

As members of the Master's Committee, we certify that we have read the thesis prepared by Joseph Blumberg, titled "Household Fruit and Vegetable Demand Estimation and Forecasting: A Revealed Preference Approach" and recommend that it be accepted as fulfilling the dissertation requirement for the Master's Degree.

Gary Thompson — Date: 5-7-19

Satheesh Aradhyula — Date: 5/7/2019

Dan Scheitrum — Date: 5/7/2019

Final approval and acceptance of this thesis is contingent upon the candidate's submission of the final copies of the thesis to the Graduate College.

I hereby certify that I have read this thesis prepared under my direction and recommend that it be accepted as fulfilling the Master's requirement.

Gary Thompson — Date: 5-7-19
Master's Thesis Committee Chair
Department of Agricultural & Resource Economics

TABLE OF CONTENTS

Tables, Figures, and Appendices

Abstract

Applying 2016-2017 household scanner data from market research firm IRI, we combine parametric and nonparametric techniques in estimating demands and forecasting consumption for six aggregated fruit and vegetable categories. The 2016 data is segmented by revealed preference (RP) such that the behavior in each subset of households is consistent with traditional utility theory, and a nonlinear Almost Ideal Demand System (NL-AIDS) model is estimated for all subsets. For comparison, demands are also estimated when the data is segmented separately on geography and household demographics. Own-price and expenditure elasticities generated across RP-consistent subsets indicate a wide range of demand responsiveness, whereas geographic and demographic subsets show similar behavior. Demand is generally more elastic for perishable goods than non-perishable. All methods of segmentation perform similarly when forecasting consumption into 2017.

1. Introduction

Fruit and vegetable consumption among adults in the United States is low, and diet-related illness is widespread. Rates of obesity, heart disease, and type-2 diabetes are significant and problematic. In 2016, 40% of U.S. adults were obese (Hales et al. 2017) and heart disease was the leading cause of death (Murphy et al. 2018). In addition to the degradation of health, diet-related illness results in financial burden from lost wages and increased costs of healthcare (WHO 2009). Consumption of fruits and vegetables, as part of a healthy diet suggested by the *2015-2020 Dietary Guidelines for Americans* (USDA 2015), is proven to reduce the risk of obesity-related diseases as well as some forms of cancer. Yet, the percentage of U.S. adults that meet the recommended daily intake of fruits and vegetables remains exceptionally low. According to the 2015 Behavioral Risk Factor Surveillance System (BRFSS) telephone survey, only 12.2% of adults met the recommended fruit intake and 9.3% the vegetable intake (Lee-Kwan et al. 2017).

Food choice is influenced by factors such as preferences, ideology, habit, access, costs, time constraints, and health considerations. Yeh et al. (2008) found that survey participants associated diets rich in fruit and vegetables with health benefits, but the prominent barriers impeding consumption were the perceived high costs of produce and a lack of preparation time due to long working hours. These findings are corroborated by Okrent & Kumcu (2016) who found that over the past 40 years demand has increased dramatically for "convenience" foods, which save time on preparation and are generally cheaper than fresh foods.[1] However, processed fruits and vegetables still satisfy federal dietary guidelines, and Stewart et al. (2016) find that the average consumer can meet federal recommendations with a limited budget, provided a smaller budget share is allocated to foods high in solid fats and added sugars.[2] So, if affordability and time constraints are not the issue, the question remains: why is fruit and vegetable consumption so low among U.S. adults? The simple, perhaps obvious, answer is that many people do not like and therefore do not eat fruits or vegetables. If that is the case, purchasing decisions will be influenced minimally by prices or income, since consumption is largely based on taste and preferences. Unfortunately, preferences cannot be quantified, which can make estimating demand a tricky affair for economists. This analysis aims to alleviate some issues introduced from unobserved preference heterogeneity via nonparametric data segmenting, and a six-good demand system is estimated using only fruit and vegetable products.

There is currently an extensive literature examining demand for both aggregate and disaggregate fruit and vegetable categories, many of which focus on differences in household income (Dong & Lin 2009; Bertail & Caillavet 2008). Segmenting consumers by observable characteristics, such as income, in microdata is commonplace in market research and demand analysis as it allows for more effective targeting for marketing, policy, and awareness initiatives. However, few studies implement a nonparametric approach to segmenting data for empirical analysis. Derived from revealed preference (RP) theory, Crawford and Pendakur (2012) developed

[1] The Consumer Price Index (CPI) estimates that average prices for fresh fruits and vegetables are consistently higher than their processed equivalents. Automation in processing and increased shelf life has led to lower prices.
[2] Stewart et al. estimate average F&V prices from retail store data and find that federal recommendations can be met for a 2,000-calorie diet on $2.10 to $2.60 a day.

an algorithm for partitioning microdata into "preference types," such that the data in each partition can be perfectly rationalized by a single utility function. Their methodology has not since been utilized, and a modified algorithm is used for this study.[3] Using 2016 household scanner data, we compare elasticities for fruit and vegetable goods generated across several group-specific nonlinear Almost Ideal Demand System (NL-AIDS) models using a sample of households segmented by two methods: on observables—demographic and geographic variables—and separately by RP. Additionally, we forecast out-of-sample to the following year for the same households to compare predication accuracy for different segmenting methodologies.

2. Literature Review

2.1. Utility Theory, Unobserved Heterogeneity, and Revealed Preferences

Neoclassical utility theory suggests that people have an underlying utility function which ranks bundles of goods given prices, product characteristics, and a budget constraint. Whichever bundle maximizes the utility of the consumer will be purchased. Therefore, choice behavior can be sufficiently explained by fitting a properly specified demand function to some observed data. Stigler and Becker (1977) postulate a rather bold *a priori* for demand analysis, that tastes are "stable over time and similar among people" (p. 76) and should not be considered as mutable. Since this is not a directly testable hypothesis, the functional form of any parametric demand system must be assumed as "true" in that it adequately approximates a consumer's utility function. In general, this traditional approach focuses purely on testing observable differences in constraints and household demographics while neglecting unobserved heterogeneity in tastes. However, as rich sources of microdata have become more readily available, heterogeneity in choice behavior is frequently observed. Consequently, the issues resulting from unobserved preference heterogeneity

[3] We are thankful to Julien Boelaert for his development of the R package "revealedPrefs," which allowed for the nonparametric segmentation of our data.

have become widely acknowledged, but to the affliction of economists, a consumer's tastes and preferences are subjective and cannot be directly measured.

The goal of RP theory, introduced by Samuelson (1938, 1948), is to provide a way to infer consumer preferences by examining observed choices. Samuelson developed algebraic conditions, consistent with traditional economic theory, that formulate a set of restrictions for a demand function in which observed behavior that satisfies those restrictions is considered maximizing behavior. In essence, if a bundle of goods $\boldsymbol{t}$ holds more overall value than bundle $\boldsymbol{v}$, and both bundles are affordable, it is said that $\boldsymbol{t}$ is revealed preferred to $\boldsymbol{v}$. By the *Weak Axiom of Revealed Preference* (WARP), $\boldsymbol{v}$ cannot simultaneously be revealed preferred to $\boldsymbol{t}$. More precisely, given some vectors of prices and purchased quantities ($\boldsymbol{p^t}, \boldsymbol{q^t}$), $\boldsymbol{q^t}$ is *directly revealed preferred* to bundle $\boldsymbol{q}$ (denoted $\boldsymbol{q^t R^0 q}$) if $\boldsymbol{p^t q^t} \geq \boldsymbol{p^t q}$ and both $\boldsymbol{q^t}$ and $\boldsymbol{q}$ are affordable given a budget constraint. If there exists multiple vectors of prices $\boldsymbol{p^r}, \boldsymbol{p^s}, \boldsymbol{p^t}, \dots, \boldsymbol{p^v}$ such that $\boldsymbol{p^r q^r} \geq \boldsymbol{p^r q^s}$, $\boldsymbol{p^s q^s} \geq \boldsymbol{p^s q^t}, \dots, \boldsymbol{p^v q^v} \geq \boldsymbol{p^v q}$, then it is said $\boldsymbol{q^t}$ is *revealed preferred* to $\boldsymbol{q}$ (denoted $\boldsymbol{q^t R q}$), where $\boldsymbol{R}$ is the "transitive closure" of $\boldsymbol{R^0}$). Later, S. N. Afriat (1967) generalized a method of nonparametrically testing consumer behavior that considered utility maximization, and *Afriat's Theorem* proves that if a set of data can be rationalized by any nontrivial utility function (i.e. satisfies "cyclical consistency") there exists a well-behaved utility function (i.e. no violations of continuity, concavity, or monotonicity) that also rationalizes that data. The tests are nonparametric in that they require no *ad hoc* assumption of the consumer's underlying utility function. The *Generalized Axiom of Revealed Preference* (GARP) (Varian 1982, 1983) condition provides a less computationally intensive, though mathematically equivalent, method for testing cyclical consistency, making empirical RP analysis practical. GARP is satisfied for a set of observed demands ($\boldsymbol{p^t}, \boldsymbol{q^t}$) if $\boldsymbol{q^t R q^v}$ implies $\boldsymbol{p^v q^v} \leq \boldsymbol{p^v q^t}$, meaning ($\boldsymbol{p^t}, \boldsymbol{q^t}$) can be rationalized by a single utility function. The benefit of the RP approach in demand analysis is that GARP provides a simple and intuitive test for rationality without assuming a functional form.

2.2. Empirical Applications and Criticisms of Revealed Preference

While the applications of RP are rooted in economic theory, the structural approach to empirical analysis is distinct from traditional econometrics. The ultimate goal in empirical econometric analysis is to gather some evidence of causal inference. An econometrician links explanatory and dependent variables by some statistical model consistent with economic theory, and an error term is introduced to explain any imperfect fit. Assumptions are made, a model is developed, and hypotheses can be tested. While perhaps an imperfect design, the econometrician is at least left with parameter estimates or elasticities with straightforward interpretations.

In contrast, the RP approach uses only price and quantity data to run expenditures through a system of inequalities to check for consistency with maximization behavior, relying neither on assumptions or model structure. If GARP is satisfied, the consumers are considered rational agents, and an economist might revel in his luck of finding a sample of *homo economicus* descendants. However, the well-behaved utility function that rationalizes these consumers' behavior remains unknown. It should also be noted that it becomes difficult to violate GARP when price data lacks variation (Varian 1982), though that should not be an issue in this study due to prices available at a unique product level (see section *3.5*). Once GARP is satisfied, economists can either then revert to the parametric approach, as we do in this analysis, or continue using nonparametric methods. Blundell et al. (2003) developed a procedure for generating nonparametric Engel curves to forecast demand within a GARP consistent sample, and Blundell (2005) further expanded this work by introducing RP bounds on demand responses. In each case, the nonparametric approach rarely produces easily interpretable results.[4] Though RP theory is well-developed, straightforward empirical applications are currently lacking (Crawford & De Rock 2014).

If GARP is violated, there are numerous possibilities to consider. Measurement error is always a plausible explanation, though Varian (1985) provides a method of measuring the magnitude of violations sufficient to restore no violations of GARP. Ruling out measurement error, let us examine the case of cross-sectional data on a sample of households. A violation indicates that the behavior of all the households in the sample cannot be rationalized by a single, well-behaved utility function. Therefore, some sub-population(s) must have a different set of preferences and multiple utility functions are necessary to explain the observed behavior. However, when testing GARP across a sample of households at a single point in time, it is assumed that all households had access to the same set of prices. It becomes ambiguous whether preferences are indeed different or if some households were purchasing under different circumstances. In this study we have aggregated product-level price data over the entire year of 2016, so it is not unreasonable to assume that households made decisions under similar prices. Average prices relative to region and county size show little variation.[5]

In the context of time-series data, let us assume GARP is violated for a single household across multiple time periods. One possible explanation is that individual preferences contain cycles (Crawford & De Rock 2014), and during one cycle a consumer may be willing to waste money even if a cheaper bundle that he has already revealed preferred is available. Whether the violation is the result of an irrational consumer or changing tastes, RP theory lacks in providing a clear

[4] For more analytic applications of RP, see Varian's "Revealed Preference and Its Applications" (2012).
[5] See Appendix A for price distributions across geographic and GARP-consistent segments.

answer. However, Blundell (2005) implements a method of characterizing changing tastes over time when behavior cannot be explained by price changes.[6]

Traditional econometrics requires some assumptions of functional form and the distribution of an error term but can provide a straightforward explanation of what occurs in the real world, conditional on those assumptions. Revealed preference theory has introduced intuitive and simple tests of rationality but struggles to provide clear answers to the questions that economists are interested in. Subsequently, this study attempts to utilize the strengths of both parametric and nonparametric approaches to empirical analysis.

3. Household Scanner Data

3.1. Nielsen/Information Resources, Inc. National Consumer Panel

The National Consumer Panel (NCP) is a joint venture by the Nielsen Company and Information Resources, Inc. (IRI) where participating members record their grocery store purchases using a barcode scanner. Households voluntarily apply to become a "panelist" by filling out an online survey, providing information such as household income, size, race, age of household head, urban/rural composition, education, and presence of children. As an incentive, households accepted into the panel can earn points for frequent reporting that are used to acquire merchandise or entry into sweepstakes. IRI strives to create a nationally representative panel, thus accepts households based on their demographic and geographic characteristics.[7] Households may be placed on a waitlist if their makeup is currently overrepresented and will receive a notification of acceptance when current panelists attrite. When accepted, households are provided with instructions for recording and transmitting their purchases, and IRI provides technical support via

[6] For our sample of NCP households, 205 violate GARP when purchase bundles for 2016 and 2017 are pairwise tested for each individual household.

[7] I personally filled out the online application; it is straightforward and takes about 10 minutes to complete. Young households are currently underrepresented, and I received an acceptance email within 48 hours.

email or hotline. NCP panelists only report purchases from grocery and some pharmacy retailers, so food away from home purchases, such as at restaurants, are not included in the data.

3.2. Reporting

After each trip to a store for food purchases a household scans every item with a barcode, either with a smartphone application or NCP-provided scanning device, and then enters the name of the store where the purchases took place. If the purchase took place at a retailer that provides point-of-sale information to IRI, the price of that product is automatically generated to reduce the burden of reporting on households. In some cases, these IRI-imputed prices are subject to measurement error (Einav et al. 2009). About 65% of transactions are assigned prices (Muth et al. 2016), which are imputed using retail-chain averages localized to 1 of 73 IRI-designated marketing areas. Otherwise, a price is manually entered, which is then checked against an IRI price-dictionary of national averages for plausibility. When an item is scanned a notification asks the panelist if the item was purchased on sale or using coupons. If so, the household inputs the type of discount (store sale, coupon, etc.), the value of the discount, and ultimately enters their total trip's discounted value. Imputed prices can be subject to error if a product was purchased using a club card or a panelist misreports the store location.

For items without a barcode, or universal product code (UPC), the panelist is instructed to select from a list of generic product codes, e.g. "lettuce" or "apples".[8] Products without UPCs are typically "random-weight" goods that are sold by the pound or count, including fresh fruits and vegetables. Once a generic code is selected the user to prompted to enter the amount paid but *not* a quantity or weight. Thus, true unit-values cannot be calculated for individual random-weight products and are imputed in this analysis (see section *4.3.*).

[8] The generic UPC for "lettuce" would be used for the purchased of any lettuce variety such as iceberg, romaine, butter, endive, etc.

3.3. "Static" Household Criteria

The current NCP is comprised of roughly 120,000 households, though not all frequently report their purchases. Muth et al. (2016) found that households with children, households in lower income brackets, and households with heads under age 35 are the least likely to report consistently. In response to underreporting, IRI recommends a subset of households, designated as "static", to be used for analysis. The static panel consists of about half of all panelists, and the following criteria must be met for a household to be considered static each year:

1. The household reports at least one purchase for 11 out of the 13 four-week reporting periods
2. The household meets a minimum weekly spending requirement depending on household size: $25/week for single-person households, $35/week for two-person households, and $45/week for households of three-persons or more

This study utilizes only a subset of the static households that have met the thresholds for two consecutive years, 2016-2017, to allow demand forecasting to the same households. This "super" static panel results in a total of 38,059 households that is less representative of the general U.S. population. Validating the findings of Muth et al. (2016), lower income households, households with heads under the age of 35, ethnically non-white households, and households with children are all underrepresented in our sample. Despite the sample not being entirely representative of the U.S., there is still enough variation, both demographically and geographically, to preserve large enough sample sizes for segmenting the data.

3.4. General Issues to Consider

As with all self-reported data, the NCP data are subject to measurement error and sample selection bias. Despite efforts to keep the panel geo-demographically representative, households may be more cognizant of prices. Lusk and Brooks (2011) find that participants in household scanning panels are slightly more price sensitive than a random U.S. sample, and Boonsaeng and

Carpio (2014) find similar results when comparing elasticities calculated from the Bureau of Labor Statistic's Consumer Expenditure Survey (CEX) to NCP. Panel participants may also be underreporting in non-random ways. All items are less likely to be scanned when a household purchases a large number of items (Einav et al. 2009), and all trips may not be recorded due to the opportunity cost of time. When compared to CEX and FoodAPS, Sweitzer et al. (2017) find lower expenditures for NCP households and underreporting as much as 45-50% for fresh vegetables. The increased burden of recording random-weight purchases may facilitate underreporting.

Despite these limitations, NCP data do have some distinct advantages to national survey data. While CEX and FoodAPS seem to have more complete record of household expenditure data, their reference period is significantly shorter. With NCP data, purchasing behavior can be analyzed for the same households for many years.[9] Highly detailed product information is available due to barcode scanning, which allows for the calculation of unit-prices as well as analysis relating to product characteristics such as branding, labeling, and nutrition.

3.5. Goods/UPCs

The USDA's MyPlate Plan, based on the *2015-2020 Dietary Guidelines for Americans*, recommends about two and a half cups of fruit and three and a half cups of vegetables daily, depending on body type, for a healthy diet. In general, one cup from the fruit group can be: (i) 1 cup raw, frozen, or canned fruit, (ii) ½ cup dried fruit, or (iii) 1 cup 100% fruit juice. For the vegetable group, one cup can be: (i) 1 cup raw, cooked, or canned vegetables, (ii) 2 cups leafy greens, or (iii) 1 cup 100% vegetable juice (https://www.choosemyplate.gov/myplateplan, accessed 03/01/2019). Typically, there is not a significant loss in nutrients when fruits and vegetables are frozen or canned, though it depends on the method of processing. Drying typically

[9] There exist 10,000 households who have remained in the "static" panel for 10 years, 2008-2017.

results in a greater loss than other methods of preservation (Rickman et al. 2007), and the health risks of added sugars to certain fruit juice and other processed fruit products can outweigh the benefits of the vitamin content (Imamura et al. 2015). Considering federal recommendations, nutritional content, and perceived barriers to consumption, we have aggregated UPCs into six broad fruit and vegetable categories:

1. Vegetables – Leafy Greens (LG_p)
2. Other Vegetables – Perishable (OV_p)
3. Other Vegetables – Non-Perishable (canned or frozen) (OV_n)
4. Fruit – Perishable (FR_p)
5. Fruit – Non-Perishable (canned or frozen) (FR_n)
6. Fruit – Dried or Juiced (DJ_n)

Categories were split by perishability due to contrasting prices, perceived cooking times, and the relative nutritional value of fresh versus processed produce. Leafy green products were aggregated into their own category due to the inclusion of pre-packaged salads that may include animal proteins or dressing. Additionally, 2 cups of leafy greens are required to satisfy one serving of vegetables. Aggregating into these categories involved the identification of 173,183 fixed-weight UPCs and 41 random-weight UPCs.

Table 1. Breakdown of Universal Product Codes (UPCs)

Fixed-Weight			Random-Weight		
Category	Count	Percentage	Category	Count	Percentage
LG_p	9,480	0.05	LG_p	5	0.12
OV_p	28,984	0.17	OV_p	16	0.39
OV_n	42,855	0.25	FR_p	20	0.49
FR_p	36,689	0.21	Total	41	1.00
FR_n	15,122	0.09			
DJ_n	40,054	0.23			
Total	173,184	1.00			

Aggregating individual products into these six categories does involve placing arguably disparate products into the same category. For example, fresh sweet potatoes and celery both reside in the same vegetable category (OV_p) despite their vastly different cooking applications and nutritional content. However, further disaggregation of products resulted in a significant number of households recording no purchases in that category, even when measuring over a year-long period.

3.6. Price Trimming and Imputation

Observations with extreme unit-prices are excluded from analysis. Households report the value of the discount if a coupon is used, so that value is subtracted from the IRI-imputed prices. This sometimes results in a negative price for a product.[10] Some households were also observed purchasing fruit and vegetable products with a price exceeding $15.00/pound or lower than $0.10/pound. We do not know if this is the result of measurement error, extreme luxury products, bulk products, or large discounts, so observations with prices at the top and bottom 1% are truncated from the data before UPCs are assigned into categories.

Recorded purchases of random-weight products require price imputation. Products with a barcode/UPC are considered “fixed-weight” and can be matched to one of two dictionaries that contain that product’s description, nutritional information, and a fixed ounce weight, which allows for the calculation of unit-prices. The “point-of-sale” dictionary holds information for all branded products, and the “perishables” dictionary holds information for perishable, typically packaged products such as produce enclosed in a bag or clamshell. To impute unit-prices for the 41 random-weight goods, average unit-prices for aggregated UPCs in the "perishables” dictionary are matched

[10] There are cases when coupons can result in a rebate for the customer, which would result in a negative unit-price for a particular product. Negative prices after discounts could also be the result of incorrect imputation on behalf of IRI.

to the closest random-weight product. For example, an average unit-price is calculated using all observations in which a fixed-weight product in the perishables dictionary under the category "lettuce" is purchased. That average unit-price for "lettuce" is then assigned to all observations in which the "lettuce" random-weight UPC is purchased. This process is repeated for all 41 random-weight goods using chain-specific data localized to marketing area.[11] Since we have the total amount paid per product in each random-weight observation, quantities are generated from the imputed unit-price. This undoubtedly introduces measurement error; however, it is more representative of actual household consumption than if random-weight products were excluded entirely.

Once the data is trimmed and random-weight quantities are imputed, UPCs are assigned into one of the six fruit and vegetable categories, and the data is aggregated across all trips over the entire year of 2016 for each household. However, price data is only available when a household records a purchase, so any household with a zero expenditure in one product category for 2016 will be missing price data for that category. In such cases, the average price for the respective category across all households within the same marketing area is assigned to that missing price.

4. Model

4.1. Segmentation

Before households were segmented by finding GARP-consistent subpopulations, a lower bound estimate on the number of subpopulations was calculated. A random permutation of all 38,059 households is drawn and each household is pairwise tested for GARP-consistency without

---

[11] Emulating IRI methodology, prices are imputed separately for each retail-chain in each of the 73 unique marketing areas. If there are no observed purchases for a product category for a specific chain within a marketing area, the average price from that marketing area is assigned. If there are no observed purchases for a product category within an entire marketing area, a total sample average is used. Total sample averages are imputed and assigned to less than 1% of the observations.

replacement.[12] If the first and second drawn households satisfy RP, the second observation is dropped and the third is tested against one. If one and three violate RP, three is retained and the fourth household tests against one and three. The fourth household is then retained only if RP is violated in both pairwise tests. After all households are drawn, we are left with a set of violators that are exhaustively pairwise GARP-inconsistent. The number of households in this set, N, is the minimal amount of unique utility functions needed to rationalize the data. Since the households are randomly permuted, this test was run 5,000 times and the largest value is reported at $N = 20$.[13]

Though the variation in choice behavior in our data should be rationalizable by 20 unique utility functions, clustering households into the minimal necessary segments is computationally infeasible.[14] Instead, households are clustered into GARP-consistent subsets using a similar method for estimating the lower bound. A random permutation of the households is drawn, the first two draws are tested, and if RP is satisfied they form the first group. The remainder of the households are tested against all current groups, creating new groups if they are not consistent with any of the former. This resulted in 107 subsets of GARP-satisfying households.[15] For this study, we only examine the largest 20 subpopulations for simplicity and complying with the lower bound estimate. This reduces our sample size from 38,059 to 26,383. Additionally, 10 households were excluded from the analysis due to reporting no purchases on any fruit or vegetable product, leaving our final sample at 26,373. Descriptive statistics and zero expenditure frequencies for the annually aggregated 2016 data are given in Table 2 and Table 3.

---

[12] For testing GARP-consistency, the method in Varian (1982) using the Floyd-Warshall algorithm is used to check for cyclical consistency. The R package 'revealedPrefs' performs this.
[13] The algorithm to estimate the lower bound runs relatively quick, averaging 10 minutes to run all 5,000 iterations.
[14] Finding the minimal, exhaustive segments would require testing RP restrictions for all possible subsets of that data at $2^{38,059}$ subsets.
[15] Due to random permutation, this algorithm should be run several times. However, each run can take up to 36 hours, so it was only repeated 5 times in this study. The clustering which resulted in the greatest sample size for the largest 20 subsets was used.

Table 2. Price, Quantity, and Expenditure Descriptive Statistics for 2016

	Price (USD/lb)				Quantity Purchased (lbs)			
Good	Mean	Median	Max	Min	Mean	Median	Max	Min
LG_p	$ 3.39	$ 3.22	$ 11.14	$ 0.46	17.5	11.9	461.3	0
OV_p	$ 1.50	$ 1.40	$ 10.61	$ 0.27	88.0	69.9	1,054.0	0
OV_n	$ 1.31	$ 1.21	$ 10.61	$ 0.26	63.1	48.8	1,346.5	0
FR_p	$ 1.61	$ 1.50	$ 10.64	$ 0.34	114.3	81.6	1,622.7	0
FR_n	$ 1.78	$ 1.68	$ 10.63	$ 0.26	17.9	9.3	875.9	0
DJ_n	$ 1.20	$ 0.94	$ 10.80	$ 0.21	91.3	52.6	2,020.6	0
	Expenditure Share				Expenditure (USD)			
LG_p	0.11	0.09	0.97	0.00	$ 56	$ 38	$ 1,424	$ -
OV_p	0.22	0.22	1.00	0.00	$ 121	$ 91	$ 1,698	$ -
OV_n	0.16	0.14	1.00	0.00	$ 74	$ 60	$ 1,183	$ -
FR_p	0.29	0.28	1.00	0.00	$ 166	$ 120	$ 2,985	$ -
FR_n	0.06	0.04	1.00	0.00	$ 30	$ 15	$ 1,361	$ -
DJ_n	0.16	0.12	1.00	0.00	$ 77	$ 51	$ 1,984	$ -

Sample Size = 26,373

Table 3. Zero Expenditures by Good

Good	LG_p	OV_p	OV_n	FR_p	FR_n	DJ_n
Frequency	1162	286	310	389	3280	775
Percentage	0.04	0.01	0.01	0.01	0.12	0.03

Sample Size = 26,373

In industry, microdata is commonly segmented by observable characteristics to account for unobserved heterogeneity. In addition to segmenting by revealed preference, our reduced sample of 26,373 is separately segmented on the observable characteristics that are commonly examined for fruit and vegetable demand in the literature. Each household is assigned into six different segments based on the following criteria: fruit and vegetable expenditure per household member (EPP) (20 groups), geography (16), yearly household income (11), average age of household head

(10), number of people in the household (6), and household composition or "type" (6).[16] Detailed information on the segments can be seen in Table 4.

Table 4. Segment Key and Sample Sizes

Segment	GARP	EPP (Quantile)		Geography		Income		Head Age		Size		Type	
1	4,292	1	1,318	Urban, NE	2,296	< $12K	857	< 35	1,422	1	6,265	Young with Children	2,273
2	2,734	2	1,319	Urban, MW	2,287	$12K - $14K	461	35-39	1,466	2	12,115	Older with Children	3,143
3	2,416	3	1,319	Urban, S	3,033	$15K - $19K	735	40-44	1,778	3	3,286	Young Singles	560
4	1,940	4	1,318	Urban, W	2,533	$20K - $24K	1,156	45-49	2,425	4	2,756	Older Singles	5,705
5	1,685	5	1,319	Suburban, NE	1,595	$25K - $34K	2,774	50-54	3,336	5	1,272	Young Couples	1,190
6	1,576	6	1,319	Suburban, MW	1,732	$35K - $44K	2,864	55-59	4,076	≥ 6	679	Older Couples	13,502
7	1,238	7	1,318	Suburban, S	3,663	$45K - $49K	1,573	60-64	4,318				
8	1,202	8	1,319	Suburban, W	1,497	$50K - $59K	2,683	65-69	3,624				
9	1,010	9	1,319	Rural, NE	584	$60K - $69K	2,264	70-74	2,205				
10	989	10	1,318	Rural, MW	1,227	$70K - $99K	5,695	≥ 75	1,723				
11	938	11	1,319	Rural, S	1,666	≥ $100K	5,311						
12	885	12	1,319	Rural, W	744								
13	827	13	1,319	Very Rural, N	346								
14	816	14	1,318	Very Rural, MW	1,435								
15	764	15	1,318	Very Rural, S	1,246								
16	675	16	1,320	Very Rural, W	489								
17	665	17	1,318										
18	587	18	1,319										
19	581	19	1,319										
20	553	20	1,318										
	26,373		26,373		26,373		26,373		26,373		26,373		26,373

Revealed preference restrictions were tested for all households within each of the 69 segments based on observables and all were in violation of GARP. This indicates that segmentation by observables is insufficient for totally mitigating unobserved heterogeneity. To check for potential patterns in observables across GARP-consistent subsets, standardized differences in proportions across observable characteristics were examined for each GARP subset against all other households in our sample. Standardized difference scores are intuitive indexes measuring an effect size between two groups. They are commonly used for comparing baseline covariates in

[16] EPP was calculated by dividing total expenditure on F&V by number of household members, and households were then ranked and segmented into 20 quantiles (highest expenditures represented by 1st quantile). For geography, counties were categorized using Nielsen Corp. ABCD criteria. Household "type" is designated by IRI.

clinical trials and studies that use propensity score matching.[17] To calculate a standardized difference metric for a categorical variable, Yang and Dalton (2012) provide a multivariate Mahalanobis distance method. A small metric indicates that the distribution of the categorical variable is similar between two groups, whereas a large metric indicates different distributions. Large metrics suggest a treatment effect, which would provide a clue to which covariate best captures unobserved heterogeneity. There is no widely accepted threshold to determine significant differences, however, Cohen (1988) suggests 0.2, 0.5, and 0.8 can represent small, medium, and large differences, respectively.

Figure 1. Standardized Differences Between GARP Subsets and Remaining Households

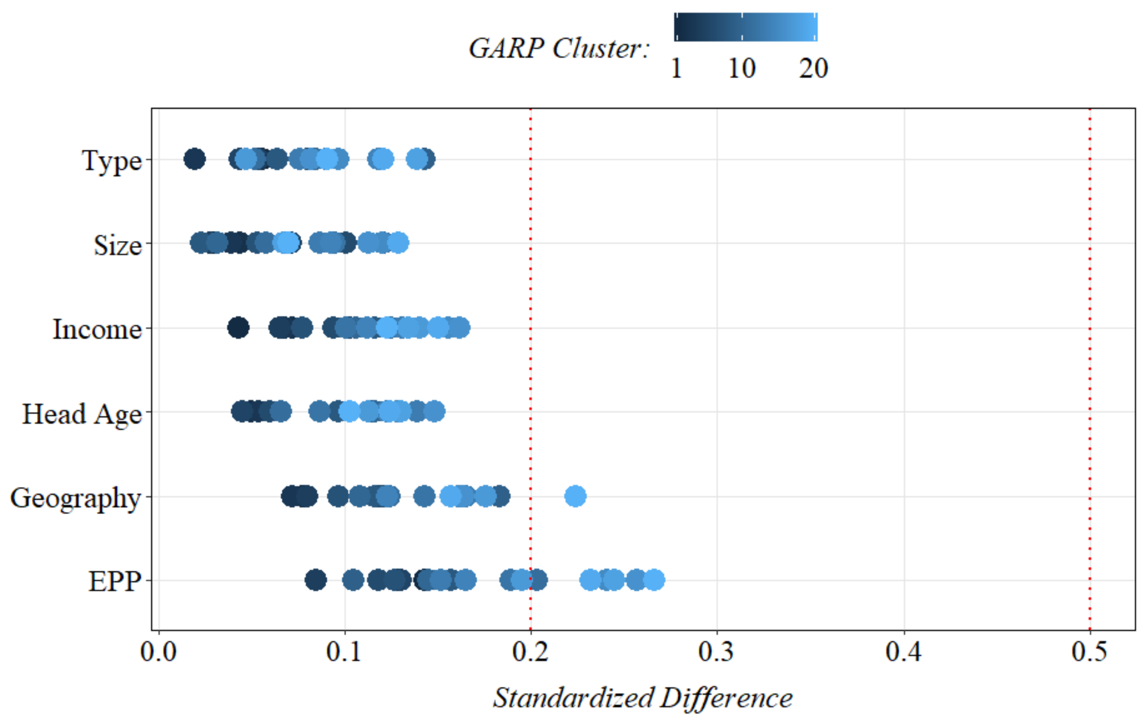

In Figure 1, each GARP subset has six standardized difference measures, one for each of the categorical variables listed in Table 3. Following Cohen's thresholds, there are only 9 out of

[17] In the context of this analysis, each individual GARP cluster is considered the treatment group, and all other households are control. A standardized difference metric is calculated for six observables for each treatment group.

120 instances when the distribution of a covariate is considered sizably different between a GARP cluster and the remaining households, all 9 of which are considered "small".[18] The prevalence of small differences indicates no treatment effect and that each GARP-consistent subset is a relatively representative sample of the remainder of the households. It is also important to notice that all differences occur in smaller sample sizes for observables with the most categorical levels.[19] Overall, there is no clear pattern in observables that can account for preference heterogeneity prior to testing RP restrictions.

4.2. Nonlinear Almost Ideal Demand System Model

In theory, our GARP-consistent subpopulations should all be rationalizable by a single integrable demand system, but the specification of that function remains unknown. At the very least, we do not need to worry about unobserved heterogeneity. We resort to using the nonlinear Almost Ideal Demand System (NL-AIDS) model (Deaton & Muellbauer 1980) due to its flexible functional form, and it is estimated once for all households and again for each unique segment. We do not treat the data for the prevalence of zero expenditures for certain product categories, so our estimation is single-stage. The data was aggregated over the entire year of 2016, so it is unlikely zero expenditures were the result from lack of access. The iterative linear least squares estimation (ILLE) is used for estimating NL-AIDS (Browning and Meghir 1991; Michalek and Keyzer 1992; Blundell and Robin 1999). The demand equations of AIDS can be simplified using expenditure shares, w_i, where i is a subscript denoting the i^{th} good in the demand system,

$$w_i = \alpha_i + \sum_j \gamma_{ij} \ln (p_j) + \beta_i \ln (x/P) \tag{1}$$

[18] Testing differences in 6 observable characteristics for 20 subsets results in 120 statistics. A 95% confidence interval was also calculated for each statistic (Yang & Dalton 2012); 0/120 lower bounds and 39/120 upper bounds exceed 0.2. The 0.5, "medium", threshold is never crossed.

[19] All differences occur in clusters less than 1,000 households, and 8/9 differences occur for the EPP observable which has 20 categorical levels.

and P is the translog price index,

$$\ln P = \alpha_0 + \sum_k a_k \ln(p_k) + \frac{1}{2} \sum_k \sum_j \gamma_{kj} \ln(p_k) \ln(p_j) \tag{2}$$

We will refer to $x \equiv \sum_j p_j q_j$ as group expenditure, which is the expenditure of all goods of interest in the demand system as opposed to total household income. Restrictions on the parameters are imposed to allow the model to conform with traditional demand theory. The "adding-up" condition ensures that expenditure shares sum up to one,

$$\sum_i \alpha_i = 1; \ \sum_i \beta_i = 0; \ \sum_i \gamma_{ij} = 0 \ \forall \, j \tag{3}$$

"homogeneity" guarantees that if all prices and income change at the same rate there will be no change in consumed quantities,

$$\sum_j \gamma_{ij} = 0 \ \forall \, i \tag{4}$$

and "symmetry" follows from simplifying AIDS to (1),

$$\gamma_{ij} = \gamma_{ji} \ \forall \, i, j \tag{5}$$

The group expenditure (income) elasticities and Marshallian (uncompensated) price elasticities can be derived from the Marshallian demand functions (1) (Anderson and Blundell 1983). The expenditure elasticities, λ_i, and Marshallian price elasticities, λ_{ij}, simplify to:

$$\lambda_i = \frac{\partial q_i}{\partial x} \frac{x}{q_i} = 1 + \frac{\beta_i}{w_i} \tag{6}$$

$$\lambda_{ij} = \frac{\partial q_i}{\partial p_i} \frac{p_i}{q_i} = -\delta_{ij} + \frac{\gamma_{ij}}{w_i} - \frac{\beta_i}{w_i} \alpha_j + \left[\sum_k \gamma_{kj} \ln(p_k) \right] \tag{7}$$

where δ_{ij} is the Kronecker delta. An approximate calculation of the covariance matrix of elasticities is calculated using the delta method to allow for hypothesis testing.

4.3. Forecasting

Demand forecasting is important to both producers and policy makers for effective planning. By limiting our sample to households that exhibit frequent purchasing for two years we can forecast purchased quantities in 2017 for every household and compare prediction accuracy across segments. Two methods of forecasting are used using 2016 estimates: direct statistical and elasticity-based. Gustavsen and Rickertsen (2003) find elasticity-based forecasting to be better than statistical forecasting. First, prices and total expenditures for 2017 are plugged into equations (1) and (2) using the unique parameter estimates from each segment to generate predicted shares and subsequently predicted quantities. Second, we use the Kastens and Brester (1996) method for forecasting using our estimated elasticities, percent changes in prices, and percent change in total expenditure:

$$q_{i,t} = \left[\sum_{j=1}^{n} \lambda_{ij} \left(\frac{p_{j,t} - p_{j,t-1}}{p_{j,t-1}}\right) + \lambda_{i} \left(\frac{x_t - x_{t-1}}{x_{t-1}}\right)\right] q_{i,t-1} + q_{i,t-1} \qquad (8)$$

Households that experience a change in an observable characteristic have their purchased quantities forecasted by the model estimates with respect to their new demographic makeup. For example, if a family moved from an urban to rural area in 2017, their purchase quantities would be predicted using one of the rural models.[20] Once quantities are predicted for each segment, root mean-squared errors (RMSE) are calculated for each good to judge prediction accuracy. The RMSE is calculated using the forecasted quantities for all households aggregated by segmenting method.[21]

---

[20] We only observe a snapshot of household demographics at the end of 2017, so we do not know exactly when a change occurs within that year. For robustness we forecasted to 2017 using both 2016 and 2017 demographically appropriate segments and there was no significant difference.

[21] For example, quantities are estimated for households within each GARP segment. The RMSE is then calculated using estimates from all 26,373 households using the GARP forecasts. This is repeated for all methods of segmenting.

5. Results

5.1. All Households

The first NL-AIDS model was estimated using all 26,373 households to create a baseline of aggregated demand responsiveness. Elasticity measures were calculated using mean prices and expenditure shares for each good. From the covariance matrix of elasticities, *t*-tests were approximated to determine marginal levels of significance (*p*-values) for each elasticity. All expenditure and own-price elasticities were statistically significant at the .001 level, see Table 5.

Table 5. Aggregate Elasticities

	Expenditure Elasticities					
Good	LG_p	OV_p	OV_n	FR_p	FR_n	DJ_n
	0.963	**1.058**	**0.769**	**1.165**	**0.942**	**0.872**
	Marshallian Own-Price Elasticities					
LG_p	**-0.941**	**0.068**	-0.026	**-0.128**	**-0.035**	**0.098**
OV_p	**0.022**	**-1.032**	**-0.168**	0.015	-0.008	**0.113**
OV_n	0.004	**-0.175**	**-0.978**	**0.289**	0.002	**0.089**
FR_p	**-0.067**	-0.013	**0.090**	**-1.210**	**-0.023**	**0.058**
FR_n	**-0.062**	-0.004	-0.021	-0.051	**-0.815**	0.010
DJ_n	**0.081**	**0.220**	**0.080**	**0.206**	0.008	**-1.467**

Bold entries indicate statistical significance at the .001 level.

Expenditure elasticities for perishable goods are systematically greater than non-perishable goods, indicating a greater increase in demand for fresh than processed foods as income increases. Similarly, own-price elasticities for perishable fruits and vegetables are greater in absolute value than their non-perishable counterparts, though the dried fruit and fruit juice category has the largest own-price elasticity. Fresh fruits and vegetables are more likely to experience fluctuations in prices due to seasonality and perishability, which may induce greater demand responses. Juice purchases often occur in large ounce-weight increments, likely causing the large demand response to changes in price. For example, a consumer may buy one gallon of orange juice on sale, which is roughly equivalent to a quantity of 8lbs (128oz). However, if fresh oranges are bought on sale, a customer may only buy a 1-2lbs (16-32oz). All cross-price elasticities are relatively close to zero, but there is evidence of a substitution effect between perishable and non-perishable vegetables due to the negative and statistically significant cross-price elasticities between those two categories.

5.2. Segmented Households

The next 89 NL-AIDS models were estimated for each unique segment, elasticities were calculated using segment-specific averages, and all elasticities were tested for statistical significance.[22] Expenditure elasticities across segments are displayed on Figures 2a and 2b, and own-price elasticities on Figures 3a and 3b. Figures 2a and 3a display segment-specific elasticities for EPP, income, and size, and 2b and 3b for geography, head age and type.[23] The baseline elasticities calculated from aggregating all households and GARP elasticities are displayed on all figures for comparison. Households partitioned by revealed preference exhibit a much larger range of expenditure and own-price elasticities than households partitioned geographically or demographically, indicating that differing tastes and preferences for fruit and vegetable products indeed exist. For every elasticity we find GARP-consistent segments ranging well-above and well-below the aggregate. On average there is a greater dispersion of elasticities for perishable goods than non-perishable for both expenditure and own-price. Elasticities generated from segments based on observables display a tight range around the aggregate, which is evidence that traditional partitioning methods are inadequate in identifying segments with significantly differing consumption behavior in the context of fruits and vegetables.

---

[22] All expenditure elasticities across all segments were found statistically different from zero at the .001 level. The own-price elasticities found *not* significant at the .001 level were: OV_n for GARP segment #19 and FR_n for GARP segments #10, 14, 15, 16, 19, 20. Insignificant own-price elasticities were excluded from figures *3a* and *3b*.
[23] See Appendix C for elasticity graphs for each segment individually.

Figure 2a. Segment-Specific Expenditure Elasticities

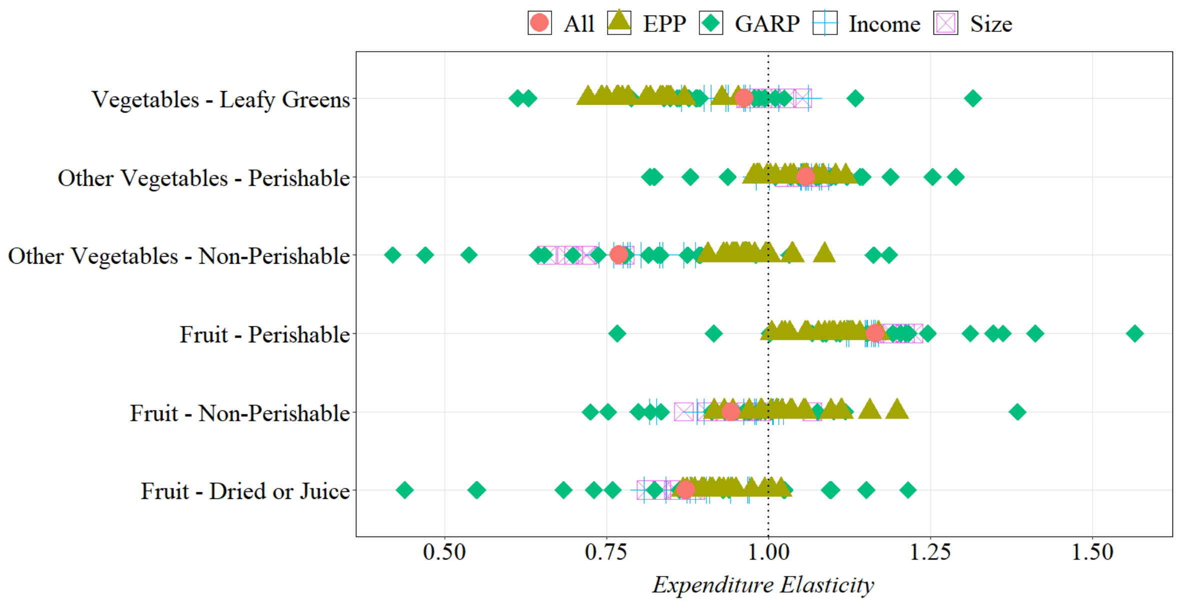

Figure 2b. Segment-Specific Expenditure Elasticities, cont.

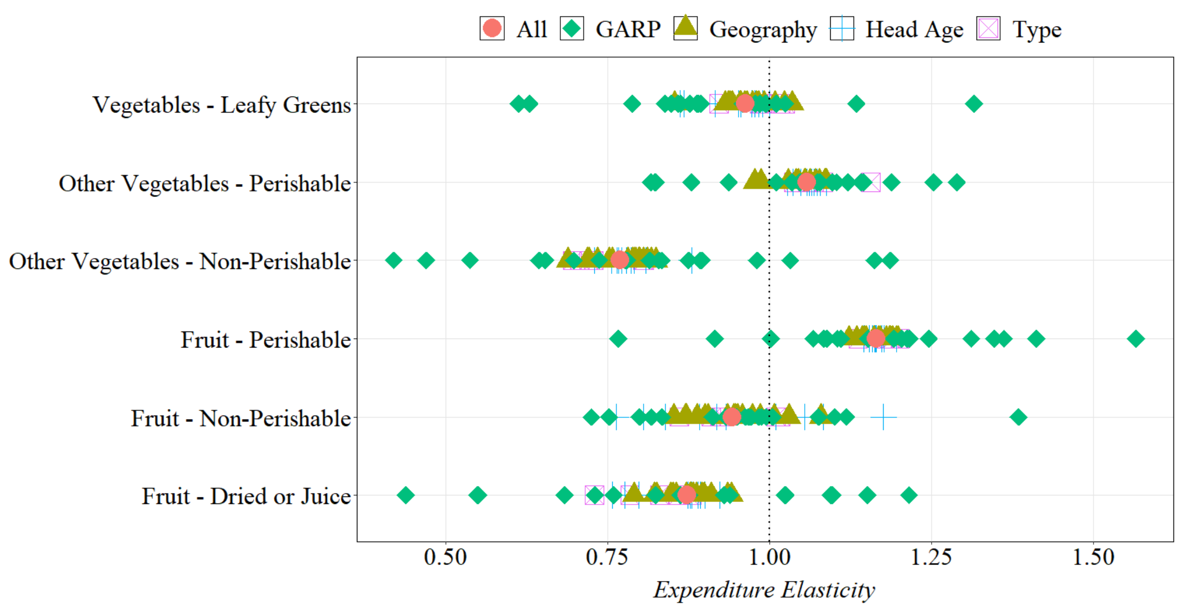

Figure 3a. Segment-Specific Marshallian Own-Price Elasticities (abs. value)

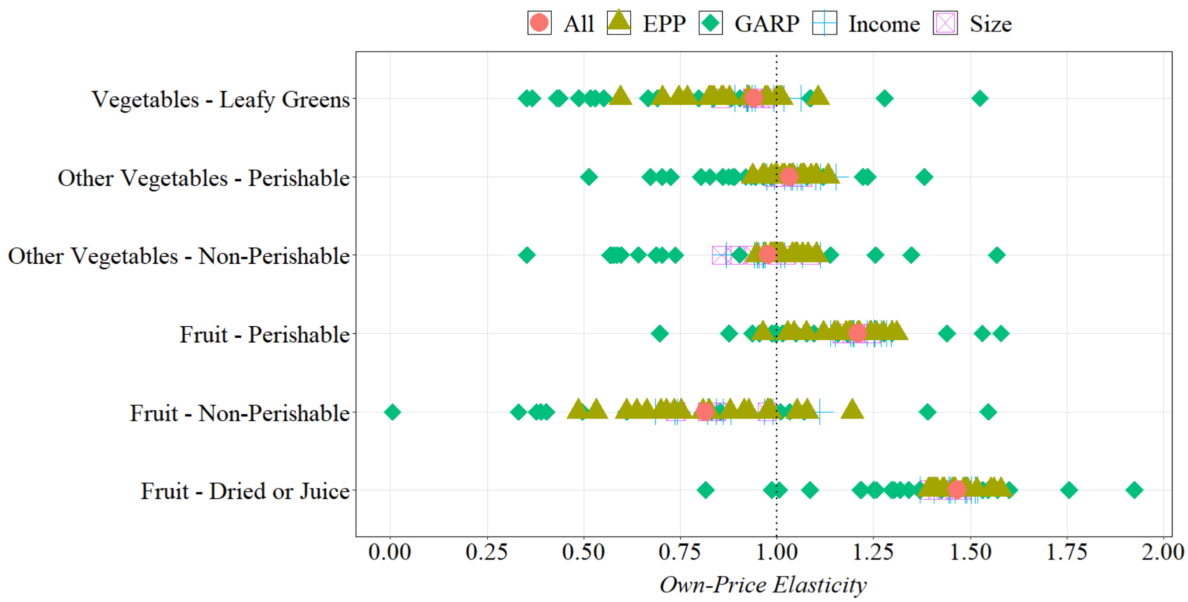

Figure 3b. Segment-Specific Marshallian Own-Price Elasticities (abs. value), cont.

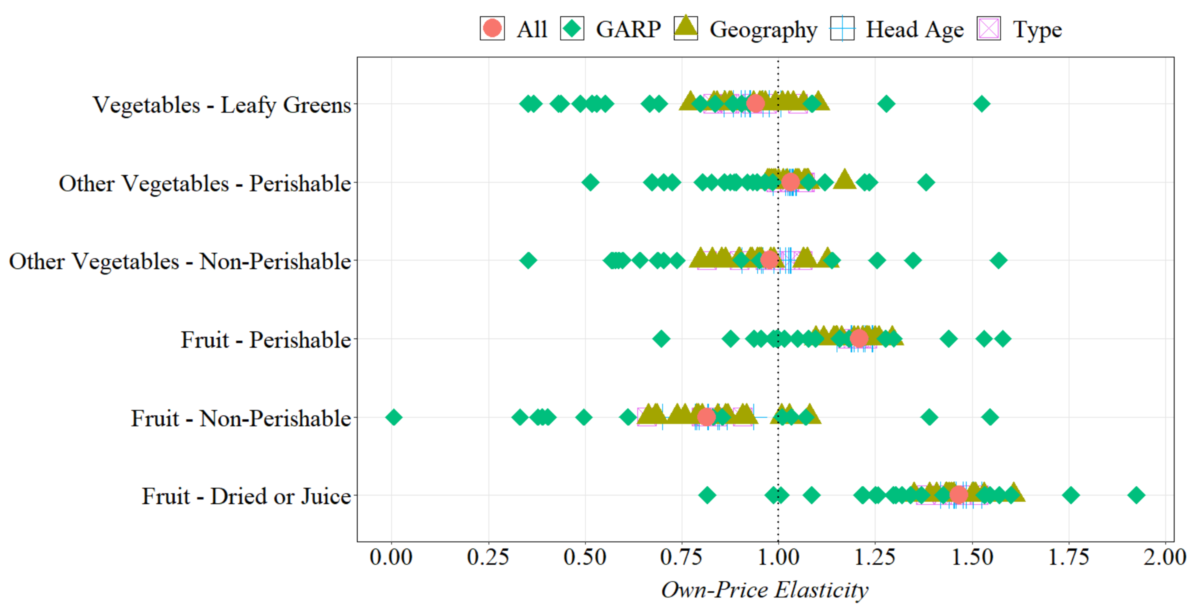

5.3. Out-of-Sample Forecasts

After generating parameter estimates and elasticities across all segments, consumption was forecasted into 2017. Elasticity-based forecasts predicted significantly better than the direct-statistical method, however, both methods of forecasting were arguably poor, see Table 6. The poor predictions may be the result of low R^2 measures for the predicted shares in each estimated demand system. Despite accounting for unobserved preference heterogeneity, predictions from segmenting by RP consistently performed the worst across all goods for both forecasting methods. This may result from consistent over- and under-estimating quantities due to "extreme" parameter and elasticity estimates generated from the GARP-consistent segments.

Table 6. Out-of-Sample Forecast RMSE, Ounces Purchased

	Direct Statistical					
Partition	LG_p	OV_p	OV_n	FR_p	FR_n	DJ_n
All	257.9	802.8	769.3	1,063.9	*427.4*	*1,415.8*
GARP	*259.5*	*806.0*	*775.9*	*1,076.8*	427.1	1,414.1
EPP	256.9	799.3	758.0	**1,035.9**	**422.5**	1,410.4
Geography	**256.0**	796.9	**756.0**	1,060.5	425.7	1,415.1
Age	256.7	802.4	765.1	1,058.0	426.2	**1,409.8**
Income	256.8	799.9	766.7	1,060.8	423.3	1,415.1
Size	257.6	**791.0**	765.7	1,066.9	426.4	1,412.1
Type	258.6	791.4	767.0	1,065.8	423.1	1,411.7
	Elasticity-Based					
All	165.0	511.4	492.2	691.0	**291.9**	936.9
GARP	*173.2*	*530.0*	*524.5*	*733.7*	294.1	*963.9*
EPP	**163.5**	508.5	501.2	**672.2**	*299.4*	**933.4**
Geography	165.3	511.3	**488.7**	689.7	295.1	942.0
Age	165.3	510.2	494.8	691.0	292.2	942.7
Income	165.4	510.4	495.4	688.5	293.6	937.6
Size	164.9	509.0	499.2	692.9	292.0	934.5
Type	165.1	**506.9**	498.1	692.9	292.8	934.5

Bold entries indicate best performing forecast for each good, *red* indicates worst performing

6. Summary and Conclusion

Revealed preference allows for an intuitive test of rationality, and segmenting microdata into GARP-consistent subsets allows for the identification of subpopulations with vastly different income and price responsiveness. A relatively low number of unique utility functions are required to perfectly rationalize fruit and vegetable purchasing behavior of 38,000 U.S. households. GARP-consistent segments display the widest range of elasticity estimates, but those elasticities result in the worst out-of-sample predictions.

Despite poor predictive power, GARP-consistent segments do provide some insight regarding low fruit and vegetable consumption that observable segments do not. Households with extreme elasticity estimates, both elastic and inelastic, may indicate extreme preferences. Incremental changes in prices and income can result in drastic demand responses for sensitive households or no response for households that simply do not like to eat fruits and vegetables. This suggests there is no "catch-all" approach in stimulating fruit and vegetable consumption for U.S. households. Nutritional assistance programs that incentivize buying fruits and vegetables by providing more credit will only benefit households who are only inhibited by purchasing power. Other households will require a change in their tastes and preferences to increase consumption, which may be attainable through education and awareness campaigns. Overall, revealed preference analysis provides key insights in understanding consumer behavior, however, it has limitations. Segmenting by GARP can be computationally intensive, and it is difficult to target consumers with different preferences, making it less attractive in industry and to policy makers.

Appendix

Appendix A. Descriptive Statistics by Segment

		Mean Price						Median Price						Max Price						Min Price					
Partition	Code	LG_p	OV_p	OV_n	FR_p	FR_n	DJ_n	LG_p	OV_p	OV_n	FR_p	FR_n	DJ_n	LG_p	OV_p	OV_n	FR_p	FR_n	DJ_n	LG_p	OV_p	OV_n	FR_p	FR_n	DJ_n
All Households	.	$ 0.21	$ 0.09	$ 0.08	$ 0.10	$ 0.11	$ 0.08	$ 0.20	$ 0.09	$ 0.08	$ 0.09	$ 0.10	$ 0.06	$ 0.70	$ 0.66	$ 0.66	$ 0.67	$ 0.66	$ 0.68	$ 0.03	$ 0.02	$ 0.02	$ 0.02	$ 0.02	$ 0.01
GARP Clusters	1	$ 0.21	$ 0.10	$ 0.08	$ 0.10	$ 0.11	$ 0.08	$ 0.20	$ 0.09	$ 0.08	$ 0.09	$ 0.11	$ 0.06	$ 0.68	$ 0.66	$ 0.51	$ 0.67	$ 0.55	$ 0.68	$ 0.06	$ 0.02	$ 0.03	$ 0.03	$ 0.02	$ 0.01
	2	$ 0.21	$ 0.09	$ 0.08	$ 0.10	$ 0.11	$ 0.08	$ 0.20	$ 0.09	$ 0.08	$ 0.09	$ 0.11	$ 0.06	$ 0.67	$ 0.40	$ 0.58	$ 0.43	$ 0.56	$ 0.61	$ 0.04	$ 0.02	$ 0.03	$ 0.02	$ 0.02	$ 0.01
	3	$ 0.21	$ 0.10	$ 0.08	$ 0.10	$ 0.11	$ 0.08	$ 0.20	$ 0.09	$ 0.08	$ 0.10	$ 0.10	$ 0.06	$ 0.69	$ 0.52	$ 0.48	$ 0.50	$ 0.66	$ 0.64	$ 0.06	$ 0.02	$ 0.03	$ 0.03	$ 0.02	$ 0.01
	4	$ 0.21	$ 0.09	$ 0.08	$ 0.10	$ 0.11	$ 0.07	$ 0.20	$ 0.09	$ 0.08	$ 0.09	$ 0.10	$ 0.06	$ 0.67	$ 0.45	$ 0.61	$ 0.59	$ 0.61	$ 0.50	$ 0.05	$ 0.02	$ 0.02	$ 0.02	$ 0.04	$ 0.02
	5	$ 0.21	$ 0.09	$ 0.08	$ 0.10	$ 0.11	$ 0.07	$ 0.20	$ 0.09	$ 0.08	$ 0.09	$ 0.11	$ 0.06	$ 0.67	$ 0.45	$ 0.47	$ 0.38	$ 0.50	$ 0.60	$ 0.06	$ 0.02	$ 0.03	$ 0.03	$ 0.03	$ 0.02
	6	$ 0.21	$ 0.10	$ 0.08	$ 0.10	$ 0.11	$ 0.08	$ 0.20	$ 0.09	$ 0.08	$ 0.09	$ 0.10	$ 0.06	$ 0.64	$ 0.44	$ 0.48	$ 0.51	$ 0.55	$ 0.52	$ 0.06	$ 0.02	$ 0.03	$ 0.03	$ 0.03	$ 0.02
	7	$ 0.21	$ 0.09	$ 0.08	$ 0.10	$ 0.11	$ 0.08	$ 0.20	$ 0.08	$ 0.08	$ 0.09	$ 0.10	$ 0.06	$ 0.63	$ 0.35	$ 0.46	$ 0.49	$ 0.41	$ 0.48	$ 0.07	$ 0.02	$ 0.03	$ 0.03	$ 0.04	$ 0.01
	8	$ 0.20	$ 0.09	$ 0.08	$ 0.10	$ 0.11	$ 0.07	$ 0.19	$ 0.08	$ 0.08	$ 0.09	$ 0.11	$ 0.06	$ 0.62	$ 0.33	$ 0.26	$ 0.41	$ 0.40	$ 0.48	$ 0.06	$ 0.02	$ 0.03	$ 0.03	$ 0.04	$ 0.01
	9	$ 0.20	$ 0.08	$ 0.08	$ 0.10	$ 0.11	$ 0.07	$ 0.19	$ 0.08	$ 0.07	$ 0.09	$ 0.11	$ 0.06	$ 0.54	$ 0.36	$ 0.33	$ 0.41	$ 0.40	$ 0.38	$ 0.06	$ 0.02	$ 0.03	$ 0.02	$ 0.02	$ 0.02
	10	$ 0.23	$ 0.10	$ 0.08	$ 0.10	$ 0.11	$ 0.07	$ 0.21	$ 0.09	$ 0.08	$ 0.10	$ 0.11	$ 0.05	$ 0.54	$ 0.27	$ 0.24	$ 0.35	$ 0.36	$ 0.55	$ 0.06	$ 0.02	$ 0.02	$ 0.03	$ 0.03	$ 0.02
	11	$ 0.21	$ 0.09	$ 0.08	$ 0.10	$ 0.11	$ 0.08	$ 0.20	$ 0.09	$ 0.08	$ 0.09	$ 0.10	$ 0.06	$ 0.69	$ 0.31	$ 0.30	$ 0.28	$ 0.41	$ 0.56	$ 0.05	$ 0.03	$ 0.03	$ 0.03	$ 0.02	$ 0.02
	12	$ 0.22	$ 0.09	$ 0.08	$ 0.10	$ 0.11	$ 0.08	$ 0.21	$ 0.08	$ 0.08	$ 0.09	$ 0.11	$ 0.06	$ 0.70	$ 0.31	$ 0.25	$ 0.37	$ 0.50	$ 0.45	$ 0.05	$ 0.02	$ 0.03	$ 0.03	$ 0.04	$ 0.02
	13	$ 0.20	$ 0.09	$ 0.08	$ 0.09	$ 0.11	$ 0.07	$ 0.20	$ 0.09	$ 0.08	$ 0.09	$ 0.11	$ 0.06	$ 0.66	$ 0.25	$ 0.53	$ 0.34	$ 0.54	$ 0.38	$ 0.03	$ 0.02	$ 0.03	$ 0.03	$ 0.04	$ 0.02
	14	$ 0.22	$ 0.09	$ 0.08	$ 0.11	$ 0.11	$ 0.07	$ 0.21	$ 0.09	$ 0.07	$ 0.10	$ 0.10	$ 0.06	$ 0.66	$ 0.51	$ 0.32	$ 0.33	$ 0.36	$ 0.51	$ 0.04	$ 0.02	$ 0.03	$ 0.03	$ 0.03	$ 0.02
	15	$ 0.22	$ 0.10	$ 0.08	$ 0.10	$ 0.11	$ 0.07	$ 0.21	$ 0.10	$ 0.08	$ 0.09	$ 0.10	$ 0.06	$ 0.68	$ 0.39	$ 0.66	$ 0.42	$ 0.44	$ 0.45	$ 0.06	$ 0.03	$ 0.03	$ 0.03	$ 0.04	$ 0.01
	16	$ 0.21	$ 0.09	$ 0.08	$ 0.11	$ 0.11	$ 0.08	$ 0.19	$ 0.08	$ 0.07	$ 0.10	$ 0.10	$ 0.06	$ 0.57	$ 0.28	$ 0.33	$ 0.37	$ 0.42	$ 0.50	$ 0.03	$ 0.02	$ 0.02	$ 0.03	$ 0.03	$ 0.02
	17	$ 0.21	$ 0.10	$ 0.08	$ 0.11	$ 0.11	$ 0.07	$ 0.20	$ 0.09	$ 0.08	$ 0.10	$ 0.11	$ 0.05	$ 0.55	$ 0.32	$ 0.23	$ 0.31	$ 0.33	$ 0.41	$ 0.06	$ 0.03	$ 0.03	$ 0.04	$ 0.04	$ 0.01
	18	$ 0.20	$ 0.09	$ 0.08	$ 0.09	$ 0.11	$ 0.07	$ 0.20	$ 0.08	$ 0.08	$ 0.09	$ 0.11	$ 0.06	$ 0.55	$ 0.35	$ 0.26	$ 0.37	$ 0.43	$ 0.40	$ 0.06	$ 0.03	$ 0.04	$ 0.02	$ 0.02	$ 0.02
	19	$ 0.21	$ 0.09	$ 0.08	$ 0.10	$ 0.11	$ 0.06	$ 0.20	$ 0.09	$ 0.08	$ 0.09	$ 0.10	$ 0.05	$ 0.65	$ 0.58	$ 0.34	$ 0.25	$ 0.34	$ 0.39	$ 0.03	$ 0.02	$ 0.04	$ 0.02	$ 0.03	$ 0.02
	20	$ 0.23	$ 0.10	$ 0.08	$ 0.10	$ 0.12	$ 0.07	$ 0.21	$ 0.09	$ 0.08	$ 0.09	$ 0.11	$ 0.06	$ 0.65	$ 0.29	$ 0.20	$ 0.28	$ 0.34	$ 0.43	$ 0.06	$ 0.02	$ 0.04	$ 0.04	$ 0.04	$ 0.01
Geography	1	$ 0.22	$ 0.10	$ 0.09	$ 0.11	$ 0.12	$ 0.08	$ 0.21	$ 0.09	$ 0.08	$ 0.10	$ 0.11	$ 0.06	$ 0.68	$ 0.52	$ 0.58	$ 0.49	$ 0.50	$ 0.62	$ 0.05	$ 0.02	$ 0.02	$ 0.03	$ 0.03	$ 0.02
	2	$ 0.21	$ 0.09	$ 0.08	$ 0.10	$ 0.11	$ 0.08	$ 0.20	$ 0.08	$ 0.08	$ 0.09	$ 0.10	$ 0.06	$ 0.66	$ 0.66	$ 0.61	$ 0.55	$ 0.56	$ 0.61	$ 0.04	$ 0.02	$ 0.03	$ 0.02	$ 0.03	$ 0.01
	3	$ 0.23	$ 0.10	$ 0.08	$ 0.10	$ 0.11	$ 0.08	$ 0.22	$ 0.10	$ 0.08	$ 0.10	$ 0.11	$ 0.06	$ 0.66	$ 0.51	$ 0.58	$ 0.38	$ 0.50	$ 0.62	$ 0.03	$ 0.02	$ 0.03	$ 0.03	$ 0.02	$ 0.02
	4	$ 0.22	$ 0.11	$ 0.09	$ 0.10	$ 0.12	$ 0.08	$ 0.21	$ 0.10	$ 0.08	$ 0.09	$ 0.12	$ 0.07	$ 0.66	$ 0.58	$ 0.66	$ 0.51	$ 0.55	$ 0.60	$ 0.04	$ 0.02	$ 0.03	$ 0.02	$ 0.02	$ 0.01
	5	$ 0.22	$ 0.09	$ 0.08	$ 0.10	$ 0.11	$ 0.08	$ 0.21	$ 0.09	$ 0.08	$ 0.10	$ 0.11	$ 0.06	$ 0.61	$ 0.33	$ 0.48	$ 0.35	$ 0.50	$ 0.60	$ 0.06	$ 0.02	$ 0.03	$ 0.03	$ 0.02	$ 0.02
	6	$ 0.20	$ 0.09	$ 0.08	$ 0.09	$ 0.11	$ 0.07	$ 0.19	$ 0.08	$ 0.07	$ 0.09	$ 0.10	$ 0.06	$ 0.67	$ 0.40	$ 0.26	$ 0.42	$ 0.40	$ 0.47	$ 0.06	$ 0.02	$ 0.02	$ 0.03	$ 0.02	$ 0.01
	7	$ 0.21	$ 0.09	$ 0.08	$ 0.10	$ 0.11	$ 0.07	$ 0.20	$ 0.09	$ 0.07	$ 0.10	$ 0.10	$ 0.06	$ 0.70	$ 0.34	$ 0.48	$ 0.67	$ 0.66	$ 0.55	$ 0.05	$ 0.02	$ 0.03	$ 0.03	$ 0.02	$ 0.01
	8	$ 0.21	$ 0.10	$ 0.09	$ 0.10	$ 0.12	$ 0.08	$ 0.21	$ 0.09	$ 0.08	$ 0.09	$ 0.11	$ 0.06	$ 0.61	$ 0.47	$ 0.57	$ 0.41	$ 0.50	$ 0.53	$ 0.05	$ 0.02	$ 0.03	$ 0.02	$ 0.04	$ 0.01
	9	$ 0.22	$ 0.09	$ 0.08	$ 0.10	$ 0.11	$ 0.07	$ 0.21	$ 0.09	$ 0.08	$ 0.09	$ 0.11	$ 0.05	$ 0.66	$ 0.36	$ 0.46	$ 0.59	$ 0.39	$ 0.50	$ 0.06	$ 0.03	$ 0.03	$ 0.03	$ 0.04	$ 0.02
	10	$ 0.20	$ 0.08	$ 0.08	$ 0.10	$ 0.11	$ 0.07	$ 0.19	$ 0.08	$ 0.07	$ 0.09	$ 0.10	$ 0.06	$ 0.69	$ 0.40	$ 0.27	$ 0.37	$ 0.61	$ 0.54	$ 0.03	$ 0.02	$ 0.03	$ 0.03	$ 0.03	$ 0.01
	11	$ 0.21	$ 0.09	$ 0.08	$ 0.10	$ 0.10	$ 0.07	$ 0.20	$ 0.08	$ 0.07	$ 0.09	$ 0.10	$ 0.06	$ 0.67	$ 0.66	$ 0.31	$ 0.50	$ 0.40	$ 0.41	$ 0.07	$ 0.02	$ 0.03	$ 0.03	$ 0.03	$ 0.01
	12	$ 0.21	$ 0.10	$ 0.08	$ 0.10	$ 0.12	$ 0.08	$ 0.20	$ 0.09	$ 0.08	$ 0.09	$ 0.11	$ 0.06	$ 0.64	$ 0.39	$ 0.33	$ 0.42	$ 0.42	$ 0.68	$ 0.07	$ 0.02	$ 0.02	$ 0.03	$ 0.03	$ 0.02
	13	$ 0.21	$ 0.08	$ 0.08	$ 0.10	$ 0.10	$ 0.07	$ 0.20	$ 0.08	$ 0.07	$ 0.09	$ 0.10	$ 0.06	$ 0.50	$ 0.23	$ 0.27	$ 0.39	$ 0.38	$ 0.64	$ 0.07	$ 0.02	$ 0.03	$ 0.03	$ 0.02	$ 0.02
	14	$ 0.19	$ 0.08	$ 0.07	$ 0.10	$ 0.10	$ 0.07	$ 0.19	$ 0.07	$ 0.07	$ 0.09	$ 0.10	$ 0.05	$ 0.62	$ 0.45	$ 0.35	$ 0.49	$ 0.32	$ 0.49	$ 0.03	$ 0.02	$ 0.03	$ 0.03	$ 0.04	$ 0.02
	15	$ 0.20	$ 0.08	$ 0.07	$ 0.10	$ 0.10	$ 0.07	$ 0.19	$ 0.08	$ 0.07	$ 0.09	$ 0.10	$ 0.06	$ 0.66	$ 0.32	$ 0.46	$ 0.40	$ 0.55	$ 0.49	$ 0.06	$ 0.02	$ 0.03	$ 0.03	$ 0.03	$ 0.02
	16	$ 0.21	$ 0.09	$ 0.08	$ 0.09	$ 0.11	$ 0.08	$ 0.20	$ 0.09	$ 0.08	$ 0.09	$ 0.11	$ 0.06	$ 0.63	$ 0.26	$ 0.47	$ 0.32	$ 0.43	$ 0.54	$ 0.07	$ 0.02	$ 0.03	$ 0.03	$ 0.04	$ 0.02

Appendix B. Individual Elasticity Graphs for Each Segment

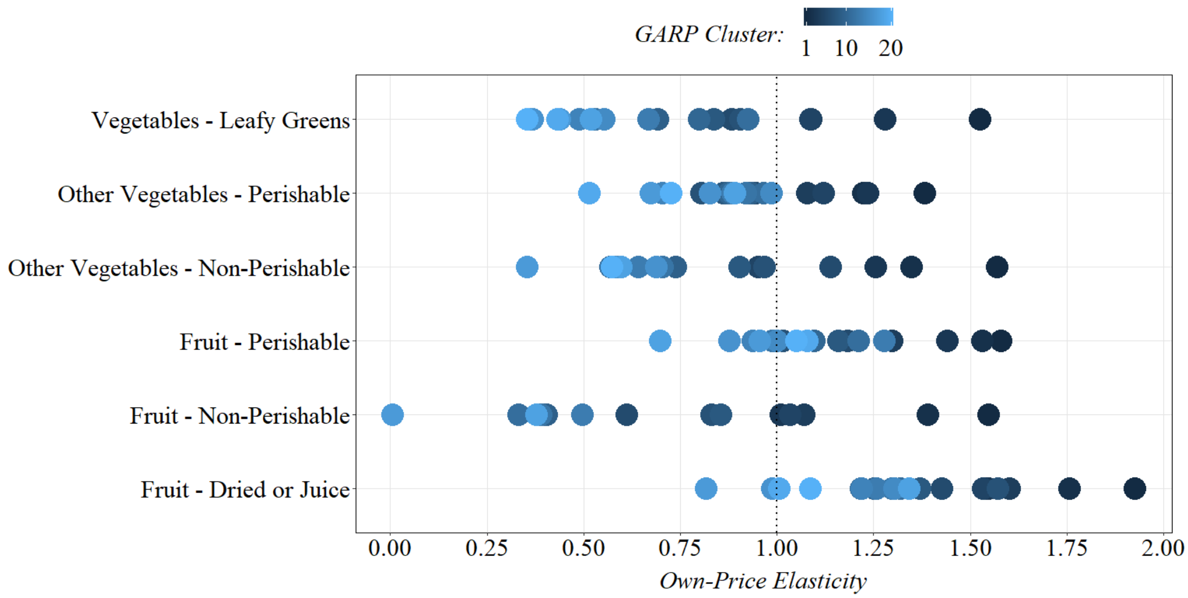

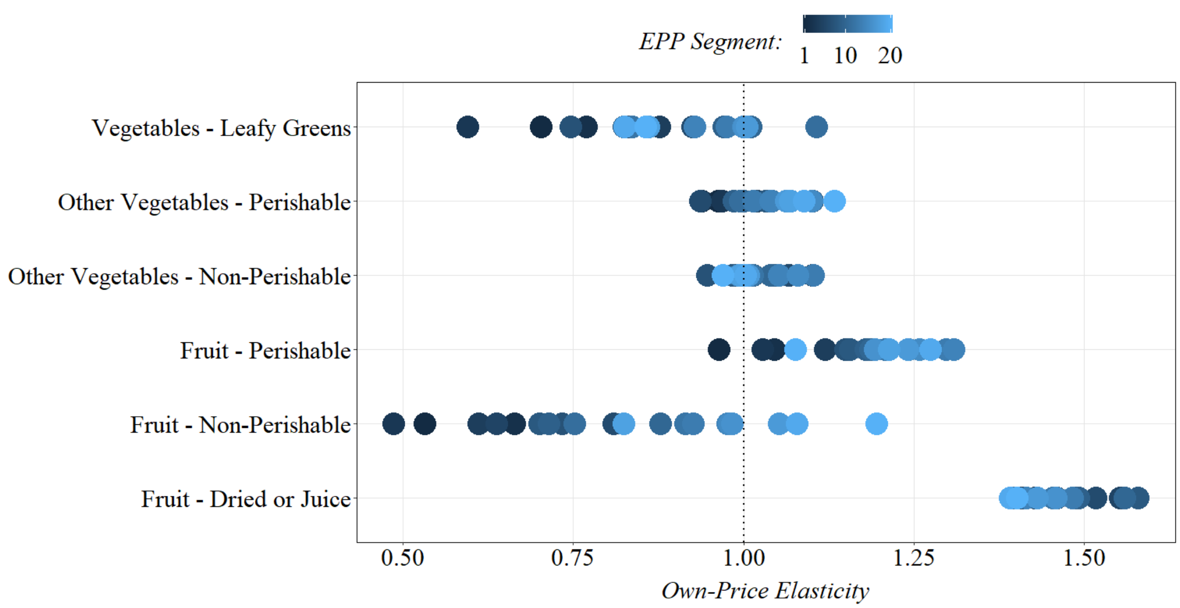

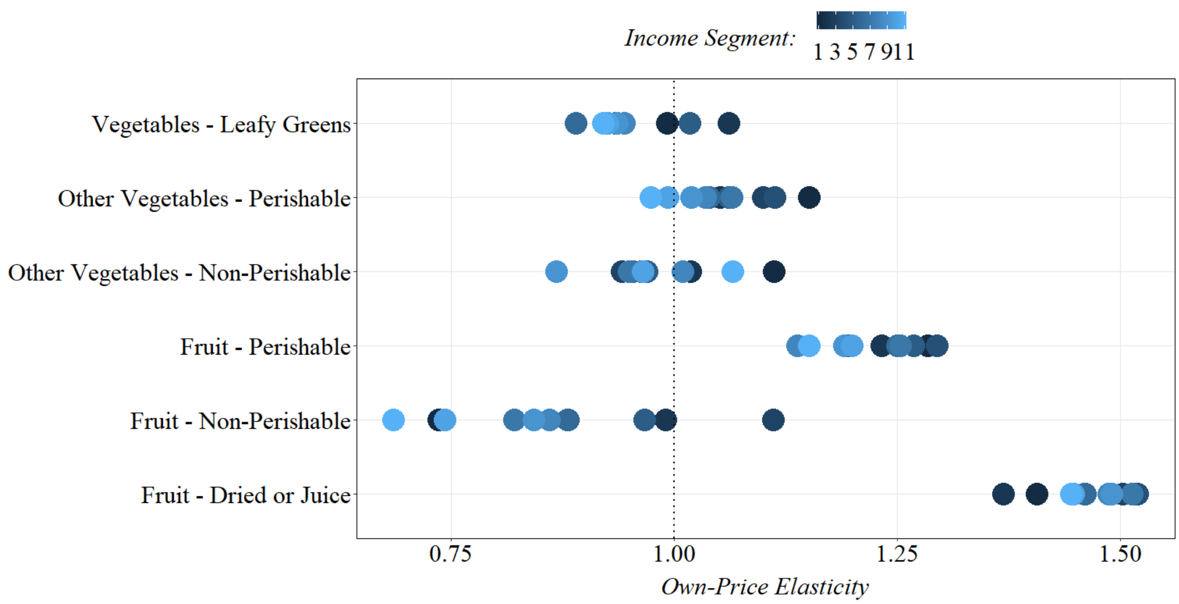
Income Segment:
1 3 5 7 9 11
Vegetables - Leafy Greens
Other Vegetables - Perishable
Other Vegetables - Non-Perishable
Fruit - Perishable
Fruit - Non-Perishable
Fruit - Dried or Juice
0.75
1.00
1.25
1.50
Own-Price Elasticity

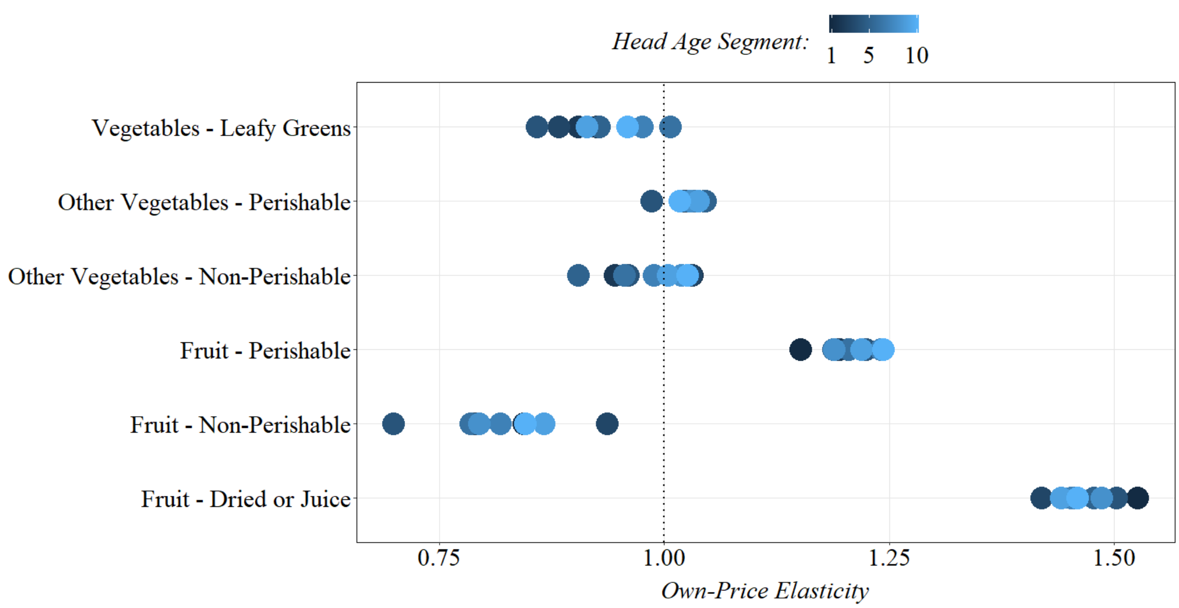
Head Age Segment:
1 5 10
Vegetables - Leafy Greens
Other Vegetables - Perishable
Other Vegetables - Non-Perishable
Fruit - Perishable
Fruit - Non-Perishable
Fruit - Dried or Juice
0.75
1.00
1.25
1.50
Own-Price Elasticity

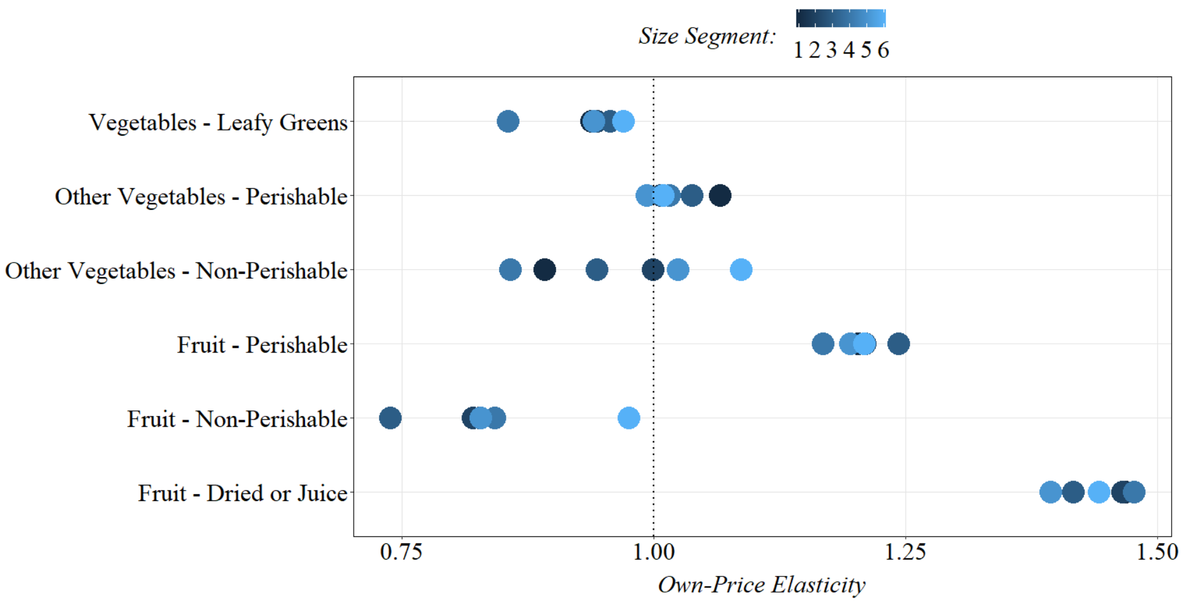
Size Segment:
1 2 3 4 5 6
Vegetables - Leafy Greens
Other Vegetables - Perishable
Other Vegetables - Non-Perishable
Fruit - Perishable
Fruit - Non-Perishable
Fruit - Dried or Juice
0.75
1.00
1.25
1.50
Own-Price Elasticity

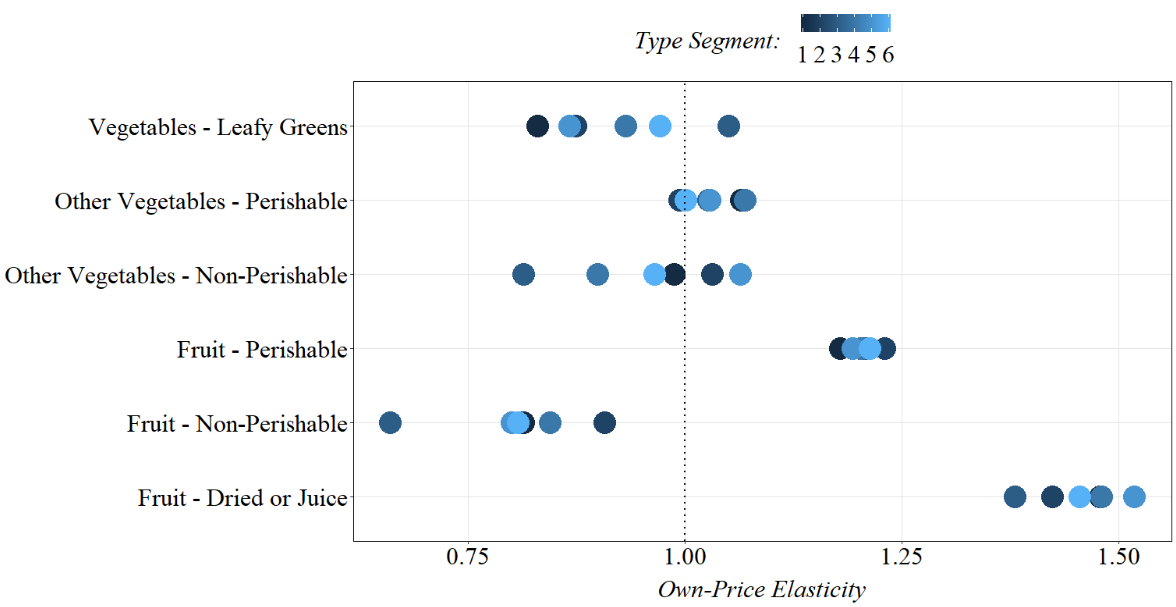
Type Segment:
1 2 3 4 5 6
Vegetables - Leafy Greens
Other Vegetables - Perishable
Other Vegetables - Non-Perishable
Fruit - Perishable
Fruit - Non-Perishable
Fruit - Dried or Juice
0.75
1.00
1.25
1.50
Own-Price Elasticity

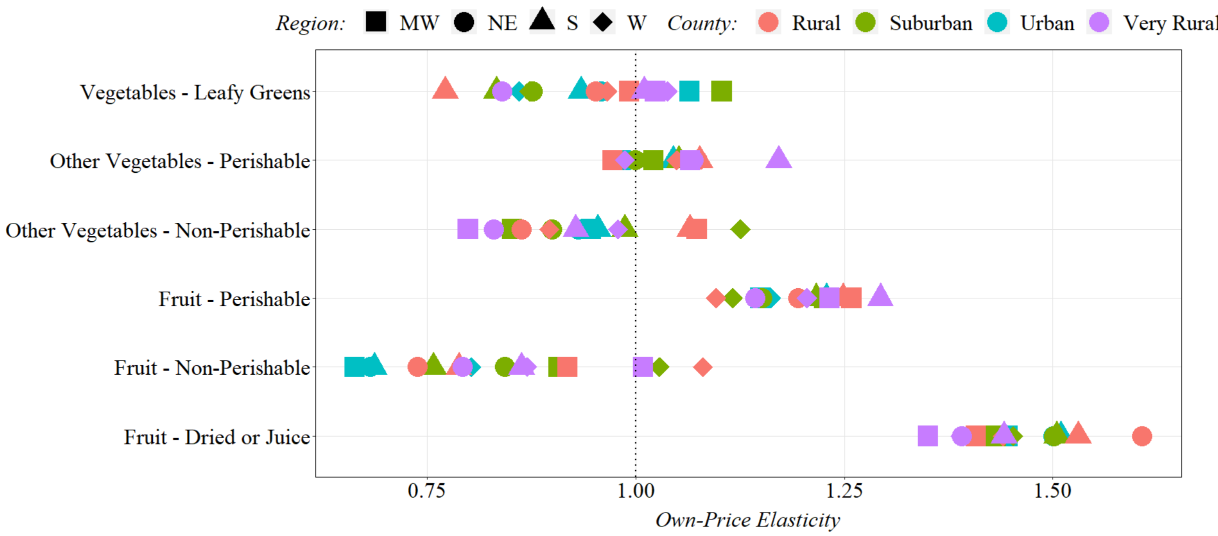
Region: MW NE S W County: Rural Suburban Urban Very Rural
Vegetables - Leafy Greens
Other Vegetables - Perishable
Other Vegetables - Non-Perishable
Fruit - Perishable
Fruit - Non-Perishable
Fruit - Dried or Juice
0.75
1.00
1.25
1.50
Own-Price Elasticity

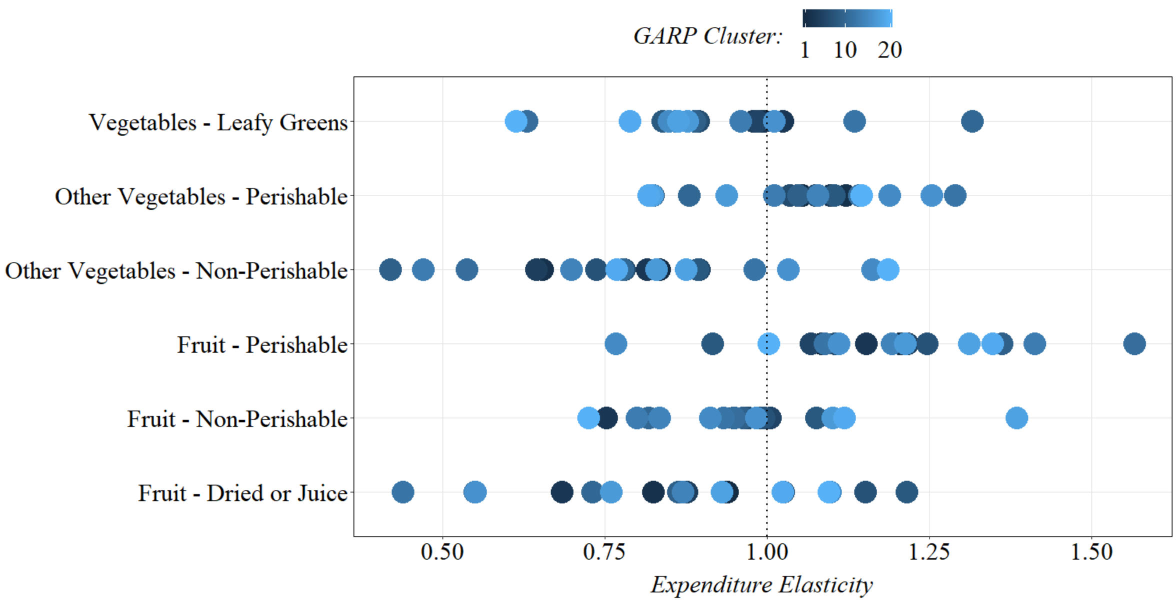
GARP Cluster: 1 10 20
Vegetables - Leafy Greens
Other Vegetables - Perishable
Other Vegetables - Non-Perishable
Fruit - Perishable
Fruit - Non-Perishable
Fruit - Dried or Juice
0.50
0.75
1.00
1.25
1.50
Expenditure Elasticity

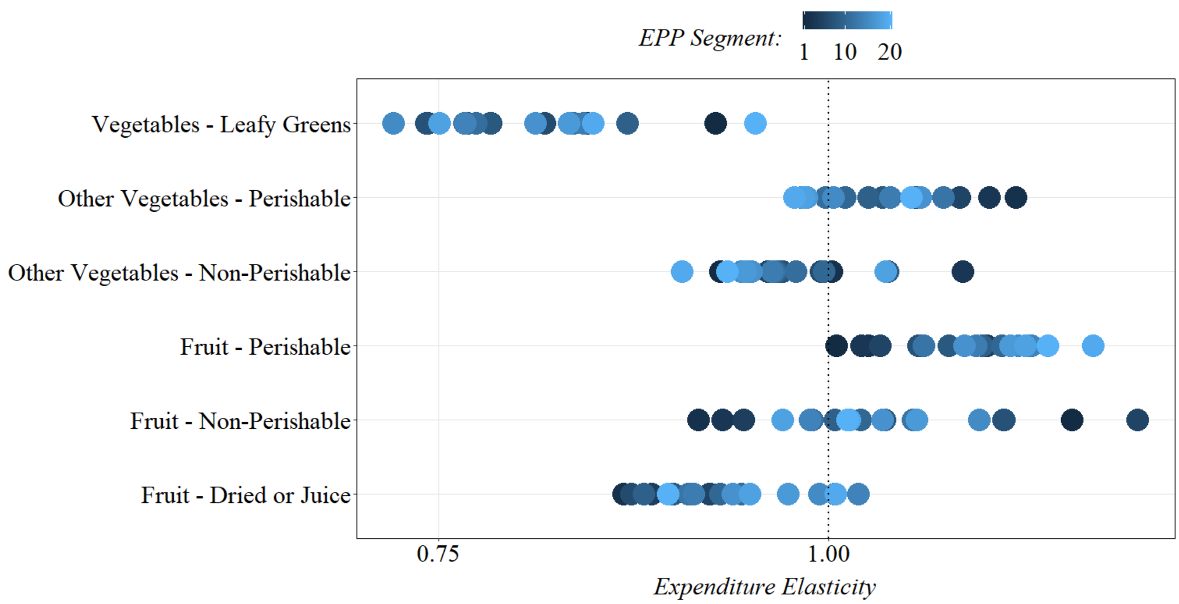

EPP Segment:
1 10 20
Vegetables - Leafy Greens
Other Vegetables - Perishable
Other Vegetables - Non-Perishable
Fruit - Perishable
Fruit - Non-Perishable
Fruit - Dried or Juice
0.75
1.00
Expenditure Elasticity

Income Segment:
1 3 5 7 9 11
Vegetables - Leafy Greens
Other Vegetables - Perishable
Other Vegetables - Non-Perishable
Fruit - Perishable
Fruit - Non-Perishable
Fruit - Dried or Juice
0.75
1.00
Expenditure Elasticity


Head Age Segment:
1 5 10
Vegetables - Leafy Greens
Other Vegetables - Perishable
Other Vegetables - Non-Perishable
Fruit - Perishable
Fruit - Non-Perishable
Fruit - Dried or Juice
0.75
1.00
Expenditure Elasticity

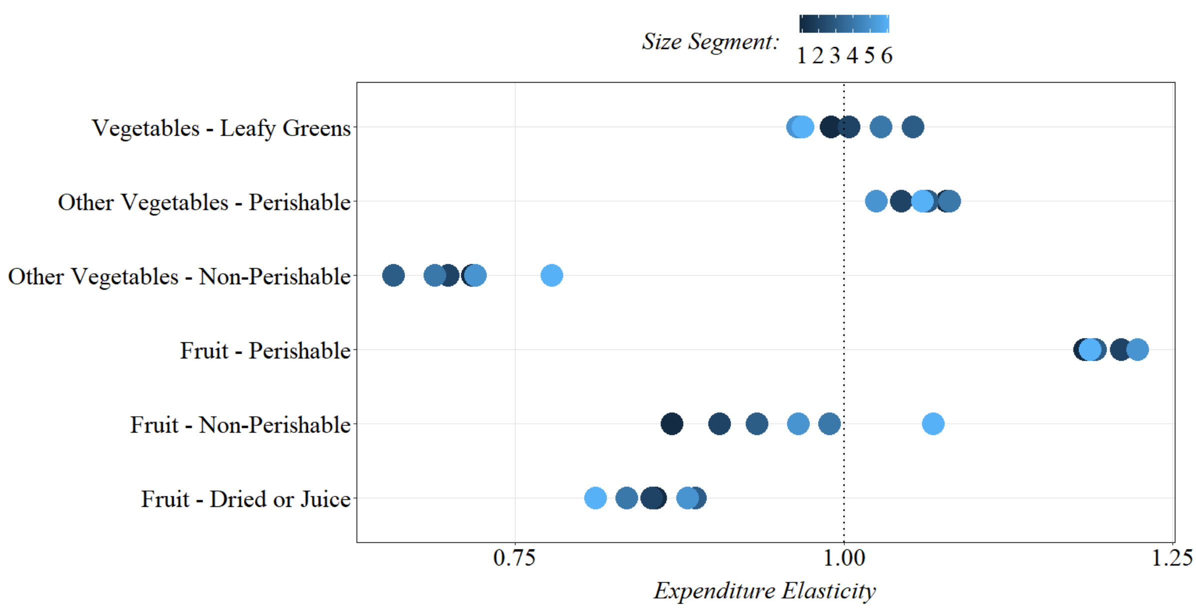
Size Segment:
1 2 3 4 5 6
Vegetables - Leafy Greens
Other Vegetables - Perishable
Other Vegetables - Non-Perishable
Fruit - Perishable
Fruit - Non-Perishable
Fruit - Dried or Juice
0.75
1.00
1.25
Expenditure Elasticity

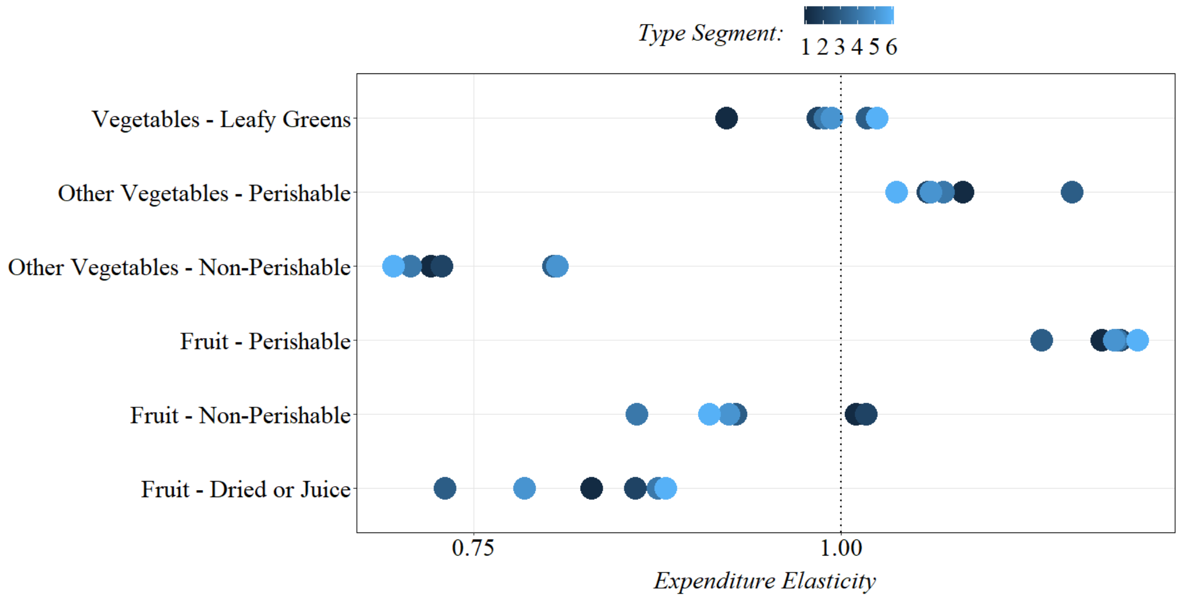

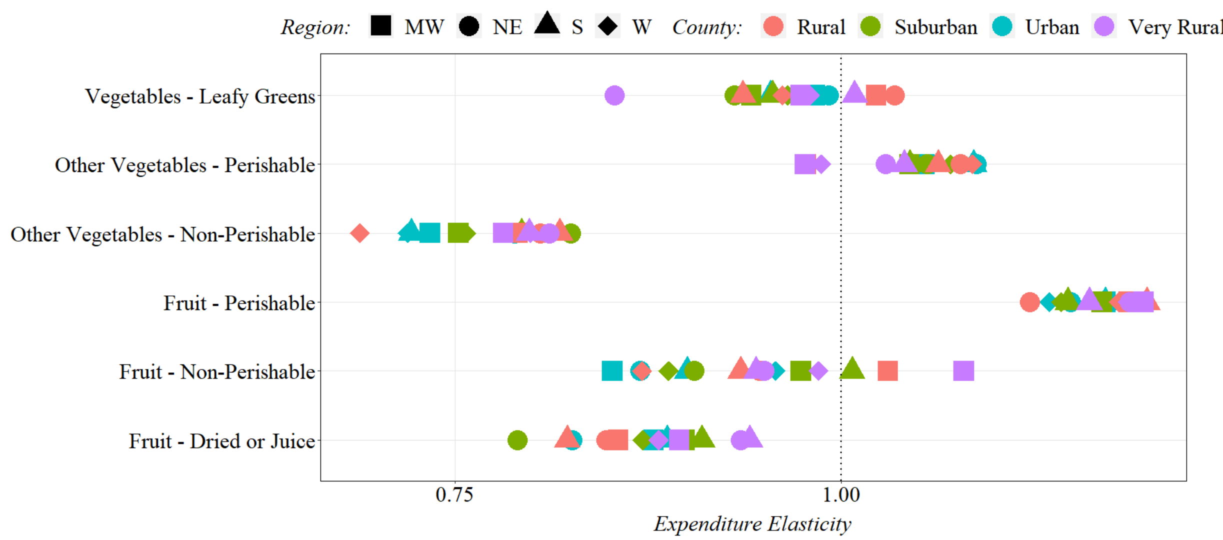

References

Afriat, S. N. "The Construction of Utility Functions from Expenditure Data." *International Economic Review* 8.1 (1967): 67.

Anderson, Gordon, and Richard Blundell. "Testing Restrictions in a Flexible Dynamic Demand System: An Application to Consumers' Expenditure in Canada." *The Review of Economic Studies* 50.3 (1983): 397.

Bertail, Patrice, and France Caillavet. "Fruit and Vegetable Consumption Patterns: A Segmentation Approach." *American Journal of Agricultural Economics* 90.3 (2008): 827-42.

Blundell, Richard W., Martin Browning, and Ian A. Crawford. "Nonparametric Engel Curves and Revealed Preference." *Econometrica* 71.1 (2003): 205-40.

Blundell, Richard, and Jean Marc Robin. "Estimation in Large and Disaggregated Demand Systems: An Estimator for Conditionally Linear Systems." *Journal of Applied Econometrics* 14.3 (1999): 209-32.

Blundell, Richard. "How Revealing Is Revealed Preference?" *Journal of the European Economic Association* 3.2-3 (2005): 211-35.

Boelaert, Julien, and Étienne Ollion. "The Great Regression." *Revue Française De Sociologie* 59.3 (2018): 475.

Boonsaeng, Tullaya, and Carlos E. Carpio. "A Comparison of Food Demand Estimation from Homescan and Consumer Expenditure Survey Data." *Selected Paper, AAEA 2014 Annual Meeting* (2014)

Browning, Martin, and Costas Meghir. "The Effects of Male and Female Labor Supply on Commodity Demands." *Econometrica* 59.4 (1991): 925.

Cohen, Jack. *Statistical Power Analysis for the Behavioral Sciences*. N.p.: L. Erlbaum Associates, 1988.

Crawford, Ian, and Bram De Rock. "Empirical Revealed Preference." *Annual Review of Economics* 6.1 (2014): 503-24.

Crawford, Ian, and Krishna Pendakur. "How Many Types Are There?" *The Economic Journal* 123.567 (2012): 77-95.

Dong, Diansheng, and Biing-Hawn Lin. "Fruit and Vegetable Consumption by Low-Income Americans." *USDA Economic Research Report Number 70* (2009): n. pag.

Einav, Liran, Ephraim Leibtag, and Aviv Nevo. "Recording Discrepancies in Nielsen Homescan Data: Are They Present and Do They Matter?" *Quantitative Marketing and Economics* 8.2 (2009): 207-39.

Gustavsen, G. W. "Forecasting Ability of Theory-constrained Two-stage Demand Systems." *European Review of Agriculture Economics* 30.4 (2003): 539-58.

Hales, Craig M., Cheryl D. Fryar, Margaret D. Carroll, and Cynthia L. Ogden. "Prevalence of Obesity Among Adults and Youth: United States, 2015–2016." *NCHS Data Brief* 288 (2017)

Imamura, Fumiaki, Laura O'Connor, Zheng Ye, Jaakko Mursu, Yasuaki Hayashino, Shilpa N. Bhupathiraju, and Nita G. Forouhi. "Consumption of Sugar Sweetened Beverages, Artificially Sweetened Beverages, and Fruit Juice and Incidence of Type 2 Diabetes: Systematic Review, Meta-analysis, and Estimation of Population Attributable Fraction." *Bmj* (2015)

Kastens, Terry L., and Gary W. Brester. "Model Selection and Forecasting Ability of Theory-Constrained Food Demand Systems." *American Journal of Agricultural Economics* 78.2 (1996): 301.

Lee-Kwan, Seung Hee, Latetia V. Moore, Heidi M. Blanck, Diane M. Harris, and Deb Galuska. "Disparities in State-Specific Adult Fruit and Vegetable Consumption — United States, 2015." *MMWR. Morbidity and Mortality Weekly Report* 66.45 (2017): 1241-247.

Lusk, Jayson L., and Kathleen Brooks. "Who Participates in Household Scanning Panels?" *American Journal of Agricultural Economics* 93.1 (2011): 226-40.

Michalek, J., and M. A. Keyzer. "Estimation of a Two-stage LES-AIDS Consumer Demand System for Eight EC Countries." *European Review of Agricultural Economics* 19.2 (1992): 137-63.

Murphy, Sherry L., Jiaquan Xu, Kenneth D. Kochanek, and Elizabeth Arias. *Mortality in the United States, 2017*. Proc. of NCHS Data Brief. U.S. Dept. of Health and Human Services, Centers for Disease Control and Prevention, National Center for Health Statistics, 2018.

Muth, Mary K., Megan Sweitzer, Derick Brown, Kristen Capogrossi, Shawn Karns, David Levin, Abigail Okrent, Peter Siegel, and Chen Zhen. "Understanding IRI Household-Based and Store-Based Scanner Data." *Economic Research Service, Technical Bulletin* 1942 (2016)

"MyPlate Plan: 2800 Calories, Age 14+." *Choose MyPlate*, 14 Aug. 2018, www.choosemyplate.gov/MyPlatePlan_2800cals_Age14plus.

Rickman, Joy C., Diane M. Barrett, and Christine M. Bruhn. "Nutritional Comparison of Fresh, Frozen and Canned Fruits and Vegetables. Part 1. Vitamins C and B and Phenolic Compounds." *Journal of the Science of Food and Agriculture* 87.6 (2007): 930-44.

Samuelson, P. A. "A Note on the Pure Theory of Consumer's Behaviour." *Economica* 5.17 (1938): 61.

Samuelson, Paul A. "Consumption Theory in Terms of Revealed Preference." *Economica* 15.60 (1948): 243.

Stewart, Hayden, Jeffrey Hyman, Andrea Carlson, and Elizabeth Frazao. "The Cost of Satisfying Fruit and Vegetable Recommendations in the Dietary Guidelines." *USDA ERS Economic Brief Number 27* (2016)

Stigler, George J., and Gary S. Becker. "De Gustibus Non Est Disputandum." *The American Economic Review* 67.2 (1977)

USDA. "Dietary Guidelines for Americans 2015-2020." N.p., Dec. 2015. Web.

Varian, Hal R. "Non-parametric Analysis of Optimizing Behavior with Measurement Error." *Journal of Econometrics* 30.1-2 (1985): 445-58.

Varian, Hal R. "Non-Parametric Tests of Consumer Behaviour." *The Review of Economic Studies* 50.1 (1983): 99.

Varian, Hal R. "Revealed Preference and Its Applications." *The Economic Journal* 122.560 (2012): 332-38.

Varian, Hal R. "The Nonparametric Approach to Demand Analysis." *Econometrica* 50.4 (1982): 945.

World Health Organization. *WHO Guide to Identifying the Economic Consequences of Disease and Injury*. 2009, www.who.int/choice/publications/d_economic_impact_guide.pdf.

Yang, Dongsheng, and Jarrod E. Dalton. "A Unified Approach to Measuring the Effect Size between Two Groups Using SAS." *SAS Global Forum 2012* Paper 335 (2012)

Yeh, M.-C., S. B. Ickes, L. M. Lowenstein, K. Shuval, A. S. Ammerman, R. Farris, and D. L. Katz. "Understanding Barriers and Facilitators of Fruit and Vegetable Consumption among a Diverse Multi-ethnic Population in the USA." *Health Promotion International* 23.1 (2008): 42-51.

Zhen, Chen, Justin L. Taylor, Mary K. Muth, and Ephraim Leibtag. "Understanding Differences in Self-Reported Expenditures between Household Scanner Data and Diary Survey Data: A Comparison of Homescan and Consumer Expenditure Survey." *Review of Agricultural Economics* 31.3 (2009): 470-92.