Substitution Patterns and Welfare Implications of Local Taxation: Empirical Analysis of a Soda Tax

Jiawei Chen Colin Reinhardt Saad Andalib Syed Shah*

January 2025

Abstract

We present a structural choice model that incorporates households' geographic and product substitution for studying the effects of localized taxation policies. Using detailed retail and household data pertaining to Philadelphia's soda tax, we estimate the choice model linking households' demographic characteristics and proximity to the city border to their tax avoidance behavior: switching from taxed to untaxed products or from Philadelphia to non-Philadelphia stores. The inclusion of travel time is vital for modeling households' heterogeneous responses, with an extra minute of travel time to reach the untaxed region equivalent to adding 28.5¢ (7%) to the product price. Compared to broader regional taxation, localized taxation proves highly inefficient: Philadelphia households on average incur a consumer surplus loss more than double their tax payment, with cross-border travel cost alone accounting for 41% of this loss.

JEL: D04, D12, H26, H71, I18, L66. *Keywords:* demand estimation, substitution patterns, consumer heterogeneity, local taxation, policy evaluation.

---

*Chen: University of California, Irvine, jiaweic@uci.edu. Reinhardt: Office of the Comptroller of the Currency, creinha1@uci.edu. Syed Shah: Wenzhou-Kean University, ssaadand@wku.edu.cn. The views expressed in this paper are the authors' and do not necessarily reflect the official position of the Office of the Comptroller of the Currency or the United States Goverment. We thank Audrey Guo, Myongjin Kim, Qihong Liu, Dongsoo Shin, Matthew Weinberg, Kyle Wilson and participants at Econometric Society Australasia Meeting, EARIE Annual Conference, AEA Annual Meeting, Santa Clara University, and University of Oklahoma for helpful comments and suggestions. Researchers' own analyses calculated or derived based in part on data from Nielsen Consumer LLC and marketing databases provided through the NielsenIQ Datasets at the Kilts Center for Marketing Data Center at The University of Chicago Booth School of Business. The conclusions drawn from the NielsenIQ data are those of the researchers and do not reflect the views of NielsenIQ. NielsenIQ is not responsible for, had no role in, and was not involved in analyzing and preparing the results reported herein. We did not generate our own data and did not register our study or write a pre-analysis plan before data collection.

1 Introduction

Governments of all types levy "sin taxes"—excise taxes imposed on certain goods deemed harmful to society and individuals—with the dual, and oftentimes competing, motives of curbing consumption and raising tax revenue. Examples include taxes on tobacco, alcohol, gambling, drugs[1], junk foods[2], etc. The present study designs a model to evaluate the effects of an increasingly popular category of sin taxes—"soda taxes," which are imposed on sweetened beverages (SBs)—by paying particular attention to both cross-border shopping (geographic substitution) and switching to alternative products (product substitution) as forms of tax avoidance.

We focus on the SB tax implemented in the US city of Philadelphia. Philadelphia provides a set of conditions that benefits researchers interested in the effects of SB taxation. First, Philadelphia is demographically diverse, particularly in terms of income and racial distributions, which allows researchers to better understand the heterogeneous effects of the taxation on the city's rich, poor, and underrepresented minority households. Second, Philadelphia is a large urban center with a substantial set of retail-level and household-level data available. Finally, the city of Philadelphia is both expansive and surrounded by a large suburban population, which provides an ideal setting for studying the effects of geographic and product substitution.

The difference between geographic and product substitution is an important one. For a local government collecting tax revenue, geographic substitution hurts local businesses and lowers tax revenue as consumers take their SB purchase—and with it their grocery shopping—to other locations, whereas product substitution leaves consumers' purchases in the same location. For public health agencies, geographic substitution defeats the purpose of the tax as consumers continue to buy unhealthy products and only change where they buy them, whereas product substitution achieves exactly the health objective of the tax by diverting consumption from unhealthy products to healthier ones. A good understanding of the relation between and the magnitudes of geographic and product substitution is then an important prerequisite for sound policymaking, for local governments and public health agencies alike.

Besides SB taxation, analogous scenarios featuring such tension between geographic and product substitution apply to many policies implemented by states, counties, or cities, including all kinds of sin taxes collected at the local level, other types of local taxes[3], and local subsidies for products such as healthy foods and gasoline[4]. By providing a structural empirical analysis

[1] Such as legal marijuana (Hollenbeck and Uetake, 2021).

[2] See for example Yazzie et al. (2020).

[3] Such as gasoline taxes at the state level and local amusement taxes (Breslow, 2019).

[4] A local subsidy not only induces local consumers to switch from unsubsidized products to subsidized ones, but also incentivizes consumers in other locations to travel to the subsidized location in pursuit of lower prices. For example, in 2022 when subsidized gasoline prices in Mexico are noticeably lower than prices in the US, many US drivers cross the border into Mexico to fill their tanks, leading to a gasoline shortage and temporary suspension of

of a local policy that decomposes consumers' heterogeneous substitution responses along the dimensions of geographic and product substitution, this paper offers a useful framework for similar studies on local taxation or subsidy policies.

To quantify the effects of Philadelphia's SB tax on consumers' product and location choices and their welfare, we construct and estimate a model of consumer demand in the random coefficients nested logit (RCNL) framework (e.g., Grigolon and Verboven (2014), Miller and Weinberg (2017), and Miravete et al. (2018)) using a combination of retail and household data. The random coefficients approach allows rich modeling of heterogeneity in consumer tastes and travel costs, while the nested structure is particularly suited to our analysis of consumers' substitution across beverage categories ("nests").[5] Aggregate-level retail data lacks the information needed to track individual households' heterogeneous responses to the tax, but measures the aggregate effect of the tax with far less noise and provides a reliable method by which one can account for endogenous variables. Micro-level household data covers only a small subset of all households, but provides an accurate measure of consumer heterogeneity and responsiveness to travel costs. Our empirical approach combines the strengths of the above elements and incorporates the two kinds of data in an internally consistent way.

In estimation, we follow an approach suggested in Grieco et al. (2022) to recover mean utility and unobserved demand shocks while accounting for heterogeneous tastes and cross-border shopping.[6] Our results include estimates of mean responses to SB taxation and travel time as well as heterogeneous parameters related to preference and substitution. Our paper is one of the first studies that estimate an RCNL model using a combination of aggregate-level and micro-level data (joining Conlon and Rao (2023)).[7]

Several key findings emerge from our analysis. (1) Travel time to the alternative region plays a key role in determining households' willingness to cross-border shop, the effectiveness of the taxation, and changes in consumer surplus. On average, an extra minute of travel time to reach a store in the alternative region is equivalent to adding 28.5¢ (7%) to the product price. (2) Compared to broader regional taxation, localized taxation proves highly inefficient: Philadelphia households on average incur a consumer surplus loss more than double their tax payment,

the gasoline subsidy in Mexico's US border region (Garrison and Barrera, 2022).

[5] An alternative to the RCNL approach is to forgo the nested structure and instead model additional random coefficients on product category dummies. In our case, there are eight beverage categories, so this alternative approach would require a much larger number of simulation draws in order to adequately approximate the distributions of the many random coefficients, which would impose a severe computational burden.

[6] Several other papers have used similar methods combining retail and household data, including Goolsbee and Petrin (2004), Chintagunta and Dubé (2005), Tuchman (2019), and Murry and Zhou (2020).

[7] In our estimation process, we found that the inclusion of household data, rather than relying solely on retail data, greatly facilitates the estimation of the RCNL model, particularly the estimation of the nesting parameter (compared to using moment conditions derived from aggregate-level data).

with cross-border travel cost alone accounting for 41% of this loss. (3) The SB tax is highly regressive and exacerbates the disparity among consumers. When measured as a percentage of annual income, low-income underrepresented minority (URM) households in Philadelphia on average pay 6.5 times more in SB taxes and incur a loss in consumer surplus 7.6 times greater than high-income non-URM households. (4) Accounting for households' heterogeneous preferences and substitution patterns, the revenue-maximizing tax rate is 2.51¢ per ounce. Compared to this rate, the actual tax rate of 1.5¢ per ounce results in 94% of the tax revenue. Philadelphia's decision to include artificially sweetened (diet) products in the tax strongly influences the revenue-maximizing tax rate, and excluding such products from the tax would result in a revenue-maximizing tax rate of 1.90¢ per ounce.

As of July 2022, excluding Cook County in the state of Illinois and the Navajo Nation, all SB taxes in the US have been implemented at the city level. Given the relatively small area of taxation, these SB taxation policies are especially vulnerable to tax avoidance behavior in the form of cross-border shopping. Roberto et al. (2019) compare pre- and post-taxation SB sales in and around Philadelphia, concluding that 24% of the decrease in Philadelphia SB sales due to the SB tax is offset by an increase in sales in the surrounding region. Similarly, Seiler et al. (2021) find evidence of cross-border shopping by Philadelphia households to the city's surrounding region, indicating that such behavior offsets 52% of the sales reduction resulting from the city's SB tax. In the general market for food products, cross-border shopping as a response to sales taxes has been observed in the District of Columbia (Fisher, 1980) and West Virginia (Tosun and Skidmore, 2007), among others.

Literature pertaining to both aggregate-level data (e.g., Thomadsen (2005), Davis (2006), and Houde (2012)) and micro-level data (e.g., McFadden et al. (1977), Capps et al. (2003), Bayer et al. (2007), and Burda et al. (2008)) finds that distance plays an important role in determining product choices. In terms of cross-border shopping, Harding et al. (2012) show that the distance to a lower-tax border affects the pass-through rates of state cigarette taxes, suggesting that consumers engage in cross-state purchasing, which pushes the burden of taxation backwards onto the factors of production. Chandra et al. (2014) find that longer driving distances strongly disincentivize shopping across the US-Canadian border in search of cheaper alternatives. Cross-border shopping as a function of geographic distance has also been identified in Denmark (Bygvrå, 2009) and Norway (Friberg et al., 2018). Our analysis builds upon the idea that distance plays a large role in inhibiting cross-border shopping, and applies it to the policy evaluation of Philadelphia's SB tax. Our modeling of *travel time* as a measure of distance within an RCNL model provides a novel approach for incorporating heterogeneous cross-locational substitution patterns into the analysis of consumer choices.

Through the inclusion of geographic and product substitution of beverages in a choice modeling structure, our paper also contributes to the expanding set of SB taxation literature. Prior works that have considered Philadelphia's SB tax as well as cross-border shopping, such as Roberto et al. (2019) and Seiler et al. (2021), have used either retail-level or household-level data but not both and have relied on reduced form estimation techniques. We complement those existing works by using both retail-level and household-level data to estimate consumer behavior and aggregate responsiveness to taxation, and by conducting counterfactual analyses made possible by the structural estimation results. In the context of structural modeling, Kifer (2015), Wang (2015), Allcott et al. (2019) and Dubois et al. (2020) have used pre-taxation data to predict the effects of hypothetical SB taxes. We take a different approach by studying the actual implementation of an SB tax, incorporating both retail-level and household-level data, and accounting for the effects of geographic substitution.

The remainder of this paper proceeds as follows. In Section 2, we introduce background information about the Philadelphia SB tax. We describe our data sources and provide information about the products and market in Section 3. Section 4 details the discrete choice model of demand that incorporates both the retail and household data. In Section 5, we discuss model identification and estimation. Section 6 presents the results of our demand estimation. We discuss the effects of the taxation on prices, market shares and consumption in Section 7. Changes in consumer surplus and the heterogeneous impact of the taxation by household demographics are discussed in Section 8. Section 9 derives the revenue-maximizing tax rate and explores the effects of alternative taxation schemes. Section 10 concludes.

2 Philadelphia Soda Tax

On June 16th, 2016, Philadelphia became the second US city to pass an SB tax, after Berkeley. Initially proposed as a 3¢-per-ounce tax on all sugar-sweetened beverages, the measure garnered widespread support. Supporters of the proposal, such as the American Medical Association, American Heart Association, and other medical groups, argued that such a tax would combat the twin epidemics of obesity and heart disease. Philadelphia ranks as one of the worst cities in the US in terms of type 2 diabetes, heart disease, and obesity. City mayor Jim Kenney predicted the tax would raise $400 million over five years, which would be used to fund universal pre-kindergarten, job creation, and development projects.

Opponents of the proposal claimed that the measure would disproportionately affect the least fortunate. The American Beverage Association, a lobbying group formed of beverage manufacturers and distributors, pushed newspaper, radio and television ads condemning the proposal as regressive—burdening URM and low-income communities with the largest share of the tax. Interest in the measure was so high that Democratic primary candidates Hillary Clinton and

Bernie Sanders weighed in with their opinions for and against the measure, respectively. After months of negotiation, a compromise was reached.

Passing with a city council vote of 13-to-4, the final draft required distributors to pay a 1.5¢-per-ounce tax on all sugar-/artificially sweetened beverages, with the law becoming effective on January 1st, 2017.[8] Thus, the tax applies to not only beverages sweetened with sugar but also diet beverages containing artificial sweeteners. While it may seem surprising to tax artificially sweetened beverages, given that artificial sweeteners have virtually no calories and that diet beverages (beverages with few or no calories) are generally considered healthier alternatives, the city council included diet beverages in the tax to make up for lost revenue as a result of decreasing the tax from the proposed 3¢ per ounce to the actual 1.5¢ per ounce. Most other soda taxes (Berkeley, CA; Boulder, CO; Seattle, WA; etc.) tax only products with added caloric sweeteners, thus excluding diet beverages.

3 Data

In this section, we describe the data used in our estimation.

3.1 Retail Data

Our retail dataset, from NielsenIQ through the Kilts Center for Marketing at The University of Chicago Booth School of Business, covers the 4-year period from January 1st, 2015 to December 31st, 2018 (Philadelphia's SB tax took effect at the midpoint of this period on January 1st, 2017). The dataset contains store-level information detailing weekly price and quantity sold at the Universal Product Code (UPC) level. For each store in the dataset, we observe a store identifier, retailer identifier, retailer type, and the store's ZIP Code prefix (a ZIP Code prefix is the first three digits of a 5-digit ZIP Code). Stores contained within the six ZIP Code prefixes in and around Philadelphia (080, 081, 189, 190, 191, 194) are considered in our analysis. We apply further restrictions by only considering stores that maintained a presence throughout the period of the dataset, whose ZIP Code could be approximated via the household-level data (as detailed later), and whose approximated ZIP Code fell within 20 minutes of the nearest ZIP Code in Philadelphia.[9]

[8] The tax is levied on distributors, and so the price increase observed by consumers is subject to a pass-through rate.

[9] NielsenIQ data provides ZIP Code information according to the United States Postal Service (USPS) designation. We match these USPS ZIP Codes to their corresponding ZIP Code Tabulation Areas (ZCTAs) as defined in 2016 according to the US Census Bureau. UDSMapper.org, funded by the American Academy of Family Physicians, provides the most up-to-date conversion of USPS ZIP Codes to their corresponding ZCTAs. ZCTA centroids and distances are provided by the NBER ZIP Code Distance Database.

Seiler et al. (2021) suggest that cross-border shopping in response to the Philadelphia SB tax occurs in the region immediately surrounding the city. They find that post SB taxation, there is a positive, statistically significant increase in SB sales in stores located 0-6 miles from Philadelphia's border, but not in stores more than 6 miles from the border. Given that the primary purpose of our work is to evaluate the effect of SB taxation on cross-border shopping and avoidance behavior, we define our market similarly. In practice, we define our market to be the collection of the ZIP Codes in Philadelphia and those within 20 (driving) minutes of the city ("city + 20 minutes"). Appendix A1 shows that sales in stores beyond the 20-minute band surrounding the city do not experience an increase in SB sales following the implementation of the SB tax. Our final retail dataset contains 196 stores: 78 stores in Philadelphia and 118 in the surrounding region.[10]

In our retail data, we observe 6,968 UPCs pertaining to eight beverage categories: Carbonated Soft Drinks, Juice, Sports Drinks, Energy Drinks, Coffee, Tea, Flavored Water, and Pure Water. For each UPC, we have information on brand, pack size, container ounces, and flavor (many UPCs relate to variations in pack size and container ounces). We rely on the USDA FoodData Central database along with several food nutrition API services[11] to collect information pertaining to ingredients, sugar content, and caloric value (grams of sugar and caloric content are reported as per ounce values). The use of both sugar and caloric content in our model is particularly important for understanding demand for coffee, tea, and some juice products where caloric content and sugar may not trend similarly due to the addition of non-sugar calorically dense ingredients (such as milk and cream). Among the UPCs we observe, we remove infrequently purchased items and consider only those available both pre- and post-tax; the remaining 2,630 UPCs account for 93.9% of total volume sales.

We then aggregate the UPCs into products, where each product is a brand/SB status/category/diet status/size combination.[12] SB status is an indicator denoting the presence of added sugar or artificial sweeteners—these products are subject to the SB tax if they are sold in Philadelphia. Diet status indicates those products marketed as "diet," "light," etc. To allow heterogeneous responsiveness to the tax by product size, we include size as a product characteristic, based on equivalent units of 12 ounces (for example, a typical 12-pack of soda contains 144 ounces and has a size of $144/12 = 12$). In total, there are 464 products, of which 329 are SBs and the other 135 are non-SBs. Prices are adjusted for inflation.[13]

[10] In our retail dataset we observe 36 grocery stores, 142 drug stores, and 18 discount stores, which comprise 60%, 25%, and 15% of our observed unit sales, respectively.

[11] world.openfoodfacts.org, chompthis.com, edamam.com, foodrepo.org and nutritionix.com.

[12] Flavor variations for the same product are aggregated together. Such variations typically have uniform price and similar sugar content and caloric values.

[13] We adjust for inflation by expressing prices as their December 2018 dollar values using the Consumer Price Index

We use the term *location* to denote Philadelphia or non-Philadelphia (the 20-minute band surrounding Philadelphia). Due to the computational constraints facing structural demand models, we aggregate our data from the store-week level to the location-month level; the aggregation over time also helps reduce the potential bias in demand estimation stemming from households' stockpiling behavior (see for example Miller and Weinberg (2017)). In our demand model, to be specified in the next section, we define an alternative in households' monthly choice set to be a product-location combination. Correspondingly, total unit sales and quantity-weighted sugar content and caloric value are considered at the product-location-month level. To avoid complications resulting from within-location store substitution, we compute product-location-month level prices as a weighted average of product-level prices, where the weights are equal to the store-level market shares pre-tax. If every product is available in every location in every month, there would be $464 \times 2 \times 48 = 44,544$ observations at the product-location-month level. In reality, not all products are available in both locations every month, and as a result our retail dataset has a smaller number of observations, at 42,782.

Finally, we obtain product-level market shares by dividing total monthly product sales by the total monthly number of store trips; the latter is found by multiplying the market household population by the monthly average number of household store trips as provided in the household-level data. Because we do not observe all the stores, we re-scale the market household population such that the observed market shares are consistent with the purchase probabilities observed in the household data, across the pre-tax time periods. Table 1 provides retail data descriptive statistics, broken down by beverage category and SB status. We note that we do not account for beverage sales at non-retailer vendors such as restaurants, fast-food outlets, and theaters, as such vendors are not covered in our data.

3.2 Household Data

NielsenIQ provides household purchase data for a sample of US households. Beverage purchases, information pertaining to the number of shopping trips, a household's ZIP Code of residence, and other household demographic data are recorded. The purchase data reports the price paid, number of units purchased, and product UPC. When available, store identifier, retailer identifier, retailer type and store location are provided. As with the retail data, store location is provided as a 3-digit ZIP Code prefix, however a household's ZIP Code of residence is provided with the full 5-digit code.

Between 2015 and 2018, there were 867 households recorded in the NielsenIQ data who lived within the 156 ZIP Codes pertaining to our market. Over the course of these 4 years, these house-

for All Urban Consumers (CPI-U).

Table 1: Retail Data Descriptive Statistics[a]

Category	Number of Products	Market Share in Beverages	Price (cents/oz)	Sugar[b] (g/oz)	Calories (Cal/oz)	Ounces (oz)
Carbonated Soft Drinks	**167**	**37.20%**				
SB	142	33.17%	3.76¢	1.93	7.29	92.55
Non-SB	25	4.03%	4.11¢	0.31	1.22	100.12
Coffee	**20**	**1.63%**				
SB	17	1.51%	15.16¢	2.29	14.75	29.58
Non-SB	3	0.12%	12.58¢	0	1.26	35.67
Energy Drinks	**36**	**4.44%**				
SB	36	4.44%	15.07¢	1.67	7.17	36.10
Non-SB	0	–	–	–	–	–
Flavored Water	**10**	**2.82%**				
SB	9	2.73%	4.93¢	0.83	3.25	65.13
Non-SB	1	0.09%	8.09¢	0	0	16
Juice	**123**	**20.49%**				
SB	58	10.72%	4.78¢	2.18	9.77	69.40
Non-SB	65	9.77%	7.86¢	2.85	13.55	55.64
Pure Water	**33**	**10.41%**				
SB	0	–	–	–	–	–
Non-SB	33	10.41%	2.30¢	0	0	236.78
Sports Drinks	**18**	**9.21%**				
SB	18	9.21%	4.10¢	1.08	4.22	92.45
Non-SB	0	–	–	–	–	–
Tea	**57**	**13.81%**				
SB	49	13.25%	3.80¢	1.67	6.99	87.14
Non-SB	8	0.56%	4.33¢	0	0	74.06

[a]Price, Sugar and Calories are presented as quantity-weighted averages.

[b]Sugar present in non-SBs is the result of natural processes and is not considered added.

holds recorded 274,686 purchase opportunities (store trips) with 120,361 beverage purchases. We differentiate between low- and high-income households and between URM and non-URM households. We create an indicator variable "low-income" for the 261 households whose annual income falls below twice the federal poverty limit for their household size; otherwise a household is labeled "high-income." Likewise, a URM indicator is created for the households that identify as African American, Hispanic American, or Native American. We focus on income and URM status as demographic variables of interest since (1) opponents of the taxation policy argued that low-income disadvantaged minorities would be most negatively affected by the SB tax, and (2) prior works suggest that low income is correlated with both a higher price sensitivity and a greater preference for sugary beverages.

3.3 ZIP Codes

We define our market as the 156 ZIP Codes either within Philadelphia or whose centroid is outside Philadelphia but within 20 minutes of the nearest Philadelphia ZIP Code centroid; 48 ZIP Codes exist within Philadelphia, while the other 108 are in the surrounding 20-minute band. ZIP Code-specific demographic data pertaining to the number of households, income, and URM status is collected from the 2018 5-Year American Community Survey (ACS).[14] Table 2 provides the household distribution of income and URM status by location. The table shows that the two locations have roughly the same number of households, with Philadelphia having significantly more low-income households and households which identify as URM.

Table 2: Household Distribution of Income and URM Status by Location

	Demographics	URM	Non-URM	**Total**
Philadelphia	Low-Income	168,027	93,284	**261,311**
	High-Income	137,395	196,187	**333,582**
	Total	**305,422**	**289,471**	**594,893**
Non-Philadelphia	Low-Income	53,873	93,524	**147,397**
	High-Income	76,195	428,564	**504,759**
	Total	**130,068**	**522,088**	**652,156**

Rather than using straight-line distance to account for location substitution in our model, we rely on travel time as provided by the Google Maps API service. For each Philadelphia ZIP Code, we find the minimum travel time to drive to a non-Philadelphia ZIP Code, and vice versa.[15] We rely on travel time rather than distance to account for location substitution for two reasons: (1) ZIP Code distances do not account for road and highway placements which can greatly alter consumers' willingness to cross-border shop, and (2) Philadelphia is home to many rivers and bridges which would remain unaccounted for if distance was the metric considered. Furthermore, driving is by far the most popular mode of transportation in and around Philadelphia (see for example Duchneskie (2016)), giving support to calculating travel time based on driving as an approximation.

[14]The Zip Code-level marginal distribution of URM status is obtained from the ACS, and the joint distribution between income and URM status is obtained from the ACS Public Use Micro Sample (PUMS). PUMS areas may overlap multiple Zip Codes; in those cases, the Zip Code-level joint distribution is assumed to be the weighted average of overlapping PUMAs using the PUMA-Zip Code crosswalk file from the Missouri Census Data Center.

[15]Travel time between two ZIP Codes is defined as the average time required to drive from one ZIP Code centroid to the other. Using ZIP Code centroids for the calculation is analogous to how ZCTA distances are calculated.

Figure 1 presents some model-free, suggestive evidence of the importance of travel time. Panel (a) shows, for each Philadelphia ZIP Code, the minimum travel time to a non-Philadelphia ZIP Code. Panel (b) shows, for each Philadelphia ZIP Code, the percentage of beverage purchases made by the ZIP Code's households that are recorded in a store within their home location (Philadelphia). A comparison of the two panels suggests these two variables are positively correlated (a longer travel time to the alternative location is associated with a higher percentage of beverage purchases in the home location), and calculation shows these two variables have a correlation coefficient of 0.51.

Figure 1: Travel Time and Beverage Purchases

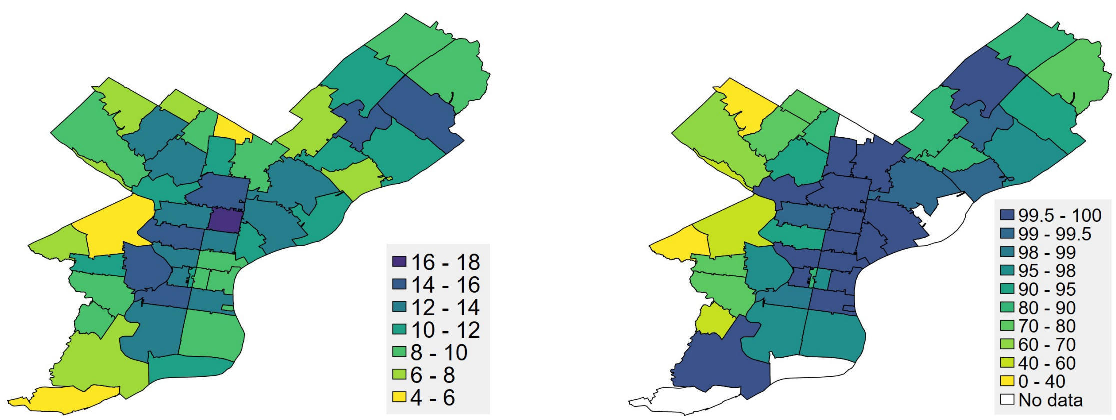

(a) Travel time (minutes) to alternative location for Phil. ZIP Codes; affected by not only distance but also roads, highways, rivers, bridges, etc.

(b) Percentage of beverage purchases in home location by Phil. households, as observed in the household dataset.

3.4 Store Location

As detailed above, the retail dataset does not provide stores' exact locations or full 5-digit ZIP Codes. Instead, we are provided with the stores' 3-digit ZIP Code prefixes (corresponding to the first three digits of the ZIP Codes). There are six ZIP Code prefixes in and around Philadelphia. Among them, two are entirely within our market: 191 is the ZIP Code prefix for Philadelphia, and 081 corresponds to a region of New Jersey that is entirely within the 20-minute band surrounding Philadelphia. Stores located within the ZIP Code prefixes of 080, 189, 190 and 194 have their locations approximated to determine whether they fall within any of the ZIP Codes pertaining to our market, as follows.

To approximate store locations, we rely on a method similar to that proposed in DellaVigna and Gentzkow (2019) and Goldin et al. (2022). For each store, we observe in the household data the ZIP Codes of residence for the households who make purchasing trips to the store. We then take the store's location to be the average of the centroids of these ZIP Codes, weighted by the total number of trips to the store originating from each of these ZIP Code during the pre-taxation period.[16] In the data, only retailers of the types "Grocery," "Discount Store," and "Drug Store" have unique identifying information that allows for this location approximation. Thus, our final retail and household dataset only considers stores of these types to remain consistent.

4 Model

In modeling the demand for beverages as a function of product and household characteristics incorporating consumer heterogeneity and demographic information, we follow the literature on discrete choice demand estimation with retail data (Berry et al. (1995) (BLP), Nevo (2000), etc.), and supplement the traditional method with household data in a process similar to that described in Goolsbee and Petrin (2004), Murry and Zhou (2020), and Grieco et al. (2022).[17] This allows us to leverage the benefits of both datasets: the retail data measures responses to the SB tax with far less noise and allows for a reliable method by which one can account for price endogeneity, while the household data provides a more accurate estimation of heterogeneous parameters, substitution patterns, and responsiveness to travel time. The model we propose utilizes the retail and household data in an internally consistent way.

4.1 Demand Specification

Consider household i in month t. The household chooses one of the available beverage options ($j = 1, \ldots, J_t$) or the outside option of no purchase ($j = 0$), where a beverage option is defined as a product-location combination.[18] Household i's indirect utility from choosing beverage option j in month t is given by

$$\begin{aligned} u_{ijt} = & x'_{jt}\beta_i + \alpha_i p_{jt} + h'_{jt}\gamma + \mathbb{1}(A_j \neq A_{z_i})(\phi_i Q_{z_i}) + \xi_{jt} + \bar{\epsilon}_{ijt}, \\ & \text{where } i = 1, \ldots, H_t,\ j = 1, \ldots, J_t,\ t = 1, \ldots, T, \text{ and } z_i = 1, \ldots, Z. \end{aligned} \tag{1}$$

[16]Centroid locations are given as latitude and longitude. We first convert the centroids to polar coordinates, calculate the weighted average, then convert back to latitude and longitude. There is a slight error introduced, as this conversion assumes a perfectly spherical earth; however given the relative closeness of locations this error is minimal.

[17]Another method is the micro-BLP estimator (Berry et al., 2004). Grieco et al. (2022) suggest that the use of micro moment conditions, as described in Berry et al. (2004), induces an additional cost in efficiency relative to a share constrained micro likelihood estimator, the type of estimator applied in this paper.

[18]Product availability varies month to month. Similar to Miravete et al. (2018), if no sales are observed for a beverage option during a specific month, then we assume that option is not present in households' choice set for that month.

x_{jt} is an $n_1 \times 1$ vector of option j's characteristics in month t, including a constant, Philadelphia dummy variable, category dummy variables, sugar content, caloric value, etc. (the full specification is given later in Section 6). p_{jt} denotes the retail price for option j in month t. The $n_2 \times 1$ vector h_{jt} contains categorical time dummies. z_i denotes household i's ZIP Code of residence (out of the 156 ZIP codes pertaining to our market). A_j and A_{z_i} are indicator variables signifying if option j and ZIP Code z_i are in the Philadelphia location, respectively. Q_{z_i} is the minimum travel time for a household living in ZIP Code z_i to drive to the alternative location (Philadelphia or non-Philadelphia). ξ_{jt} denotes unobserved quality, and $\bar{\epsilon}_{ijt}$ denotes unobserved idiosyncratic preferences. The indirect utility from choosing the outside option excluding $\bar{\epsilon}_{i0t}$ is normalized to 0.

We characterize household i by a $d \times 1$ vector of demographic attributes D_i, including low-income (below twice the federal poverty limit), URM (African American, Hispanic American, and Native American) and location (non-Philadelphia). We model unobserved household preference heterogeneity through the use of the multivariate normal distribution. Households' preferences for price, beverage option characteristics, and travel time are as follows:

$$\begin{pmatrix} \alpha_i \\ \beta_i \\ \phi_i \end{pmatrix} = \begin{pmatrix} \alpha \\ \beta \\ \phi \end{pmatrix} + \Pi D_i + \Sigma v_i, \quad v_i \sim \mathcal{N}(\mathbf{0}, \mathbf{I}_{n_1+2}), \tag{2}$$

where Π is an $(n_1 + 2) \times d$ matrix that measures the impact of observable demographic attributes on preferences, and Σ is an $(n_1 + 2) \times (n_1 + 2)$ matrix that captures the covariance of unobserved household preferences. In our study we estimate only the variance of unobserved household preferences and restrict $\Sigma_{hk} = 0 \ \forall h \neq k$.

Given the specification in Eq. (2), the indirect utility in Eq. (1) excluding $\bar{\epsilon}_{ijt}$ can be decomposed into its common and idiosyncratic components, δ_{jt} and μ_{ijt}, respectively, where

$$\begin{aligned} &\delta_{jt} = x'_{jt}\beta + \alpha p_{jt} + h'_{jt}\gamma + \xi_{jt}, \text{ and} \\ &\mu_{ijt} = \left[p_{jt}, x'_{jt}, \mathbb{1}(A_j \neq A_{z_i})Q_{z_i}\right](\Pi D_i + \Sigma v_i) + \mathbb{1}(A_j \neq A_{z_i})(\phi Q_{z_i}). \end{aligned} \tag{3}$$

We assume that unobserved idiosyncratic preferences for beverage options, $\bar{\epsilon}_{ijt}$, are correlated within the same beverage category. In our data we observe eight beverage categories (coffee, carbonated soft drinks, energy drinks, flavored water, juice, pure water, sports drinks, and tea), and the outside option of no purchase is defined to be category zero. $\bar{\epsilon}_{ijt}$ follows the distributional assumption of a nested logit model and can be decomposed into

$$\bar{\epsilon}_{ijt} = \zeta_{igt} + (1 - \rho)\epsilon_{ijt}, \tag{4}$$

where ϵ_{ijt} is i.i.d. type I extreme value, $\rho \in [0, 1]$ is the nesting parameter, $g \in \{0, 1, \ldots, 8\}$ is the category that option j belongs to, and ζ_{igt} has a (unique) distribution such that $\bar{\epsilon}_{ijt}$ is distributed extreme value. The nesting parameter ρ measures the correlation in preferences across beverages within the same category. Perfect within-nest substitution is obtained if ρ equals one, while as ρ goes to zero, the model reduces to the standard random coefficients logit specification. The probability of household i choosing option j belonging to category g in month t is then

$$\pi_{ijt} = \frac{\exp\Big(\big(\delta_{jt} + \mu_{ijt}\big)/(1-\rho)\Big)}{\exp\Big(I_{igt}/(1-\rho)\Big)} \times \frac{\exp\big(I_{igt}\big)}{\exp\big(I_{it}\big)}, \tag{5}$$

where the "inclusive values" I_{igt} and I_{it} are given by

$$I_{igt} = (1-\rho)\log \sum_{j \in \mathcal{J}_{gt}} \exp\Big(\frac{\delta_{jt} + \mu_{ijt}}{1-\rho}\Big) \tag{6}$$

with $\mathcal{J}_{gt}$ denoting the set of beverage options in category g in month t, and

$$I_{it} = \log\Big(1 + \sum_{g=1}^{8} \exp\Big(I_{igt}\Big)\Big). \tag{7}$$

4.2 Household Choice Probabilities

In the household dataset, for each household i and each month $t \in \mathcal{T}_i$ during which household i is in the data, we observe the household's O_{it} purchase opportunities (store trips). During each opportunity, the household chooses one of the available beverage options or the outside option of no purchase.[19] Integrating over the distribution of unobserved household attributes, denoted $F_v(v_i)$, the density of household i's observed sequence of choices is given by

$$L_i(Y_i|x,p,h,Q_{z_i},D_i;\delta,\Theta) = \int \prod_{t\in\mathcal{T}_i}\prod_{o=1}^{O_{it}}\prod_{j=0}^{J_t}[\pi_{ijt}(x_t,p_t,h_t,Q_{z_i},D_i,\delta_t,\Theta,v_i)]^{y_{ijot}}\, dF_v(v_i), \tag{8}$$
$$\text{where } \delta_t = (\delta_{1t},\ldots,\delta_{J_t t})',\ x_t = (x'_{1t},\ldots,x'_{J_t t})',\ p_t = (p_{1t},\ldots,p_{J_t t})', \text{ and } h_t = (h'_{1t},\ldots,h'_{J_t t})'.$$

We summarize the model's heterogeneous taste, travel time, and nesting parameters as $\Theta = (\Pi,\Sigma,\phi,\rho)$, and use Y_i to denote the observed sequence of household i's choices, where $y_{ijot} = 1$ if household i chooses beverage option j during purchase opportunity o in month t.

4.3 Retail Market Shares

At the retail level, we use M_t to denote the market size in month t, i.e., the total number of purchase opportunities experienced that month, obtained as the total number of households in

[19] A household's number of store trips per month does not show statistically significant correlation with our demographic variables of interest, and we assume that the number of purchase opportunities is independent of unobservable individual characteristics. Such an assumption is necessary for our estimation to be tractable under the BLP framework, and is one innately imposed by researchers working solely with retail data (i.e., Berry et al. (1995), Nevo (2000), etc.).

the market multiplied by the average number of grocery store trips per household in that month as observed in the household data. We assume a continuum of purchase opportunities of mass M_t, and the household data is assumed to be a finite sample drawn from it.[20]

Consider the set of household-specific characteristics that lead to the purchase of beverage option j in month t, $\{(D_i, z_i, v_i, \bar{\epsilon}_{ijt})|u_{ijt} > u_{ikt}\ \forall k = 0, 1, \ldots, J_t\}$. The distribution of $\bar{\epsilon}_{ijt}$ is extreme value as given in Eq. (4), which leads to household choice probabilities π_{ijt} given in Eq. (5). The distribution of v_i is multivariate normal as given in Eq. (2), and the distributions of z_i and $D_i|z_i$ are obtained from the ACS. Integrating over the distributions of v_i, z_i, and $D_i|z_i$, we obtain the predicted market share for beverage option j in month t as

$$s_{jt} = \int_{v_i} \int_{z_i} \int_{D_i} \pi_{ijt}(x_t, p_t, h_t, Q_{z_i}, D_i, \delta_t, \Theta, v_i) dF_D(D_i|z_i) dF_z(z_i) dF_v(v_i). \tag{9}$$

In assuming a continuum of households, as is routine in the literature, and conditioning on ξ, through δ, the market share in Eq. (9) is deterministic, and the aggregate demand for beverage option j is obtained as $M_t s_{jt}$.

5 Identification and Estimation

Our objective is to estimate the parameters α, β, γ, Π, Σ, ϕ, and ρ. While we are not necessarily interested in the value of δ per se, it is required to recover the mean taste parameters α, β, and γ. Thus, our estimation proceeds with two steps. First, we maximize a likelihood function using the retail and household data. This identifies all the parameters except those derived from the mean utility. Next, to estimate α, β, and γ, we use a two-stage least squares (TSLS) regression and instrument p_{jt} with a Hausman style instrument (as seen in Nevo (2001)) to control for correlation with the error term ξ_{jt}.[21]

5.1 Maximum Likelihood

In the first stage of our estimation, for any candidate values of Θ and δ, the density of a household's choice history is given by Eq. (8), and the corresponding log-likelihood of the household data is

$$\mathcal{L}(Y; \delta, \Theta) = \sum_{i=1}^{H} \log[L_i(Y_i|x, p, h, Q_{z_i}, D_i; \delta, \Theta)]. \tag{10}$$

[20] Appendix A2 provides details about the case of multiple purchases during a single trip.

[21] We calculate the average price of each product across all US stores in the NielsenIQ data, excluding those in the Philadelphia designated market area (DMA) which contains the market of our demand model, and use this average to instrument the price in our model.

In theory it is possible to estimate δ directly via maximum likelihood solely with the household-level data; practically, however, this is computationally infeasible considering the large number of beverage options available. Instead, we rely upon the work of Berry (1994) who shows that for any given value of Θ, there exists a unique vector of δ such that the predicted market shares from Eq. (9), s_{jt}, exactly match those observed in the retail data, S_{jt}. Consequently, given the retail market shares, we can treat δ as a known function of Θ.[22] Appendix A3 shows in more detail how a unique vector of δ is obtained from our retail data.

Thereby, the log-likelihood of the household-level data shown in Eq. (10) can be re-written as

$$\mathcal{L}(Y;\Theta) = \sum_{i=1}^{H} \log[L_i(Y_i|x, p, h, Q_{z_i}, D_i, \delta(\Theta);\Theta)], \tag{11}$$

where $\delta(\Theta)$ is given by the one-to-one contraction mapping from the retail market share constraint. In performing the contraction mapping, we evaluate the integrals of Eq. (9) by Monte Carlo simulation with 4000 Halton draws, per month, from the distributions of v, z, and $D|z$ (i.e., 4000 simulated households). Similarly, we use a separate set of 200 Halton draws from the distribution of v when evaluating the integral in Eq. (8).[23] Our estimation proceeds by searching for the value of Θ that maximizes Eq. (11).[24] Finally, we obtain robust standard errors for Θ by sandwiching the covariance of the household-level gradient between the inverted Hessian at the optimum of the likelihood function.[25]

5.2 Mean Utility Coefficients

Given $\hat{\delta}$ resulting from the optimal $\hat{\Theta}$ in the maximum likelihood step, we use the relation expressed Eq. (3) to determine our mean utility parameters. In practice, we rewrite each product's

[22] By assuming the aggregate market shares are derived from a continuum of households, the asymptotic variance of the shares is zero. Grieco et al. (2022) shows that this assumption has a cost in terms of both efficiency and inference, unless the household sample size is negligibly small when compared to the size of the market population. This is similar to the efficiency loss of the standard micro-BLP (Berry et al., 2004). In our model $H/N = 0.0007$, where $H = 867$ is the size of the household dataset and $N = 1,247,049$ is the population of households in and around Philadelphia from which those 867 households were drawn; accordingly, the efficiency loss should be minimal. Furthermore, to use a mixed data likelihood estimator as suggested in Grieco et al. (2022) would be too computationally burdensome, as each δ_{jt} must be treated as a parameter of interest in the likelihood estimation.

[23] Results from Train (1999) show simulation variance with 100 Halton draws to be lower than 1000 random draws in a mixed logit application.

[24] Our tolerance for the contraction mapping step is set to $.5e^{-12}$. For the likelihood maximization algorithm, we set a tolerance of $2e^{-10}$ and provide computed numerical gradients. We consider several randomized starting values when proceeding with the maximization algorithm to rule out local minima.

[25] See Train (2009), p. 201.

mean utility from Eq. (3) as:

$$\begin{aligned}\delta_{jt} &= x_{jt}'\beta + \alpha p_{jt} + h_{jt}'\gamma + \xi_{jt} \\ &= u_j + \widetilde{x}_{jt}'\widetilde{\beta} + \alpha p_{jt} + h_{jt}'\gamma + \Delta\xi_{jt}.\end{aligned} \tag{12}$$

ere, we introduce a time- and location-invariant fixed effect, u_j, to account for systematic variations in preferences across products.[26] With this decomposition, $\widetilde{x}_{jt}$ now includes the subset of time- and/or location-variant characteristics originally in x_{jt}. The covariates represented in $\widetilde{x}_{jt}$, p_{jt}, and h_{jt} therefore control for the effects of time- and/or location-variant preferences, while deviations from these are captured by $\Delta\xi_{jt}$. To estimate mean utility associated with time- and location-invariant characteristics, such as product categories and size, we project the estimated product fixed effects onto the remaining observable product characteristics.

To assess the relationship expressed in Eq. (12), we conduct a Two-Stage Least Squares (TSLS) regression using a Hausman-style instrument for price, given the likelihood that p_{jt} may be correlated with the error term ξ_{jt}. Specifically, we use the average price of the product across all regions outside of Philadelphia and its surrounding areas as an instrument for the price variable in our model. This approach significantly strengthens the estimated price response, more than doubling its magnitude, which is consistent with the expectation when accounting for endogenous pricing behavior.

Standard errors for $(\hat{\alpha}, \hat{\beta}, \hat{\gamma})$ are calculated using a two-stage bootstrap procedure, where the first stage captures the estimation error from the maximum likelihood step and the second stage captures the typical sampling error. Specifically, we begin by first taking 1000 draws from the asymptotic distribution of Θ. Next, for each draw Θ_d, we find its corresponding vector $\delta(\Theta_d)$. We then draw with replacement from the sample $\{(\delta_{11}(\Theta_d), x_{11}, p_{11}, h_{11}), \ldots, (\delta_{J_T T}(\Theta_d), x_{J_T T}, p_{J_T T}, h_{J_T T})\}$ to create a bootstrapped dataset (of a size equal to the original sample). Given this bootstrapped sample, we then perform the TSLS regression to estimate $(\alpha_d^*, \beta_d^*, \gamma_d^*)$. From the distribution of $(\alpha_d^*, \beta_d^*, \gamma_d^*)$, we find the standard errors of our mean utility parameters.

6 Demand Estimates

Table 3 presents the demand estimates of our preferred specification of the RCNL model using the two-step procedure outlined above.[27] To avoid perfect collinearity, we have dropped the

[26] An alternative approach would be to allow the product fixed effect to vary over time, effectively turning our estimation into a difference-in-differences model. However, this approach is not feasible because our product set includes items that are not available in both locations.

[27] We considered an alternative nested logit model with the choice between beverages and the outside option at the highest level and the choices of beverage category and then beverage option at subsequent levels; however such a model did not improve model fit.

category pure water. On average, consumer valuation for beverages increases with calories but decreases with sugar content, and SBs display a lower valuation relative to non-SBs. Additionally, the negative interaction term for "Philadelphia" may indicate that consumers perceive Philadelphia stores as less desirable shopping locations, potentially due to factors like smaller store sizes, limited non-beverage offerings, inadequate parking, and higher levels of crowdedness, all of which can reduce consumer utility. Supporting this, the City of Philadelphia's Department of Public Health has released a research report highlighting the scarcity of healthy foods and produce in many Philadelphia retail food stores, which encourages residents to shop elsewhere and rely on automobiles for grocery shopping (Hillengas et al., 2019).

Taxation Bias Following the taxation, a change in consumer valuation for Philadelphia SBs arising from tax-induced price changes should be absorbed by the Price response, but public awareness of the taxation policy, which results in the same SB being more expensive in Philadelphia than outside, has the potential to impact consumer valuation for products beyond the Price response, for example due to psychological reasons such as perceived inequality/unfairness (Weber et al. (2014), Xia et al. (2004), Campbell (1999)). We find that, after the taxation, valuation for Philadelphia SBs declines as demonstrated by the negative coefficient on the indicator variable Tax Bias, which equals 1 for Philadelphia SBs in post-tax months. However, this decline in valuation only holds for smaller-size products; we find that Philadelphia SBs with a size greater than about 10 standard beverage cans (120 ounces) see an increase in average valuation, as shown by the Tax Bias $\times$ Size interaction.

This result is intuitive. Despite larger-size SBs seeing the greatest level of taxation and the greatest increase in product price, larger-size products are generally cheaper per ounce. Therefore, we find that a well-publicized local taxation policy shifts consumer preference towards products that offer the greatest cost-saving benefits.

Demographic Interactions We allow for variation in consumer valuations across observed demographic characteristics, including income status, URM status, and location. The results are presented in columns 4, 5, and 6 of Table 3. The estimation of Π reveals significant differences in consumer valuations for beverage options. Compared to non-URM households, URM households have higher valuations for all inside options (based on the estimates for URM interactions with the constant and category fixed effects). We also find low-income households have statistically significant reduction in valuation for flavored water, juice, sports drinks and tea. Travel time and its interaction with demographic characteristics are discussed later in this section.

Furthermore, through examination of estimation results from model specifications with and without the URM interaction terms, we find that much of the effect of the low-income status diminishes (many interaction terms become insignificant) with the addition of the URM inter-

Table 3: RCNL Demand Estimates[a]

	Mean	Standard	Demographic Interactions (Π)		
	(α, β, ϕ)	**Deviation (Σ)**	**Low-Income**	**URM**	**Non-Phil.**
Constant	-5.876*** (0.270)	1.669*** (0.097)	0.400 (0.279)	2.127*** (0.360)	
Price	-0.427*** (0.034)	0.080*** (0.009)	0.020 (0.014)	-0.052*** (0.017)	
Taxation Bias	-0.242*** (0.021)				
Taxation Bias $\times$ Size	0.026*** (0.005)				
Calories	0.025** (0.010)				
Sugar	-0.157*** (0.025)	0.212*** (0.016)	0.020 (0.017)	0.150*** (0.022)	
Diet	-0.116 (0.085)				
Size	0.152*** (0.014)	0.060*** (0.004)	-0.005 (0.006)	-0.060*** (0.009)	
Size^2	-0.003*** (<0.001)				
SB	-0.368*** (0.055)	0.613*** (0.044)	0.036 (0.074)	0.347*** (0.093)	
Carb. Soft Drinks	2.243*** (0.244)		-0.354 (0.295)	-1.415*** (0.326)	
Coffee	0.300 (0.361)		-0.710 (0.472)	-0.880** (0.436)	
Energy Drinks	0.702 (0.485)		0.721 (0.446)	-0.617 (0.593)	
Flavored Water	0.239 (0.349)		-0.710* (0.389)	-0.204 (0.442)	
Juice	1.981*** (0.254)		-0.895*** (0.290)	-0.956*** (0.283)	
Sports Drinks	1.745*** (0.275)		-0.618* (0.324)	-1.870*** (0.356)	
Tea	1.687*** (0.239)		-0.923*** (0.301)	-1.306*** (0.296)	
Philadelphia	-0.931*** (0.084)	0.949*** (0.054)			
Travel Time	-0.168*** (0.012)	0.011*** (0.005)	0.019* (0.010)	0.018* (0.010)	0.067*** (0.020)
Category Nesting (ρ)	0.614*** (0.024)				

***p<.01, **p<.05, *p<.1

[a]Standard errors are reported in parentheses. Estimates for the nesting parameter, travel time, random coefficients, and demographic interactions are obtained from the maximum likelihood estimation with 867 households and 273,943 household observations over 48 months. Parameter values for price, tax amount, calories, and sugar are obtained from the projection of δ onto these characteristics, product fixed effects, and time-category fixed effects. The remaining mean utility estimates are obtained from the projection of estimated product fixed effects onto the observable characteristics.

action terms. Given that low-income and URM are correlated, this finding suggests that works which fail to account for the URM status might incorrectly attribute the effect of the URM status to the low-income status, due to omitted variable bias. By including URM interactions in our work, we help to better capture the complexities and nuances of consumer choice, particularly in contexts where URM status may significantly influence choice and outcomes.

Travel Time In measuring responsiveness to travel time, we allow for a rich set of heterogeneous model parameters. Disutility from travel time is greater for high-income households, in line with prior transportation research (e.g., Hymel et al. (2010)) which finds that high-income households have a higher valuation of their time. In addition, we find URM households have lower disutility for travel time.

Taking the average of the ratio of travel time responsiveness to price responsiveness across all simulated households, we find that on average an extra minute of travel time to reach the store is equivalent to adding 28.5¢ (7%) to the product price (averaging \$3.88 per product), suggesting travel time plays a significant role in households' willingness to cross-border shop.[28]

Additionally, the interaction term between travel time and non-Philadelphia households shows that non-Philadelphia households have a lower disutility for travel time than Philadelphia households. While Philadelphia households live in an urban setting, non-Philadelphia households are largely suburban or rural, and the finding that their responsiveness to travel time differs noticeably is intuitive. Through the inclusion of this non-Philadelphia interaction term, we are accounting for how differences in urbanization, population density, job location, and store accessibility may impact a household's willingness to travel.

Random Coefficients and the Nesting Parameter We include in our model a rich set of random coefficient parameters (Σ), all of which exhibit statistical significance and sensible results. For instance, the standard deviation of the random coefficient on sugar content suggests that 49% of households who identify as URM experience an increase in utility from higher sugar content, whereas only 23% of high-income, non-URM households experience the same. Regarding the distribution of household disutility for travel time, we see that households hold a negative valuation of travel time, regardless of their demographic status. This result fits with what prior transportation research suggests. Finally, the nesting parameter ρ is estimated very precisely, and implies that households show a strong correlation in preferences across beverages within the same category. To corroborate this point, consider the price elasticity of demand.

Price Elasticities Table 4 provides the price elasticity of demand by household location, reporting own- and cross-elasticities averaged at the location level and category-location level.

[28] Our travel time variable measures the time needed to travel *to* a store in the alternative location, so for example purchasing at a store 10 minutes away would involve a 20-minute round trip.

Cross-elasticities of demand are reported for beverage options from the same category, same category and same location, same category and different location, and all beverage options.

For own-elasticity of demand averaged at the category-location level, estimates show that households have elastic demand for beverages in their home location, with the elasticity ranging from -1.40 to -2.16. When considering a household's non-home location, the own-elasticity of demand varies from -0.10 to -0.62. Differences in own-elasticities between home and alternative locations make intuitive sense: individuals purchasing beverages while traveling outside their home location are already engaging in a costly activity, and thus are likely less sensitive to increased product prices in the alternative location.

Further, we find Philadelphia households are about four times more sensitive than non-Philadelphia households to own-price changes in their non-home location. This is primarily a result of household-specific travel times. Nearly half of all Philadelphia households live within a 10-minute drive from the border, whereas only a quarter of non-Philadelphia households live within a 10-minute drive from the border. Differences in own-elasticities of demand for the non-home location thus reflect the differences in travel times and the overall willingness to cross-border shop.

Finally, considering cross-elasticity of demand, we see that it is higher between beverages in the same category, and furthermore it is generally higher between beverages in the same category and same location when compared to beverages in the same category but different locations. These results delineate a clear order of preference in terms of substitution.

7 Pass-Through, Substitution Patterns, and Consumption Changes

In the remainder of this paper, we study the effects of Philadelphia's SB tax by using our demand estimates to evaluate various counterfactual scenarios. Comparing the outcome under taxation to the counterfactual scenario of no tax, this section examines households' substitution patterns and consumption changes brought about by the tax, while the next section analyzes the welfare implications and regressivity of the tax. Then in Section 9 we consider the effects of alternative tax rates, alternative tax coverages, and changes in travel time.

A major consideration in these subsequent analyses is the responsiveness of beverage companies to our counterfactual taxation schemes and other counterfactual scenarios. Consequently, we specify a simple supply-side pricing model and use it to predict how prices would change at the product level. We assume a competitive Bertrand model where each month beverage companies strategically set prices across all products within their brand portfolio (e.g. the Coca-Cola company also controls Sprite, Fanta, Dasani, Powerade, Honest Tea, among others). This assumption allows coordination in pricing across products in various categories, all falling under

Table 4: Price Elasticity of Demand by Household Location

	Average Level	**Own-Elasticity**	**Cross-Elasticity**			
			Same Category			All Bev. Options
			All Bev. Options	Same Location	Different Location	
Philadelphia Households	**Phil. Bev. Options**	-2.1137	0.0123	0.0214	0.0033	0.0020
	Carbonated Soft Drinks	-2.1153	0.0046	0.0080	0.0013	0.0020
	Coffee	-2.0855	0.0379	0.0670	0.0089	0.0020
	Energy Drinks	-2.0716	0.0193	0.0331	0.0055	0.0020
	Flavored Water	-2.0684	0.0751	0.1239	0.0263	0.0021
	Juice	-2.1618	0.0060	0.0106	0.0015	0.0020
	Pure Water	-2.1269	0.0228	0.0402	0.0056	0.0020
	Sports Drinks	-1.9567	0.0351	0.0610	0.0093	0.0019
	Tea	-2.0922	0.0121	0.0212	0.0031	0.0020
	Non-Phil. Bev. Options	-0.4067	0.0025	0.0032	0.0018	0.0004
	Carbonated Soft Drinks	-0.3910	0.0008	0.0011	0.0006	0.0004
	Coffee	-0.4253	0.0077	0.0094	0.0059	0.0004
	Energy Drinks	-0.6197	0.0058	0.0075	0.0040	0.0006
	Flavored Water	-0.5072	0.0181	0.0250	0.0113	0.0005
	Juice	-0.3957	0.0011	0.0014	0.0008	0.0004
	Pure Water	-0.4569	0.0048	0.0060	0.0036	0.0004
	Sports Drinks	-0.2729	0.0048	0.0062	0.0034	0.0003
	Tea	-0.338	0.0019	0.0025	0.0014	0.0003
Non-Philadelphia Households	**Phil. Bev. Options**	-0.1324	0.0008	0.0005	0.0010	0.0001
	Carbonated Soft Drinks	-0.1287	0.0003	0.0002	0.0004	0.0001
	Coffee	-0.1365	0.0025	0.0019	0.0030	0.0001
	Energy Drinks	-0.1051	0.0010	0.0007	0.0012	0.0001
	Flavored Water	-0.1014	0.0036	0.0022	0.0050	0.0001
	Juice	-0.1395	0.0004	0.0003	0.0005	0.0001
	Pure Water	-0.1377	0.0015	0.0011	0.0019	0.0001
	Sports Drinks	-0.1382	0.0025	0.0017	0.0033	0.0001
	Tea	-0.1419	0.0008	0.0006	0.0011	0.0001
	Non-Phil. Bev. Options	-1.6751	0.0096	0.0185	0.0005	0.0016
	Carbonated Soft Drinks	-1.7026	0.0038	0.0072	0.0002	0.0016
	Coffee	-1.6146	0.0293	0.0566	0.0020	0.0016
	Energy Drinks	-1.4019	0.0129	0.0249	0.0009	0.0014
	Flavored Water	-1.4807	0.0541	0.1060	0.0021	0.0015
	Juice	-1.7247	0.0048	0.0093	0.0003	0.0016
	Pure Water	-1.6202	0.0173	0.0334	0.0012	0.0016
	Sports Drinks	-1.6651	0.0302	0.0591	0.0013	0.0016
	Tea	-1.7409	0.0101	0.0197	0.0005	0.0017

brands controlled by the same beverage company. Beverage companies then compete to maximize their respective profits, as consumers face various counterfactual scenarios influencing their demand for beverage products. Unlike prior works exploring SB taxation, our counterfactual estimates incorporate consumer heterogeneity in beverage preference, travel costs, and locational and categorical substitution. For further information, Appendix A4 details how firms' marginal costs and counterfactual prices are obtained.

As a first step in our counterfactual analyses, we explore product prices under a no-tax simulation. Specifically, we remove the per-product tax from each product's corresponding marginal cost, and remove the disutility associated with taxation from the consumer's utility function. We then obtain each firm's set of profit-maximizing counterfactual product/month-level prices.

7.1 Pass-Through Rate

After obtaining counterfactual prices under the no-tax scenario, we first explore the pass-through rate of the tax, based on the difference between observed and counterfactual prices at the per-ounce level.

Studying Philadelphia's SB tax, several authors have conducted analyses on this topic. Cawley et al. (2018) and Roberto et al. (2019) find pass-through rates of 55% and 68%, respectively, while Bleich et al. (2020) and Cawley et al. (2020) find higher pass-through rates of 120% and 105%, respectively. More recently, Seiler et al. (2021) find a pass-through rate of 97%, relying upon their finding that the region within the four 3-digit ZIP Code prefixes surrounding Philadelphia but more than 6 miles away from Philadelphia does not exhibit an increase in SB sales in response to Philadelphia's SB tax. They proceed by treating that region as their control: it is close enough to Philadelphia to experience similar marketing and demand shocks while uninfluenced by cross-border shopping.

However, unlike the reduced form methods used in the above studies, we generate our pass-through rate through simulating, based on the structural demand and marginal costs estimates, firms' counterfactual prices in the no-tax scenario. Thus, our findings offer a unique perspective in the analysis of the Philadelphia SB tax's pass-through rate, one that adds to the growing body of literature that analyzes SB taxation.

Category-specific pass-through rates are presented in Table 5. Recall the tax is equal to 1.5¢ per ounce. Thus, a Price Difference of 1.5 would denote a full pass-through, and dividing the Price Difference by 1.5 provides the pass-through rate in percentage terms. Results in Table 5 are based on average monthly price for all taxed products by category. Across all taxed products, the average pass-through rate is 100%, while category-level estimates range from 99% to 107%, falling within the range of pass-through rates seen in prior research (Cawley et al. (2018), Roberto et al. (2019), Bleich et al. (2020), Cawley et al. (2020), and Seiler et al. (2021)).

Table 5: Pass-Through Rate of SB Tax, by Category

	SB Category						
	Carb. Soft Dr.	Coffee	Energy Dr.	Flav. Water	Juice	Sports Dr.	Tea
Price Difference (in Cents)	1.483	1.607	1.52	1.485	1.501	1.502	1.505
Pass-Through Rate	99%	107%	101%	99%	100%	100%	100%

Prices without the presence of taxation are the result of a no-tax counterfactual analysis. Average differences between observed and counterfactual product prices represent the impact of SB taxation; from this, we can obtain category-specific average price increases and pass-through rates. The category pure water is excluded as that category contains no taxed products.

7.2 Substitution Patterns

Next, we examine how the SB tax induces categorical and locational substitution. To perform this analysis, we simulate 100 draws from the distribution of $\bar{\epsilon}_{ijt}$ for each combination of simulated household i, beverage option j, and post-taxation month $t = 25, \ldots, 48$ (January 2017 to December 2018).[29] We then determine product-level utility with and without the SB tax holding the $\bar{\epsilon}_{ijt}$ draws constant. That is, product-level utility takes the form

$$u_{ijt}^{\text{with tax}} = \delta_{jt}^{\text{with tax}} + \mu_{ijt}^{\text{with tax}} + \bar{\epsilon}_{ijt}, \text{ and} \tag{13}$$

$$u_{ijt}^{\text{without tax}} = \delta_{jt}^{\text{without tax}} + \mu_{ijt}^{\text{without tax}} + \bar{\epsilon}_{ijt}. \tag{14}$$

Thus, beverage choice with and without the tax is given by the maximal value of the utilities found in Eqs. (13) and (14), respectively.

In both equations, the coefficients are the estimated coefficients from our demand estimation, and the household and beverage option characteristics are the observed characteristics. In Eq. (13), the prices are the observed prices, while in Eq. (14), the prices are obtained according to the no-tax counterfactual simulation.

Holding $\bar{\epsilon}_{ijt}$ to be the same between the two equations when examining beverage choices with and without the tax allows us to isolate the effects of the tax on households' beverage choices. In comparison, tracking how the households in our household dataset actually change their choices from the pre-taxation period to the post-taxation period would not paint an accurate picture of the tax-induced substitution patterns, because households' idiosyncratic preferences $\bar{\epsilon}_{ijt}$, product availability, and demand shocks all have changed between the two periods. Likewise, relying on the retail data would not allow us to track how households switch from one category to another and/or from one location to the other as a result of the tax.

We report our findings in Table 6. The first column of the table provides the category market shares for Philadelphia SBs under the counterfactual scenario of no taxation, averaged across

[29]In accordance with the methodology presented in Wang and Ye (2024), we simulate draws from a nested logit distribution.

Table 6: Model Predicted Substitution Patterns

Choice w/o Tax		Top Four Choices with Tax							
Phil. SBs		**1st Choice**		**2nd Choice**		**3rd Choice**		**4th Choice**	
Carb.	(41.38%)	P Carb. S	(59.91%)	NP Carb. S	(16.41%)	Outside Op.	(13.09%)	P Carb. NS	(6.02%)
Coffee	(2.36%)	P Coffee S	(74.97%)	Outside Op.	(9.92%)	NP Coffee S	(9.72%)	P Coffee NS	(3.19%)
Energy	(7.09%)	P Energy S	(77.65%)	NP Energy S	(10.86%)	Outside Op.	(9.61%)	P Water NS	(1.15%)
Flav.	(3.14%)	P Flav. S	(70.91%)	NP Flav. S	(14.06%)	Outside Op.	(10.30%)	P Flav. NS	(1.93%)
Juice	(19.05%)	P Juice S	(55.00%)	P Juice NS	(18.73%)	Outside Op.	(10.33%)	NP Juice S	(10.16%)
Sports	(8.23%)	P Sports S	(66.36%)	NP Sports S	(14.90%)	Outside Op.	(14.69%)	P Water NS	(1.82%)
Tea	(18.75%)	P Tea S	(62.9%)	Outside Op.	(15.06%)	NP Tea S	(15.05%)	P Tea NS	(1.95%)

P, NP, S, and NS denote Philadelphia, non-Philadelphia, SB, and non-SB, respectively.

the 24 post-taxation months. The next four columns provide the first, second, third and fourth location $\times$ category $\times$ SB status choices with taxation, given the household would have chosen the leftmost item of that row without taxation. For example, without taxation, sweetened carbonated soft drinks would have made up 41.38% of the market share for Philadelphia SBs; with taxation, 59.91% of the households who would have chosen sweetened Philadelphia carbonated soft drinks continue to choose the same (no substitution), 16.41% choose sweetened carbonated soft drinks in the non-Philadelphia location (geographic substitution), 13.09% choose the outside option (consumption reduction), and 6.02% choose non-sweetened carbonated soft drinks in Philadelphia (product substitution). Information like this can be particularly useful to policymakers for understanding people's behavior patterns in response to the implementation of a public policy.

As expected, for all SB categories, the primary choice with taxation remains the same as that without. We observe that the categories of Philadelphia SBs that are the most responsive to taxation are juice, carbonated soft drinks, tea, and sports drinks, as measured by the proportion of households who switch away. Excluding tea, juice, and coffee, the primary choice of substitution is the same category of SBs in the alternative location, followed by the outside option. For juice, the primary choice of substitution is non-SB juice in the same location; this is an intuitive result, as many non-SB juice products are similar in terms of sugar content through the use of juice concentrates. For Philadelphia SBs, the proportion of households who transfer their consumption to the same category of SBs in the alternative location is either larger than (carbonated soft drinks, energy drinks, flavored water, and sports drinks) or very close to the proportion of those who switch to the outside option. This provides clear evidence towards consumers' willingness to cross-border shop in the presence of an SB tax.

The Philadelphia SB categories of coffee and energy drinks retain the greatest proportion of original consumers. We hypothesize that this pattern is due to the heterogeneous interaction of taxation policy with product size. Coffee and energy drink products are primarily sold in

Table 7: Simulated Market Shares by SB Status, Size, and Location

SB Status × Size × Bev. Location	Without Tax	With Tax	Difference	% Change
Philadelphia Bev. Options				
Non-SB × Small	0.37%	0.51%	+0.14	26.97%
Non-SB × Medium	1.01%	1.48%	+0.47	25.27%
Non-SB × Large	0.76%	0.91%	+0.15	12.46%
SB × Small	3.46%	3.01%	-0.45	-13.09%
SB × Medium	4.38%	2.60%	-1.78	-40.70%
SB × Large	1.94%	0.58%	-1.36	-70.25%
Non-Philadelphia Bev. Options				
Non-SB × Small	0.61%	0.63%	+0.02	3.29%
Non-SB × Medium	2.77%	2.86%	+0.09	3.20%
Non-SB × Large	1.65%	1.68%	+0.03	1.33%
SB × Small	3.90%	4.23%	+0.34	8.67%
SB × Medium	7.68%	8.28%	+0.59	7.73%
SB × Large	3.91%	4.18%	+0.27	7.00%
Outside Option	67.56%	69.07%	+1.51	2.24%

small, single-serving containers with relatively high per-ounce prices. Thus, the price increase due to the tax is proportionally smaller than those observed in other categories, where products on average come in larger sizes with lower per-ounce prices.

Supportive evidence is provided in Table 7, which displays simulated market shares for SBs and non-SBs by size categories (small, medium, and large) and location with and without taxation. We define small products as those less than 24 ounces, medium products as those greater than or equal to 24 ounces but less than 120 ounces, and large products as those greater than or equal to 120 ounces. Unlike Table 6, for Table 7 we do not need to keep track of how each simulated household switches from one choice to another in response to the tax, and so the market shares reported in Table 7 are found by averaging the choice probabilities of the original Halton draws across the 24 post-taxation months without directly simulating product choices. The "without tax" counterfactual is conducted with the effect of taxation removed from the individual-level utility.

From Table 7 we observe that the effect of the SB tax is heterogeneously distributed among differently sized SBs. The tax decreases the market shares of small, medium, and large Philadelphia SBs, with large Philadelphia SBs seeing the biggest drop followed by medium then small. SBs in the non-Philadelphia location experience an increase in market share regardless of size; so do non-SBs in Philadelphia.

These are intuitive results. Consider Philadelphia SBs, which are subject to the SB tax. Compared to small products, large products are typically sold at a "quantity discount" and have a lower per-ounce price. Consequently, the SB tax—levied at 1.5¢ per ounce—results in proportionally larger price increases for large products, thereby having a more negative impact on their market shares. Some of the market share that leaves large Philadelphia SBs goes to small Philadelphia SBs due to their proportionally smaller price increases and relatively high substitutability, leading to a significantly smaller decrease in the market share of small Philadelphia SBs.

7.3 Effects of SB Tax on Beverage Consumption

We now consider the effects of Philadelphia's SB tax on households' beverage consumption, as well as their cross-border shopping and tax avoidance behavior. For each simulated household in each post-taxation month, we compute the household's expected consumption (in ounces) of Philadelphia SBs, Philadelphia non-SBs, non-Philadelphia SBs, and non-Philadelphia non-SBs, respectively, based on the model predicted choice probabilities and adjusting the amounts to account for the expected numbers of products and units purchased per trip and the expected number of trips in that month. We then sum over the 24 post-taxation months and compute the average per household over all households, Philadelphia households, and non-Philadelphia households, respectively.[30] We do this twice, without tax and with tax, then calculate the differences. The results are reported in Table 8.

One may notice that our counterfactual "Without Tax" simulation suggests that compared to non-Philadelphia households, (1) Philadelphia households consume a smaller amount of beverages overall, and (2) Philadelphia households are significantly more likely to cross-border shop. First, we hypothesize that differences in consumption result from Philadelphia households having greater access to products at non-retailer vendors such as restaurants, fast-food outlets, bodegas, etc. Considering fast food, for instance, there should be no surprise that inner-city consumers exhibit differences in consumption behavior when compared to those living in more rural regions. One primary reason for such differences arises from consumers' proximity to fast-food establishments. Rahkovsky et al. (2018) find that urban regions have 10 times more fast-food outlets within a one-mile radius than rural regions, and urban households consume 20% more fast food each week than rural households. These differences in consumption habits lead to differences in purchasing behavior and highlight a limitation of our analysis (common to prior works). Namely, we only measure consumption at grocery stores, discount stores, and drug stores, and cannot account for consumption at non-retailer vendors such as restaurants, fast-food outlets, bodegas, etc.

Next, we find that Philadelphia households are significantly more likely to cross-border

[30]The same procedure for computing the expected amount for each simulated household and then averaging across simulated households is used in subsequent analyses when we compute the average amount of tax paid, loss in consumer surplus, and sugar and caloric consumption.

Table 8: Average Beverage Consumption per Household[a]

Bev. Location × SB Status	Without Tax	With Tax	Difference	% Change
All Households				
Philadelphia Non-SBs	2,168	2,653	+485	22.37%
Philadelphia SBs	3,940	1,840	-2,100	-53.30%
Non-Philadelphia Non-SBs	4,833	4,910	+77	1.59%
Non-Philadelphia SBs	7,285	7,825	+540	7.41%
Philadelphia Households				
Philadelphia Non-SBs	4,194	5,184	+990	23.61%
Philadelphia SBs	7,624	3,673	-3,951	-51.82%
Non-Philadelphia Non-SBs	1,352	1,472	+120	8.88%
Non-Philadelphia SBs	2,187	3,030	+843	38.55%
Non-Philadelphia Households				
Philadelphia Non-SBs	321	344	+23	7.17%
Philadelphia SBs	578	167	-411	-71.11%
Non-Philadelphia Non-SBs	8,009	8,047	+38	0.47%
Non-Philadelphia SBs	11,937	12,199	+262	2.19%

[a]In ounces; aggregate amount over the post-taxation period January 2017 to December 2018.

shop. This is supported by our household-level data, which suggests that compared to non-Philadelphia households who live within a 10-minute drive of Philadelphia, Philadelphia households who live within a 10-minute drive of the surrounding region are about twice as likely to crosser-border shop. Further, nearly half of Philadelphia households live within a 10-minute drive of the surrounding region, whereas only a quarter of non-Philadelphia households live within a 10-minute drive of Philadelphia. Finally, we hypothesize that the greater willingness to cross-border shop by Philadelphia households is driven by the greater availability of stores in the non-Philadelphia region. Compared to Philadelphia, the non-Philadelphia region contains 2.5 times more grocery stores, twice as many discount stores, and 30% more drug stores despite the number of households living in Philadelphia and in the surrounding region being roughly the same. This greater willingness to cross-border shop by Philadelphia households has a direct impact on the effectiveness of Philadelphia's SB taxation policy.

Turning now to the percentage change in consumption with and without taxation, our counterfactual simulation shows that Philadelphia's SB tax reduces an average household's purchase of Philadelphia SBs by 53%. 26% ($= 540/2,100$) of this reduction is offset by an increase in the purchase of non-Philadelphia SBs, leading to a net reduction equivalent to 40% of the purchase of Philadelphia SBs in the no-tax scenario. These findings differ from that of Seiler et al. (2021), however they reflect the outcomes seen in Roberto et al. (2019) which we discuss in greater detail later in this section.

Of course, considering only the average household does not provide a full picture. Instead, a primary benefit of our structural estimation using a combination of retail and household data is the ability to explore how the taxation policy affects households' behavior conditional on the location of their residence.

As expected, Philadelphia households on the whole favor Philadelphia beverage options. In the case without taxation, 78% of Philadelphia households' SB purchase is for SBs sold within the city limits. The implementation of the SB tax reduces their purchase of Philadelphia SBs by 52% and increases their purchase of non-Philadelphia SBs by 39%. Since Philadelphia households' purchase of non-Philadelphia SBs without taxation is relatively small, a 39% increase in their non-Philadelphia purchase offsets only 21% of the reduction in their Philadelphia purchase. When considering the change in SB purchase in the two locations combined, Philadelphia households experience an average reduction of 32%.

Non-Philadelphia households also prefer beverage options in their home location. In the case without taxation, the purchase of Philadelphia SBs accounts for only 5% of non-Philadelphia households' SB purchase. Furthermore, compared to Philadelphia households, non-Philadelphia households' purchase of Philadelphia SBs is more responsive to the SB tax: they reduce their purchase of Philadelphia SBs by 71% and offset 63% of this reduction through an increase in non-Philadelphia SB purchase. This is an intuitive result, as non-Philadelphia households already live in a region without taxation and travel carries an inherent cost. When considering non-Philadelphia households' SB purchase in the two locations combined, we find that the tax leads to a drop of only 1.2%.

Finally, from Table 2 we know that non-Philadelphia households comprise 52.3% of all households in our market, and from Table 8 we find that relative to Philadelphia households, non-Philadelphia households display a greater tendency to transfer their SB purchase from Philadelphia to the surrounding region in response to the SB tax. It is then not surprising that 25% of the increase in the purchase of non-Philadelphia SBs comes from non-Philadelphia households avoiding the taxed region, rather than cross-border shopping by Philadelphia households. Prior studies of SB taxation consider the increase in SB sales in the surrounding untaxed region to be a result of cross-border shopping by residents of the taxed region. Our results highlight the multiple sources of such an increase and suggest that SB taxation may be more effective than previously thought, if we consider that the tax's intended target is those households residing within the city limits.

Two prior papers, Roberto et al. (2019) and Seiler et al. (2021), also use retail scanner data to examine the Philadelphia SB tax. There exist several similarities and dissimilarities between our works. In particular, our counterfactual simulation finds a decrease in volume sales of Philadelphia SBs greater than that suggested by either prior paper. Roberto et al. (2019) find that volume sales of Philadelphia taxed beverages decline by 51% after the taxation policy and that 24% of this reduction is offset by an increase in volume sales in the surrounding region for a net reduction of 38%, while Seiler et al. (2021) find a decrease of 46% in volume sales of Philadelphia taxed beverages, with 52% of this reduction offset by an increase in volume sales in the surrounding region for a net reduction of 22%. In comparison, we find that the SB tax results in a 53% reduction in volume sales of Philadelphia taxed beverages, with an increase in volume sales in the surrounding region offsetting 26% of this reduction for a net reduction of 40%.

Differences in the estimated impact of the SB tax can result from a multitude of factors. Firstly, to the best of our knowledge, our paper is the first to analyze the effects of an SB tax in a structural context where geographic substitution plays a primary role in determining consumers' choices. The works of Roberto et al. (2019), Cawley et al. (2020), and Seiler et al. (2021), among others, employ reduced form estimations that consider the change from pre-taxation to post-taxation SB volume sales. Using a structural model, we complement prior works by forming our counterfactual estimation directly on the post-taxation months and incorporating the presence of shocks unrelated to changes in tax policy; thus, we model purchase as it would have been in the post-taxation period barring the presence of taxation. Secondly, both Roberto et al. (2019) and Seiler et al. (2021) use data obtained from IRI whereas our data is provided by NielsenIQ; differences in the retail stores covered by the different data sources can contribute to differences in the expected outcome. Differences in weighting procedures, controls, sample distributions, and cleaning methods may also contribute to overall discrepancies in findings. Finally, to more accurately account for households' heterogeneous responsiveness, we rely upon both retail and household data, which is another potential source for differing results between our work and those of others.

8 Welfare Implications

In this section, we examine the welfare effects of the soda tax on consumers shopping at grocery stores, drug stores, and discount stores, focusing on both the tax burden on households and the subsequent changes in their consumer surplus. To provide a nuanced understanding of the welfare impacts of localized soda taxation, we decompose these changes across multiple dimensions, exploring the broader and more comprehensive effects of such fiscal measures on overall welfare.

Following this, we delve into the implications of deadweight loss resulting from localized taxation, comparing it to a counterfactual scenario where all regions are uniformly taxed. This comparison highlights the significant role of localized taxation and cross-border shopping in generating economic distortions. Finally, we revisit the changes in consumer surplus, but this time decompose them according to household demographic characteristics rather than location.

Table 9: Loss in Consumer Surplus per Household[a]

	All Households	Phil. Households	Non-Phil. Households
Expected CS Loss	**\$56.59**	**\$112.06**	**\$5.98**
- Tax Paid	\$27.80	\$55.52	\$2.51
- Travel Cost	\$16.21	\$46.34	-\$11.28
- Product Substitution	\$3.24	\$6.61	\$0.16
- Consumption Reduction	\$6.85	\$11.31	\$2.79
- Residual Loss	\$2.49	-\$7.72	\$11.80

[a] Aggregate amount over the post-taxation period January 2017 to December 2018.

This allows us to examine the differential effects of the tax on various demographic groups, providing a more granular understanding of the impact of the tax and shedding light on the varying consequences of localized taxation in different segments of the population.

8.1 Consumer Surplus Changes by Household Location

We begin by evaluating the average amount of tax paid and loss in consumer surplus per household during the 24 post-taxation months. Under a nested-logit framework, loss in consumer surplus is the difference between the expected utility (the "inclusive value" I_{it} in Eq. (7)) without and with taxation divided by the household's marginal utility of money α_i.

To decompose the loss of consumer surplus across multiple dimensions, we employ the Monte Carlo simulation described in Subsection 7.2 to evaluate changes in expected choice utility. Specifically, we simulate draws from the distribution of $\bar{\epsilon}_{ijt}$, holding them constant while assessing product-level utilities with and without taxation. The maximum value of the product-level utilities in each scenario determines the beverage choice, allowing us to calculate the corresponding changes in consumer surplus. The Monte Carlo simulation error is minimal, with the derived change in consumer surplus demonstrating a difference of less than 1% relative to the change in consumer surplus obtained from the standard two-level nested logit expression.

Table 9 presents our findings averaged across all households, Philadelphia households, and non-Philadelphia households, respectively. We decompose the loss of expected consumer surplus into five primary dimensions: tax paid, travel cost, product substitution, consumption reduction, and residual loss. The sum of these dimensions equals the total loss in expected consumer surplus associated with Philadelphia's tax policy from January 2017 through December 2018.

Expected Consumer Surplus Loss Across all households identified as living in or within the region substitutable with Philadelphia, the expected loss in consumer surplus is \$56.59. This loss represents the dollar amount the government needs to pay a household for them to reach their pretax utility level. As expected with localized taxation, households with the greatest loss in surplus are those living within the city limits, with an average Philadelphia household's reduction in consumer surplus 19 times that of an average non-Philadelphia household. This difference follows from non-Philadelphia households' lower demand for Philadelphia SBs (as discussed in Subsection 7.3) and the fact that they can purchase in the untaxed location without incurring travel costs. For all consumers, this loss in expected surplus is decomposed along multiple dimensions.

Tax Paid First, we examined the loss in consumer surplus associated with the dollar amount paid in taxes by households. As Philadelphia's soda tax is applied to retailers, the amount households pay is dependent on the pass-through rate. The pass-through rate for most beverage options is about 100%, as such this measure is also similar to the total tax generated per household.

We find that about half of the expected loss in consumer surplus is generated from the tax paid. For example, for Philadelphia households, tax paid makes up about 49% of the expected loss in consumer surplus. Although this loss of consumer surplus due to the tax paid is less than that suggested by Wang (2015), it is intuitive given the localized nature of Philadelphia's taxation policy. Demand for SBs remains high despite taxation, and to purchase these products, consumers face two choices: pay the post-tax price or travel. Many choose to pay and some choose to travel.

Travel Cost For those who choose to change their travel due to the post-tax price, they incur a change in their expected surplus. This can either be a loss in surplus from Philadelphia households incentivized to purchase non-Philadelphia SBs post-tax, or a gain in surplus from non-Philadelphia households avoiding a more costly trip. Focusing first on Philadelphia households, we find that loss in surplus from post-tax travel cost is nearly as large as that due to the tax paid. Whereas, for non-Philadelphia consumers, there is a gain in surplus associated with travel. Again, these results are intuitive.

High demand for SBs encourages those who wish to avoid the tax to purchase in the untaxed region. For Philadelphia households, this avoidance of taxation, yet continued desire for SBs, leads some to incur costly travel in favor of costly taxation. However, for non-Philadelphia households who wish to avoid taxes while still consuming SBs, their solution is simple: avoid Philadelphia and shop near home. Thus, non-Philadelphia households experience a gain in surplus from reducing their travel.

Product Substitution Next, we consider the loss in surplus from those who would forgo taxation and travel changes in favor of substitution from taxed SBs to untaxed beverages within Philadelphia. In this, our analysis is unique when compared to similar works, as Philadelphia's soda tax extends to both sugar-sweetened *and* diet products.

For example, Wang (2015) decomposes the loss in utility associated with soda taxation, and finds that the loss in consumer surplus from product switching due to a penny-per-ounce tax is about 21% to 29% of the total reduction in utility, dependent on income status. However, our work suggests that the loss in surplus due to product switching is on average $3.24, or about 6% of the total loss in consumer surplus. This difference arises due to the differences in tax policies between the hypothetical one explored in Wang (2015) and that realized in Philadelphia.

Extending the soda tax to both sugar-sweetened and artificially-sweetened (diet) products eliminates the primary alternative consumers would consider under a typical soda taxation policy. With high demand for SBs and few close substitutes available, it is no surprise that consumers would view product substitution as an undesirable alternative. Hence, they do not incur large deadweight losses along this channel. This is further supported by the fact that surplus loss associated with a reduction in consumption is greater than that of product substitution.

Consumption Reduction After travel costs, consumption reduction—where households switch from purchasing Philadelphia sugary beverages to opting out post-tax—emerges as the third-largest contributor to the decline in consumer surplus. This outcome is intuitive. As previously mentioned, the inclusion of diet products within the sugary beverage taxation policy leaves few close substitutes for most taxed beverages. Consequently, those who choose neither to travel nor to pay the tax are more likely to opt for no purchase at all rather than substitute with less desirable alternatives.

Residual Loss We define all changes in consumer surplus not captured by the previously defined channels as "residual loss." This encompasses, but is not limited to, changes in surplus due to substitution between taxed products, utility changes resulting from counterfactual price adjustments, and utility shifts beyond travel costs from switching to untaxed products in alternative regions.

8.2 Deadweight Loss

Overall, our findings suggest that Philadelphia's localized taxation policy is highly inefficient. The reductions in consumer surplus beyond the tax paid represents consumer *deadweight loss*, with over half of the total loss in consumer surplus attributable to factors other than the tax itself. The localized nature of the policy significantly contributes to this distortion, as travel costs alone account for 29% of the total consumer surplus loss across all households, rising to 41% when considering only those residing within the city limits.

Table 10 examines the counterfactual scenario where all locations are taxed. In this scenario, deadweight loss decreases significantly while the share of Consumer Surplus Loss attributed to the Tax Paid dimension increases from 49% to 71%. These findings reinforce the idea that

Table 10: Loss in Consumer Surplus per Household: All Regions Taxed Counterfactual[a]

	All Households	Phil. Households	Non-Phil. Households
Expected CS Loss	**$123.95**	**$109.85**	**$136.82**
- Tax Paid	$87.53	$79.33	$95.01
- Travel Cost	-$8.84	-$15.75	-$2.54
- Product Substitution	$9.71	$5.64	$13.43
- Consumption Reduction	$21.44	$27.00	$16.37
- Residual Loss	$14.11	$13.63	$14.55

[a] Aggregate amount over the post-taxation period January 2017 to December 2018.

localized taxation is highly inefficient when travel time to untaxed regions is small, as cross-border shopping contributes significantly to deadweight loss.

This result is consistent with existing findings. Agarwal (2012) argues that uniform tax rates within a region can lead to substantial distortions along borders, especially when neighboring regions have significantly lower prices. To mitigate the resulting deadweight loss, he proposes that state planners could implement tax rates that decrease as they approach the border. However, Philadelphia's relatively small geographic size likely limits the feasibility of such a strategy, and instead faces significant reductions in deadweight loss.

8.3 Sugar and Caloric Consumption by Household Location

All is not bad for consumers of SBs. We next consider how the taxation policy reduces sugar and caloric consumption—a "side benefit" of the SB tax, whose stated primary goal is to generate tax revenue. According to the US Center for Disease Control and Prevention, SBs are the leading source of added sugars in the American diet, and frequent consumption of sugary drinks is associated with obesity, type 2 diabetes, heart disease, and kidney diseases, among a plethora of other negative health effects.[31] Table 11 presents the changes in sugar and caloric consumption from beverages during the 24 post-taxation months, averaged across all households, Philadelphia households, and non-Philadelphia households, respectively.

We find that, for an average household living in Philadelphia and the surrounding region, there is an expected reduction in the consumption of sugar by 12%. This effect is strongest for Philadelphia households, who have an average reduction of 26%, whereas non-Philadelphia households—who are not the targeted population of the SB tax—experience an average reduction of 1%. To put this reduction in context, we consider the expected caloric reduction. For Philadelphia households, the implementation of the SB tax translates to a decrease in caloric intake equal to 23,609 calories—approximately 12 days' worth of caloric intake (under a 2,000-calories-a-day diet). The sizable reduction in sugar and caloric consumption among Philadelphia

[31] See https://www.cdc.gov/nutrition/data-statistics/sugar-sweetened-beverages-intake.html.

Table 11: Average Sugar and Caloric Consumption from Beverages per Household[a]

	Without Tax	With Tax	Difference	% Change
All Households				
Sugar (g)	25,395	22,380	-3,015	-11.87%
Calories (Cal)	103,984	92,162	-11,822	-11.37%
Philadelphia Households				
Sugar (g)	23,216	17,198	-6,018	-25.92%
Calories (Cal)	94,069	70,460	-23,609	-25.10%
Non-Philadelphia Households				
Sugar (g)	27,383	27,108	-275	-1.00%
Calories (Cal)	113,032	111,962	-1,070	-0.95%

[a] Aggregate amount over the post-taxation period January 2017 to December 2018.

households attests to the substantial public health benefits of the SB tax. Note that our results only consider the decrease in sugar and caloric consumption from beverages purchased at grocery stores, discount stores, and drug stores; overall reduction will be larger when considering other avenues of purchase. Also note that our analysis does not consider substitution to sugary non-beverage alternatives.

8.4 Differences between High- and Low-Income Households

We now consider to what extent households with different income status differ in their amount of tax paid, loss of consumer surplus, and reduction in sugar and caloric consumption. This will in turn inform us about the degree to which the taxation policy exhibits regressive tendencies, which is particularly relevant in this context, as a primary concern for opponents of Philadelphia's SB tax was its potential impact on the city's poor—households who, as found in past studies, generally display a greater demand for SBs, the products to be taxed.

From our structural setup, there are several mechanisms by which low-income households may react differently to the implementation of an SB tax. This includes differences in preference for product categories, price, travel time, etc.; a full accounting is found in Table 3. For example, our results suggest that low-income households incur less disutility from travel time, which may result in a greater willingness to cross-border shop. Additionally, price sensitivity may differ between those with means and those without.

Price Elasticity of Demand We begin by considering price elasticity of demand by income status. Table 12 presents our findings. Unlike in Table 4, here we consider own- and cross-elasticities of demand averaged only at the "all beverage options" level to highlight the differences between high- and low-income households in terms of their responsiveness to price increases. As before, cross-elasticities of demand are reported for beverage options from the same category, same

Table 12: Price Elasticity of Demand by Income Status

Average Level	**Own-Elasticity**	**Cross-Elasticity**			
		Same Category			All Bev. Options
		All Bev. Options	Same Location	Different Location	
High-Income					
All Bev. Options	-1.0276	0.0058	0.0103	0.0013	0.0010
Low-Income					
All Bev. Options	-1.1686	0.0072	0.0120	0.0023	0.0011

Table 13: Average SB Consumption per Household, by Location and Income Status[a]

Income Status × SB Location	**Without Tax**	**With Tax**	**Difference**	**% Change**
Philadelphia Households				
High-Income				
Philadelphia SBs	7,951	3,763	-4,188	-52.67%
Non-Philadelphia SBs	1,928	2,735	+807	41.86%
Low-Income				
Philadelphia SBs	7,208	3,558	-3,650	-50.64%
Non-Philadelphia SBs	2,517	3,406	+889	35.32%
Non-Philadelphia Households				
High-Income				
Philadelphia SBs	510	141	-369	-72.35%
Non-Philadelphia SBs	12,022	12,259	+237	1.97%
Low-Income				
Philadelphia SBs	811	258	-553	-68.19%
Non-Philadelphia SBs	11,649	12,000	+351	3.01%

[a] In ounces; aggregate amount over the post-taxation period January 2017 to December 2018.

category and same location, same category and different location, and all beverage options.

We find that low-income households display a greater price sensitivity than high-income households, with respect to both own-elasticity and cross-elasticities. For example, low-income households' own-elasticity of demand is 14% greater than high-income households'. This finding is intuitive, as one would expect those with less income to display a greater sensitivity to changes in price. However, despite their greater price elasticity, regardless of location low-income households still display less sensitivity to SB taxation, as demonstrated in Table 13.

SB Consumption Table 13 considers SB consumption by location and income status with and without taxation during the post-taxation period. Interestingly, we find that high-income households consume more SBs than low-income households. In data we observe that high-income households drink substantially more tea and juice, including the SB varieties, while low-income households drink considerably more sweetened carbonated soft drinks.

Table 13 shows that low-income households are less responsive to the taxation policy than high-income households. Among Philadelphia households, low-income households reduce their consumption of Philadelphia SBs by 51% in response to the tax, 2 percentage points lower than their high-income counterparts, while among non-Philadelphia households, low-income households reduce their consumption of Philadelphia SBs by 68% in response to the tax, 4 percentage points lower than their high-income counterparts.

When considering the net reduction in Philadelphia SB consumption (i.e., after accounting for the offset by cross-border shopping) among Philadelphia households, we see that low-income households experience a net reduction of 38% (with the increase in cross-border shopping offsetting 24% of the reduction in Philadelphia SB consumption), whereas high-income households experience a bigger net reduction at 43% (with the increase in cross-border shopping offsetting only 19% of the reduction in Philadelphia SB consumption).

Table 13 also shows that among both Philadelphia households and non-Philadelphia households, and both without tax and with tax, low-income households exhibit a greater tendency to cross-border shop than high-income households. These two types of households' geographic distribution may go towards explaining the discrepancies between their cross-border shopping behavior. Figure 2 shows the percentage of high-income households for each ZIP Code in and around Philadelphia. We observe that within Philadelphia, low-income households tend to live along the city's western edge, through the city center, and up to the northern city limits. Outside Philadelphia, we see that low-income households tend to live near the western portion of the city border. It is along these western and northern city limits—with high population density and where some of Philadelphia's most vulnerable low-income households reside—that we observe some of the highest rates of cross-border shopping. In comparison, especially among non-Philadelphia households, high-income households find it more beneficial to avoid cross-border shopping since their saving of travel costs would be higher, both because they have greater disutility from travel time and because they tend to live farther from the border in less population-dense regions. Such a pattern therefore offers an explanation for the greater tendency among low-income households to cross-border shop.

Amount of Tax Paid and Loss in Consumer Surplus Following directly from the differences in purchasing behavior, we consider the differences between high- and low-income households in the amount of tax paid and loss in consumer surplus. Table 14 presents these results.

Within Philadelphia, high-income households pay 6% more taxes than their low-income counterparts, while outside Philadelphia, low-income households pay 83% more taxes than their high-income counterparts. Among Philadelphia households, the difference in tax paid arises primarily from the pattern that high-income households have a greater preference for tea and

Figure 2: Percentage of High-Income Households by ZIP Code[a]

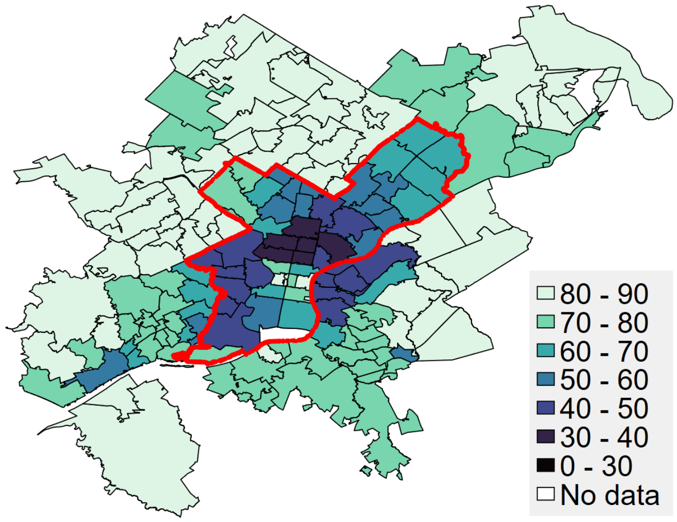

[a]The city of Philadelphia is the area outlined in red. Total map area consists of those ZIP codes included in model estimation.

Table 14: Average Tax Paid and Loss in Consumer Surplus per Household, by Location and Income Status[a]

	All Households	Phil. Households	Non-Phil. Households
High-Income			
Tax Paid	\$23.71	\$56.44	\$2.11
ΔCS	-\$47.83	-\$112.78	-\$4.96
Low-Income			
Tax Paid	\$35.58	\$53.37	\$3.87
ΔCS	-\$-75.83	-\$112.93	-\$9.72

[a]Aggregate amount over the post-taxation period January 2017 to December 2018.

juice SBs and tend to purchase more of these SBs with or without taxation, thus driving up their overall consumption of taxed products. Among non-Philadelphia households, in addition to the greater preference for SBs displayed by low-income households, another factor that contributes to the difference in tax paid is a household's home location. As shown in Figure 2, outside Philadelphia, low-income households tend to live close to the city border. Their proximity to Philadelphia coupled with their lower disutility from travel time contributes to their much larger purchase of Philadelphia SBs, with or without taxation, than their high-income counterparts (as shown in Table 13). This, in turn, is the primary driver behind the difference in the amount of tax paid between high- and low-income non-Philadelphia households.

This border proximity also helps explain the significantly larger loss in consumer surplus observed for non-Philadelphia households who are low-income. For them, compared to their high-income counterparts who tend to live farther away from the city and have a lower preference for SBs, Philadelphia SBs are more likely to be the most preferred among all options in their choice set when there is no tax, and therefore the imposition of a tax on Philadelphia SBs has a more negative impact on their consumer surplus. Finally, considering households living within Philadelphia, we find that the loss in consumer surplus from taxation is nearly equal between high- and low-income households.

Regressivity of the SB Tax by Income Status Although low-income Philadelphia households on average incur a loss of consumer surplus similar to that for their high-income counterparts, and their average taxation amount is slightly less, the large income difference between these two groups of households needs to be taken into account when assessing the regressivity of the SB tax.[32] According to data from the 2018 ACS, the average annual income for low-income Philadelphia households is $17,330, whereas their high-income counterparts have a much higher average of $93,245. These numbers, together with the loss-in-consumer-surplus numbers reported in Table 14, show that when measured as a percentage of annual income, low-income Philadelphia households on average incur a loss of consumer surplus 5.39 times as large as their high-income counterparts', suggesting that the tax is highly regressive. Similarly, among those living outside the city limits, low-income households have an average annual income of $20,264, whereas high-income households have a much higher average of $125,060. Therefore, when measured as a percentage of annual income, low-income non-Philadelphia households again incur a much larger loss of consumer surplus than their high-income counterparts. We will explore the regressivity of SB taxation in greater detail in Subsection 8.5, where we consider the joint distribution of income and representation status.

Changes in Sugar and Caloric Consumption Lastly, we consider changes in sugar and caloric consumption from beverages for high- and low-income households by home location. We find that among Philadelphia households, high-income households on average consume less sugar and fewer calories and experience a greater percentage reduction in their consumption, while among non-Philadelphia households, high-income households on average consume less sugar and fewer calories but experience a smaller percentage reduction in their consumption. Detailed results are reported and discussed in Appendix A5.

[32]Joint distribution of average income and low-income status is obtained from the PUMA ACS for those ZIP Codes pertaining to our sample of interest.

Table 15: Price Elasticity of Demand by Representation Status

Average Level	**Own-Elasticity**	**Cross-Elasticity**			
		Same Category			All Bev. Options
		All Bev. Options	Same Location	Different Location	
Non-URM					
All Bev. Options	-0.9026	0.0049	0.0089	0.0009	0.0008
URM					
All Bev. Options	-1.3928	0.0087	0.0145	0.0028	0.0013

8.5 Differences between URM and Non-URM Households

Similar to above, we now consider to what extent households with different representation status differ in their amount of tax paid, loss of consumer surplus, and reduction in sugar and caloric consumption. This will in turn inform us about the degree to which the taxation policy exhibits regressive tendencies among some of the most vulnerable communities. Again, this is particularly relevant in this context, as a primary concern for opponents of Philadelphia's SB tax was its potential impact on the city's most vulnerable.

From our structural setup, there are several mechanisms by which URM households may react differently to the implementation of an SB tax. This includes differences in preference for product categories, price, travel time, etc.; a full accounting is found in Table 3. For example, our results suggest that URM households incur less disutility from travel time, which may result in a greater willingness to cross-border shop. Additionally, the price sensitivity of URM households is greater than that of non-URM households.

Price Elasticity of Demand We begin by considering price elasticity of demand by representation status. Table 15 presents our findings. Unlike in Table 4, here we consider own- and cross-elasticities of demand averaged only at the "all beverage options" level to highlight the differences between URM and non-URM households in terms of their responsiveness to price increases. As before, cross-elasticities of demand are reported for beverage options from the same category, same category and same location, same category and different location, and all beverage options.

We find that, with respect to both own-elasticity and cross-elasticities, URM households display a noticeably greater price sensitivity than non-URM households. For example, URM households' own-elasticity of demand is 54% greater than non-URM households'. This difference is significantly larger than that between low- and high-income households (Table 12), suggesting that representation status plays a much larger role in choice behavior than income status.

Similar to low-income households, despite their greater price sensitivity, regardless of loca-

Table 16: Average SB Consumption per Household, by Location and Representation Status[a]

Rep. Status × SB Location	Without Tax	With Tax	Difference	% Change
Philadelphia Households				
Non-URM				
Philadelphia SBs	7,592	3,464	-4,128	-54.37%
Non-Philadelphia SBs	1,673	2,390	+717	42.86%
URM				
Philadelphia SBs	7,656	3,871	-3,785	-49.44%
Non-Philadelphia SBs	2,674	3,637	+963	36.01%
Non-Philadelphia Households				
Non-URM				
Philadelphia SBs	467	124	-343	-73.45%
Non-Philadelphia SBs	11,886	12,103	+217	1.83%
URM				
Philadelphia SBs	1,025	342	-683	-66.63%
Non-Philadelphia SBs	12,139	12,585	+446	3.67%

[a] In ounces; aggregate amount over the post-taxation period January 2017 to December 2018.

tion URM households display less sensitivity to SB taxation, as demonstrated by Table 16.

SB Consumption Table 16 considers SB consumption by location and representation status with and without taxation during the post-taxation period. We find that URM households consume more SBs than non-URM households. Considering consumption at the category level, regardless of location, URM households consume less water, less tea, and significantly more juice. Further, considering juice consumption, non-URM households drink a nearly even ratio of SB and non-SB juice, whereas URM households prefer SB juice—these products make up two-thirds of their total juice consumption.

Regardless of household location, URM households reduce their consumption of Philadelphia SBs by a smaller proportion in response to the tax. Among Philadelphia households, URM households reduce their consumption of Philadelphia SBs by 49% in response to the tax, 5 percentage points lower than their non-URM counterparts. Among non-Philadelphia households, the two types of households exhibit an even greater discrepancy, with URM households reducing their consumption of Philadelphia SBs by 67% in response to the tax, 6 percentage points lower than their non-URM counterparts. As with the price elasticity of demand, we observe that the difference in behavior between URM and non-URM households is greater than that between low- and high-income households, again suggesting that representation status plays a stronger role in determining consumer behavior than income status.

When considering the net reduction in Philadelphia SB consumption (i.e., after accounting for the offset by cross-border shopping) among Philadelphia consumers, we see that URM households experience a net reduction of 37% (with the increase in cross-border shopping offsetting 25% of the reduction in Philadelphia SB consumption), whereas non-URM households experience a bigger net reduction at 45% (with the increase in cross-border shopping offsetting only 17% of the reduction in Philadelphia SB consumption). Overall, URM households are less responsive to the tax and are more likely to cross-border shop than non-URM households.

Geographic distribution may go towards explaining the differences between URM and non-URM households' cross-border shopping behavior. Figure 3 shows the percentage of URM households for each ZIP Code in and around Philadelphia. We observe that within Philadelphia, URM households tend to live along the city's western edge, through the city center, and up to the northern city limits. Outside Philadelphia, we see that URM households tend to live near the western portion of the city border. It is along these western and northern city limits—with high population density and where some of Philadelphia's most vulnerable groups reside—that we observe some of the highest rates of cross-border shopping.

Comparing Figure 3 to Figure 2, which presents the geographic distribution of high-income households, we observe that the concentration of URM households along the Philadelphia city border is greater than that of low-income households. This, in turn, suggests that URM households may display a greater willingness to cross-border shop as compared to low-income households. This hypothesis is confirmed by comparing Table 16 to Table 13, both with and without taxation.

Considering non-Philadelphia households, given the geographic distribution of households observed in Figure 3, it is of no surprise that URM households are more likely to shop within Philadelphia than non-URM households, both with and without taxation. To a large extent, the majority of URM households live either within Philadelphia or close to the city limits. This raises concerns that vulnerable communities such as URM households may bear the brunt of Philadelphia's SB taxation, which we examine next.

Amount of Tax Paid and Loss in Consumer Surplus Following directly from the differences in purchasing behavior, we consider the differences between URM and non-URM households in the amount of tax paid and loss in consumer surplus. Table 17 presents these results.

We find that URM households bear the largest tax burden. Within Philadelphia, URM households pay 12% more taxes than their non-URM counterparts, while outside Philadelphia, URM households pay an astounding 176% more taxes than their non-URM counterparts. Among Philadelphia households, the difference in tax paid arises primarily from URM households' stronger preference for SBs and hence their tendency to purchase more of these products with or without taxation. Among non-Philadelphia households, in addition to the greater preference for SBs displayed by URM households, another factor that contributes to the difference in tax paid

Figure 3: Percentage of URM Households by ZIP Code[a]

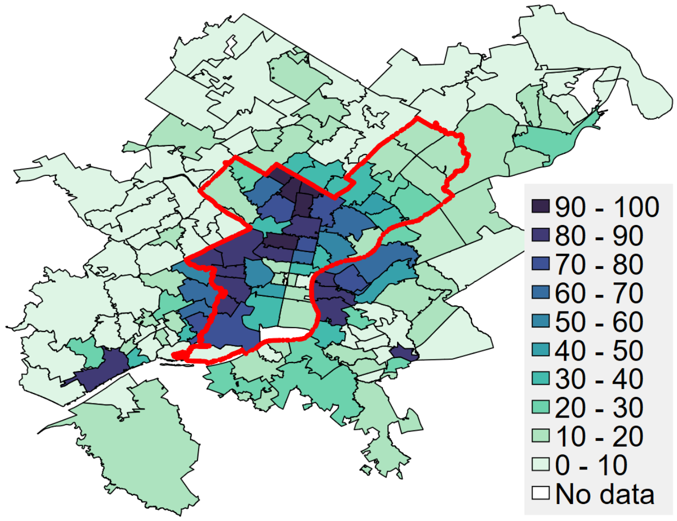

[a]The city of Philadelphia is the area outlined in red. Total map area consists of those ZIP codes included in model estimation.

Table 17: Average Tax Paid and Loss in Consumer Surplus per Household, by Location and Representation Status[a]

	All Households	Phil. Households	Non-Phil. Households
Non-URM			
Tax Paid	\$19.73	\$51.95	\$1.86
ΔCS	-\$37.89	-\$98.56	-\$4.25
URM			
Tax Paid	\$42.25	\$58.06	\$5.13
ΔCS	-\$92.59	-\$126.37	-\$13.30

[a]Aggregate amount over the post-taxation period January 2017 to December 2018.

is home location.

As shown in Figure 3, outside Philadelphia, URM households tend to live close to the city border. Their proximity to Philadelphia coupled with their lower disutility from travel time contributes to their much larger purchase of Philadelphia SBs, with or without taxation (as shown in Table 16). This, in turn, is the primary driver behind the difference in the amount of tax paid between URM and non-URM non-Philadelphia households.

This border proximity also helps explain the significantly larger loss in consumer surplus observed for URM non-Philadelphia households. For them, compared to their non-URM counterparts who tend to live farther away from the city and have a lower preference for SBs, Philadelphia SBs are more likely to be the most preferred among all options in their choice set when there is no tax, and therefore the imposition of a tax on Philadelphia SBs has a more negative impact on their consumer surplus.

Finally, considering consumers living within Philadelphia, we find that the loss in consumer surplus is greater for URM households, primarily driven by their stronger preference for SBs.

Burden of the SB Tax Within Philadelphia, URM households comprise 52% of the households and pay 12% more taxes than their non-URM counterparts. Outside Philadelphia, URM households comprise only 20% of the households and pay 176% more taxes than their non-URM counterparts. Taking into consideration total taxes paid and household distributions, we find that URM households, despite comprising only 35% of the households in and around Philadelphia, generate 54% of the tax revenue. This disparity is even worse when we consider that URM households on average earn \$49,454 a year, whereas non-URM households on average earn \$98,886 a year—a 100% increase. These results suggest that disadvantaged and underrepresented communities pay the greatest share of the tax revenue generated, exacerbating the existing racial disparity among households.

Regressivity of the SB Tax by Income and Representation Status Examination of the tax revenue by income and representation status groups further demonstrates an amplification of disparity resulting from the adverse consequences of taxation. Specifically, low-income URM households (which are the lowest income group within our research and on average earn \$1,000 less than low-income non-URM households) comprise 18% of the households in and around Philadelphia but generate 28% of the total tax revenue. In contrast, low-income non-URM households make up 15% of the households and generate only 14% of the total tax revenue.

Focusing solely on Philadelphia, URM households on average earn \$43,663 a year, 42% less than non-URM households, whose average annual income is \$75,748. The regressivity of SB taxation is exacerbated when considering the joint distribution of income and representation status: within Philadelphia, low-income URM households on average earn \$17,129 a year and pay \$55.64 in SB taxes, whereas high-income non-URM households on average earn \$105,819 a year and pay \$53.23 in SB taxes. Taken as a whole, the consideration of representation status serves to widen the disparity in outcome observed in Subsection 8.4 between high- and low-income households. When measured as a proportion of their respective annual incomes, low-income URM Philadelphia households pay 6.5 times more in SB taxes and see a reduction in consumer surplus 7.6 times greater than their high-income non-URM counterparts. These findings highlight the regressive nature of the SB tax: those who bear the greatest burden from the tax are underrepresented communities with the least means.

Changes in Sugar and Caloric Consumption Lastly, we consider changes in sugar and caloric consumption from beverages by home location and representation status. We find that among

Philadelphia households, non-URM households on average consume less sugar and fewer calories and experience a greater percentage reduction in their consumption, while among non-Philadelphia households, non-URM households on average consume less sugar and fewer calories but experience a smaller percentage reduction in their consumption. Detailed results are reported and discussed in Appendix A6.

9 Alternative Scenarios

In this section, we examine several alternative scenarios to further our understanding of Philadelphia's SB tax. We first vary the tax rate to identify the one that maximizes the tax revenue. We then consider the changes in sugar and caloric consumption, consumer surplus, and the revenue-maximizing tax rate if diet products are not subject to taxation (as in the original proposal of the Philadelphia SB tax). We also consider the impact of taxation on SB consumption and consumer surplus if not only Philadelphia but also its surrounding region are subject to the same tax (as would be the case if the tax is implemented in a broader region, for instance at the state or federal level). Lastly, we consider how reduction in travel time (resulting from improved roads, for example) would affect SB consumption and cross-border shopping behavior.

9.1 Revenue-Maximizing Tax Rate

Here, we use our estimates of beverage demand and taxation responses to predict outcomes under counterfactual tax rates. Counterfactual prices under hypothetical tax rates are obtained based on the supply-side model described in Section 7. As mentioned earlier, unlike prior works, our estimates of demand responsiveness to taxation account for consumer heterogeneity in terms of beverage preferences, travel costs, and locational and categorical substitution.

Given any counterfactual tax rate and the corresponding counterfactual beverage prices, we calculate the average amount of SB tax payment per household during the 24 post-taxation months for each income status/representation status/location combination. We use these averages and the demographic distribution of households provided in Table 2 to obtain the total tax revenue for each combination. Summing over all such combinations provides the total tax revenue for the given tax rate.

We obtain a revenue-maximizing tax rate of 2.51¢ per ounce—much closer to the initially proposed tax rate of 3¢ per ounce than the actual tax rate of 1.5¢ per ounce. We find that the actual tax rate generates a revenue of $34.4 million during the 24 post-taxation months, which equals 94% of the $36.7 million that would be generated at the revenue-maximizing tax rate, whereas the initially proposed tax rate of 3¢ per ounce would generate a revenue more than 99% of the maximal revenue. We note that these revenue figures account for revenue generated at the stores in our sample, namely grocery stores, discount stores, and drug stores; we do not consider revenue from other sources such as supercenters, gas stations, dollar stores, and non-retailer vendors such as restaurants, fast-food outlets, and theaters.

Furthermore, we find that the revenue-maximizing tax rate demonstrates even bigger regressive tendencies, with 30.4% of the tax revenue coming from low-income URM households, compared to 28.2% under the actual tax rate. Additionally, the revenue-maximizing tax rate would lead to an average consumer surplus loss of $72.14, constituting a 27% increase when compared to the actual tax rate's average loss of $56.99.

Differences between the revenue-maximizing tax rate found in our work and those found in other research likely arise from differences in the structure of the demand curve. Of particular note is Seiler et al. (2021), who find a revenue-maximizing tax rate of 1.63¢ per ounce for Philadelphia's SB tax, when assuming a linear demand curve. Our findings, in comparison, are derived from our demand estimates based on the RCNL modeling structure. We contend that our higher revenue-maximizing tax rate is driven primarily by the persistent consumption of Philadelphia SBs by a subset of Philadelphia households who lack inexpensive substitutes for SBs within their home region, many of these households facing large travel costs and high SB preferences. This pattern leads to such households' low price sensitivity with respect to Philadelphia SBs, which in turn gives rise to a higher figure for the revenue-maximizing tax rate.

In terms of structural modeling, Allcott et al. (2019) consider the optimal SB tax rate for a government with preferences for wealth redistribution. They determine an optimal tax rate between 1¢ and 2.1¢ per ounce. Differing from their work, our analysis is concerned with the revenue-maximizing tax rate rather than the socially optimal tax rate. Furthermore, they focus on a national SB tax imposed on sugar-sweetened beverages only, whereas our analysis is concerned with the city of Philadelphia and includes diet products containing artificial sweeteners. A more appropriate comparison between our work and Allcott et al. (2019) is found in the next subsection, where we report a revenue-maximizing tax rate of 1.90¢ per ounce under the assumption that diet products are excluded from the tax.

9.2 Additional Counterfactuals

Next we consider three additional counterfactual scenarios. Detailed results are reported in Table 18; here we summarize the main findings.

We find that removing diet products from the tax (column 4 of Table 18) induces a greater reduction in households' sugar and caloric consumption and reduces households' tax paid and loss in consumer surplus, both of which are beneficial for households. The main intuition here is that sugary beverages and their diet counterparts are good substitutes for some households, therefore when diet products are excluded from the tax, these households are able to switch from sugary beverages to their diet counterparts in order to avoid the tax, rather than having to travel for cross-border shopping or switch to less substitutable products. Correspondingly, we see that cross-border shopping is reduced, illustrating the inverse relationship between product substitution and geographic substitution. Additionally, the relatively high substitutability between taxed products and their untaxed diet counterparts results in a significantly lower revenue-maximizing tax rate of 1.90¢ per ounce, compared to 2.51¢ per ounce when diet products are also taxed.

If the SB tax is levied upon both Philadelphia and its surrounding region (column 5 of Table 18), there is a more significant reduction in households' sugar and caloric consumption, and households experience a larger amount of tax paid and a bigger loss in consumer surplus. Interestingly but intuitively, the consumption of Philadelphia SBs is reduced by a noticeably smaller proportion in response to the SB tax. The removal of the option to buy untaxed SBs in the non-Philadelphia location results in some households choosing to pay the heightened prices and continue buying Philadelphia SBs.

If the travel time experienced by all households is reduced by 25% of the baseline (column 6 of Table 18), cross-border shopping is much more active, offsetting a significantly larger proportion of the reduction in the consumption of Philadelphia SBs. The reduction in sugar and caloric consumption is diminished, and the amount of tax paid and the loss in consumer surplus are reduced.

Together, our counterfactual analyses show that when designing local taxation policies, policymakers need to pay careful attention to the scope of the tax—in terms of product and geographic coverage—as well as households' ability and tendency for cross-border shopping, as they are shown to have significant impact on households' responses and the consumption and welfare consequences of the tax. Furthermore, the structural approach that we use in our analyses provides a useful framework for policymakers in their policy development and evaluation.

10 Conclusion

In this work, we employ a structural modeling framework that combines both retail and household data to study the relationship between local taxation and households' tax avoidance behavior including cross-border shopping and product substitution, focusing on Philadelphia's SB tax.

We find that travel time to the untaxed region surrounding Philadelphia plays an important role in determining households' substitution patterns. In response to the implementation of an SB tax, our results quantify households' reduction in the consumption of taxed beverages in Philadelphia and their willingness to seek untaxed products in locales beyond the city border. Accounting for household location, we find that 26% of the rise in SB sales in the surrounding

Table 18: Additional Counterfactuals

		Actual	No Diet[a]	Both Loc. Taxed	Travel Time Reduced
%Δ Phil. SBs	**All Households**	**-53.30%**	**-60.19%**	**-45.20%**	**-58.71%**
	URM	-50.36%	-56.31%	-41.69%	-57.09%
	Non-URM	-56.26%	-64.76%	-48.79%	-60.35%
	Low-Income	-51.68%	-58.16%	-43.06%	-58.00%
	High-Income	-54.42%	-61.64%	-46.70%	-59.21%
% Offset by Cross-Border Shopping	**All Households**	**25.72%**	**17.54%**	**-**	**50.89%**
	URM	28.28%	20.64%	-	56.23%
	Non-URM	23.39%	14.42%	-	45.75%
	Low-Income	27.44%	19.34%	-	55.22%
	High-Income	24.57%	9.85%	-	47.95%
%Δ Sugar	**All Households**	**-11.87%**	**-14.76%**	**-37.04%**	**-8.16%**
	URM	-14.97%	-18.35%	-34.07%	-9.63%
	Non-URM	-9.82%	-12.36%	-39.01%	-7.18%
	Low-Income	-14.96%	-18.56%	-36.97%	-9.85%
	High-Income	-10.36%	-12.87%	-37.08%	-7.32%
%Δ Calories	**All Households**	**-11.37%**	**-13.75%**	**-35.81%**	**-7.68%**
	URM	-14.38%	-17.26%	-32.69%	-9.07%
	Non-URM	-9.42%	-11.48%	-37.83%	-6.79%
	Low-Income	-14.45%	-17.63%	-35.78%	-9.36%
	High-Income	-9.90%	-11.90%	-35.82%	-6.88%
Tax Paid	**All Households**	**-$27.60**	**-$18.67**	**-$89.21**	**-$24.40**
	URM	-$42.25	-$31.64	-$96.23	-$36.53
	Non-URM	-$19.73	-$11.72	-$85.43	-$17.90
	Low-Income	-$35.58	-$25.00	-$89.40	-$30.93
	High-Income	-$23.71	-$15.59	-$89.11	-$21.23
Δ CS	**All Households**	**-$56.99**	**-$42.67**	**-$164.16**	**-$25.87**
	URM	-$92.59	-$74.68	-$190.24	-$39.25
	Non-URM	-$37.89	-$25.48	-$150.16	-$18.69
	Low-Income	-$75.83	-$58.01	-$169.31	-$30.94
	High-Income	-$47.83	-$35.20	-$161.66	-$23.41

[a]The No Diet scenario removes diet beverages from being taxed, and the reported relative changes in consumption and cross-border shopping pertain to sugar-sweetened beverages only.

region is due to an avoidance of Philadelphia SBs by those residing in the surrounding region, rather than cross-border shopping by Philadelphia households.

Our model and estimation allow for heterogeneous consumer behavior based on their demographic characteristics and proximity to the city border. Taking into account consumers' heterogeneous responses to the tax, we show that the actual tax rate of 1.5¢ per ounce is well below the revenue-maximizing tax rate of 2.51¢ per ounce. Our results suggest that, without readily available substitutes and facing large travel costs associated with cross-border shopping, a subset of Philadelphia households are persistent in their consumption of Philadelphia SBs, willing to pay the higher prices resulting from the tax. Their low price sensitivity with respect to Philadelphia SBs is a main factor behind the high revenue-maximizing tax rate.

Based on our demand estimates, we calculate the average amount of tax paid and loss in consumer surplus for households at different locations and with different income and representation status. Taking into account travel costs and the switch to less preferred products, Philadelphia households on average incur a loss in consumer surplus more than twice the amount of tax paid, with low-income URM households bearing the largest burden. When measured as a percentage of annual income, low-income URM Philadelphia households on average pay 6.5 times more in SB taxes and incur a loss in consumer surplus 7.6 times greater than their high-income non-URM counterparts, suggesting that the tax is highly regressive.

These findings are especially relevant for governmental entities weighing the benefits of a revenue-generating, healthy-habit-inducing tax against the drawbacks of a strongly regressive taxation policy. Additionally, through counterfactual analyses in which we vary the geographic coverage of the tax as well as travel times to the alternative region, we provide supportive evidence for the notion that policymakers must carefully consider geographic coverage and geographic substitution when designing and evaluating local policies.

Lastly, our model's applicability extends beyond the context studied in this work. Any local tax or subsidy susceptible to cross-border shopping offers an opportunity for study under our framework, which facilitates rich modeling and sensible estimation of individuals' heterogeneous travel costs and substitution patterns, as well as the policy's potentially vastly different impact on different individuals.

References

Agarwal, David R. 2012. "Games within borders: are geographically differentiated taxes optimal?" *Journal of Marketing Research* 19 (2): 574–597.

Allcott, Hunt, Benjamin B Lockwood, and Dmitry Taubinsky. 2019. "Regressive sin taxes, with an application to the optimal soda tax." *The Quarterly Journal of Economics* 134 (3): 1557–1626.

Bayer, Patrick, Fernando Ferreira, and Robert McMillan. 2007. "A Unified Framework for Measuring Preferences for Schools and Neighborhoods." *Journal of Political Economy* 115 (4): 588–638.

Berry, Steven, James Levinsohn, and Ariel Pakes. 1995. "Automobile Prices in Market Equilibrium." *Econometrica* 63 (4): 841–890.

Berry, Steven, James Levinsohn, and Ariel Pakes. 2004. "Differentiated Products Demand Systems from a Combination of Micro and Macro Data: The New Car Market." *Journal of Political Economy* 112 (1): 68–105.

Berry, Steven T. 1994. "Estimating Discrete-Choice Models of Product Differentiation." *The RAND Journal of Economics* 25 (2): 242–262.

Bleich, Sara N, Hannah G Lawman, Michael T LeVasseur, Jiali Yan, Nandita Mitra, Caitlin M Lowery, Ana Peterhans, Sophia Hua, Laura A Gibson, and Christina A Roberto. 2020. "The Association of a Sweetened Beverage Tax with Changes in Beverage Prices and Purchases at Independent Stores." *Health Affairs* 39 (7): 1130–1139.

Breslow, Samantha K. 2019. "Chicago and Cook County Amusement Tax." *The National Law Review* IX (114). https://www.natlawreview.com/article/chicago-and-cook-county-amusement-tax.

Burda, Martin, Matthew Harding, and Jerry Hausman. 2008. "A Bayesian mixed logit–probit model for multinomial choice." *Journal of Econometrics* 147 (2): 232–246, Estimating demand systems and measuring consumer preferences.

Bygvrå, Susanne. 2009. "Distance and cross-border shopping for alcohol – evidence from Danes' cross-border shopping 1986–2003." *Nordic Studies on Alcohol and Drugs* 26 (2): 141–163.

Campbell, Margaret C. 1999. "Perceptions of Price Unfairness: Antecedents and Consequences." *Journal of Marketing Research* 36 (2): 187–199.

Capps, Cory, David Dranove, and Mark Satterthwaite. 2003. "Competition and Market Power in Option Demand Markets." *The RAND Journal of Economics* 34 (4): 737–763.

Cawley, John, David Frisvold, Anna Hill, and David Jones. 2020. "The impact of the Philadelphia beverage tax on prices and product availability." *Journal of Policy Analysis and Management* 39 (3): 605–628.

Cawley, John, Barton Willage, and David Frisvold. 2018. "Pass-through of a tax on sugar-sweetened beverages at the Philadelphia International Airport." *JAMA* 319 (3): 305–306.

Chandra, Ambarish, Keith Head, and Mariano Tappata. 2014. "The Economics of Cross-Border Travel." *The Review of Economics and Statistics* 96 (4): 648–661.

Chintagunta, Pradeep K., and Jean-Pierre Dubé. 2005. "Estimating a Stockkeeping-Unit-Level Brand Choice Model That Combines Household Panel Data and Store Data." *Journal of Marketing Research* 42 (3): 368–379.

Conlon, Christopher, and Nirupama Rao. 2023. "The Cost of Curbing Externalities with Market Power: Alcohol Regulations and Tax Alternatives." working paper, New York University.

Davis, Peter. 2006. "Spatial Competition in Retail Markets: Movie Theaters." *The RAND Journal of Economics* 37 (4): 964–982.

DellaVigna, Stefano, and Matthew Gentzkow. 2019. "Uniform Pricing in U.S. Retail Chains." *The Quarterly Journal of Economics* 134 (4): 2011–2084.

Dubois, Pierre, Rachel Griffith, and Martin O'Connell. 2020. "How Well Targeted Are Soda Taxes?" *American Economic Review* 110 (11): 3661–3704.

Duchneskie, John. 2016. "Philly's daily commute: How do your habits compare?." *The Philadelphia Inquirer*. https://www.inquirer.com/philly/business/transportation/405506645.html.

Fisher, Ronald C. 1980. "Local Sales Taxes: Tax Rate Differentials, Sales Loss, and Revenue Estimation." *Public Finance Quarterly* 8 (2): 171–188.

Friberg, Richard, Frode Steen, and Simen A. Ulsaker. 2018. "Hump-shaped cross-price effects and the extensive margin in cross-border shopping." Discussion Paper Series in Economics 20/2019, Norwegian School of Economics.

Garrison, Cassandra, and Adriana Barrera. 2022. "Mexico suspends gasoline subsidy as U.S. drivers cross border for deals." *Reuters*. https://www.reuters.com/world/americas/mexico-suspends-subsidy-gasoline-border-region-with-us-2022-04-02/.

Goldin, Jacob, Tatiana Homonoff, and Katherine Meckel. 2022. "Issuance and Incidence: SNAP Benefit Cycles and Grocery Prices." *American Economic Journal: Economic Policy* 14 (1): 152–78.

Goolsbee, Austan, and Amil Petrin. 2004. "The Consumer Gains from Direct Broadcast Satellites and the Competition with Cable TV." *Econometrica* 72 (2): 351–381.

Grieco, Paul LE, Charles Murry, Joris Pinkse, and Stephan Sagl. 2022. "Conformant and Efficient Estimation of Discrete Choice Demand Models." working paper, The Pennsylvania State University.

Grigolon, Laura, and Frank Verboven. 2014. "Nested Logit or Random Coefficients Logit? A Comparison of Alternative Discrete Choice Models of Product Differentiation." *The Review of Economics and Statistics* 96 (5): 916–935.

Harding, Matthew, Ephraim Leibtag, and Michael F. Lovenheim. 2012. "The Heterogeneous Geographic and Socioeconomic Incidence of Cigarette Taxes: Evidence from Nielsen Homescan Data." *American Economic Journal: Economic Policy* 4 (4): 169–98.

Hillengas, Amory, Cheryl Bettigole, Amanda Wagner, and Hannah Lawman. 2019. "Neighborhood Food Retail in Philadelphia." research report, Philadelphia Department of Public Health, Philadelphia, PA, USA, https://www.phila.gov/media/20190923114738/GHP_FoodAccessRpt_Final_wDate.pdf.

Hollenbeck, Brett, and Kosuke Uetake. 2021. "Taxation and Market Power in the Legal Marijuana Industry." *The RAND Journal of Economics* 52 (3): 559–595.

Houde, Jean-François. 2012. "Spatial Differentiation and Vertical Mergers in Retail Markets for Gasoline." *American Economic Review* 102 (5): 2147–82.

Hymel, Kent, Kenneth Small, and Kurt Van Dender. 2010. "Induced demand and rebound effects in road transport." *Transportation Research Part B: Methodological* 44 1220–1241.

Kifer, Avigail. 2015. "The Incidence of a Soda Tax, in Pennies and Pounds." Working Paper 1-017, Kilts Center for Marketing at Chicago Booth – Nielsen Dataset Paper Series.

McFadden, Daniel, Antti Talvitie, and Associates. 1977. *Demand Model Estimation and Validation*. Institute of Transportation Studies, Berkeley CA.

Miller, Nathan H., and Matthew C. Weinberg. 2017. "Understanding the Price Effects of the MillerCoors Joint Venture." *Econometrica* 85 (6): 1763–1791.

Miravete, Eugenio J., Katja Seim, and Jeff Thurk. 2018. "Market Power and the Laffer Curve." *Econometrica* 86 (5): 1651–1687.

Murry, Charles, and Yiyi Zhou. 2020. "Consumer Search and Automobile Dealer Colocation." *Management Science* 66 (5): 1909–1934.

Nevo, Aviv. 2000. "A Practitioner's Guide to Estimation of Random-Coefficients Logit Models of Demand." *Journal of Economics & Management Strategy* 9 (4): 513–548.

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

Rahkovsky, Ilya, Lisa Mancino, and Young Jo. 2018. "Chapter 8: How Food Environment and Proximity to Restaurants Affect Nutritional Quality." *America's Eating Habits: Food Away from Home. M. Saksena and A. Okrent editors, Economic Research Service, U.S. Department of Agriculture,* (2018), 109–122.

Roberto, Christina A., Hannah G. Lawman, Michael T. LeVasseur, Nandita Mitra, Ana Peterhans, Bradley Herring, and Sara N. Bleich. 2019. "Association of a Beverage Tax on Sugar-Sweetened and Artificially Sweetened Beverages with Changes in Beverage Prices and Sales at Chain Retailers in a Large Urban Setting." *JAMA* 321 (18): 1799–1810.

Seiler, Stephan, Anna Tuchman, and Song Yao. 2021. "The Impact of Soda Taxes: Pass-Through, Tax Avoidance, and Nutritional Effects." *Journal of Marketing Research* 58 (1): 22–49.

Thomadsen, Raphael. 2005. "The Effect of Ownership Structure on Price in Geographically Differentiated Industries." *RAND Journal of Economics* 36 908–929.

Tosun, Mehmet, and Mark Skidmore. 2007. "Cross-Border Shopping and the Sales Tax: An Examination of Food Purchases in West Virginia." *Topics in Economic Analysis & Policy* 7 1877–1877.

Train, Kenneth. 1999. "Halton sequences for mixed logit." working paper, UC Berkeley.

Train, Kenneth E. 2009. *Discrete choice methods with simulation*. Cambridge University Press.

Tuchman, Anna E. 2019. "Advertising and Demand for Addictive Goods: The Effects of E-Cigarette Advertising." *Marketing Science* 38 (6): 994–1022.

Wang, Ao. 2021. "A BLP Demand Model of Product-Level Market Shares with Complementarity." The Warwick Economics Research Paper Series (TWERPS) 1351, University of Warwick.

Wang, Emily Yucai. 2015. "The impact of soda taxes on consumer welfare: implications of storability and taste heterogeneity." *The RAND Journal of Economics* 46 (2): 409–441.

Wang, Ke, and Xin Ye. 2024. "A practical method to draw from multivariate extreme value distributions." *Journal of Choice Modelling* 52 100506.

Weber, Till Olaf, Jonas Fooken, and Benedikt Herrmann. 2014. "Behavioural Economics and Taxation." Working Paper N.41 - 2014, European Commission Taxation Papers.

Xia, Lan, Kent B. Monroe, and Jennifer L. Cox. 2004. "The Price Is Unfair! A Conceptual Framework of Price Fairness Perceptions." *Journal of Marketing* 68 (4): 1–15.

Yazzie, Del, Kristen Tallis, Caleigh Curley, Priscilla R. Sanderson, Regina Eddie, Timothy K. Behrens et al. 2020. "The Navajo Nation Healthy Diné Nation Act: A Two Percent Tax on Foods of Minimal-to-No Nutritious Value, 2015–2019." *Preventing Chronic Disease* 17:200038.

Online Appendix

A1 SB Sales in the Region Surrounding Philadelphia

To assess the exclusion of stores beyond the 20-minute band surrounding Philadelphia, we examine SB sales in stores within the 3-digit ZIP Code prefixes (080, 181, 189, 190, 191, 192 and 194) pertaining to Philadelphia and its surrounding region. Specifically, we estimate the impact of the SB tax on SB sales within Philadelphia and the surrounding region by travel time. In addition to the 196 stores in the "city + 20 minutes" region referenced in Section 3.1, we also observe the sales of 57 stores 20–30 minutes from the city, and 49 stores 30+ minutes from the city. SB sales are aggregated at the store-week level for estimation purposes.

Results in Table A1 provide evidence that stores 20+ minutes from the city border are not affected by cross-border shopping, as their SB sales do not demonstrate a positive response to the Philadelphia SB tax. The first regression, Column (1), treats stores 20+ minutes from Philadelphia as our control group. We find a positive and significant change in SB ounces sold among stores within the 20-minute region surrounding Philadelphia, suggesting the presence of cross-border shopping. Column (2) acts as a test of the validity of our "city + 20 minutes" cutoff. Specifically, we treat stores 30+ minutes from Philadelphia as our control group, and find no significant response in sales among stores 20–30 minutes from Philadelphia.

Table A1: Regression of SB Volume Sales[a]

Dependent Variable: Weekly SB Sales (in Ounces)	**(1)**	**(2)**
Post-Tax × Philadelphia	-57109.64***	-56072.88***
	(1795.76)	(2194.39)
Post-Tax × (0–10 minutes from Philadelphia)	63887.86***	64924.60***
	(2876.15)	(3140.52)
Post-Tax × (10–15 minutes from Philadelphia)	16812.08***	17848.84***
	(2158.95)	(2500.33)
Post-Tax × (15–20 minutes from Philadelphia)	21004.13***	22040.88***
	(2025.41)	(2385.98)
Post-Tax × (20–25 minutes from Philadelphia)	–	3369.40
		(2762.45)
Post-Tax × (25–30 minutes from Philadelphia)	–	208.57
		(2920.85)
Store FEs	Y	Y
Week FEs	Y	Y
Observations	63,067	63,067
Weeks	209	209

***p<.01, **p<.05, *p<.1

[a]Robust standard errors reported in parentheses. SB sales aggregated at the store-week level.

A2 Multiple Purchases During a Single Trip

During some observed purchase opportunities, households buy multiple units of the same product or buy multiple different products. However, in our retail data, information pertaining to individual-level purchase variety and amounts is unavailable—we only observe aggregate store sales. To make our model tractable under a discrete choice framework, and to reconcile with the retail data, assumptions are required. In cases where we observe the purchasing of multiple units or multiple products during a single trip, we treat them as arising from multiple purchase opportunities. Finally, we test for and do not find multiple purchase occurrences to be correlated with our demographic variables of interest.

Current literature involving the purchase of multiple units or multiple products under the BLP framework considers bundling units of the same or different goods together as a sort of composite product (e.g., Wang (2021)). This approach would be computationally infeasible in our case given the large number of beverage products. As such, our approach described above (1) simplifies our estimation, (2) makes the model tractable under the BLP framework, and (3) is simply following the assumptions innately made by researchers working solely with retail data (i.e., Berry et al. (1995), Nevo (2000), etc.).

A3 Estimation Procedure During the Retail Data Step

Provided a candidate draw of Θ, for each month $t = 1, \ldots, T$ we need to solve for $\delta_t = (\delta_{1t}, \ldots, \delta_{J_t t})'$ such that

$$s_{jt}(\delta_t; \Theta) = S_{jt}, \text{ for } j = 1, \ldots, J_t, \tag{A1}$$

where $s_{jt}(\cdot)$ is the predicted retail market shares from Eq. (9) and S_{jt} is the observed retail market shares. In solving this system of equations, we require two steps.

We start by calculating the left-hand side of (A1). In practice, we rely upon Monte Carlo integration where Equation (9) is approximated by

$$s_{jt}(\delta_t; \Theta) = \frac{1}{R} \sum_{r=1}^{R} \pi_{rjt}(x_t, p_t, h_t, Q_{z_r}, D_r, \delta_t, \Theta, v_r). \tag{A2}$$

Each simulated household $r = 1, \ldots, R$ is represented by Halton draws of v_r, z_r, and D_r from the distributions of v, z, and $D|z$, respectively. We draw $R = 4000$ simulated households per month to compute Eq. (A2).

Next, to obtain δ_t, we need to invert our system of equations (A1). For the RCNL model, this system of equations is non-linear and is solved numerically. Grigolon and Verboven (2014) provides the contraction mapping algorithm for the random coefficients logit model with nesting parameters. In the case of a one-level nested model, the algorithm iteratively computes

$$\begin{aligned}&\delta_t^{k+1} \equiv \delta_t^k + (1-\rho)[\ln(S_t) - \ln(s_t(\delta_t^k;\Theta))],\ k = 1,2,\ldots,\\ &\text{where } S_t = (S_{1t},\ldots,S_{J_t t})' \text{ and } s_t = (s_{1t},\ldots,s_{J_t t})',\end{aligned} \tag{A3}$$

until the relative difference between δ_t^{k+1} and δ_t^k is less than our tolerance of .5e^{-12}. Once the inversion has been completed for each $t = 1,\ldots,T$, a unique $\delta(\Theta)$ has been obtained, and we proceed to the evaluation of our household-level log-likelihood.

A4 Supply-Side Model

In this appendix, we detail how we obtain counterfactual prices based on the estimates found in Section 6. First, under the assumption that product prices are set optimally at the firm level, product marginal costs are inferred from observed prices and market shares and expected price sensitivity. Specifically, we assume that, independently across each month, firms set their prices to maximize their expected monthly profits. The first-order condition with respect to the price of product j in the set of products $\mathcal{J}_{ft}$ sold by firm f in month t is

$$0 = \frac{\partial \pi_{ft}}{\partial p_{jt}} = \frac{\partial}{\partial p_{jt}} \sum_{n\in\mathcal{J}_{ft}} S_{nt}(p_{nt} - mc_{nt}) = S_{jt} + \sum_{n\in\mathcal{J}_{ft}} \frac{\partial S_{nt}}{\partial p_{jt}}(p_{nt} - mc_{nt}),$$

which can be rewritten in vector form, for each month t, as

$$0 = S_t + \Delta_t'(p_t - mc_t), \tag{A4}$$

where $S_t = (S_{1t},\ldots,S_{jt},\ldots,S_{J_t t})'$, $p_t = (p_{1t},\ldots,p_{jt},\ldots,p_{J_t t})'$, $mc_t = (mc_{1t},\ldots,mc_{jt},\ldots,mc_{J_t t})'$, and Δ_t is a $J_t \times J_t$ matrix with the (n,j) element equal to $\frac{\partial S_{nt}}{\partial p_{jt}}$ if n and j are owned by the same firm, and zero otherwise. Thus, the vector of marginal costs for all products for month t is

$$mc_t = (\Delta_t')^{-1} S_t + p_t. \tag{A5}$$

Once the vector of marginal costs has been obtained, we can proceed to predict the impact of our counterfactual scenarios of interest. We do so by incorporating the counterfactual changes into our model of demand and supply. Take, for instance, a counterfactual increase in the tax rate. An increase in the tax rate increases the marginal costs for taxed products, which induces firms to re-optimize the prices that they set, which in turn changes the equilibrium market outcome.

Provided a gradient vector comprising the first-order conditions for profit maximization, we find the profit-maximizing price vector $\hat{p}_{ft}$ for each firm f in each month t. In application, for each month t, we iterate over the firms, maximizing each firm's profits given the other firms' choice of prices. We continue iterating until $\hat{p}_{ft}$ converges for each firm.[A1]

[A1] Our tolerance for convergence is set to 1e-7.

Table A2: Average Sugar and Caloric Consumption from Beverages per Household, by Location and Income Status[a]

	Without Tax	**With Tax**	**Difference**	**% Change**
Philadelphia Households				
High-Income				
Sugar (g)	22,745	16,504	-6,241	-27.44%
Calories (Cal)	93,137	68,444	-24,693	-26.51%
Low-Income				
Sugar (g)	23,817	18,083	-5,734	-24.08%
Calories (Cal)	95,260	73,034	-22,226	-23.33%
Non-Philadelphia Households				
High-Income				
Sugar (g)	27,035	26,798	-237	-0.88%
Calories (Cal)	112,433	111,509	-924	-0.52%
Low-Income				
Sugar (g)	28,603	28,196	-407	-1.42%
Calories (Cal)	115,169	113,600	-1,569	-1.36%

[a] Aggregate amount over the post-taxation period January 2017 to December 2018.

A5 Changes in Sugar and Caloric Consumption for Households by Income Status

Table A2 reports changes in sugar and caloric consumption from beverages for households by home location and income status. We find that among Philadelphia households, high-income households on average consume less sugar and fewer calories and experience a greater percentage reduction in their consumption. The pattern is different for non-Philadelphia households, where high-income households on average consume less sugar and fewer calories but experience a smaller percentage reduction in their consumption. Differences in these outcomes in response to the taxation are best explained by Table 13, where we observe that, in terms of the total volume of Philadelphia and non-Philadelphia SBs consumed, low-income Philadelphia households are less responsive to the tax compared to their high-income counterparts, but the opposite is true for low-income non-Philadelphia households, who experience a larger reduction in SB consumption, and therefore a larger reduction in sugar and caloric consumption, compared to their high-income counterparts.

Table A3: Average Sugar and Caloric Consumption from Beverages per Household, by Location and Representation Status[a]

	Without Tax	**With Tax**	**Difference**	**% Change**
Philadelphia Households				
Non-URM				
Sugar (g)	19,564	13,502	-6,061	-30.99%
Calories (Cal)	79,807	55,790	-24,017	-30.09%
URM				
Sugar (g)	26,679	20,701	-5,978	-22.41%
Calories (Cal)	107,588	84,366	-23,222	-21.58%
Non-Philadelphia Households				
Non-URM				
Sugar (g)	25,658	25,433	-225	-0.88%
Calories (Cal)	106,449	105,572	-877	-0.82%
URM				
Sugar (g)	34,305	33,824	-481	-1.40%
Calories (Cal)	139,426	137,577	-1,849	-1.33%

[a] Aggregate amount over the post-taxation period January 2017 to December 2018.

A6 Changes in Sugar and Caloric Consumption for Households by Representation Status

Table A3 reports changes in sugar and caloric consumption from beverages for households by home location and representation status. We find that among Philadelphia households, non-URM households on average consume less sugar and fewer calories and experience a greater percentage reduction in their consumption. The pattern is different for non-Philadelphia households, where non-URM households on average consume less sugar and fewer calories but experience a smaller percentage reduction in their consumption. Differences in these outcomes in response to the taxation are best explained by Table 16, where we observe that, in terms of the total volume of Philadelphia and non-Philadelphia SBs consumed, URM Philadelphia households are less responsive to the tax compared to their non-URM counterparts, but the opposite is true for URM non-Philadelphia households, who experience a larger reduction in SB consumption, and therefore a larger reduction in sugar and caloric consumption, compared to their non-URM counterparts.