Market Heterogeneity and the Distributional Incidence of Soft-drink Taxes: Evidence from France

Fabrice Etilé, Sebastien Lecocq, Christine Boizot-Szantai

WORKING PAPER N° 2019 – 17

Market Heterogeneity and the Distributional Incidence of Soft-drink Taxes: Evidence from France

Fabrice Etilé
Sébastien Lecocq
Christine Boizot-Szantaï

JEL Codes: D12, E31, H22, I18
Keywords : Soft-drink tax; Nutrition; Tax incidence; Inequality; Market Structure; Consumer Price Index

Market Heterogeneity and the Distributional Incidence of Soft-drink Taxes: Evidence from France*

Fabrice Etilé[a,†], Sébastien Lecocq[b] and Christine Boizot-Szantaï[b]

[a]Paris School of Economics, F-75014 Paris, France, and INRA

[b]ALISS UR1303, INRA, Université Paris-Saclay, F-94200 Ivry-sur-Seine, France

26th March 2019

Abstract

Market heterogeneity may affect the distributional incidence of soft-drink taxes if households sort by income across markets with different characteristics. We use the Kantar Worldpanel homescan data to analyse the distributional incidence of the 2012 French soda tax on Exact Price Indices (EPIs) that measure consumer welfare from the price, availability and consumption of Sugar-Sweetened Beverages (SSBs) at a local market level. After correcting prices for consumer heterogeneity in preferences, we find that the soda tax had a significant but small national average impact corresponding to a pass-through of approximately 40%. Producers and retailers set significantly higher pass-throughs in low-income, less-competitive and smaller markets and for cheaper but less popular brands. Market heterogeneity ultimately has substantial distributional effects, as it accounts for approximately 35% of the difference in welfare variation between low- and high-income consumers.

Keywords: Soft-drink tax; Nutrition; Tax incidence; Inequality; Market Structure; Consumer Price Index.

JEL Classification: D12, E31, H22, I18.

* We thank Andrew Clark, Jessie Handbury, Miren Lafourcade, Lorenzo Rotunno, Judith Valls and seminar participants at Bocconi University (CERGAS), the EuHEA 2018 Conference, IRDES, Leeds University (AUHE), Monash University (CHE), Pompeu Fabra University (CRES), the Toulouse School of Economics (Food Seminar) and the University of York (YSHE). We gratefully acknowledge research support from the INRA Metaprogramme DiD'IT and the European Research Network DEDIPAC.

† Corresponding author. Paris School of Economics, 48, Boulevard Jourdan, 75014 Paris. fabrice.etile@psemail.eu.

1 Introduction

The worldwide rise in obesity and diabetes has prompted public health officials to devote particular attention to sugar intake from Sugar-Sweetened Beverages (SSBs). Taxing these beverages is considered a means of decreasing their consumption by increasing prices, at zero cost to public finances.[1] SSB taxes and, more generally, nutritional taxes are often criticised on the basis of their regressivity. As the poor tend to allocate a larger budget share to unhealthy food and beverages, they may mechanically face a higher tax burden.[2] However, the welfare consequences of any consumer tax depend not only on initial prices and quantities consumed but also on the incidence of the tax on consumer prices. A tax is unlikely to be shifted 1:1 into market prices due to changes in behaviour on the demand and supply sides of markets. The distributional impacts of a tax will thus depend not only on consumer preferences but also on the market characteristics driving producers and retailers' decisions. It is possible that low-income households face a higher tax burden partly because they are more likely to reside in markets with characteristics (*e.g.,* fewer retailers) conducive to higher pass-through of the tax to consumer prices.

The main purpose of the present study is therefore to demonstrate how heterogeneity in market characteristics contributes to the distributional impact of soft-drink taxes. We use homescan panel data to estimate the incidence of the French soda tax on soft-drink prices and consumer welfare, with a particular focus on heterogeneity across local markets. The French soda tax was passed in November 2011 and introduced on 1 of January 2012. Until 2014, it consisted of a unit excise tax of 0.0716 euro/litre on the producer price. It is levied on manufacturers or importers of SSBs (soft drinks and nectars) and Non-Calorically Sweetened Beverages (NCSBs). For space limitations, however, this study focusses on SSBs. We report additional results for NCSBs in the discussion section.

We examine the incidence of the tax on SSB prices using six years of nationally representative homescan data provided by Kantar Worldpanel (KWP) (2008-2013). This unique and detailed dataset covers 75% of SSB purchases in France and contains information on household purchases at the

[1]See the Harvard School of Public Health: https://www.hsph.harvard.edu/nutritionsource/sugary-drinks-fact-sheet/; and the World Health Organization: http://www.who.int/elena/titles/ssbs_childhood_obesity/en/. Epidemiological analyses clearly show that high SSB consumption is associated with greater risks of obesity and diabetes, especially for children (Malik, Pan, Willett, & Hu, 2013).

[2]While a large body of ex ante evaluation studies has focused on evaluating the potential aggregate health benefits of nutritional taxes, few have analysed their distributional impacts. Recent efforts using scanner data to analyse soft-drink taxes include Finkelstein et al. (2013), Wang (2015), Tiffin, Kehlbacher, and Salois (2015), Sharma, Hauck, Hollingsworth, and Siciliani (2014) and Etilé and Sharma (2015) for the U.S., U.K., and Australia. Madden (2015) and Tiffin and Salois (2014) use food expenditure surveys in Ireland and the UK to examine the distributional effects of revenue-neutral fiscal policies combining taxes on unhealthy food and subsidies for healthier food. As wealthy households spend relatively more on healthy food, such fiscal mixes tend to increase the relative burden on the poor. However, Madden (2015) show that they might be neutral with respect to poverty. Muller, Lacroix, Lusk, and Ruffieux (2017) validate these findings with incentivised framed field experiments, wherein subjects had to select an entire day's worth of food from a large set of food products, the prices of which varied substantially (±30%) across tax-subsidy treatments.

product level. Our dependent variable is a theoretically rigorous nested-CES price index that exactly measures variations in the utility from one unit of SSB consumption across local markets. This Exact Price Index (EPI) is tailored to provide measures of tax incidence that account for consumer substitution across products and for variations in their price and availability across locations and over time. It is constructed from local transaction prices and purchase quantities following recent methodological advances in trade and spatial economics (see, e.g., Handbury & Weinstein, 2014; Redding & Weinstein, 2016).

We have two motivations for working on a price index rather than on separate price series of product varieties. First, the welfare incidence of the tax depends essentially on household preferences for quantity and on the pass-through of the tax to the EPI for SSBs. For small taxes such as the French soda tax, welfare variations can be measured using compensating variation, which is approximately equal to the initial SSB quantity consumed times the variation in the price index. Hence, the distributional effects crucially depend on the tax incidence on the EPI, which varies across households as a function of their preferences for products and their place of residence. Second, the EPI can be adjusted for consumer and retailer heterogeneity to abstract from welfare changes reflecting variations in preferences for products and store formats across households within a population. To examine the tax incidence on SSB prices, we construct a global EPI with full adjustment for heterogeneity in household preferences, as well as separate EPIs for low- and high-income households. The difference in national tax incidence between the income group-specific EPIs allows us to test whether variations in preferences across income groups produce differences in the incidence on aggregate prices. The global EPI allows us to abstract from income-related preference heterogeneity to specifically identify the impact of market characteristics on tax incidence.[3]

We estimate the tax incidence with a before-after approach that controls for the rise in the cost of sugar after EU sugar quota policy was revised in October 2011. As other unobserved macro-shocks might still affect our estimates, we apply a difference-in-difference (DiD) design that uses changes in the EPI of water as a counterfactual. Our two identification strategies eventually produce the same results. Taking the before-after estimates, the tax increased the price of SSBs by approximately 4.1% on average, corresponding to a tax incidence of 39.1%. We find evidence that the average effect of the tax was similar for low- and high-income households. This indicates that heterogeneity in preferences for products and stores across income groups did not produce significant distributional effects. We then consider tax incidence across markets and find significant spatial heterogeneity. As expected, tax incidence decreases in retailer competition and market size. In addition, conditional on local competition, tax incidence is higher in low-income markets. Finally, using compensating variation, we find that and market heterogeneity accounts for at least 36% of the difference in yearly

[3]Further, the EPI is the relevant price statistic for evaluating the impact of the tax on sugar intake: following standard consumer theory, aggregate demand depends on aggregate price levels; previous research has shown that the crucial behavioural margin is not SSB quality but aggregate SSB quantity (Bonnet & Réquillart, 2013b).

welfare loss per capita between income groups. Sorting of households by income across markets is a significant determinant of the distributional effects of the tax.

Overall, this article complements the literature studying the distributional effects of nutritional taxes (Tiffin & Salois, 2014; Madden, 2015; Muller et al., 2017). Our findings emphasise the importance of accounting for supply-side reactions and market structure when simulating the impact of nutritional taxes and assessing their distributional impacts. We also revise downward previous estimates from ex ante and ex post studies, which both concluded that the French soda tax would have been over-shifted (Bonnet & Réquillart, 2013b; Berardi, Sevestre, Tepaut, & Vigneron, 2016). Our national average pass-through rate is similar to ex post estimates for the Berkeley soda tax, which resulted in incidence rates of between 22% and 47% (Falbe, Rojas, Grummon, & Madsen, 2015; Cawley & Frisvold, 2017).

The remainder of the paper is organised as follows. Section 2 describes the data, presents the nested-CES EPI for aggregate SSB consumption, and analyses its evolution over time. The details of the index construction appear in supplementary appendices. Section 3 sets out the identification strategies and examines the role of income-related preference heterogeneity. Section 4 analyses the heterogeneity of tax incidence across markets, as a function of their characteristics. Section 5 discusses the results, and Section 6 concludes the paper.

2 Derivation of local price indices

We construct local monthly price indices from homescan data collected by Kantar Worldpanel (KWP) over the 2008-2013 period. KWP follows a nationally representative sample of more than 21,000 French households, which record every purchase they make, including online purchases. We implement a methodology proposed by Handbury and Weinstein (2014) to derive market-level nested-CES EPIs for each of the following four product groups: SSBs, NCSBs, Unsweetened Beverages (USBs) and Water. This will allow us to compare price trends in the taxed (SSBs, NCSBs) vs. untaxed (USBs, Water) categories. The EPIs account for the spatial and time variability in prices and product availability, which are largely related to the heterogeneity in the distribution of retailers across markets. This will be used for identifying the heterogeneity of tax effects as a function of market characteristics and the tax.

2.1 Data

Each observation in the KWP data represents the purchase of a unique product variety in a particular store by a particular household on a given day. Households use handheld scanners to register the quantity, the expenditure, and the Universal Product Code (UPC) of the purchase or a set of product descriptors when there is no UPC. KWP has not provided us with the UPC but with a broad set of product attributes: flavour, brand, volume, type of packaging, type of beverage (family), whether it is carbonated, whether it is light, and whether it has been sweetened using caloric or non-caloric sweeteners. We use these attributes to define a set of 526 distinct products, belonging to 14 different families of beverages: colas, carbonated fruit drinks, non-carbonated fruit drinks, fruit nectars, lemonades, iced teas, tonics, energy drinks, flavoured water, natural water, fruit juices without added sugar, syrups (cordials/squash), pulps and milk-based fruit juices (for further information, see the Supplementary Appendix, Section A.1). Following Bonnet and Réquillart (2013b, 2013a), we also define 10 homogeneous categories of retailer stores according to the company name and the store format (hard discount, supermarket, hypermarket).[4] These two criteria are significant determinants of retailers' price-quality marketing mix.

We finally apply a three-tiered nomenclature to define and classify household purchases. In the upper tier, all purchases are sorted into one of the four following groups: SSBs, NCSBs, USBs, and Water. The middle tier consists of 81 brand-modules defined by interacting the four groups, the 14 beverage families and the brand names, *e.g.,* Coca-Cola Classic (group = SSBs, family = Colas, brand = Coca-Cola), Diet Coke (group = NCSB, family = Colas, brand = Coca-Cola). The lower tier consists of "artificial" UPCs, defined by the interaction of products with retailer categories (e.g., a 1-litre plastic bottle of Coca Cola Classic sold in a Carrefour hypermarket).

We end up with a total of 2,770 UPCs. Defining UPCs as product-retailer pairs captures that (*i*) the utility obtained from purchasing a product may vary from one store to another, as stores offer different levels of amenities; and (*ii*) beverage price and promotion policies are retailer-specific, as they are a means to attract or retain customers (Handbury & Weinstein, 2014; Bonnet & Réquillart, 2013b).

We define a local market as a "living zone" in a given month. The French National Statistics Office (INSEE) delineates a living zone as the smallest territory where inhabitants have access to everyday facilities and services, including stores.[5] From a retailer's perspective, these living zones represent consumer catchment areas. We assign each household to a living zone according to the city code of its residence. The purchase data are then matched by living zone to the TradeDimensions panel provided by Nielsen, and to INSEE census and fiscal data. The TradeDimensions panel provides exhaustive information about the presence of retailer stores in any given living zone in each month. These information will be used in Section 4 to characterise market heterogeneity in terms of local competition, affluence and size.

[4] The categories are: Auchan (Atac, Maximarché); Carrefour (Stock, Shopi, Proxi); Intermarché; Leclerc; a grouping of Casino (Monoprix, EcoService, PetitCasino, Spar, and Maxicoop), Cora, U and others (cheesemongers, grocery stores); subsidiary hard discount stores (Ed-Dia, Franprix, Leader Price); and independent hard discount (Lidl, Aldi).

[5] "Bassin de vie" in French; see https://www.insee.fr/en/metadonnees/definition/c2060.

2.2 Sample selection and characteristics of the national market

To ensure the statistical representativeness of prices, we retain living zones where at least 10 households are observed each year over the whole period. This leaves us with 263 living zones, out of a total of 1,633. Although we lose rural areas, this selection does not alter the distribution of other household characteristics (Supplementary Appendix A.2). We also select the 1,891 UPCs that are purchased at least 100 times over the 2008-2013 period and retain from these the 995 UPCs that are purchased at least once in each of the 72 months. There are 400 UPCs in the SSB group, 127 in the NCSB group, 338 in the USB group, and 130 in Water. Our final sample therefore consists of 30,254 distinct households over the six-year period (roughly 15,000 households are observed each year) and over four million purchases. We observe at least 35 households in 90% of the living zones over the period, and the median number of households per local market (living zone $\times$ month) is 100. For each UPC, household, month and retailer, we calculate the mean expenditure and mean quantity. Dividing mean expenditures by mean quantities produces mean unit prices that we further deflate by the general French Consumer Price Index (CPI).

Table 1 reports selected market statistics by beverage family, for each of the four groups. The last line indicates that SSBs represent 25.9% of the total volume of non-alcoholic beverages purchased for at-home consumption in France. This is much larger than the NCSB figure (only 8.3%) but smaller than that for USBs and Water (34.7% and 31.0%, respectively). Colas are dominant in the SSB and NCSB groups but face many competitors in the SSB category. Table 1 also shows the average unit value in each segment. Interestingly, there is not a particularly large price premium for NCSB products compared to SSB products within the same beverage family. The average unit value of non-calorically sweetened colas is even lower than that of sugar-sweetened colas.

Table 1: Beverage groups – Descriptive statistics

	SSB			NCSB			USB			Water		
	UPC #	Market share	Unit value Mean (SD)	UPC #	Market share	Unit value Mean (SD)	UPC #	Market share	Unit value Mean (SD)	UPC #	Market share	Unit value Mean (SD)
Colas	61	11.51	0.97 (0.49)	67	6.47	0.93 (0.46)						
Carbonated fruit drinks	73	3.82	1.13 (0.82)	24	0.69	1.05 (0.34)						
Non-carbonated fruit drinks	63	3.18	0.98 (0.45)									
Nectars	64	3.18	1.26 (0.60)	5	0.20	1.34 (0.65)						
Lemonades	40	1.11	0.57 (0.53)	5	0.09	0.59 (0.23)						
Iced teas	41	1.56	0.76 (0.39)	8	0.13	0.84 (0.30)						
Tonics	28	0.72	1.04 (0.57)	3	0.04	1.15 (0.07)						
Energy drinks	12	0.33	2.88 (1.66)									
Flavoured water	18	0.51	0.89 (0.39)	15	0.67	0.96 (0.17)				13	0.32	0.79 (0.25)
Natural water										117	30.72	0.37 (0.26)
Juices (no added sugar)							221	29.00	1.51 (0.91)			
Syrups							94	4.57	2.86 (2.73)			
Pulps							13	0.68	3.56 (0.56)			
Milk-based fruit juices							10	0.50	1.96 (0.35)			
Total	400	25.92	1.02 (0.66)	127	8.29	0.95 (0.43)	338	34.75	1.77 (1.47)	130	31.04	0.38 (0.27)

Notes: Kantar Worldpanel data 2008-2013. Unit values are deflated by the Consumer Price Index for consumer goods (Base: 2011) and are expressed in euros/litre. Market shares are defined by the volume of transactions over total non-alcoholic beverage transactions observed in the estimation sample (weighted by household sample weights).

2.3 Exact Price Indices for aggregate SSB consumption

An Exact Price Index (EPI) measures the change in expenditure required to hold utility constant as the prices of product varieties vary. It is therefore an index of consumer welfare. This price index can be formally defined for a representative household of population $\mathcal{P}$, product group g and supply available in market c as

$$EPI_{gc}^{\mathcal{P}} = \frac{C(V, \boldsymbol{p_{gc}}; \mathcal{P})}{C(V, \boldsymbol{p_{g\mathcal{R}}}; \mathcal{P})}, \tag{1}$$

where $C(V, \boldsymbol{p_{gc}}; \mathcal{P})$ is the cost of attaining utility V for a household that is endowed with representative preferences of $\mathcal{P}$ and faces prices $\boldsymbol{p_{gc}}$; $\boldsymbol{p_{g\mathcal{R}}}$ is a vector of reference prices.

The construction of any EPI relies on structural assumptions regarding the household choice problem to adjust for preference heterogeneity in the population, substitutions across products, and variations in product price and availability (Triplett, 2001). Following Handbury and Weinstein (2014), we assume the weak separability of non-alcoholic beverages from other food and beverages, and we impose nested-CES preferences for consumer utility over brand-modules and UPCs. Households allocate their non-alcoholic beverage budget among the four beverage groups g, then among the "brand-modules" b within each beverage group (e.g., Coca-Cola Classic, Pepsi-Cola (regular)). Finally, it is divided among UPCs u within each brand-module: UPCs are products purchased from a specific retailer. This multi-stage budgeting process thus mirrors the three-tiered nomenclature of purchases presented in Section 2.1.

As UPCs are not all available in every market, the EPI for g is the product of a "conventional" nested-CES Exact Price Index (CEPI) and an adjustment coefficient for Variety Availability (VA):

$$EPI_{gc}^{\mathcal{P}} = CEPI_{gc}^{\mathcal{P}} VA_{gc}^{\mathcal{P}}. \tag{2}$$

$CEPI_{gc}^{\mathcal{P}}$ is the EPI obtained under the assumption that the choice set in every market c is the same as that in the reference market $\mathcal{R}$ chosen to calculate the reference prices. $VA_{gc}^{\mathcal{P}}$ is an adjustment for differences in the available choice sets between markets c and $\mathcal{R}$. We here define the reference market as the "national market" (*i.e.,* the union of all living zones) in 2011, the pre-tax year. The Supplementary Appendix B details the formula and its derivation. The Supplementary Appendix C provides exhaustive details on the construction of the EPI.

We now explain the EPI in intuitive terms. The conventional price index $CEPI_{gc}^{\mathcal{P}}$ is a sales-weighted average of the local prices of products purchased by households of population $\mathcal{P}$ living in c. Any rise in the price of a UPC increases the CEPI. However, since more popular products have larger market shares, they also have higher weights in the CEPI and larger impacts on consumer welfare. The CEPI is therefore adjusted for consumer preferences over products and for conventional substitution effects. The variety-adjustment term $VA_{gc}^{\mathcal{P}}$ is determined by the local availability of products and their popularity in population $\mathcal{P}$ at the national level. The availability of products will vary across markets as a function of the localisation of retailers, and with entries and exits of products. The loss of welfare due to locally missing varieties translates into a higher price index. The welfare loss is unimportant for varieties that have a very small share of the national market, since they are not very popular among consumers. The welfare loss from a lack of variety also decreases with an increase in the elasticities of substitution across brand-modules and across products.

Since we wish to identify the specific role of market heterogeneity in tax incidence, we need to account for the impact of within-market preference heterogeneity on observed prices and sales. In a given market c, the observed unit prices for a UPC are likely to vary from one household to another for two reasons. First, households choose to shop in specific stores, which may differ in terms of amenities. Stores adjust their prices as a function of the amenities they provide. Second, stores also adjust their prices as a function of customer demand and characteristics. In addition, households may differ in, among other characteristics, their shopping behaviour and sensitivity to sales promotions. To abstract from preference heterogeneity, we follow Handbury and Weinstein (2014) by constructing the CEPI and VA terms from unit prices and shares adjusted for within-market variations in household and retailer heterogeneity. This has two desirable implications. First, the EPI will measure spatial and time variations in the welfare of a representative consumer endowed with identical preferences and shopping in homogeneous stores. These variations will be caused primarily by shocks to production, logistic and retailing costs, and variations in market structure. Second, it makes the UPC and brand-modules homogeneous, in terms of subjective quality. This renders plausible the assumption of constant elasticities of substitution, which underlies the use of CES preferences.

We construct a global EPI for the entire household population and specific EPIs for low- and high-income households. To define the income groups, we consider an equal division of the population using the median real household equivalent income, *i.e.,* adjusting household income for inflation (via the CPI) and units of consumption (via the OECD scale). We employ an extensive list of variables to adjust for household and retailer heterogeneity within these three populations: household equivalent income, age and gender of the main shopper, household structure, education, type of residential area, some interactions between income and product characteristics, and the name and format of the retailer. The global EPI thus measures the welfare variations of the representative French household across space and time. We leverage these variations to identify the impact of market characteristics on tax incidence. We use the income group-specific EPIs to compare the average tax incidence between low- and high-income households. This will reveal the importance of income-related preference heterogeneity in the distributional effects of the tax.

Figure 1 presents the evolution of the average global EPI for the four product categories. For the four groups, the EPI shows a slight increase up to mid-2009, followed by a decline until 2012. There is then a steep increase for SSBs, NCSBs and USBs (*i.e.,* all soft drinks) in 2012-2013, while the price of Water fell. Interestingly, the absence of a steep price increase before January 2012—the month that the tax was implemented— shows that producers and retailers did not pass the tax on to consumers in advance, although the soda-tax project was announced in late August.[6]

Figure 2 illustrates the spatial heterogeneity in prices, which motivates our focus on market heterogeneity. It displays the histogram of the EPI (left panel) and a quantile-quantile (Q-Q) dot plot comparison of the EPI and the CEPI (right panel) for SSBs in 2013. The histograms demonstrate the importance of spatial price variations, despite that the prices have been adjusted for retailer and consumer heterogeneity. In the Q-Q plots, the departures from the 45-degree line indicate the effect of VA on local prices. The dispersion of dots is explained by the spread in the distribution of VA: local prices can be up to 50% higher in some areas due to the absence of subsets of products. The comparison of the local EPI and local CEPI, ranked by percentiles, shows that the VA factor substantially affects the price ranking of markets, as shown by the dots far from the 45-degree line. The VA factor also inflates heterogeneity in local prices, as the distribution of the CEPI has a standard deviation of 11.7% vs. 15.9% for the EPI.

[6]A simple event analysis reveals that SSB prices in August, September, October and November were on average 1.7% higher, 1.7% lower, 0.3% lower and 0.2% higher, respectively, than in December. These differences are not significant in October and November. This lack of anticipation can be explained by the existence of annual contracts between manufacturers and retailers (renewed in February-March) and by the uncertainty surrounding the legislative process, as the tax was eventually adopted in Parliament on 21 December 2011, after intense lobbying and debate (Le Bodo, Etilé, Gagnon, & de Wals, 2017). See https://lexpansion.lexpress.fr/actualite-economique/taxe-sodas-light-comment-coca-cola-a-perdu-la-bataille_1440607.html and https://www.legifrance.gouv.fr/, "LOI n°2011-1906 du 21 décembre 2011 de financement de la sécurité sociale pour 2012".

Figure 1: Exact Price Index, monthly average, 2008-2013

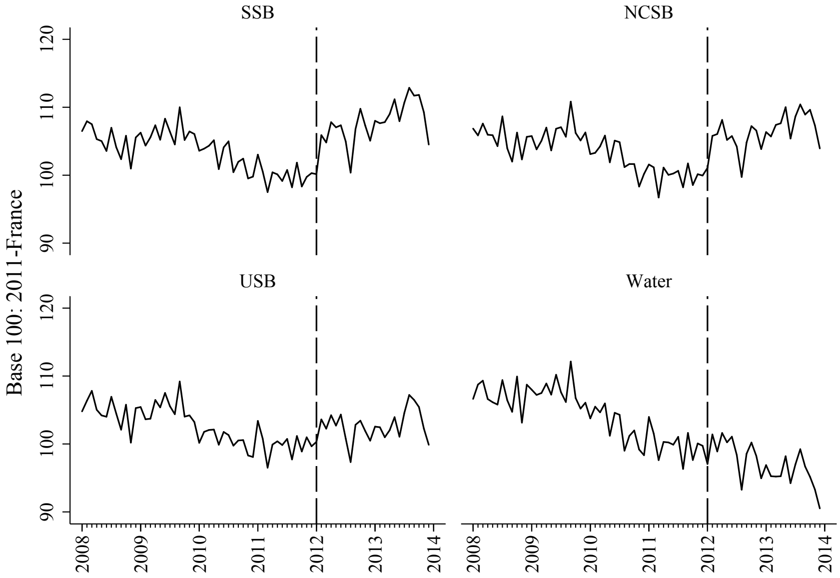

Notes: Kantar Worldpanel data 2008-2013. This figure shows the changes over time in the national Exact Price Indices (EPIs) of the four product categories. The national EPI is a weighted average of local EPIs for the general population, where the weights adjust for the share of living zones in national sales in 2011. The reference market is the union of all markets in 2011.

Figure 2: Spatial distribution of SSB prices in 2013

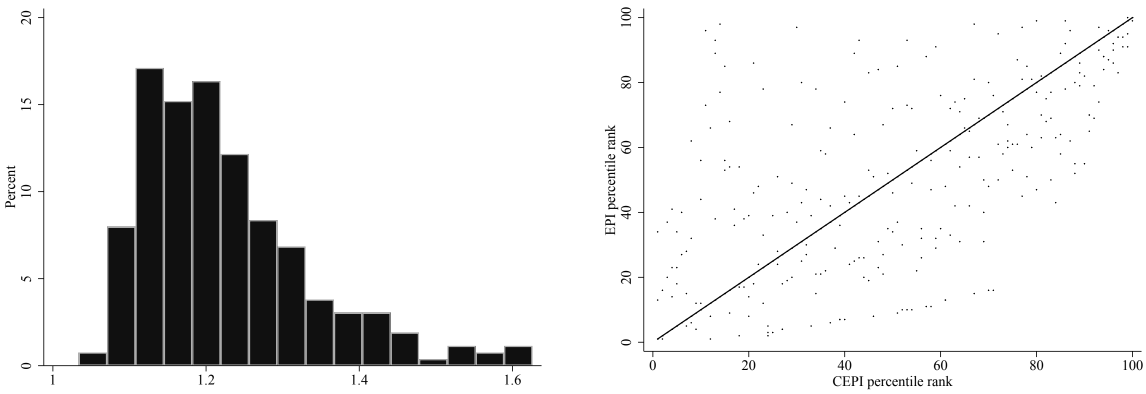

Distribution of EPI

EPI vs. CEPI across markets (Q-Q plot)

Notes: Kantar Worldpanel data 2013. This Figure shows for SSBs (i) the distribution of the Exact Price Index (EPI) in the left panel and (ii) a quantile-quantile plot of the CEPI against the EPI in the right panel, with departures from the 45-degree line indicating the effect of VA on prices. Average yearly statistics are calculated for each living zone.

3 Tax incidence and preference heterogeneity across income groups

We first estimate the average incidence of the tax at a national level. The comparison of results for the global EPI and for the income group-specific EPIs allows us to assess the importance of income-related preferences in tax incidence. We begin by describing the empirical design that we exploit to obtain these baseline results.

3.1 Empirical design

We identify the tax incidence at the national level by examining year-to-year changes. This is similar to empirical pass-through specifications that regress the annual change in prices on the annual changes in costs and identify the pass-through by comparing the change in prices to the change in costs across equilibrium situations (Hong & Li, 2017; Amiti, Itskhoki, & Konings, 2014).

We have two empirical strategies. Given the absence of anticipated responses on the supply side, we first carry out a before-after estimation, which compares the average 2012 price to that in 2011. The following equation is estimated on our sample of local EPIs, where each observation is the local price of SSBs observed in market c (living zone a × month t)

$$\ln(\boldsymbol{P}_{SSB,c}) = \delta Post_{t\geq 2012} + \delta_y + \delta_m + \gamma C_t + \delta_a + \epsilon_{SSB,c}. \tag{3}$$

In this equation, the before-after estimate of the tax effect is given by δ. The equation compares the average EPI in 2012 (after: $Post_{t\geq 2012}$) to that in 2011 (before), adjusting for year effects (δ_y: 2008, 2009, 2010 and 2013), month effects (δ_m), which are restricted to be the same in all years (allowing for year-specific month effects yields the same results), living-zone fixed effects (δ_a), and input costs C_t. We add living zone-specific fixed-effects to increase statistical efficiency.

Input costs may have played a role in the evolution of SSB prices over the period. In particular, the Producer Price Index for sugar increased following the revision of the EU sugar quota policy in September 2011, which was politically unrelated to the soda tax (Supplementary Appendix A.3). As sugar is an important input in the production of soft drinks, this shock is a potential confounding factor in the evaluation of tax incidence. We therefore control for the cost of sugar in C_t. The main identifying assumption is then that the remaining variation is entirely attributable to the tax.

The before-after regression results may be driven by movements in other supply-side costs. We cannot add more input prices, as this produces considerable multicollinearity in the regressions.[7] We cannot control for other manufacturing and retailing costs. Therefore, as a robustness check, we

[7] We have time series on many inputs, such as oil, metal, plastic, glass, paper, electricity, gas, and sugar. When we include some of them in C_t, the associated Variance Inflation Factors are over 20, which is clear evidence of multicollinearity. One likely explanation is that most input costs are indexed on the price of oil and/or are driven by similar macroeconomic shocks. Introducing non-linear functions of sugar costs (higher-order polynomials or piece-wise functions) also produces a considerable amount of multicollinearity, with exploding Variance Inflation Factors.

adopt a DiD approach, with Water as the control, and focus on the 2010-2012 period to ensure that the common-trend assumption holds. We estimate the following model for the comparison between SSBs and Water (group-index g)

$$\ln\left(\boldsymbol{P}_{gc}\right) = \delta Post_{t\geq 2012} + \delta_{SSB} + \delta D_{SSB,t\geq 2012} + \delta_y + \delta_m + \delta_{SSB,m} + \gamma\delta_{SSB} \times C_{t,sugar} + \delta_a + \epsilon_{gc} \tag{4}$$

where the tax effect is given by the coefficient δ, $D_{SSB,t\geq 2012}$ is a dummy for SSB prices observed after December 2011, δ_{SSB} is an SSB fixed effect, $\delta_{SSB,m}$ are group-specific month effects for differences in seasonality between SSB and Water consumption. We control for the cost of sugar for SSBs only, since sugar is not an input for Water, and we want to avoid multicollinearity problems.

We choose Water as the control group for four reasons. First, Water was obviously not targeted by the soda tax. Second, apart from sugar, the inputs used in the supply of Water are similar to those for SSBs, and they are also similar in terms of cost structure: plastic, glass and aluminium for packaging; natural water; and marketing, logistic and retailing costs. Third, while the companies owning SSBs (and NCSBs) have some important USB brands, they have zero or very small market shares for Water.[8] This limits any firm strategic reactions producing changes in the supply price of Water. Fourth, we estimated an AIDS demand system for the four groups of beverages to identify the price substitutions across the four markets. Our results show that the market for Water is largely disconnected from that for soft drinks (Supplementary Appendix D.1).

To check whether the common-trends assumption holds in the pre-policy period, Figure 3 plots the annual average of the EPI, compared to Water. As in all of our results, each observation is weighted by the share of national sales in each local market in 2011. Although the trends in SSB and Water prices differ slightly before 2010, the common-trends assumption holds for 2010-2011.

3.2 Results

The upper panel of Table 2 presents the baseline results, which are obtained with EPI constructed from the full sample. The observations are weighted by the share of national sales in the living zone in 2011. The estimates thus represent average welfare variations for a representative French household. Column (1) displays the estimation of a before-after specification with month and living-zone fixed effects (δ_m and δ_a in equation 3). The estimated tax impact on the EPI is significant at the 1% level. The average price of SSBs in 2012 was approximately 5.4% higher than in 2011. Column (2) shows that this impact is smaller when we additionally control for the cost of sugar, as it declines to 4.1%. This adjustment is in line with available evidence regarding the pass-through of variations in sugar

[8] Coca-Cola, PepsiCo and Orangina-Suntory are the main owners of the national SSB brands. PepsiCo owns Tropicana, which is the leading national brand in the USB market. Danone and Nestlé own the most popular national brands of Water.

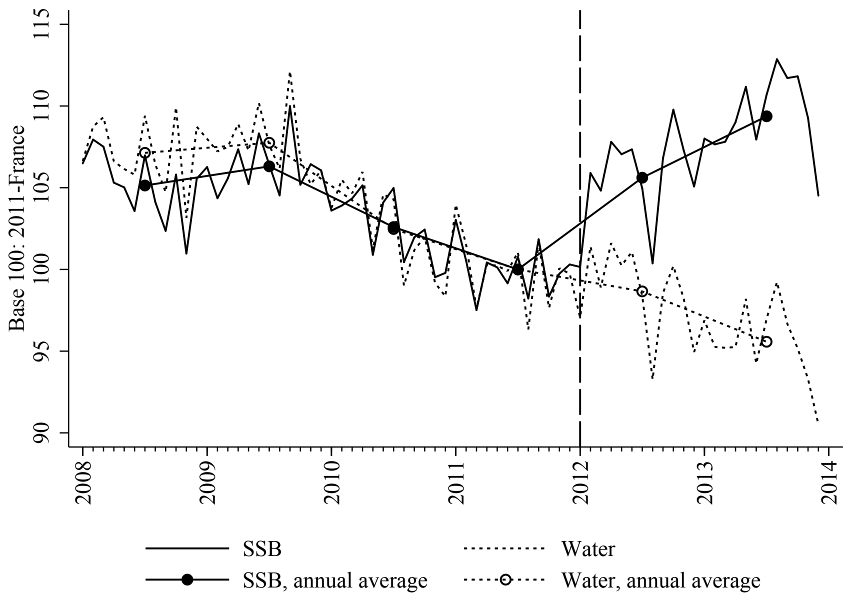

Figure 3: Common trends, difference-in-difference estimation

Notes: Kantar Worldpanel data 2008-2013. Each point represents the value of the average EPI in a given year, while the oscillating lines shows the monthly variations in the indices around their yearly trends. Each average price figure is calculated by taking the weighted mean of local values, using market sales as weights.

prices into consumer SSB prices in France.[9]

The third column of Table 2 reports the DiD estimates. The estimated impact for SSBs is very similar to that from the before-after estimation: 4.2% vs. 4.1%. Column (4) in Table 2 provides a very conservative test of the common-trends assumption, using a placebo policy change on 1 January 2011 (one year before). The estimated placebo impact for SSBs, although significant, is more than seven times smaller.[10]

Tax incidence is likely to vary across income groups, depending on consumer preferences and residential sorting of households across local markets with varying characteristics. To uncover specifically

[9]Using an empirical IO model, Bonnet and Réquillart (2011) finds that a 36% decrease in the sugar price leads to an average decrease in SSB prices by 3.4% (see their simulations). This corresponds to an elasticity of approximately 0.1, which implies that the increase in sugar prices between 2011 and 2012 (approximately +18%) would correspond to a +1.8% increase in consumer prices. The difference between the estimates in columns (1) and (2) provides a close result ($5.426\% - 4.144\% = 1.283\%$)

[10]More generally, taking any placebo date before January 2012 for the implementation of the tax produces an estimated impact that is much lower than the estimate in column (3). This can readily be seen in Figure 3. We also constructed a placebo distribution treatment by permuting SSBs (treated product) and Water (control) in randomly drawn living zones. This permutation procedure assesses the uncertainty regarding the absence of policy effect for Water. The DiD effect in column (3) is significantly higher than the placebo effects at any significance level.

Table 2: Tax incidence: Price-index variation (% points) – national average

	Before-after		DiD	
Population	(1)	(2)	2012	2011
All households	5.426***	4.144***	4.261***	0.570**
	(0.171)	(0.232)	(0.402)	(0.214)
Low-income households	5.965***	4.614***	4.273***	0.341
	(0.185)	(0.485)	(0.578)	(0.315)
High-income households	4.989***	3.600***	4.482***	0.402
	(0.204)	(0.346)	(0.549)	(0.289)
Differential incidence (Low Inc. - High Inc.)	0.920***	1.047	-0.053	-0.269
	(0.314)	(0.590)	(0.806)	(0.430)
Additional controls				
Sugar price (in log) for SSBs and NCSBs	No	Yes	Yes	Yes
Group-specific month effects	No	No	Yes	Yes
Period	2008-2013	2008-2013	2010-2012	2009-2011

Notes: The dependent variable is the log of EPI or, for differential incidence, the log-difference in EPI between low- and high-income households. The EPI is estimated from Kantar Worldpanel data 2008-2013 using market-level observations (living zone-month). These estimates represent changes in the EPI, in % points, between 2011 and 2012 (before-after columns: δ in Equation (3)) and the difference in the changes between SSBs/NCSBs and Water (DiD: difference-in-difference columns: δ in Equation (4)). The DiD-2011 column is a placebo test, focussing on the 2010-2011 change. Living zone fixed effects are included. Each observation is weighted by the population-specific share of national sales in the living zone in 2011. For estimating the differential incidence, the weights are the share of national sales of low-income households. Standard errors are in parentheses; ***, ** and * indicate significance at the 1%, 5% and 10% levels.

the role of income-related preferences, we examine the differences between low- and high-income households at the national level. We use local EPIs for each income group as defined in Section 2.3. As the EPIs are still corrected for within-group consumer and retailer heterogeneity, they measure the welfare variations of representative households in the low- and high-income groups. The observations are weighted by the income group-specific share of national sales in the living zone in 2011.

The results appear in the second panel of Table 2. The tax incidence is slightly higher for low-income households in the before-after estimates (4.6% vs. 3.6% for high-income households in specification (2)). However, the DiD estimates in column 3 show minor differences only (+4.5% for the high-income vs. +4.3% for the low-income households). This indicates that income-related preference heterogeneity did not cause low-income households to be significantly more impacted by the tax than high-income households. These national-level results might still be driven by residential sorting by income across living zones, because the regression weights depend on the purchase volume of each income group in each market in 2011. Therefore, the last line of the lower panel estimates the differential in tax incidence between low- and high-income households, by replacing the log of EPI with the log-difference in EPI observed in each market between low- and high-income households. Each observation is weighted by the share of national sales observed for low-income households in the living zone in 2011. The estimated differential incidence thus reflects only the role of income-related preference heterogeneity. The DiD results confirm that there are no significant differences between the two income groups. Socio-economic variations in household preferences over product varieties do not produce important distributional effects.

4 Tax incidence and heterogeneity across markets

We now analyse the impact of local market characteristics to reveal their contributions to the distributional effects of the tax. We work with the global EPI, which is constructed from the full sample of households. We can isolate the effect of local cost and market structures because we retained market fixed effects in our quality-adjusted prices, while we purged the average effects of retailer and consumer heterogeneity. The local variations in EPI therefore reflect neither preference heterogeneity nor national-level variations in costs or strategies across retailers, nor their interactions with retailer localisation.

4.1 Market structure: affluence, size and competition

We first analyse the association between tax incidence and market size and affluence, with the former proxied by the number of consumption units and the latter by median fiscal income.[11] We then add an indicator for the degree of local-market competition. It is measured by the Herfindahl-Hirschman Index (HHI) of sales area per capita, dichotomised using a threshold of $2,000$ (a value of $10,000$ corresponds to a monopoly). This threshold is used by the European Commission to reflect a lack of horizontal competition (Official Journal C 31 of 05/02/2004). As the competition variable is strongly correlated with market size, we cannot simultaneously control for market size and competition. As the DiD analysis confirmed the before-after results, we implement the before-after strategy.[12] Interaction effects between $Post_{t\geq 2012}$ and the market-level characteristics are added to equation (3). As we focus on market heterogeneity, we do not weight market-level observations in the regressions.

Table 3 reports the results. Specification 1 replicates our earlier estimates, except that now all markets are given the same weight. The key estimates are thus slightly different, 4.94% against 4.14% in column 2 of Table 2. Specification 2 adds the interactions between $Post_{t\geq 2012}$ and the logarithm of median income and the number of consumption units. Median income positively affects price (+8.90 percentage points), and it has a negative effect on tax incidence. We have centred the log-income variable on its mean, so that the estimated coefficient (-3.08) implies that the tax incidence

[11] These variables are provided in the annual Census and fiscal data from INSEE. Postcode-level statistics were aggregated to the level of living zones using population weights.

[12] In addition, introducing the market variables and their interactions with $Post_{t\geq 2012}$ in the DiD analysis produces multicollinearity problems.

is approximately 25% lower when median income is 50% above average. The tax reduced the price gap between less-affluent and more-affluent markets: it is thus regressive. The coefficient on N_{cu} shows that doubling the market size also reduces tax incidence by 0.23 percentage points, *i.e.,* 5% of the baseline effect. This is in line with the theoretical predictions from New Economic Geography models that competition should be more intense in larger markets, reducing the possibilities of firm mark-up adjustments (see Handbury & Weinstein, 2014).

Specifications (3) to (5) specifically address the role of competition. In specification (3), the HHI dummy is positively correlated with prices (+1.14 percentage points in concentrated markets). It also has a large, significant and positive effect on tax incidence, which is approximately 12% higher in concentrated markets (0.556/4.799). Specifications (4) and (5) compare the impact of market competition between the full sample and the sub-sample of living zones with median fiscal incomes under the national median figure. Average tax incidence is 12% higher in low-income markets (5.555/4.954). We also find a stronger effect of competition in poorer areas. In low-income, high-HHI markets, tax incidence is 20% higher (1.130/5.555) than in low-income low-HHI markets, and 35% higher than the national average. In other words, tax incidence is similar in non-competitive, high-income and competitive, low-income markets. This illustrates that taking competition into account can significantly moderate our conclusions regarding the distributional effects of taxes.

4.2 Price-setting vs. assortment strategies

The spatial heterogeneity in retailer aggregate price responses to taxes is driven by their price-setting and choice of assortments, *i.e.,* the number of products they offer to consumers. We investigate these mechanisms separately in Table 4, which reports the effect of the tax on the conventional EPI (CEPI: upper panel) and the variability-adjustment factor (VA: right panel), for specifications (1), (3) and (5). The comparison of the estimates for CEPI and VA reveals that it is the former rather than the latter that drives heterogeneity and the level of the tax incidence. The CEPI increased more in poorer areas, and competition significantly reduces the tax burden for consumers, especially in poorer areas (Table 4, left panel).

Regarding variety adjustment, we find some evidence of interaction effects between the tax and affluence and market competition. At the baseline, a lack of horizontal competition increases the impact of variety adjustment: VA is 0.80 percentage points higher for SSBs in high-HHI markets. In Specification (3), income also has a negative effect on VA, albeit not significant, suggesting that wealthier areas tend to have access to more varieties. The effect of the tax on VA does not seem to vary with competition. However, when we focus on low-income areas (specification (5)), the competition effect becomes significant for SSBs: the tax incidence is +0.26 percentage points higher in high-HHI areas. One potential explanation is that, to adjust to the tax, retailers have changed their SSB assortments to reduce price competition (Hamilton, 2009). To test this argument, we

Table 3: The heterogeneity of tax incidence across markets (% points)

	(1)	(2)	(3)	(4)	(5)
Post	4.940***	4.722***	4.799***	4.954***	5.555***
	(0.387)	(0.457)	(0.449)	(0.387)	(0.561)
$\times \ln(Income)$		-3.077*	-3.932**		
		(1.761)	(1.656)		
$\times \ln(N_{cu})$		-0.233*			
		(0.122)			
$\times 1_{HHI>2000}$			0.556**	0.551**	1.130**
			(0.249)	(0.244)	(0.488)
$\ln(Income)$		8.901*	8.614*		
		(5.014)	(4.895)		
$\ln(N_{cu})$		-2.727			
		(4.932)			
$1_{HHI>2000}$			1.142***	1.148***	1.234**
			(0.322)	(0.321)	(0.492)
Sample	Full	Full	Full	Full	Inc<Q(50)

Notes: The dependent variable is the log of EPI. The EPI is estimated from Kantar Worldpanel data 2008-2013. Estimated impacts in % points from a before-after specification (observations are *not* weighted by market-specific sales). N_{cu}: number of consumption units (cu) in each market (INSEE census data). $Income$: market average of the median real equivalent income in the market's postcodes (INSEE fiscal data). HHI is a Herfindahl-Hirschman index based on the sales area of retailers (TradeDimensions data). An HHI greater than 2000 reflects horizontal competition concerns. All of these variables vary across markets c, *i.e.,* across living zones a and periods t. All estimates include area, month and year fixed effects. Full sample: $N = 18,927$ living zone-month observations. The sample Inc<Q(50) contains only markets where the median income is below the median figure ($N = 9,466$ living zone-month observations). Standard errors are clustered at the area-level in parentheses; ***, ** and * indicate significance at the 1%, 5% and 10% levels.

compared the number of UPCs available before and after the tax. It turns out that retailers have *not* changed the breadth of their product varieties on offer.

The moderating effect of competition on changes in variety adjustment can then only be explained by changes in the national share of UPCs available in each specific market, *i.e.,* changes in the popularity of products among consumers. The estimated VA effects thus reflect a decline in the popularity of some UPCs that were specifically available in less-competitive, low-income markets. Tax incidence differs not only across markets but also across UPCs, as different retailers have different degrees of bargaining power with respect to producers. We use the before-after design of Section 3.1 to estimate national average pass-through rates for UPCs (See Supplementary Appendix D.2, Table D.2). The pass-through rates are, on average, lower for the UPCs corresponding to top national brands (19.2%) than for other national, retailer and hard-discount products (48.5%, 47.4%, and

38.5%). These higher pass-through rates eventually contributed to the spatial variations in tax incidence, as hard-discount and retailer brands are more likely to be available and purchased in low-income markets. Products with higher pass-through rates lost market shares, and this decline in popularity explains our estimates of the effect of the tax on VA in low-income markets.

Table 4: The heterogeneity of tax incidence across markets – CEPI and VA (% points)

	CEPI			VA		
	(1)	(3)	(5)	(1)	(3)	(5)
Post	4.956***	4.739***	5.331***	-0.016	0.060	0.224
	(0.370)	(0.430)	(0.537)	(0.108)	(0.125)	(0.141)
$\times \ln(Income)$		-2.960*			-0.973**	
		(1.584)			(0.459)	
$\times 1_{HHI>2000}$		0.534**	0.871*		0.021	0.259**
		(0.238)	(0.467)		(0.069)	(0.123)
$\ln(Income)$		9.410**			-0.796	
		(4.682)			(1.357)	
$1_{HHI>2000}$		0.341	0.688		0.801***	0.546***
		(0.308)	(0.471)		(0.089)	(0.124)
Sample	Full	Full	Inc<Q(50)	Full	Full	Inc<Q(50)

Notes: The dependent variable is the log of CEPI or VA. CEPI and VA are estimated from Kantar Worldpanel data 2008-2013. Other details are as in Table 3. Standard errors are clustered at the area level in parentheses; ***, ** and * indicate significance at the 1%, 5% and 10% levels.

5 Discussion

This study provides evidence on the role of market structure in the distributional incidence of the 2012 French soft-drink tax. The tax incidence was 12% higher than the average in low-income markets and 35% higher in low-income less-competitive markets (Table 3). To illustrate the magnitude of this market effect, we use compensating variation measures of welfare loss by income group. Low-income households lost on average 0.52€/capita/year as against 0.34€/capita/year for high-income households. The difference (0.18€/year) can be decomposed into a market effect that reflects residential sorting by income across markets (we control for heterogeneity in preferences) and a preference effect that is produced by differences in preferences for quantity and quality between low- and high-income households living in the same market. This exercise reveals that market heterogeneity accounts for 35.8% of the difference in welfare loss between income groups (Supplementary Appendix, D.3, Table D.3). Residential sorting across markets is an important determinant of the distributional effects of soft-drink taxes.

The market heterogeneity in tax incidence is essentially explained by variations in product-level pass-through rates, differences in product market shares across markets, and spatial variations in pre-tax market structure. Conditional on local competition, initial prices were lower in low-income markets, but the tax incidence is ultimately higher. This effect does not reflect differences in household preferences across markets, as the price indices are adjusted for product availability and consumer and retailer heterogeneity. In our set-up, there is a single demand curve corresponding to a representative consumer. The likely explanation is therefore that markets were initially at different equilibrium positions along the demand curve because retailers face higher operating and rental costs in more affluent markets.[13] As the supply curve is shifted by the cost differential, markets differ in the slope and curvature of the demand initially faced by retailers. These two characteristics of the demand curve are crucial determinants of the pass-through of taxes to consumer prices (see, e.g., Weyl & Fabinger, 2013).[14]

Our results also reveal that the soda tax increased the prices of SSBs by approximately 4.1% on average at the national level. On the basis of the estimates, we can calculate the aggregate pass-through of the tax into the EPI, which provides a measure of the distribution of the tax burden between consumers and suppliers (See Supplementary Appendix D.4). The estimated national pass-through to consumer prices is 39.1% for SSBs. This is higher than the average product-level pass-through, because the EPI assigns relatively less importance to the pass-through of national brands. As the quality of national brands is higher than that of retailer and hard-discount brands, their quality-adjusted market shares are lower than their observed market shares. Hence, after adjusting for consumer and retailer heterogeneity, the aggregate tax incidence gives more weight to the pass-through rates of retailer and hard-discount brands.

At the national level, the average impact of the tax on consumer prices does not differ between low- and high-income households. This result demonstrates that income-related heterogeneity in the preference for quality is not an essential driver of SSB tax regressivity. The negative welfare impact of the tax was indeed larger for a wealthy household living in a poor neighbourhood with few retailers than for a poor household living in a wealthy neighbourhood with many retailers. However, income-related heterogeneity in the preference for quantity plays an important role, as demonstrated by our decomposition of welfare losses (Supplementary Appendix, Table D.3). For reasons of statistical power, we were unable to specifically consider the households in the bottom of the income distribution. Future ex post evaluation studies could attempt to obtain more specific results, as experimental works have found significant distributional effects for poor households (see, e.g., Muller et al., 2017).

The 2012 French soft-drink tax reform also implemented a specific excise tax of 0.0716€/L for

[13]The role of local costs in reducing pass-through is well-documented in work on empirical trade; see Nakamura and Zerom (2010).

[14]The intuition is that, following a taxation shock, a profit-maximising firm has to increase its prices to maintain the equality between marginal revenue and marginal cost. The optimal price-setting strategy will eventually depend on the rate at which demand falls with the mark-up adjustment on each unit sold.

NCSBs. We therefore replicated our empirical analysis for this group (see Supplementary Appendix D.5). The estimated average tax incidence, +4.2%, corresponds to a pass-through rate of 39.0%, which is similar to the pass-through for SSBs. The likely explanation is that producers manage product portfolios that include both SSB and NCSB products. When we analyse the data at the product level, we find clear evidence that producers and retailers tied NCSB prices to their twin-variety SSB prices, with a correlation in price variations of approximately 0.75. The analysis of market heterogeneity also shows a strong competition effect. Although the impact of median market income is imprecisely estimated, it is similar in sign to that observed for SSBs. We also find a similar hierarchy of product-level pass-through rates. Unlike SSBs, however, we find some evidence of differential tax incidence across income groups at the national level. Although the estimated effect is not significantly different from zero, this suggests that the excise tax on NCSBs is regressive for low-income households, partly because they have a specific preference for quality over NCSB products. Nevertheless, when we consider both SSBs and NCSBs, the decomposition of the compensating variation reveals that market heterogeneity explains more than half of the difference in welfare losses between low- and high-income households (Supplementary Appendix, Table D.3).

As a final evaluation exercise, we can predict the effect of the policy on consumption. In the Supplementary Appendix D.6, we provide evidence that the SSB tax did not change the distribution of sugar content within the group of SSB products. The key health benefits can only be obtained through substitutions towards outside options, such as NCSBs or USBs, which is consistent with the estimates of Bonnet and Réquillart (2013b). By combining estimates of price elasticities and tax incidence, we then find that the tax reduced purchases by approximately 1L/cap/year in low-income households vs. 0.5L/cap/year in high-income households. This implies that regressivity in consumer welfare may be partially offset by progressivity in health benefits, although empirical findings from the US suggest that small SSB taxes are ineffective in the short-term at reducing obesity (Fletcher, Frisvold, & Tefft, 2015, 2010).

The estimated pass-through is in line with the ex post estimates in Cawley and Frisvold (2017) and Falbe et al. (2015), who adopt a DiD approach with geographic control groups to estimate the incidence of the Berkeley tax and find pass-through rates of between 22% and 47%. However, different markets can yield different pass-through estimates. For instance, Colchero et al. (2015) and Grogger (2017) examine the impact of an excise tax in Mexico and find that the tax was over-shifted for carbonated varieties of SSBs.

Our findings lead to an unexpected revision of previous ex post evaluation results by Berardi and colleagues for France (Berardi et al., 2016). They concluded that the tax was fully shifted into SSB prices after six months (a 100% pass-through). There are however four important differences between their study and ours. First, they exploited extracts of shopping prices collected between August 2011 and June 2012 from the online sites of approximately 1,800 drive-through outlets. We exploit KWP data, which cover all outlet formats and provide a representative sample of purchases. Second, they work with store-level product prices weighted by national-level fixed market shares, i.e., the Laspeyres index. We construct theoretically founded price indices that control for preference heterogeneity, product availability and substitution. Third, their empirical before-after analysis does control neither for variations in sugar prices nor for month-of-the-year (seasonal) effects. As we have six years of data, we are able to control for changes in the cost of sugar and for month-of-the-year fixed effects. Fourth, they assume that the pass-through has to be evaluated by comparing the price levels at the end of their observation period (June 2012) with those observed in December 2011. We base the current analysis on year-to-year comparisons, following an approach that is widely used in the literature (e.g., Hong & Li, 2017).

An alternative approach in pass-through analysis is to track the monthly changes in price resulting from the taxation shock to costs in January 2012, the effect of which may be felt with some lags (Gopinath & Itskhoki, 2010; Nakamura & Zerom, 2010). The results from this event study suggest that the tax was passed on quite rapidly to consumer prices, after one quarter (Supplementary Appendix D.7). This is unsurprising given that, over the period 2008-2013, the contractual framework between manufacturers and retailers was regulated, with annual negotiations that had to be resolved by the end of March. The price levels reached in March-April 2012 are similar to our earlier results in the before-after specification. The rise in prices over the second quarter of 2012 was due to a seasonal effect, which Berardi et al. (2016) could not control for due to the limited time window covered by their data.

6 Conclusion

This study of the French soft-drink tax shows that aggregate market characteristics alter the tax incidence and distributional effects of soft-drink taxes. On average, at the national level, the tax burden was not significantly higher for low-income households. However, households living in low-income markets with few retailers faced considerable price increases. Residential sorting of households across markets with varying structure is an important determinant of the effectiveness and equity of soft-drink taxes.

While this paper has demonstrated the importance of accounting for market heterogeneity in tax incidence analysis, our approach is based on a theoretical framework that sets aside concerns about consumer behavioural biases. Allcott, Lockwood, and Taubinsky (2017) propose a method for estimating optimal sin taxes when direct measures of bias-proneness are available. Therefore, it would be interesting to replicate our analysis with exact price indices adjusted for behavioural biases. We leave this for future research.

References

Allcott, H., Lockwood, B. B., & Taubinsky, D. (2017). A theory of regressive sin taxes, with an application to the optimal soda tax. *NBER Working Paper*, *23085*.

Amiti, M., Itskhoki, O., & Konings, J. (2014). Importers, exporters, and exchange rate disconnect. *American Economic Review*, *104*(7), 1942–1978.

Berardi, N., Sevestre, P., Tepaut, M., & Vigneron, A. (2016). The impact of a soda tax on prices: evidence from french micro data. *Applied Economics*, *48*(41), 3976–3994.

Bonnet, C. & Réquillart, V. (2011). Does the eu sugar policy reform increase added sugar consumption? an empirical evidence on the soft drink market. *Health Economics*, *20*(9), 1012–1024.

Bonnet, C. & Réquillart, V. (2013a). Impact of cost shocks on consumer prices in vertically-related markets: the case of the french soft drink market. *American Journal of Agricultural Economics*, *95*(5), 1088–1108.

Bonnet, C. & Réquillart, V. (2013b). Tax incidence with strategic firms in the soft drink market. *Journal of Public Economics*, *106*, 77–88.

Cawley, J. & Frisvold, D. E. (2017). The pass-through of taxes on sugar-sweetened beverages to retail prices: the case of berkeley, california. *Journal of Policy Analysis and Management*, *36*(2), 303–326.

Colchero, M. A., Salgado, J. C., Unar-Munguía, M., Molina, M., Ng, S., & Rivera-Dommarco, J. A. (2015). Changes in prices after an excise tax to sweetened sugar beverages was implemented in mexico: evidence from urban areas. *PLoS One*, *10*(12), e0144408.

Etilé, F. & Sharma, A. (2015). Do high consumers of sugar-sweetened beverages respond differently to price changes? a finite mixture iv-tobit approach. *Health Economics*, *24*(9), 1147–1163.

Falbe, J., Rojas, N., Grummon, A. H., & Madsen, K. A. (2015). Higher retail prices of sugar-sweetened beverages 3 months after implementation of an excise tax in berkeley, california. *American Journal of Public Health*, *105*(11), 2194–2201.

Finkelstein, E. A., Zhen, C., Bilger, M., Nonnemaker, J., Farooqui, A. M., & Todd, J. E. (2013). Implications of a sugar-sweetened beverage (ssb) tax when substitutions to non-beverage items are considered. *Journal of Health Economics*, *32*(1), 219–239.

Fletcher, J. M., Frisvold, D., & Tefft, N. (2010). Can soft drink taxes reduce population weight? *Contemporary economic policy*, *28*(1), 23–35.

Fletcher, J. M., Frisvold, D., & Tefft, N. (2015). Non-linear effects of soda taxes on consumption and weight outcomes. *Health economics*, *24*(5), 566–582.

Gopinath, G. & Itskhoki, O. (2010). Frequency of price adjustment and pass-through. *Quarterly Journal of Economics*, *125*(2), 675–727.

Grogger, J. (2017). Soda taxes and the prices of sodas and other drinks: evidence from mexico. *American Journal of Agricultural Economics*, *99*(2), 481–498.

Hamilton, S. F. (2009). Excise taxes with multiproduct transactions. *American Economic Review*, *99*(1), 458–471.

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

Hong, G. H. & Li, N. (2017). Market structure and cost pass-through in retail. *Review of Economics and Statistics*, *99*(1), 151–166.

Le Bodo, Y., Etilé, F., Gagnon, F., & de Wals, P. (2017, January). Conditions influencing the adoption of a soda tax: analysis of the french case. Conference presentation: The role of taxation and other economic incentives in promoting healthier lifestyles: a focus on tobacco, alcohol and unhealthy diets (A. Alemanno and A. Garde). London.

Madden, D. (2015). The poverty effects of a 'fat-tax'in ireland. *Health economics*, *24*(1), 104–121.

Malik, V. S., Pan, A., Willett, W. C., & Hu, F. B. (2013). Sugar-sweetened beverages and weight gain in children and adults: a systematic review and meta-analysis. *American Journal of Clinical Nutrition*, *98*(4), 1084–1102.

Muller, L., Lacroix, A., Lusk, J. L., & Ruffieux, B. (2017). Distributional impacts of fat taxes and thin subsidies. *The Economic Journal*, *127*(604), 2066–2092.

Nakamura, E. & Zerom, D. (2010). Accounting for incomplete pass-through. *Review of Economic Studies*, *77*(3), 1192–1230.

Redding, S. J. & Weinstein, D. E. (2016). *A unified approach to estimating demand and welfare*. National Bureau of Economic Research.

Sharma, A., Hauck, K., Hollingsworth, B., & Siciliani, L. (2014). The effects of taxing sugar-sweetened beverages across different income groups. *Health Economics*, *23*(9), 1159–1184.

Tiffin, R., Kehlbacher, A., & Salois, M. (2015). The effects of a soft drink tax in the uk. *Health Economics*, *24*(5), 583–600.

Tiffin, R. & Salois, M. (2014). The distributional consequences of a fiscal food policy: evidence from the uk. *European Review of Agricultural Economics*, *42*(3), 397–417.

Triplett, J. E. (2001). Should the cost-of-living index provide the conceptual framework for a consumer price index? *The Economic Journal*, *111*(472), 311–334.

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

Weyl, E. G. & Fabinger, M. (2013). Pass-through as an economic tool: principles of incidence under imperfect competition. *Journal of Political Economy*, *121*(3), 528–583.

Market Heterogeneity and the Distributional Incidence of Soft-Drink Taxes: Evidence from France.

Supplementary Appendix

A Additional descriptive statistics

A.1 Products

Our choice of attributes to define the product varieties is based on nomenclatures used by the beverage industry (*Syndicat National des Boissons Rafraichissantes*). Table A.1 shows the distribution of attributes in each group of beverages in our data. It reveals that the SSB market is not dominated by carbonated drinks, while this is the case for NCSB, reflecting the innovative path taken by the market leaders (Coca and Pepsi). This is also reflected in the large share of top national brands in NCSB products (68%, as against 34% for SSB). Interestingly, the SSB and NCSB markets have lower unit values than USBs and larger unit values than Water. Last, the carbohydrate content in the SSB category is, as expected, much higher than that in NCSB, and equal to that in USB.

A.2 Household sample

Table A.2 sets out some household descriptive statistics for the original KWP household sample and the final sample used in the construction of the price indices. These are very similar, except for the type of residential area. As we drop living zones with under 10 households, the countryside and small towns are under-represented, while larger cities are over-represented. We observe 43,379 households over the whole period, with each household remaining in the sample for three years on average.

Table A.1: **Product varieties - Descriptive statistics**

	SSB		NCSB		USB		Water	
UPC #	400		127		338		130	
Carbonated (%)								
No	47.00		14.96		100.00		64.62	
Yes	53.00		85.04				35.38	
Carbohydrates (SD)	8.96	(1.98)	0.65	(1.25)	10.08	(3.74)	0.09	(0.41)
Light (%)								
No	100.00				97.93		90.00	
Yes			100.00		2.03		10.00	
Packaging (%)								
Plastic	63.50		77.95		23.37		100.00	
Carton	14.25		0.79		34.32			
Metal	17.25		21.26		15.68			
Glass	5.00				26.63			
Flavour (%)								
Citrus	5.94		8.40					
Plain cola	19.06		56.30					
Multifruit	9.06		0.84					
Peach	9.06		6.72					
Orange	22.81		11.76		31.07			
Plain							90.00	
Grenadine					8.88			
Mint					7.10			
Apricot-peach					1.18			
Lemon-lime	0.63		3.36		2.07		6.15	
Other	33.44		12.61		49.70		3.85	
Brand (%)								
Top national	34.00		67.72		31.95		40.00	
Other national	28.25		13.39		23.08		33.85	
Retailer	27.00		14.17		32.54		18.46	
Hard discount	10.75		4.72		12.43		7.69	

Notes: Kantar Worldpanel data 2008-2013. Carbohydrates are expressed in grams per 100 ml. The top national brand segment includes from one to six brands, depending on the product family. These are unweighted product-level statistics.

Table A.2: Households - Descriptive statistics

	Final sample		Full sample	
Monthly household income (SD)	1,589	(1,056)	1,521	(992)
Household income class (%)				
Rich	16.86		14.69	
Mid-rich	30.37		29.26	
Mid-poor	37.43		39.75	
Poor	15.34		16.30	
Household size (SD)	2.26	(1.35)	2.34	(1.31)
Household structure (%)				
Single	22.02		19.76	
Old	22.60		22.31	
Couple without children	22.29		22.59	
Couple with children	33.09		35.34	
Main shopper				
Age (SD)	48.79	(17.14)	48.86	(1.92)
Gender (%)	12.02		11.58	
Highest education level (%)				
Primary	5.41		5.89	
High school	21.83		23.34	
Baccalauréat	23.60		24.76	
2 years, technical/university	21.26		21.49	
3 years and more, university	27.90		24.52	
Residential area (%)				
Countryside	11.19		24.15	
Small town	4.37		11.86	
Town	9.49		10.80	
Large town	16.90		12.02	
City	58.05		41.17	
Households	30,254		43,379	
Years per household, on average (SD)	2.97	(1.90)	2.99	(1.90)
Observations (households x years)	89,930		129,911	

Notes: Kantar Worldpanel data 2008-2013. Household income is in 2011 Euros, per consumption unit (OECD scale). All statistics are weighted using the survey sample weights.

A.3 Trends in sugar and beverage prices: series from the National Statistics office

The National Statistics office (INSEE) provides a national Consumer Price Index (CPI) and a national Producer Price Index (PPI) for sugar, water and an aggregated beverage category. These price indices are annually chained Laspeyres price indices. The beverage category includes both taxed products (SSBs and NCSBs) and untaxed products (Unsweetened Beverages: USBs). Hence, it can not be used for studying the incidence of the soda tax.

The revision of the EU sugar quota policy in September 2011 may confound the impact of the soda tax on prices, because the soft drinks produced in France do not contain high-fructose corn syrup but standard sucrose (see the analysis published by the consumer news magazine *60 millions de consommateurs* in July 2012). The potential role of the cost of sugar is illustrated in Figure A.1, which plots national statistics data on the PPI for sugar (on the right panel), the CPI (on the left panel) and the PPI for all non-alcoholic beverages and its two main components, soft drinks, juices and syrups on the one hand (SSBs+NCSBs+USBs) and Water on the other.

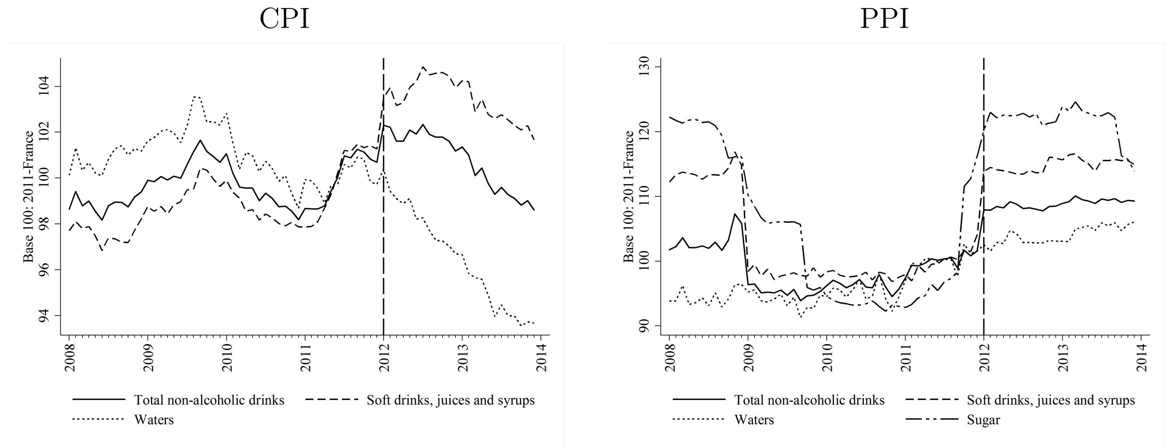

Figure A.1: Consumer and Producer Price Indices (INSEE)

Notes: This Figure shows the changes over time in Consumer Price Indices (CPI: left panel) and Producer Price Indices (PPI: right panel) reported by the French National Statistics Office (*INSEE: Institut National de la Statistique et des Etudes Economiques*) and Eurostat. Producer Price Indices measure monthly changes in the trading price of products from the producers' perspective. Consumer Price Indices measure changes in the prices of consumer goods.

The CPI of both components move similarly until 2011, with a positive trend from mid-2008 to mid-2009, and then a negative trend until the end of 2011. The CPIs then diverge until mid-2012, with a much larger and longer CPI increase for soft drinks, juices and syrups than for water, before going back to a common downward-sloping trend. The steep 2012 increase in the CPI of soft drinks, juices and syrups series may reflect (at least partly) the rise in the PPI for sugar following the revision of the EU sugar quota policy.

B Derivation of the Exact Price Index (EPI)

While the EPI has long been central in the measurement of the true cost-of-living using household budget data, a recent literature has proposed its adaptation to account for the spatial and temporal variability in product availability, which is largely related to the heterogeneity in the spatial distribution of retailers across markets. It is made possible by the availability of scanner data, which contain almost exhaustive information on markets. In this paper, we use the nested-CES price index, following the theoretical and empirical works of Handbury and Weinstein (2014), Feenstra (1994), Broda and Weinstein (2010, 2006).

B.1 Theoretical framework

An Exact Price Index (EPI) measures the change in expenditure required to keep utility constant as the prices of product varieties vary. It is therefore an index of consumer welfare for a given population $\mathcal{P}$ (index omitted in the formula). It can be formally defined for product group g and a representative consumer in market c as

$$\boldsymbol{P}_{gc} = \frac{C(V, \boldsymbol{p_{gc}})}{C(V, \boldsymbol{p_{g.}})}, \tag{B.1}$$

where $C(V, \boldsymbol{p_{gc}})$ is the cost of attaining utility V when facing prices $\boldsymbol{p_{gc}}$, and $\boldsymbol{p_{g.}}$ is a vector of reference prices.

We assume that households take a four-stage budgeting approach to decide their beverage consumptions (Deaton & Muellbauer, 1980b). They first allocate their consumption budget between broad food groups, here alcoholic and non-alcoholic beverages. The non-alcoholic beverage budget is then allocated between the four beverage groups g: (1) SSBs (sodas and fruit drinks essentially); (2) Non-Calorically Sweetened Beverages (NCSBs); (3) Unsweetened Beverages (USBs, mainly fruit juices without added sweeteners); and (4) Water. The budget is then allocated between "brand-modules" b within each beverage group (Coca-Cola regular, Pepsi-Cola regular, Diet Coke, etc.). Last it is split up between UPCs u within each brand-module: UPCs are the products purchased from a specific retailer.

This multi-stage budgeting process thus mirrors the three-tiers nomenclature of purchases presented in Section 2.2. Purchases are classified into product groups g (the upper level), product groups are made up of brand-modules b (the intermediate level), and brand-modules include a number of distinct UPCs u. For example, g = SSB; b = Coca Cola regular, and u = a 1-liter plastic bottle of Coca Cola regular sold in a Carrefour hypermarket. This classification also matches the business nomenclature used by producers and retailers.

Households purchase on disjoint markets c that are unique combinations of living zones a and time periods t. They purchase from retailers r, who may or may not be present in the market c. They hence have access to a market-specific set of brand-modules, B_{gc}, and a market-specific set of

UPCs, U_{bgc}, depending on their residential location and the period. This set-up allows variability in the availability of UPCs across markets. For instance, a retailer may launch its own private-label cola in a given year, test it, and withdraw it if it does not attract a profitable market share. We denote $\mathcal{R}$ the reference market. It is the "national market" (*i.e.* the union of all living zones) in 2011, the pre-tax year. The reference set of brand modules is $B_g = \bigcup_{c \in \mathcal{R}} B_{gc}$, and $U_{bg} = \bigcup_{c \in \mathcal{R}} U_{bgc}$ similarly for the reference set of UPCs within a brand-module.

To be consistent with multi-stage budgeting, and given our focus on aggregate consumption, we make the following weak-separability assumption regarding household preferences.

Assumption 1 (Weak Separability):

1. *The set of products is partitioned into G mutually-exclusive groups. We note $\boldsymbol{q}_g$ the vector of goods in group g, with the related price vector $\boldsymbol{p}_{gc}$ in market c, and q_k an elementary product priced at p_k.*

2. *Household preferences are separable, so that each household h living in market c solves the following maximisation program*

$$Max_{\boldsymbol{q_g}, \forall g=1,\ldots,G} F_h \left(U_{h1}(\boldsymbol{q_1}), \ldots, U_{hG}(\boldsymbol{q_G}) \right),$$
$$s.t. \sum_{g=1}^{g=G} \boldsymbol{p_{gc}}' \boldsymbol{q_g} = X_h,$$
$$\forall i, q_i \geq 0.$$

We then assume nested-CES functional forms for consumer preferences over brand-modules and UPCs. This assumption yields two benefits. First, the derivation of the EPI is relatively straightforward. Second, nested-CES utility functions represent the behaviour of a household that would be representative of a population having nested-logit preferences over brand-modules (a nest) and UPCs (Anderson, de Palma, & Thisse, 1988; Redding & Weinstein, 2016).

Assumption 2 (nested-CES subutility functions): *Consumer preferences over brand-modules and UPCs are represented by a two-level CES utility function $U_{hg}(q_g)$:*

- *Upper-level:*

$$\boldsymbol{Q_{hg}} = U_{hg}(\boldsymbol{q_g}) = \left[\sum_{b \in B_{gc}} \boldsymbol{Q_{bg}}^{\frac{\sigma_{hg}-1}{\sigma_{hg}}} \right]^{\frac{\sigma_{hg}}{\sigma_{hg}-1}},$$

- *Lower-level:*

$$\boldsymbol{Q_{bg}} = \left[\sum_{u \in U_{bgc}} \left(\varphi_{hug} q_{ug} \right)^{\frac{\sigma_{hb}-1}{\sigma_{hb}}} \right]^{\frac{\sigma_{hb}}{\sigma_{hb}-1}},$$

where φ_{hug} is the household-specific (subjective) quality of variety u.

It is important to note here that the aggregation of consumer behaviours is made possible and plausible by the introduction of a household-UPC specific quality φ_{hug}, which adjusts the quantity purchased by consumer heterogeneity in preference over quality. Adjusting for household and retailer heterogeneity makes the UPCs and brand-modules homogeneous, in terms of subjective quality. As such, it renders plausible the assumption of a constant elasticity of substitution.

Quality-adjustment is all the more necessary that homescan data do not provide retailer prices, but rather unit prices (or unit values). In a given market c, the observed unit prices for an UPC are likely to vary from one household to another for three reasons. First, households choose to shop in specific stores, which may differ in terms of amenities. Stores adjust their prices as a function of the amenities they provide. Second, stores also adjust their prices as a function of customer demand and characteristics. In addition, households may differ in their shopping behaviour, sensitivity to sales promotions, etc. Third, as we define UPCs from a restricted set of product characteristics, purchases of a given UPC may be heterogeneous in terms of unobserved attributes related to consumer preferences (*e.g.* a particular flavour). In practice, household unit prices will be adjusted for retailer fixed effects (retailer heterogeneity), household characteristics and some interactions between UPC characteristics and household characteristics (household heterogeneity) - see Appendix C.

Given our Assumption (2), the utility-maximization program at the lower level of UPCs for household h purchasing in market c is

$$\begin{aligned} Max_{q_{ug}, \forall u \in U_{bgc}} \boldsymbol{Q_{bg}}, \\ s.t. \sum_{u \in U_{bgc}} p_{ubgc} q_{ug} = \nu_{bgc}, \\ q_{ug} \geq 0, \end{aligned} \tag{B.2}$$

where ν_{bgc} is the budget constraint and the quantity index Q_{bg} is a direct measure of consumer utility

$$\boldsymbol{Q_{bg}} = \left[\sum_{u \in U_{bgc}} \left(\varphi_{hug} q_{ug} \right)^{\frac{\sigma_{hb}-1}{\sigma_{hb}}} \right]^{\frac{\sigma_{hb}}{\sigma_{hb}-1}}. \tag{B.3}$$

Solving the dual cost-minimization problem, we obtain the following equality

$$P_{bgc}\boldsymbol{Q_{hbg}} = \nu_{bgc}, \tag{B.4}$$

where P_{bgc} is the unit cost function (*i.e.* the expenditure required to obtain one unit of utility Q_{hbg})

$$P_{bgc} = \left(\sum\nolimits_{u\in U_{bgc}} \left(\frac{p_{ubgc}}{\varphi_{hug}}\right)^{1-\sigma_b}\right)^{\frac{1}{1-\sigma_b}}. \tag{B.5}$$

We can similarly solve the household utility-maximization problem at the upper-level of brand modules in order to obtain a unit cost function P_{gc} measuring the cost of one unit of utility U_{hg} from the consumption of products in group g.

Since we have assumed that $U_{hg}(\boldsymbol{q_g}) = \left[\sum\nolimits_{b\in B_{gc}} \boldsymbol{Q_{bg}}^{\frac{\sigma_{hg}-1}{\sigma_{hg}}}\right]^{\frac{\sigma_{hg}}{\sigma_{hg}-1}}$, we have

$$P_{gc} = \left(\sum\nolimits_{b\in B_{gc}} (P_{bgc})^{1-\sigma_g}\right)^{\frac{1}{1-\sigma_g}}. \tag{B.6}$$

B.2 Adjusting the EPI for product availability

Suppose that the prices are adjusted for differences in subjective quality (tastes) φ_{hug}, so that we can calculate representative prices from observed transaction prices (Appendix C explains the procedure). These representative prices are adjusted for household and retailer heterogeneity and are denoted $\tilde{p}_{ubgc}$. All quality-adjusted variables are indicated by a tilda.

We assume that $\forall h, \sigma_{hg} = \sigma_g$ and $\sigma_{hb} = \sigma_b$ for both theoretical and empirical reasons. First, elasticities do not vary across markets and do not vary across households within markets, as we want to construct a local price index for a representative consumer. Second, elasticities do not vary over time, because we do not have enough observations to estimate them separately for each month.

The reference market $\mathcal{R}$ is the "national market" (*i.e.* the union of all living zones) in 2011, and we assume that preferences do not vary between living zones. In this case, the price index for the product category g in market c can be written as

$$I_{gc} = \tilde{P}_{gc}/\tilde{P}_{g\mathcal{R}}, \tag{B.7}$$

where

$$\tilde{P}_{gc} = \left(\sum\nolimits_{b\in B_{gc}} \left(\tilde{P}_{bgc}\right)^{1-\sigma_g}\right)^{\frac{1}{1-\sigma_g}}, \tag{B.8}$$

with $\tilde{P}_{g\mathcal{R}}$ the price of c in the reference market $\mathcal{R}$

$$\tilde{P}_{g\mathcal{R}} = \left(\sum\nolimits_{b \in B_{g\mathcal{R}}} (\tilde{P}_{bg\mathcal{R}})^{1-\sigma_g}\right)^{\frac{1}{1-\sigma_g}}, \tag{B.9}$$

where $\tilde{P}_{bg\mathcal{R}}$ is the "national price"of brand-module b in 2011. The price indices will thus reflect deviations from the 2011 national price.

We have for the (taste-adjusted) share of any specific brand module b within product group g

$$\tilde{S}_{bgc} = \left(\frac{\tilde{P}_{bgc}}{\tilde{P}_{gc}}\right)^{1-\sigma_g}, \tag{B.10}$$

and therefore

$$\ln(\tilde{P}_{gc}) = \ln(\tilde{P}_{bgc}) - \frac{\ln(\tilde{S}_{bgc})}{1-\sigma_g}, \qquad \forall b \in B_{gc}. \tag{B.11}$$

For the reference market $\mathcal{R}$, *i.e.* for all brand-modules in $B_{g\mathcal{R}} = \bigcup_{c' \in \mathcal{R}} B_{gc'}$, we have

$$\tilde{P}_{g\mathcal{R}} = \tilde{P}_{bg\mathcal{R}}(\tilde{S}_{bg\mathcal{R}})^{\frac{1}{1-\sigma_g}} \Longrightarrow \ln(\tilde{P}_{g\mathcal{R}}) = \ln(\tilde{P}_{bg\mathcal{R}}) - \frac{\ln(\tilde{S}_{bg\mathcal{R}})}{1-\sigma_g}, \qquad \forall b \in B_{g\mathcal{R}}. \tag{B.12}$$

Hence

$$I_{gc} = \tilde{P}_{gc}/\tilde{P}_{g\mathcal{R}} = \frac{\tilde{P}_{bgc}}{\tilde{P}_{bg\mathcal{R}}}\left(\frac{\tilde{S}_{bgc}}{\tilde{S}_{bg\mathcal{R}}}\right)^{-\frac{1}{1-\sigma_g}}, \qquad \forall b \in B_{gc} \cap B_{g\mathcal{R}}. \tag{B.13}$$

$B_{gc} \cap B_{g\mathcal{R}} = B_{gc\mathcal{R}}$ is the set of brand-modules available both on c and on $\mathcal{R}$. Now, let $U_{b\mathcal{R}} = \bigcup_{c' \in \mathcal{R}} U_{bc'}$, and let $\tilde{\nu}_{ubgc}$ be the (taste-adjusted) expenditure on u in market c, and note that

$$\tilde{S}_{bgc} = \frac{\sum\limits_{u \in U_{bc}} \tilde{\nu}_{ubgc}}{\sum\limits_{b' \in B_{gc}} \sum\limits_{u \in U_{b'c}} \tilde{\nu}_{ub'gc}} = \underbrace{\frac{\sum\limits_{u \in U_{bc}} \tilde{\nu}_{ubgc}}{\sum\limits_{b' \in B_{gc\mathcal{R}}} \sum\limits_{u \in U_{b'c}} \tilde{\nu}_{ub'gc}}}_{\tilde{S}^{c\mathcal{R}}_{bgc}} \times \underbrace{\frac{\sum\limits_{b' \in B_{gc\mathcal{R}}} \sum\limits_{u \in U_{b'c}} \tilde{\nu}_{ub'gc}}{\sum\limits_{b' \in B_{gc}} \sum\limits_{u \in U_{b'c}} \tilde{\nu}_{ub'gc}}}_{\tilde{s}^{c\mathcal{R}}_{gc}}. \tag{B.14}$$

$$\tilde{S}_{bg\mathcal{R}} = \frac{\sum\limits_{u \in U_{b\mathcal{R}}} \sum\limits_{c' \in \mathcal{R}} \tilde{\nu}_{ubgc'}}{\sum\limits_{b' \in B_{g\mathcal{R}}} \sum\limits_{u \in U_{b'\mathcal{R}}} \sum\limits_{c' \in \mathcal{R}} \tilde{\nu}_{ub'gc'}} = \underbrace{\frac{\sum\limits_{u \in U_{b\mathcal{R}}} \sum\limits_{c' \in \mathcal{R}} \tilde{\nu}_{ubgc'}}{\sum\limits_{b' \in B_{gc\mathcal{R}}} \sum\limits_{u \in U'_{b\mathcal{R}}} \sum\limits_{c' \in \mathcal{R}} \tilde{\nu}_{ub'gc'}}}_{\tilde{S}^{c\mathcal{R}}_{bg\mathcal{R}}} \times \underbrace{\frac{\sum\limits_{b' \in B_{gc\mathcal{R}}} \sum\limits_{u \in U_{b'\mathcal{R}}} \sum\limits_{c' \in \mathcal{R}} \tilde{\nu}_{ub'gc'}}{\sum\limits_{b' \in B_{g\mathcal{R}}} \sum\limits_{u \in U_{b'\mathcal{R}}} \sum\limits_{c' \in \mathcal{R}} \tilde{\nu}_{ub'gc'}}}_{\tilde{s}^{c\mathcal{R}}_{g\mathcal{R}}}. \tag{B.15}$$

Define the variety-adjusted Sato-Vartia weights W_{bc} only on the set of UPCs available in both markets

c and $\mathcal{R}$ as follows

$$W_{bc} = \left(\frac{\tilde{S}_{bgc}^{c\mathcal{R}} - \tilde{S}_{bg\mathcal{R}}^{c\mathcal{R}}}{\ln(\tilde{S}_{bgc}^{c\mathcal{R}}) - \ln(\tilde{S}_{bg\mathcal{R}}^{c\mathcal{R}})} \right) \Bigg/ \sum\nolimits_{b' \in B_{gc\mathcal{R}}} \left(\frac{\tilde{S}_{b'gc}^{c\mathcal{R}} - \tilde{S}_{b'g\mathcal{R}}^{c\mathcal{R}}}{\ln(\tilde{S}_{b'gc}^{c\mathcal{R}}) - \ln(\tilde{S}_{b'g\mathcal{R}}^{c\mathcal{R}})} \right). \tag{B.16}$$

As these weights sum up to one, we can take the geometric mean of the log of the price index accross the varieties in $B_{gc\mathcal{R}}$

$$\begin{aligned} \ln(\tilde{P}_{gc}) - \ln(\tilde{P}_{g\mathcal{R}}) &= \sum_{b \in B_{gc\mathcal{R}}} W_{bc} \left(\ln(\tilde{P}_{gc}) - \ln(\tilde{P}_{g\mathcal{R}}) \right) \\ &= \sum_{b \in B_{gc\mathcal{R}}} W_{bc} \left(\ln(\tilde{P}_{bgc}) - \ln(\tilde{P}_{bg\mathcal{R}}) \right) - \sum_{b \in B_{gc\mathcal{R}}} W_{bc} \frac{\left(\ln(\tilde{S}_{bgc}) - \ln(\tilde{S}_{bg\mathcal{R}}) \right)}{1 - \sigma_g}, \end{aligned} \tag{B.17}$$

and

$$\begin{aligned} \sum_{b \in B_{gc\mathcal{R}}} W_{bc} \frac{\left(\ln(\tilde{S}_{bgc}) - \ln(\tilde{S}_{bg\mathcal{R}}) \right)}{1 - \sigma_g} &= \sum_{b \in B_{gc\mathcal{R}}} W_{bc} \frac{\ln(\tilde{S}_{bgc}^{c\mathcal{R}}) - \ln(\tilde{S}_{bg\mathcal{R}}^{c\mathcal{R}}) + \ln\left(\tilde{s}_{gc}^{c\mathcal{R}}\right) - \ln\left(\tilde{s}_{g\mathcal{R}}^{c\mathcal{R}}\right)}{1 - \sigma_g} \\ &= \frac{\ln\left(\tilde{s}_{gc}^{c\mathcal{R}}\right) - \ln\left(\tilde{s}_{g\mathcal{R}}^{c\mathcal{R}}\right)}{1 - \sigma_g}. \end{aligned} \tag{B.18}$$

This implies that

$$I_{gc} = \left\{ \prod\nolimits_{b \in B_{gc\mathcal{R}}} \left(\frac{\tilde{P}_{bgc}}{\tilde{P}_{bg\mathcal{R}}} \right)^{W_{bc}} \right\} \left(\frac{\tilde{s}_{g\mathcal{R}}^{c\mathcal{R}}}{\tilde{s}_{gc}^{c\mathcal{R}}} \right)^{\frac{1}{1-\sigma_g}}. \tag{B.19}$$

For each brand-module, we have similarly

$$I_{bgc} = \frac{\tilde{P}_{bgc}}{\tilde{P}_{bg.}} = \left\{ \prod\nolimits_{u \in U_{bc\mathcal{R}}} \left(\frac{\tilde{p}_{ubgc}}{\tilde{p}_{ubg.}} \right)^{w_{ubc}} \right\} \left(\frac{\tilde{s}_{b\mathcal{R}}^{c\mathcal{R}}}{\tilde{s}_{bc}^{c\mathcal{R}}} \right)^{\frac{1}{1-\sigma_b}}, \tag{B.20}$$

with $U_{bc\mathcal{R}} = U_{bc} \cap U_{b\mathcal{R}}, \forall (b, c)$, and

$$\tilde{s}_{bc}^{c\mathcal{R}} = \frac{\sum\limits_{u \in U_{bc\mathcal{R}}} \tilde{\nu}_{ubgc}}{\sum\limits_{u \in U_{bc}} \tilde{\nu}_{ubgc}}, \tag{B.21}$$

$$\tilde{s}_{b\mathcal{R}}^{c\mathcal{R}} = \frac{\sum\limits_{u \in U_{bc\mathcal{R}}} \sum\limits_{c' \in \mathcal{R}} \tilde{\nu}_{ubgc'}}{\sum\limits_{u \in U_{b\mathcal{R}}} \sum\limits_{c' \in \mathcal{R}} \tilde{\nu}_{ubgc'}}, \tag{B.22}$$

and the Sato-Vartia weights

$$w_{ubc} = \left(\frac{\tilde{s}^{c\mathcal{R}}_{ubgc} - \tilde{s}^{c\mathcal{R}}_{ubg\mathcal{R}}}{\ln(\tilde{s}^{c\mathcal{R}}_{ubgc}) - \ln(\tilde{s}^{c\mathcal{R}}_{ubg\mathcal{R}})} \right) \Big/ \sum\nolimits_{u' \in U_{bc}} \left(\frac{\tilde{s}^{c\mathcal{R}}_{u'bgc} - \tilde{s}^{c\mathcal{R}}_{u'bg\mathcal{R}}}{\ln(\tilde{s}^{c\mathcal{R}}_{u'bgc}) - \ln(\tilde{s}^{c\mathcal{R}}_{u'bg\mathcal{R}})} \right), \tag{B.23}$$

with

$$\tilde{s}^{c\mathcal{R}}_{ubgc} = \frac{\tilde{\nu}_{ubgc}}{\sum\limits_{u \in U_{bc\mathcal{R}}} \tilde{\nu}_{ubgc}}. \tag{B.24}$$

$$\tilde{s}^{c\mathcal{R}}_{ubg\mathcal{R}} = \frac{\sum_{c'} \tilde{\nu}_{ubgc'}}{\sum\limits_{u \in U_{bc\mathcal{R}}} \sum\limits_{c'} \tilde{\nu}_{ubgc'}}. \tag{B.25}$$

We end up with the exact price index for product group g in market c

$$EPI_{gc} = CEPI_{gc} VA_{gc}, \tag{B.26}$$

where

$$CEPI_{gc} = \prod\nolimits_{b \in B_{gc\mathcal{R}}, u \in U_{bc\mathcal{R}}} \left(\frac{\tilde{p}_{ubgc}}{\tilde{p}_{ubg\mathcal{R}}} \right)^{w_{ubc} W_{bc}}, \tag{B.27}$$

$$VA_{gc} = \left(\frac{\tilde{s}^{c\mathcal{R}}_{g\mathcal{R}}}{\tilde{s}^{c\mathcal{R}}_{gc}} \right)^{\frac{1}{1-\sigma_g}} \prod\nolimits_{b \in B_{gc}} \left(\frac{\tilde{s}^{c\mathcal{R}}_{b\mathcal{R}}}{\tilde{s}^{c\mathcal{R}}_{bc}} \right)^{\frac{W_{bc}}{1-\sigma_b}}, \tag{B.28}$$

where $CEPI_{gc}$ is the EPI obtained under the assumption that the same choice set is observed in market c and in the reference market, and VA_{gc} is an adjustment for the differences in the available choice sets. In the formula for $CEPI_{gc}$, $\tilde{p}_{ubgc}$ and $\tilde{p}_{ubg\mathcal{R}}$ are respectively the quality-adjusted prices of u in market c and in the reference market. $\tilde{p}_{ubgc} = p_{ubgc}/\varphi_{hug}$ are thus the unit prices adjusted for within-market variations in household tastes and retailer heterogeneity. W_{bc} and w_{ubc} are Sato-Vartia weights that reflect the relative importance of brand-modules b and UPCs u in market c as compared to the reference market.

In the formula for VA_{gc}, the taste-adjusted shares $\tilde{s}^{c\mathcal{R}}_{gc}$ and $\tilde{s}^{c\mathcal{R}}_{bc}$ are the expenditure shares for market c of category g's brand-modules and brand-module b's UPCs that are available in both c and $\mathcal{R}$. The shares $\tilde{s}^{c\mathcal{R}}_{g\mathcal{R}}$ and $\tilde{s}^{c\mathcal{R}}_{b\mathcal{R}}$ are the corresponding expenditure shares for the reference market $\mathcal{R}$. σ_g is the elasticity of substitution across brand-modules in product group g, and σ_b is the elasticity of substitution across UPCs within a brand-module.

The variety-adjustment term VA_{gc} is determined by the relative availability of UPCs and their relative popularity in market c as compared to the reference market $\mathcal{R}$. Given our choice of a very large reference market, it turns out that, *in our data*, $\tilde{s}^{c\mathcal{R}}_{gc} = \tilde{s}^{c\mathcal{R}}_{bc} = 1$: all of the UPCs that are observed in market c are always available in the reference market. Then, the variety-adjustment term will first vary with the quality-adjusted shares of the available UPCs $\tilde{s}^{c\mathcal{R}}_{b\mathcal{R}}$ in a brand-module b, and the quality-adjusted shares of available brand-modules $\tilde{s}^{c\mathcal{R}}_{g\mathcal{R}}$ in product category g in market c. These shares do not reflect consumer choices in market c but rather the availability of UPCs, as $\tilde{s}^{c\mathcal{R}}_{b\mathcal{R}}$ is defined as the ratio of total expenditure in $\mathcal{R}$ on UPCs u in brand-module b available in both markets c and $\mathcal{R}$, to the total expenditure on all UPCs u in brand-module b available in $\mathcal{R}$. This ratio is therefore below 1 whenever $U_{bc\mathcal{R}}$ is smaller than $U_{b\mathcal{R}}$, *i.e.* when a UPC in brand-module b is unavailable in market c (which is always the case in our data). Now suppose that there are many UPCs that are not available in market c, so that $\tilde{s}^{c\mathcal{R}}_{b\mathcal{R}}$ falls. As $1/(1-\sigma_b)$ is negative ($\sigma_b > 1$), VA_{gc} will increase. The loss of welfare due to the absence of some UPCs translates into a higher price index. For UPCs within brand-modules, this is unimportant if the brand-module has a low Sato-Vartia weight W_{bc}, *i.e.* if it is not very popular among consumers. Similarly, $\tilde{s}^{c\mathcal{R}}_{gc}$ is the ratio of the total expenditure on all brand-modules b available in market c, to the total expenditure on all brand-modules in g. This is lower than 1, and will produce a rise in VA_{gc} whenever $B_{gc\mathcal{R}}$ is smaller than $B_{g\mathcal{R}}$. Entries of UPCs already available in $\mathcal{R}$ will on the contrary produce a drop in VA_{gc}, corresponding to an increase in consumer welfare.

C Construction of the EPI

We here present the empirical steps taken to construct the nested-CES price index. The method closely follows Handbury and Weinstein (2014).

C.1 Overview of the procedure

In the EPI formula (Supplementary Appendix B, Equations B.27 and B.28), prices, expenditures and quantities are adjusted for differences in subjective quality between UPCs. These differences in subjective quality are related to consumer and retailer heterogeneity. We purge these by adjusting household unit prices via period specific regressions, where we control for UPC and household characteristics, retailer and store-format fixed effects, and interactions between product and household attributes. These estimates are then used to construct market-specific prices adjusted for retailer and store period-specific fixed effects (time-varying retailer heterogeneity) and period-specific household characteristics (time-varying household heterogeneity) – See Handbury and Weinstein (2014, subsections 3.2. and 5.1). We use sampling weights in all of the calculations, to ensure that the prices are economically and demographically representative.

To construct the set of available products in each market, we match our data to the Nielsen TradeDimensions panel, by living zone and time period. This panel provides exhaustive information on the set of retailers operating in each living zone a at each period t. We construct the set of UPCs u available in market c by assuming that each retailer proposes the same UPCs in all of the living zones in which it operates at t. This procedure may overestimate the number of UPCs actually available in market c. However, this is a minor issue as we focus only on the most-populated living zones, in which the local assortment of UPCs proposed by a retailer likely corresponds to its national assortment. This in addition is not a concern for the evaluation of the variety of brand-modules, as all brand-modules appear on (or disappear from) all of the local markets at the same time. For example, the introduction of Coca Zero was national, and even if all package sizes were not available from every retailer, this was unlikely to have had a large effect on consumer welfare: variety/innovation biases should primarily be corrected at the brand-module level.

The computation of the variety-adjustment terms requires the computation of elasticities of substitutions within and across brand-modules. These are estimated through systems of CES demand *and* supply equations, following the method in Feenstra (1994) and extended by Broda and Weinstein (2006). We assume that elasticities do not vary across across households, as we want to construct a local price index for a representative consumer. The estimated median substitution elasticities among varieties within brand-modules are almost the same for SSBs and NCSBs (5.48 and 5.39, respectively). This figure is larger for USBs (9.59) and smaller for Water (4.57). The across-brand-module elasticities are fairly large for SSBs and NCSBs (6.04 and 6.69, respectively) and lower for USB (3.35) and Water (3.13). For USBs, the low elasticity figure is explained by the small number of brand modules and the considerable differences between them (juices are very different from pulps). For Water, we observe strong brand loyalty in our data (see Table C.1). We have tested the robustness of our main findings to the use of separate elasticities for 2008-2011 and 2012-2013, as the soft-drink tax may also have altered consumer preferences. The elasticities are fairly similar between the two periods, and the results presented in the next sections are therefore unaffected.

C.2 Adjusting unit prices, quantities and expenditures: details

Let p_{ucrh} be the "unadjusted"average price that a household h paid for UPC u in retailer r in market c. We construct a quality-adjusted average price by estimating the following OLS regressions

$$\ln(p_{ucrh}) = \alpha_c + \alpha_r + X_u\beta_u + X_h\beta_h + X_{uh}\beta_{uh} + \epsilon_{ucrh}, \tag{C.1}$$

where α_c are market fixed-effects, α_r is a vector of dummy variables indicating the retailer name (seven dummies) and format (three dummies), X_u, X_h and X_{uh} are vectors of UPC characteristics, household attributes and interactions, respectively, and β_u, β_h and β_{uh} the corresponding parameters. In the empirical application (after a specification search), X_u includes carbohydrate density and dummy variables for the group (four dummies), beverage family (14 dummies), brand (26 dummies), flavor (11 dummies), packaging (4 dummies) and volume (up to five dummies per family); X_h contains the age and gender of the main shopper, household structure (four dummies), equivalent income class (four dummies), education (five dummies) and type of residential area (five dummies), and X_{uh} interacts the low-income class with carbohydrate density, and the beverage-family and volume dummies.

The adjustment regressions are performed month by month, with 263 fixed-effects for living zones. This allows to better control for variations in retailers' amenities over time and implies that there is no redundancy between the inclusion of retailer fixed effects and the definition of an UPC as a given variety purchased at a given retailer. Each observation is weighted by the transaction value (household expenditures, v_{ucrh}) multiplied by the Kantar sample weight for the household (ω_{ch}). This gives more weight to varieties that attract higher national expenditure shares. The equations are estimated on data pooled over the four groups of products. Most R^2s range between 0.8 and 0.9. Adding UPC fixed effects instead of a large set of product characteristics only slightly increases the fit, while it greatly weakens the identification of the impact of household characteristics as it then essentially relies on households purchasing different UPCs within a month. In addition, separate regressions by product category, albeit theoretically preferable, produce lower R^2s (around 0.7-0.8) and the market fixed effects were not well-identified: the final EPI exhibited large and implausible monthly changes.

The average price corrected for retailer and household heterogeneity is finally

$$\tilde{p}_{ucrh} = \exp\left[\ln(p_{ucrh}) - (\hat{\alpha}_r + X_h\hat{\beta}_h + X_{uh}\hat{\beta}_{uh})\right]. \tag{C.2}$$

We use these adjusted prices to calculate quality-adjusted market-specific expenditures

$$\tilde{v}_{ubgc} = N_c \sum_{h\in H_c} \left\{ \frac{\omega_{ch}}{\sum_{h\in H_c}\omega_{ch}} \sum_{r\in R_c} \tilde{p}_{ucrh} \frac{q_{ucrh}}{SIZE_{hc}} \right\}, \tag{C.3}$$

where H_c is the set of households observed in market c, N_c is the population in the market, $SIZE_{hc}$ is the number of members in household h, and R_c is the set of retailers operating in market c (see Handbury & Weinstein, 2014, footnote 28). We first take a weighted average (with $\omega_{ch}/\sum_{h\in H_c}\omega_{ch}$ as the relative weights) of per capita household expenditures, and then multiply the result by the population size in c to obtain total expenditure in market c that can be compared to the total expenditure observed in other markets. In addition, we observe the following market-specific quantities

$$q_{ubgc} = N_c \sum_{h\in H_c} \left\{ \frac{\omega_{ch}}{\sum_{h\in H_c}\omega_{ch}} \sum_{r\in R_c} \frac{q_{ucrh}}{SIZE_{hc}} \right\}. \tag{C.4}$$

The prices, adjusted for consumer and retailer heterogeneity, and accounting for quantities, can then be calculated as

$$\tilde{p}_{ubgc} = \tilde{v}_{ubgc}/q_{ubgc}. \tag{C.5}$$

C.3 Estimating CES elasticities for VA: details

The CES model implies that, within each market, the difference between the quality-adjusted demand shares is proportional to the difference between the quality-adjusted prices. This is used to derive substitution elasticities as in Feenstra (1994), Broda and Weinstein (2006). Since suppliers also react to demand, a structural CES model is specified for the supply-side. The elasticity is obtained by solving for the equilibrium, under the identifying assumption that, within each market, within-brand-module unobserved shocks to demand are unrelated to within-brand-module unobserved shocks to supply. A key argument in favour of this assumption is that the demand and supply functions here are estimated from the price and market share data adjusted for consumer and retailer heterogeneity. A second assumption is that the elasticities do not vary across markets, *i.e.* over time and across living zones.

We first consider the estimation of within brand-module elasticities. Given the assumption regarding the equality of within-brand elasticities in a product category, for all UPCs u in brand-module

b in product category g we can write

$$\ln(\tilde{p}_{ubgc}) = \frac{\ln(\tilde{s}_{ubgc})}{1-\sigma_b} + \ln(\tilde{P}_{bgc}), \qquad \forall b \in B_{gc}, \forall u \in U_{bc}. \tag{C.6}$$

Let k_{bg} be one of the UPCs in the set of varieties U_{bc}; we then have the following demand equation

$$\Delta^{k_{bg}} \tilde{s}_{ubgc} = (1-\sigma_b)\Delta^{k_{bg}} \tilde{p}_{ubgc}, \qquad \forall b \in B_{gc}, \forall u \in U_{bc}, \tag{C.7}$$

where $\Delta^{k_{bg}} \tilde{s}_{ubgc} = \ln(\tilde{s}_{ubgc}) - \ln(\tilde{s}_{k_{bg}bgc})$ and $\Delta^{k_{bg}} \tilde{p}_{ubgc} = \ln(\tilde{p}_{ubgc}) - \ln\left(\tilde{p}_{k_{bg}bgc}\right)$.[15] To estimate the demand equation, we jointly specify a CES supply equation, so that we have the following system describing market equilibrium

$$\begin{aligned} \mathit{Demand}: \Delta^{k_{bg}} \tilde{s}_{ubgc} &= (1-\sigma_b)\Delta^{k_{bg}} \tilde{p}_{ubgc} + \varepsilon^{k_{bg}}_{ubgc}, \\ \mathit{Supply}: \Delta^{k_{bg}} \tilde{p}_{ubgc} &= \frac{\omega_b}{1+\omega_b}\Delta^{k_{bg}} \tilde{s}_{ubgc} + \delta^{k_{bg}}_{ubgc}, \end{aligned} \tag{C.8}$$

where ω_b is the supply elasticity, and $\varepsilon^{k_{bg}}_{ubgc}$ and $\delta^{k_{bg}}_{ubgc}$ are two error terms capturing the impact of random shocks on demand and supply respectively. Note that (i) as for the demand elasticity, the supply elasticity is assumed to be the same within all brand modules, and (ii) $\delta^{k_{bg}}_{ubgc}$ captures for instance the impact of assembly-line shocks that affect some UPCs within a brand module but not others (Broda and Weinstein, 2010).

The within-brand differentiation eliminates all brand-specific shocks. One credible identification restriction is then that the within-brand shocks to demand and supply are unrelated whatever the market: $\mathbb{E}(\varepsilon^{k_{bg}}_{ubgc}\delta^{k_{bg}}_{ubgc}|c) = 0$. To see this, multiply the demand and supply equations

$$\begin{aligned} \varepsilon^{k_{bg}}_{ubgc}\delta^{k_{bg}}_{ubgc} &= (\Delta^{k_{bg}} \tilde{s}_{ubgc} - (1-\sigma_b)\Delta^{k_{bg}} \tilde{p}_{ubgc})(\Delta^{k_{bg}} \tilde{p}_{ubgc} - \frac{\omega_b}{1+\omega_b}\Delta^{k_{bg}} \tilde{s}_{ubgc}) \\ &= \Delta^{k_{bg}} \tilde{s}_{ubgc}\Delta^{k_{bg}} \tilde{p}_{ubgc}\left(1 + \frac{(1-\sigma_b)\omega_b}{1+\omega_b}\right) - (1-\sigma_b)(\Delta^{k_{bg}} \tilde{p}_{ubgc})^2 \\ &- \frac{\omega_b}{1+\omega_b}(\Delta^{k_{bg}} \tilde{s}_{ubgc})^2. \end{aligned} \tag{C.9}$$

[15] Handbury and Weinstein (2015) use a double-differentiation, that is $\Delta^{k_{bg}} x_{ubgc} = [\ln(x_{ubgc}) - \ln(x_{k_{bg}bgc})] - [\ln(x_{ubg.}) - \ln(x_{k_{bg}bg.})]$ for any variable x_{ubgc}. The reason is that they consider a slightly different specification based on Broda and Weinstein (2010), where prices and shares are not "quality-adjusted".

Rearranging, we have

$$\underbrace{(\Delta^{k_{bg}}\tilde{p}_{ubgc})^2}_{Y_{ubgc}} = -\frac{\omega_b}{(1+\omega_b)(1-\sigma_b)}\underbrace{(\Delta^{k_{bg}}\tilde{s}_{ubgc})^2}_{X^1_{ubgc}}$$
$$+\frac{1+2\omega_b-\sigma_b\omega_b}{(1+\omega_b)(1-\sigma_b)}\underbrace{\Delta^{k_{bg}}\tilde{s}_{ubgc}\Delta^{k_{bg}}\tilde{p}_{ubgc}}_{X^2_{ubgc}}\underbrace{-\varepsilon^{k_{bg}}_{ubgc}\delta^{k_{bg}}_{ubgc}/(1-\sigma_b)}_{\upsilon_{ubgc}}. \tag{C.10}$$

Broda and Weinstein (2010, footnote 28) use the following reparameterization

$$\omega_b = \gamma_b/(\sigma_b(1-\gamma_b)-1), \tag{C.11}$$

so that we have

$$Y_{ubgc} = \underbrace{\frac{\gamma_b}{(\sigma_b-1)^2(1-\gamma_b)}}_{\theta_1}X^1_{ubgc} + \underbrace{\frac{(2\gamma_b-1)}{(\sigma_b-1)(1-\gamma_b)}}_{\theta_2}X^2_{ubgc} + \upsilon_{ubgc}. \tag{C.12}$$

Since prices and shares are correlated with the errors $\varepsilon^{k_{bg}}_{ubgc}$ and $\delta^{k_{bg}}_{ubgc}$, X^1_{ubgc} and X^2_{ubgc} are correlated with υ_{ubgc}. Feenstra (1994) shows that a consistent estimator can be obtained by averaging (C.12) over time. Removing the variations within each living zone, we have $\mathbb{E}(\overline{X}^1_{ubga}\overline{\upsilon}_{ubga}) = 0$ and $\mathbb{E}(\overline{X}^2_{ubga}\overline{\upsilon}_{ubga}) = 0$, where the upper bar denotes the sample mean. Then, assuming that $\mathbb{E}(\overline{\upsilon}_{ubga}) = 0$ implies that the between estimator of (C.12) provides consistent estimates of θ_1 and θ_2. Let $\widehat{\theta}_1$ and $\widehat{\theta}_2$ denote these estimates, which can be obtained by applying the Weighted Least Squares (WLS) estimator to the transformed equation

$$\overline{Y}_{ubga} = \theta_1\overline{X}^1_{ubga} + \theta_2\overline{X}^2_{ubga} + \overline{\upsilon}_{ubga}, \tag{C.13}$$

where the share of expenditures on u in brand module b and living zone a, $\overline{\tilde{s}}_{ubga}$, is used as the weight. If $\overline{X}^1_{ubga}$ and $\overline{X}^2_{ubga}$ are not asymptotically collinear, then θ_1 and θ_2 are separately identified from (C.13). Moreover, adding a constant term to the regression renders these estimates consistent even when the unit values are measured with errors. This estimator corresponds to Hansen's (1982) Generalized Method of Moments (GMM) estimator, where the moment condition $\mathbb{E}(\upsilon_{ubgc}) = 0$ is approximated by choosing θ_1 and θ_2 to minimize the weighted sum of squared sample moments $\overline{\upsilon}_{ubga}$. It is also equivalent to applying an Instrumental Variable (IV) estimator to equation (C.12), assuming $\mathbb{E}(\upsilon_{ubgc}|a) = 0$ and therefore using all living zone fixed effects to instrument X^1_{ubgc} and X^2_{ubgc} (see Feenstra, 1994).

It is then possible to recover σ_b and γ_b from $\widehat{\theta}_1$ and $\widehat{\theta}_2$. Feenstra (1994) shows in his Proposition

2 that as long as $\widehat{\theta}_1 > 0$ the estimates of σ_b and γ_b are as follows

$$\widehat{\sigma}_b = 1 + \left(\frac{2\widehat{\gamma}_b - 1}{1 - \widehat{\gamma}_b}\right)\frac{1}{\widehat{\theta}_2}$$
$$\widehat{\gamma}_b = \frac{1}{2} \pm \left(\frac{1}{4} - \frac{1}{4 + \left(\widehat{\theta}_2^2/\widehat{\theta}_1\right)}\right)^{1/2}, \tag{C.14}$$

the plus and minus signs in the last expression applying for $\widehat{\theta}_2 > 0$ and $\widehat{\theta}_2 < 0$, respectively. As $\widehat{\theta}_2 \to 0$, then $\widehat{\gamma}_b \to 1/2$ and $\widehat{\sigma}_b \to 1 + \widehat{\theta}_1^{-1/2}$. For all brand-modules but three, $\widehat{\theta}_1 > 0$ so that we can use Feenstra's (1994) formulae. In the remaining brand modules, $\widehat{\theta}_1 < 0$ and we follow Broda and Weinstein (2006, 2010): we perform a grid-search over values of $\sigma_b > 1$ and $\gamma_b > 0$, and retain the values minimizing the GMM objective function, where the residuals, $\overline{v}_{ubga}$ for WLS and v_{ubgc} for GMM, are weighted by their corresponding shares, $\overline{\tilde{s}}_{ubga}$ and $\tilde{s}_{ubgc}$ respectively. The objective function is evaluated for $\sigma_b \in [1.05; 131.5]$ at intervals 0.05 apart, and for $\gamma_b \in [0.01; 1]$ at intervals 0.01 apart. Only the combinations of σ_b and γ_b that imply $\sigma_b > 1$ and $\omega_b > 0$ (where ω_b is given by (C.11)) are used. The grid-search and Feenstra's original method lead to very similar results when $\widehat{\theta}_1 > 0$. The standard errors in all cases can be obtained by bootstrapping.

We can then calculate the price indices for each brand module, $\tilde{P}_{bgc}$, which are given by the formula

$$\tilde{P}_{bgc} = \left(\sum\nolimits_{u \in U_{bc}} (\tilde{p}_{ubgc})^{1-\sigma_b}\right)^{\frac{1}{1-\sigma_b}}, \tag{C.15}$$

and apply the same procedure to estimate the across brand-module elasticities, σ_g.

Table C.3 describes the distribution of the estimates for the within parameters (especially elasticities, σ_b) and the across parameters (especially elasticities, σ_g) (the estimates obtained for each brand module are available from the authors). These are obtained using all available data (2008-2013). Out of the 81 brand-modules, 8 include only one UPC. For these singletons, we cannot obtain estimates of σ_b and γ_b. These are set to zero, so that the prices of the corresponding brand modules do not affect the EPI of the group. Over all four product groups, the median within-brand-module elasticity is 5.85, so that a 1% increase in the price of a UPC within a brand module reduces its sales on average by 5.85%. The larger the elasticity, the more substitutes are the UPCs within a brand module. It is hence not surprising to find a positive correlation between the within-elasticity of a brand module and the number of distinct UPCs in that brand module: more available alternatives yield higher elasticities of substitution. Omitting both the zeroes and the six largest values (over 20), the distribution of σ_b looks log-normal, as shown in Figure C.1. The median within brand-module elasticity is almost the same for SSBs and NCSBs (5.48 and 5.39, respectively); it is larger for USBs (9.59) and smaller for Water (4.57). The values that are used below for σ_b are those obtained over the whole period but it is worth noting that there is not much change in the values estimated before and after the tax: the distribution is log-normal in both cases and the quartiles are 4.41 (4.23), 6.25 (6.27) and 10.31 (11.03) in 2008-2011 (2012-2013) – see Figure C.2.

Regarding the across-brand-module elasticities, σ_g, the larger the elasticity, the closer substitutes are brand modules within a group. As can be seen at the bottom of Table 1, a distinction can be made between groups. SSBs and NCSBs are characterized by rather large elasticities (6.04 and 6.69, respectively), showing that they are both composed of brand modules that are highly substitutable, at least more than those composing USBs and Water (3.35 and 3.13, respectively). These low elasticities may be explained by the smaller number of brand modules in USBs and Water than in SSBs and NCSBs. In addition, the USB group groups together very heterogeneous beverages (juices, syrups, pulps and milk-based drinks). The Water group is apparently more homogeneous, but there is still differentiation between sparkling and still waters and we also observe that brand loyalty is high.

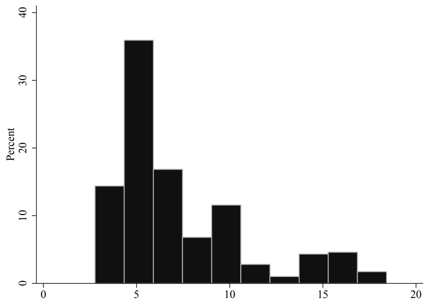

Figure C.1: **Distribution of σ_b, 2008-2013**

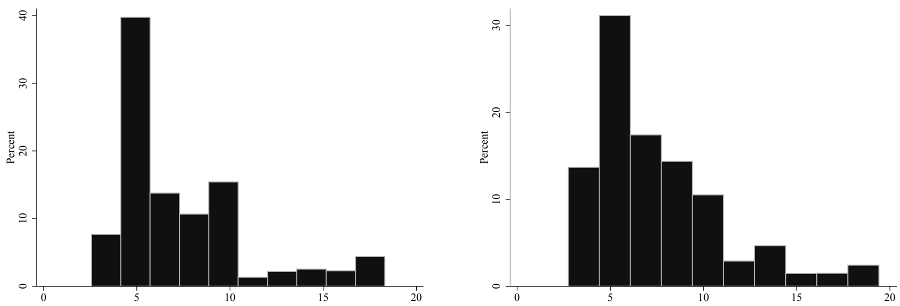

Figure C.2: **Distribution of σ_b, 2008-2011 (left) vs. 2012-2013 (right)**

Table C.1: CES elasticity estimates

	SSB			NCSB			USB			Water		
Within (81 brand-modules)	#UPC	σ_b	γ_b	#UPC	σ_b	γ_b	#UPC	σ_b	γ_b	#UPC	σ_b	γ_b
Percentile 1	1	0	0	1	0	0	1	0	0	12.0	3.30	0.01
Percentile 5	1	0	0	1	0	0	1	0	0	12.0	3.30	0.01
Percentile 10	2	3.07	0.27	1	0	0	20.5	5.89	0.26	12.0	3.36	0.04
Percentile 25	9.6	4.52	0.39	1.3	3.04	0.06	33.0	7.17	0.53	12.0	4.15	0.21
Percentile 50	13.4	5.48	0.59	5.9	5.39	0.30	27.7	9.59	0.57	12.0	4.57	0.41
Percentile 75	12.4	9.81	0.68	6.2	10.61	0.59	29.3	13.73	0.69	14.9	8.37	0.58
Percentile 90	11.7	14.90	0.84	5.8	34.37	0.91	26.1	16.84	0.87	13.4	14.81	0.75
Percentile 95	11.6	17.93	0.89	5.7	37.19	0.92	26.0	34.37	0.88	13.0	18.42	0.90
Percentile 99	11.4	186.55	0.98	5.8	122.35	0.96	26.0	34.37	0.88	13.0	18.42	0.90
Average	11.1	12.02	0.54	5.8	13.76	0.35	26.0	11.49	0.57	13	6.80	0.40
Across (four product groups)	#B-M	σ_g	γ_g	#B-M	σ_g	γ_g	#B-M	σ_g	γ_g	#B-M	σ_g	γ_g
	36	6.04	0.45	22	6.69	0.51	13	3.35	0.01	10	3.13	0.11

Notes: These are within brand-module elasticities (σ_b) and between brand-module elasticities for each product group (σ_g). Kantar Worldpanel data 2008-2013. #UPC is the average number of distinct UPCs for brand modules at a given percentile of the distribution of within brand-module elasticities.

D Additional results

D.1 Estimating the price response of demand

Our difference-in-difference estimation strategy uses Water as a counterfactual for SSB. This is possible insofar as Water is not a substitute for SSB since, otherwise, the impact of taxation on equilibrium SSB prices would depend on trends in Water prices. To delineate the boundaries of substitution in the market for SSB, we specify an Almost Ideal Demand System (AIDS) for the four groups of non-alcoholic beverages (Deaton & Muellbauer, 1980a). The dependent variables are the (market average) budget shares of SSBs, NCSBs, USBs and Water, and the explanatory variables are the logarithms of EPI, the logarithm of total expenditure on non-alcoholic beverages deflated by the AIDS aggregated price index, and controls for macro shocks (year and month dummies) and demographics across markets. The logarithm of real total expenditure is instrumented by the logarithm of average real household income, thus allowing for income effects. Homogeneity and symmetry constraints are imposed (Lecocq & Robin, 2015). The model is estimated using the pre-tax years only (2008-2011).

We can safely assume that prices are exogenous here as the local price indices are adjusted for consumer, product and retailer heterogeneity. In principle, one could test this assumption through an instrumental variable strategy. The literature has used the prices in adjacent locations as instruments (Hausman, Leonard, & Zona, 1994; Zhen, Finkelstein, Nonnemaker, Karns, & Todd, 2013) or time variations in production input costs (Bonnet & Réquillart, 2013a). But as expected, since we have four prices to instrument and the correlations among these instruments are high, these instrument sets did not pass standard weak instrument tests.

The upper panel of Table D.1 lists the estimated coefficients for budget shares, while the lower panel shows the corresponding Marshallian elasticities for quantities. The own-price elasticities of SSBs and NCSBs are large and significant, -0.87 and -0.85 respectively. Interestingly, the Marshallian cross-price elasticities between SSBs and NCSBs are negative and marginally significant. An increase in SSBs price lowers NCSBs consumption. A change in soft-drink prices has no impact on the consumption of Water, so that the relevant market for soft-drinks includes USBs but not Water. They also imply that the soft-drink tax had a large negative effect on soft-drink consumption, with beneficial health consequences in terms of sugar intake.

Table D.1: The response of quantity demanded to price

	SSB	NCSB	USB	Water
Price effects				
SSB	0.040***	-0.013	-0.016**	-0.011
	(0.009)	(0.010)	(0.008)	(0.008)
NCSB	-0.013**	0.028***	-0.008	-0.008
	(0.006)	(0.007)	(0.006)	(0.005)
USB	-0.016**	-0.008	0.061***	-0.038***
	(0.008)	(0.009)	(0.006)	(0.007)
Water	-0.011	-0.008	-0.038***	0.057***
	(0.008)	(0.009)	(0.007)	(0.007)
Budget effects	0.043***	0.069***	-0.079***	-0.033***
	(0.010)	(0.011)	(0.009)	(0.010)
Price elasticities				
SSB	-0.872***	-0.112*	-0.019	-0.027
	(0.034)	(0.051)	(0.027)	(0.031)
NCSB	-0.043*	-0.846***	-0.035*	-0.031
	(0.020)	(0.031)	(0.016)	(0.019)
USB	-0.144***	-0.228***	-0.625***	-0.078*
	(0.037)	(0.056)	(0.030)	(0.034)
Water	-0.103***	-0.171***	-0.031	-0.742***
	(0.028)	(0.043)	(0.023)	(0.026)

Notes: These results come from the estimation of an Almost Ideal Demand System (Deaton and Muellbauer, 1980). Kantar Worldpanel data 2008-2011. The observation unit is a market (a living zone in a month); there are 11,779 observations. The dependent variables are the (market-average) budget shares on SSBs, NCSBs, USBs and Water in the upper panel, and the corresponding quantities in the lower panel. The independent variables are the logarithms of the Exact Price Indices, the logarithm of total expenditure on non-alcoholic beverages deflated by the AI aggregated price index, and control variables for macro shocks (year and month dummies) and market demographics (average household size, average age of the main shopper, proportion of households where the main shopper is a male, and the proportion of households in four household structures). The logarithm of real total expenditure is instrumented by the logarithm of average real household income. Homogeneity and symmetry constraints are imposed. Standard errors are in parentheses; ***, ** and * indicate significance at the 1%, 5% and 10% levels.

D.2 Product-level pass-through rates

The product-level pass-through rates are constructed from estimates of UPC-level pass-throughs. We first fit a before-after model separately for each UPC in order to identify the impact of the tax T

$$\ln(p_{ubgc}) = \delta_u Post_{t\geq 2012} + \delta_{y,u} + \delta_{m,u} + \gamma_u C_t + \delta_{a,u} + \epsilon_{u,c}, \tag{D.1}$$

where UPC-market observations are weighted by market sales and C_t is the retail price of sugar. To avoid the influence of outliers, the dependent variable is not the log-mean but the log of the median unadjusted UPC price observed in each cluster. Sample weights are taken into account.

From an ex-ante perspective, the pass-through is defined relative to the expected impact of the tax on the prices observed in 2011, so that we have

$$\rho_{u,2011} = \frac{\overline{p_{ubgc}}^{y=2011}\left[exp(\delta_u) - 1\right]}{T} \approx \frac{\overline{p_{ubgc}}^{y=2011}\delta_u}{T}, \tag{D.2}$$

where $\overline{p_{ubgc}}^{y=2011}$ is the average of cluster UPC prices observed in 2011. From an ex-post perspective, the pass-through is defined relative to the prices observed in 2011 purged of the specific impact of the tax

$$\rho_{u,2012} = \frac{\overline{p_{ubgc}}^{y=2012}\left[1 - \frac{1}{exp(\delta_u)}\right]}{T} \approx \frac{\overline{p_{ubgc}}^{y=2012}\delta_u}{exp(\delta_u) \times T}, \tag{D.3}$$

where $\overline{p_{ubgc}}^{y=2012}$ is the average of cluster UPC price in 2012, which takes into account all the factors that affected the change in prices between 2011 and 2012. The excise tax of 0.0716 Euro/Liter applies to producer prices; it should be multiplied by 1.055 (5.5% being the VAT rate) to obtain the tax passed into consumer prices with 100% pass-through. Hence $T = 0.075538$.

In practice $\rho_{u,2011}$ and $\rho_{u,2012}$ are very similar. We choose to take the ex-ante perspective.

Table D.2 reports the pass-through estimated at the UPC level. The results show that the pass-through were, on average, slightly higher for SSB (36.4%) than for NCSB (32.0%). However, the ranking of pass-throughs across brands is similar, with higher pass-through for retailer brands.

D.3 Decomposition of the distributional incidence of the tax

The full incidence of the soft-drink tax policy can be measured through the associated compensating variation (CV), which is the amount of additional income that is needed to keep utility constant after the passing of the tax to consumer prices. Given that the tax had small estimated effects on purchases (-4.5% and -3.9% for SSB purchases of low- and high-income households respectively), we can use the standard first-order approximation CV formula to compare the differential in welfare variation between a representative household of population $\mathcal{P}_1$, and a representative household of population $\mathcal{P}_2$. This difference will depend on three key elements: the residential sorting of the population across markets; the average quantities purchased by each population on each market ; the incidence of the

Table D.2: **Tax incidence: pass-throughs at the UPC level (% points)**

	SSB		NCSB	
	#UPC	Pass-through (%)	#UPC	Pass-through (%)
All	400	36.4	127	32.0
Top national	136	19.2	86	23.9
Other national	113	48.5	17	66.0
Retailer	108	47.4	18	38.8
Hard discount	43	33.5	6	35.4
Standard Coca-Cola	31	38.5		
Diet Coke/Coke Light			28	35.3
Coca-Cola Zero			17	12.4

Notes: This table reports average pass-through rates that are calculated from UPC-specific pass-through rates for the set of UPCs indicated in the first column. Each UPC-specific pass-through is estimated using a before-after specification similar to specification (2) in Table 2, where the dependent variable is the median unadjusted UPC price observed in each market, and observations are weighted by market-specific sales.

tax on consumer EPI in each market. To compare the relative importance of these three elements, we now derive a decomposition of the difference in tax incidence. We let $\omega_c^{\mathcal{P}}$ be the probability that a representative household of population $\mathcal{P}$ resides in market c. This household purchases quantities $Q_{gc}^{\mathcal{P}}$ of product category g, with the corresponding aggregate price being denoted by $\mathbb{P}_{gc}^{\mathcal{P}}$. For this household, the approximate compensating variation associated to the tax can thus be written as:

$$CV_c^{\mathcal{P}} \simeq \sum_g Q_{gc}^{\mathcal{P}} \Delta \mathbb{P}_{gc}^{\mathcal{P}}$$

where $\Delta \mathbb{P}_{gc}^{\mathcal{P}}$ measures the incidence of the tax on the aggregate price. This incidence is equal to the variation in the EPI times the reference price.

We can now write an approximate decomposition of the differential in tax incidence between populations $\mathcal{P}_1$ and $\mathcal{P}_2$:

$$\begin{aligned}
CV^{\mathcal{P}_2} - CV^{\mathcal{P}_1} &= \sum_c \omega_c^{\mathcal{P}_2} CV_c^{\mathcal{P}_2} - \sum_c \omega_c^{\mathcal{P}_1} CV_c^{\mathcal{P}_1} \\
&= \sum_c \left(\omega_c^{\mathcal{P}_2} - \omega_c^{\mathcal{P}_1}\right) CV_c^{\mathcal{P}_2} + \sum_c \omega_c^{\mathcal{P}_1} \left(CV_c^{\mathcal{P}_2} - CV_c^{\mathcal{P}_1}\right) \\
&\simeq \sum_c \left[\left(\omega_c^{\mathcal{P}_2} - \omega_c^{\mathcal{P}_1}\right) \sum_g Q_{gc}^{\mathcal{P}_2} \Delta\mathbb{P}_{gc}^{\mathcal{P}_2}\right] + \sum_c \omega_c^{\mathcal{P}_1} \left(\sum_g Q_{gc}^{\mathcal{P}_2} \Delta\mathbb{P}_{gc}^{\mathcal{P}_2} - \sum_g Q_{gc}^{\mathcal{P}_1} \Delta\mathbb{P}_{gc}^{\mathcal{P}_1}\right) \\
&\simeq \sum_c \left[\left(\omega_c^{\mathcal{P}_2} - \omega_c^{\mathcal{P}_1}\right) \sum_g Q_{gc}^{\mathcal{P}_2} \Delta\mathbb{P}_{gc}^{\mathcal{P}_2}\right] + \sum_c \omega_c^{\mathcal{P}_1} \left(\sum_g \left(Q_{gc}^{\mathcal{P}_2} - Q_{gc}^{\mathcal{P}_1}\right) \Delta\mathbb{P}_{gc}^{\mathcal{P}_2}\right) \\
&\qquad + \sum_c \omega_c^{\mathcal{P}_1} \left(\sum_g Q_{gc}^{\mathcal{P}_1} \left(\Delta\mathbb{P}_{gc}^{\mathcal{P}_2} - \Delta\mathbb{P}_{gc}^{\mathcal{P}_1}\right)\right)
\end{aligned}$$

This expression is the sum of three elements. The first element,

$$\sum_c \left[\left(\omega_c^{\mathcal{P}_2} - \omega_c^{\mathcal{P}_1}\right) \sum_g Q_{gc}^{\mathcal{P}_2} \Delta\mathbb{P}_{gc}^{\mathcal{P}_2}\right] \tag{D.4}$$

is produced by the residential sorting of the populations across living zones, through the difference in residence probabilities $\left(\omega_c^{\mathcal{P}_2} - \omega_c^{\mathcal{P}_1}\right)$. The second element,

$$\sum_c \omega_c^{\mathcal{P}_1} \left(\sum_g \left(Q_{gc}^{\mathcal{P}_2} - Q_{gc}^{\mathcal{P}_1}\right) \Delta\mathbb{P}_{gc}^{\mathcal{P}_2}\right) \tag{D.5}$$

depends on the difference in purchases quantities between the two populations. This is likely to be driven primarily by differences in preferences for quantities; yet, it remains possible that the two populations face different aggregate price indices because they have different preferences for products within beverage groups (different tastes). Such differences in preferences for products will partly translate into differences in price indices, and therefore quantities. The third element,

$$\sum_c \omega_c^{\mathcal{P}_1} \left(\sum_g Q_{gc}^{\mathcal{P}_1} \left(\Delta\mathbb{P}_{gc}^{\mathcal{P}_2} - \Delta\mathbb{P}_{gc}^{\mathcal{P}_1}\right)\right) \tag{D.6}$$

depends on the differential in incidence of the tax on consumer prices between a low- and high-income households, consuming the same quantity and living in the same market. It is driven by population differences in preferences for products within beverage groups, i.e. preference for quality.

Table D.3 reports the results of the decomposition of the difference in welfare loss between low- and high-income households ($\mathcal{P}_1$: low-income households; $\mathcal{P}_2$: high-income households). To produce these results, we have estimated models for market heterogeneity (as in Table D.5, specification 3) by income groups. The estimates thus provides tax incidence by market for each income group, as a function of market affluence and concentration. The market effect is significant, especially when we also include NCSB in the computation.

Table D.3: Decomposition of the difference in welfare loss across income groups

	SSB		SSB & NCSB	
Compensating variations	cts of €/capita/year			
Low-income households	52.04		49.29	
High-income households	33.55		65.57	
Difference Low/High	+18.49		+16.28	
Decomposition	cts of €/cap/year	% of the difference	cts of €/cap/year	% of the difference
Market	6.63	35.8%	8.41	51.7%
Quantity	8.95	48.5%	5.69	34.9%
Quality	2.91	15.7%	2.18	13.4%

Notes: This table reports the estimated welfare losses for the low-income and high-income households, in cts of €/cap/year, using a compensating variation welfare measure. In the lower panel, we report the results of the decomposition of the difference between income groups. There are three components: a market heterogeneity effect corresponding to equation D.4; a quantity effect, corresponding to equation D.5; a quality effect corresponding to D.6.

D.4 Group level pass-through rate: formula

At the level of the SSB group g, we define the pass-through ρ_{gc} as the ratio of the estimated change in quality-adjusted unit cost $\tilde{P}_{gc}$ produced by the tax to the average change that would have been observed had the tax been fully shifted into UPC prices. Given the relationship between the quality-adjusted unit cost and the price index, we have

$$\rho_{gc} = \frac{EPI_{gc}^{y=2011}\left[exp\left(\delta\right)-1\right]}{EPI_{gc}^{y=2012,*}-EPI_{gc}^{y=2011}}. \tag{D.7}$$

In the numerator, δ is the estimated impact of the tax on a price index observed in 2011, $EPI_{gc}^{y=2011}$ (in percentage points); $EPI_{gc}^{y=2011} \times exp\left(\delta\right)$ is the EPI that would have been observed in 2012 in the same living zone and same month, had nothing other than the tax policy occurred. In the denominator, $EPI_{gc}^{y=2012,*}$ is the EPI that would have been observed in 2012 had the tax been fully shifted into UPC prices, no behavioural response had happened, and no other changes had occurred. We can rewrite the pass-through as

$$\rho_{gc} = \frac{\left[exp\left(\delta\right)-1\right]}{\frac{EPI_{gc}^{y=2012,*}}{EPI_{gc}^{y=2011}}-1}. \tag{D.8}$$

To construct $EPI_{gc}^{y=2012,*}$, counterfactual household specific prices p^*_{ucrh} are calculated under the following assumptions: (i) a 100% pass-through, *i.e.* $p^*_{ucrh} = p_{ucrh} + T$, where p_{ucrh} is predicted from equation D.1 (setting $\delta_u = 0$) ; (ii) the average subjective quality on each market is constant, so that the ratio of adjusted to observed UPC prices does not change. We then obtain the counterfactual quality-adjusted prices as $\tilde{p}^*_{ucrh} = (p^*_{ucrh}/p_{ucrh}) \times \tilde{p}_{ucrh}$. We can finally construct the desired counterfactual price index.

D.5 Results for Non-Calorically Sweetened Beverages (NCSB)

The health risks and benefits of artificially-sweetened beverages are still debated in the public-health literature (see, e.g. Borges et al., 2017). In the French case, the decision to create a twin tax on NCSBs was not motivated by public health reasons (Le Bodo, Etilé, Gagnon, & de Wals, 2017). NCSBs were included at the end of a political process that started in August 2011, when the government announced the creation of a SSB tax to fight children obesity. After several rounds of discussions between the government, the parliament and the industry, an agreement was reached. The original public-health motivation for the tax — fighting obesity — became secondary, and the legal text focussed on a fiscal motivation: raising revenue for Social Security and the farming sector. NCSB were included in the fiscal basis as a "voluntary" contribution of the beverage industry to Social Security. In July 2018, the tax schemes changed. The unit tax on SSBs is now increasing with their sugar content, with a floor rate of 0.03 Euro/Litre that applies to all soft-drinks with less than 1 kg of added sugar per hectolitre (including NCSBs).

The following tables provide our estimates of the tax incidence on NCSB prices. We apply exactly the same method as for SSBs. To ease comparisons, results for both groups are reported.

Table D.4 displays estimates of the national average tax incidence, for all households, low-income households, and high-income households. The upper panel shows that the average tax incidence for NCSBs (+4.2%) is similar to the incidence for NCSBs. The medium and lower panels show however differences in incidence for NCSBs between low- and high-income groups. The incidence of the tax on NCSB prices was higher for low-income households, both in the before-after and difference-in-difference regressions. This indicates that income-related preference heterogeneity across income groups have significant distributional consequences in terms of welfare for NCSB consumption.

Table D.5 reports the estimates of the moderating impact of market heterogeneity. We find no significant direct or interaction effects of income and market size on NCSB prices. However, in specification (3), the HHI dummy has a large, significant and positive effect on tax incidence. In concentrated markets, tax incidence is about 33% higher for NCSBs, a larger effect than that found for SSBs (+12%). Specifications (5) reveals that the average tax incidence is higher in low-income markets both for SSBs and NCSBs (5.55% for SSBs and 5.15% for NCSBs). We do not find a stronger effect of competition in poorer areas for NCSBs.

Table D.4: **Tax incidence: Price-index variation (% points) - national average**

	Before-after			DiD	
	(0)	(1)	(2)	2012	2011
All households					
SSB	4.430***	5.426***	4.144***	4.261***	0.570**
	(0.263)	(0.171)	(0.232)	(0.402)	(0.214)
NCSB	4.154***	5.217***	4.248***	3.162***	0.579
	(0.261)	(0.366)	(0.343)	(0.493)	(0.303)
Low-income households					
SSB	5.042***	5.946***	4.575***	4.174***	0.350
	(0.923)	(0.182)	(0.466)	(0.581)	(0.319)
NCSB	3.815***	6.087***	4.683***	3.831***	0.449
	(0.340)	(0.321)	(0.413)	(0.826)	(0.398)
High-income households					
SSB	3.912***	5.049***	3.597***	4.383***	0.500*
	(0.293)	(0.217)	(0.344)	(0.539)	(0.280)
NCSB	3.815***	4.928***	3.801***	2.646***	0.173
	(0.340)	(0.427)	(0.431)	(0.641)	(0.405)
Differential incidence: low-income EPI minus high-income EPI (log)					
SSB		0.919***	1.001*	-0.271	-0.152
		(0.302)	(0.574)	(0.810)	(0.437)
NCSB		1.145**	0.904	1.208	0.464
		(0.475)	(0.680)	(1.114)	(0.558)
Adjustment for preference heterogeneity	No	Yes	Yes	Yes	Yes
Additional controls					
Sugar price (in log) for SSB and NCSB	No	No	Yes	Yes	Yes
Group-specific month effects	No	No	No	Yes	Yes
Period	2008-2013	2008-2013	2008-2013	2010-2012	2009-2011

Notes: The dependent variable is the log of EPI. The EPI is estimated from Kantar Worldpanel data 2008-2013 using market-level observations (living zone-month). In column 1, it is not adjusted for preference heterogeneity. In columns (2-5), it is adjusted for within-population preference heterogeneity. These estimates represent changes in % points, between 2011 and 2012 (before-after columns), and the difference in the changes between SSB/NCSB and Water (DiD: Difference-in-Difference columns). The DiD-2011 column is a placebo test, focussing on the 2010-2011 change. Each observation is weighted by the population-specific share of the national sales in the market in 2011. For estimating the differential incidence, the weights correspond to the low-income population. Living zone fixed effects are included. Standard errors are in parentheses; ***, ** and * indicate significance at the 1%, 5% and 10% levels.

Table D.5: The heterogeneity of tax incidence across markets (% points)

	SSB					NCSB				
	(1)	(2)	(3)	(4)	(5)	(1)	(2)	(3)	(4)	(5)
Post	4.940***	4.722***	4.799***	4.954***	5.555***	4.046***	3.980***	4.120***	4.088***	5.146***
	(0.387)	(0.457)	(0.449)	(0.387)	(0.561)	(0.722)	(0.849)	(0.835)	(0.722)	(1.151)
$\times \ln(Income)$		-3.077*	-3.932**				-2.162	-3.659		
		(1.761)	(1.656)				(3.243)	(3.046)		
$\times \ln(N_{cu})$		-0.233*					-0.365			
		(0.121)					(0.226)			
$\times 1_{HHI>2000}$			0.556**	0.551**	1.130**			1.372***	1.338***	-0.121
			(0.249)	(0.244)	(0.488)			(0.458)	(0.450)	(0.979)
$\ln(Income)$		8.901*	8.614*				4.910	3.456		
		(5.014)	(4.895)				(9.240)	(9.037)		
$\ln(N_{cu})$		-2.727					-0.771			
		(4.932)					(9.207)			
$1_{HHI>2000}$			1.142***	1.148***	1.234**			-0.074	-0.058	-0.481
			(0.322)	(0.321)	(0.492)			(0.598)	(0.597)	(1.009)
Sample	Full	Full	Full	Full	Inc<Q(50)	Full	Full	Full	Full	Inc<Q(50)

Notes: The dependent variable is the log of EPI. The EPI is estimated from Kantar Worldpanel data 2008-2013. The estimated impacts in % points come from a before-after specification (observations are *not* weighted by market-specific sales). N_{cu}: number of Consumption Units (cu) in each market (INSEE census data). $Income$: market average of the median real equivalent income in the market's postcodes (INSEE fiscal data). HHI is a Herfindahl-Hirschman index based on the sales area of retailers (TradeDimensions data). The European Commission considers that a HHI greater than 2000 reflects horizontal-competition concerns (Official Journal C 31 of 05/02/2004). All of these variables vary across markets c, *i.e.* across areas a and periods t. All estimates include area, month and year fixed effects. Full sample: $N = 18,927$ (SSB), 17,453 (NCSB) living zone-month observations. The sample Inc<Q(50) contains only markets where the median income is below the median figure ($N = 9,466$ (SSB), 8,761 (NCSB) living zone-month observations). Standard errors are clustered at the area-level in parentheses; ***, ** and * indicate significance at the 1%, 5% and 10% levels.

D.6 Intra-group heterogeneity in sugar content

One important policy issue is whether the SSB group is heterogeneous in terms of sugar content. If this were the case, then a soda tax may induce a considerable amount of substitution from expensive products (top-national brands) to less expensive, but perhaps more sugary, products (retailer own brands).

The left panel in Figure D.1 displays the distribution of the sugar density of SSB products, both weighted by market sales (black) and unweighted (grey). The X-axis is in g/L. The grey distribution shows that most products have a sugar density of between 5 g/L and 11 g/L; the black distribution shows that households purchase products that are very homogeneous in terms of sugar content, with about 80% of purchases being concentrated at around 10-11 g/L. The right panel in Figure D.1 plots the concentration curve of SSB market shares against their sugar density. There is one curve for the pre-policy year (2011) and one for the post-policy year (2012). These coincide, so that the policy did not affect the distribution of the sugar density of SSB purchases: households did not switch to more or less sugary SSB products.

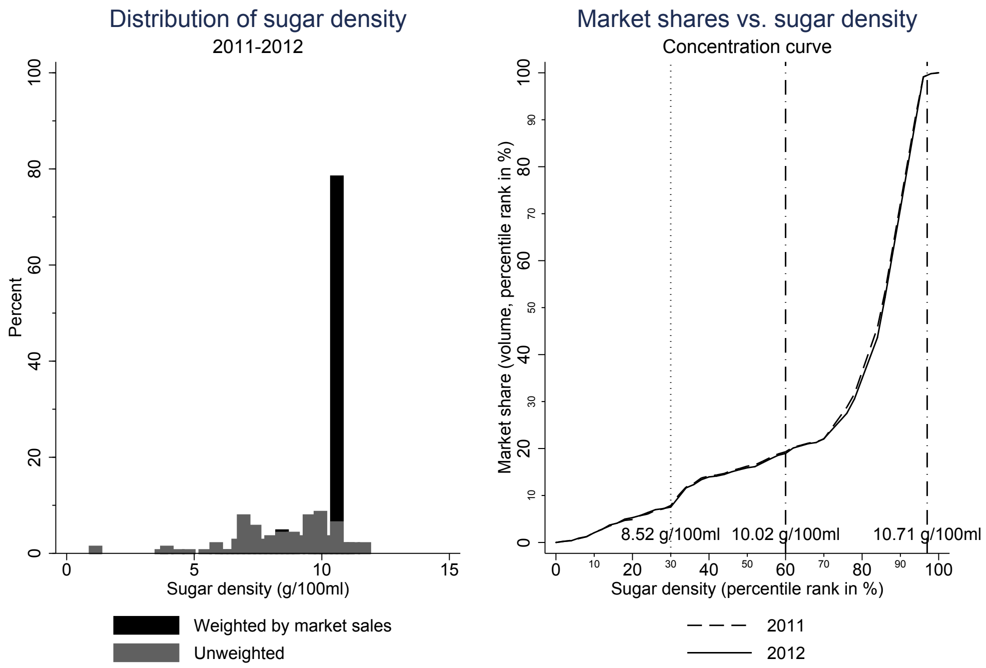

Figure D.1: Distribution of the sugar density of SSB products, 2011-2012

Notes: Kantar Worldpanel data 2011-2012. The left panel shows the distribution of the sugar density of products in 2011. The histogram in grey shows the unweighted data, while that in black weights the products by market sales. The right panel shows the concentration curve of the market shares of products (Y-axis) ranked by their sugar density (X-axis). The curves are shown for 2011 and 2012.

This finding is in line with the ex-ante evaluation results in Bonnet and Réquillart (2013b), who use a mixed multinomial logit model explicitly taking into account all substitutions between all SSB and NCSB product varieties. Their simulation results show substantial effects of tax policies, but these are explained uniquely by substitutions from SSBs to NCSBs and USBs (the outside option in their model). Substitutions within SSBs plays no role. More specifically, their model predicts that a VAT increase from 5.5% to 19.6% yields an average reduction of added sugar intake of 352 g/capita/year (-21%). An excise tax of 0.09 Euro cents per 100g of sugar produces an average fall of 629 g/capita/year (-38%). However, the ex-post average sugar density of SSB varieties predicted for these two policies are respectively 92.0 g/L and 92.6 g/L, as against 92.1 g/L before the policies (see their Appendix).

One straightforward consequence is that the effectiveness of soda taxes in France depends on substitution between groups, *i.e.* from SSBs to NCSBs, and USBs. As such, the pass-through of the tax should be measured at the aggregate level of the group, with a price measure that accurately reflects the impact of the tax on consumer utility from purchasing SSBs.

D.7 The timing of tax shifting

The timing of the tax incidence can be analysed via an event study, by adding particular month effects for 2012 to the second specification in Table 2. The implicit concept of pass-through here is the change in price resulting from the taxation shock to costs in January 2012, the effect of which may be felt with some lags (Gopinath & Itskhoki, 2010; Nakamura & Zerom, 2010). The estimated coefficients appear in Figure D.2. Each point here is the observed gap in 2012 from the usual month-of-the-year effect, with December 2011 being the absolute reference. The horizontal line represents the effect estimated in the second column of Table 2, *i.e.* the yearly average for 2012. The average prices in January are similar to those observed usually in Januaries. This is as expected, and is explained by the fall in the value of Christmas inventories owned by retailers that leads them to propose "clearance prices" (sales) to consumers (Smith & Achabal, 1998; Gupta, Hill, & Bouzdine-Chameeva, 2006). The prices of both SSBs and NCSBs then increase, but do not significantly vary between February and April, increase again in May, and then return to the 2012 average. The subsequent increases observed in October and November cannot be attributed to the tax.

This analysis suggests that the tax was passed on quite rapidly to consumer prices, after one quarter. This is unsurprising given that, over 2008-2013, the contractual framework between manufacturers and retailers was regulated, with annual negotiations that had to be resolved by the end of March. The price levels reached in March-April 2012 are similar to our earlier results in the before-after specification.

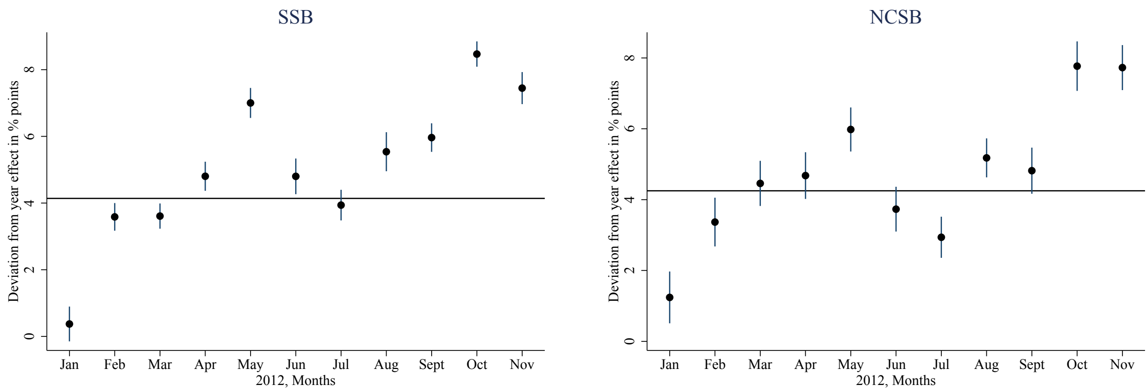

Average tax incidence: 4.14% | Average tax incidence: 4.25%

Figure D.2: The timing of tax shifting: event study

Notes: Kantar Worldpanel data 2008-2013. Each point represents the estimated coefficient on the respective month-to-month tax change indicator variable in 2012, *i.e.* EPIs relative to the EPI in December 2011. The horizontal lines represent the before-after effects estimated in Table 2, specification (2), i.e. the average effect in 2012. The bars extending from each point represent the bounds of the 95 percent confidence interval calculated from standard errors that are clustered at the area level. The control variables are as in Table 2, specification (2).

References

Anderson, S. P., de Palma, A., & Thisse, J.-F. (1988). The ces and the logit: two related models of heterogeneity. *Regional Science and Urban Economics*, *18*(1), 155–164.

Bonnet, C. & Réquillart, V. (2013a). Impact of cost shocks on consumer prices in vertically-related markets: the case of the french soft drink market. *American Journal of Agricultural Economics*, *95*(5), 1088–1108.

Bonnet, C. & Réquillart, V. (2013b). Tax incidence with strategic firms in the soft drink market. *Journal of Public Economics*, *106*, 77–88.

Borges, M. C., Louzada, M. L., de Sá, T. H., Laverty, A. A., Parra, D. C., Garzillo, J. M. F., ... Millett, C. (2017). Artificially sweetened beverages and the response to the global obesity crisis. *PLoS medicine*, *14*(1), e1002195.

Broda, C. & Weinstein, D. E. (2006). Globalization and the gains from variety. *Quarterly Journal of Economics*, *121*(2), 541–585.

Broda, C. & Weinstein, D. E. (2010). Product creation and destruction: evidence and price implications. *American Economic Review*, *100*(3), 691–723.

Deaton, A. & Muellbauer, J. (1980a). An almost ideal demand system. *American Economic Review*, *70*(3), 312–326.

Deaton, A. & Muellbauer, J. (1980b). *Economics and consumer behavior*. Cambridge university press.

Feenstra, R. C. (1994). New product varieties and the measurement of international prices. *American Economic Review*, *84*(1), 157–177.

Gopinath, G. & Itskhoki, O. (2010). Frequency of price adjustment and pass-through. *Quarterly Journal of Economics*, *125*(2), 675–727.

Gupta, D., Hill, A. V., & Bouzdine-Chameeva, T. (2006). A pricing model for clearing end-of-season retail inventory. *European Journal of Operational Research*, *170*(2), 518–540.

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

Hausman, J., Leonard, G., & Zona, J. D. (1994). Competitive analysis with differenciated products. *Annales d'Economie et de Statistique*, *34*, 159–180.

Le Bodo, Y., Etilé, F., Gagnon, F., & de Wals, P. (2017, January). Conditions influencing the adoption of a soda tax: analysis of the french case. Conference presentation: The role of taxation and other economic incentives in promoting healthier lifestyles: a focus on tobacco, alcohol and unhealthy diets (A. Alemanno and A. Garde). London.

Lecocq, S. & Robin, J.-M. (2015). Estimating almost-ideal demand systems with endogenous regressors. *Stata Journal*, *15*(2), 554–573.

Nakamura, E. & Zerom, D. (2010). Accounting for incomplete pass-through. *Review of Economic Studies*, *77*(3), 1192–1230.

Redding, S. J. & Weinstein, D. E. (2016). *A unified approach to estimating demand and welfare*. National Bureau of Economic Research.

Smith, S. A. & Achabal, D. D. (1998). Clearance pricing and inventory policies for retail chains. *Management Science*, *44*(3), 285–300.

Zhen, C., Finkelstein, E. A., Nonnemaker, J. M., Karns, S. A., & Todd, J. E. (2013). Predicting the effects of sugar-sweetened beverage taxes on food and beverage demand in a large demand system. *American Journal of Agricultural Economics*, *96*(1), 1–25.