NBER WORKING PAPER SERIES

A SIN BAN, NOT A SIN TAX? CONSUMPTION AND INCIDENCE OF SNAP SODA RESTRICTIONS

Crossan Cooper
Katja Seim
Jintaek Song

Working Paper 35739
http://www.nber.org/papers/w35739

NATIONAL BUREAU OF ECONOMIC RESEARCH
1050 Massachusetts Avenue
Cambridge, MA 02138
September 2026

Katja Seim has no relevant or material financial interests that relate to the research described in this paper. Crossan Cooper has no relevant or material financial interests that relate to the research described in this paper. Jintaek Song gratefully acknowledges financial support from Policy Impacts at MIT and has no relevant or material financial interests that relate to the research described in this paper. The views expressed herein are those of the authors and do not necessarily reflect the views of the National Bureau of Economic Research.

NBER working papers are circulated for discussion and comment purposes. They have not been peer-reviewed or been subject to the review by the NBER Board of Directors that accompanies official NBER publications.

© 2026 by Crossan Cooper, Katja Seim, and Jintaek Song. All rights reserved. Short sections of text, not to exceed two paragraphs, may be quoted without explicit permission provided that full credit, including © notice, is given to the source.

A Sin Ban, Not a Sin Tax? Consumption and Incidence of SNAP Soda Restrictions
Crossan Cooper, Katja Seim, and Jintaek Song
NBER Working Paper No. 35739
September 2026
JEL No. H30, I18, I38

ABSTRACT

We study the first statewide SNAP restrictions on sweetened beverages and candy. Restricted spending falls 11 percent; calories and sugar fall 5–8 percent. Substitution is asymmetric: eligible fruit drinks partly offset soda declines, while candy restrictions reduce purchases of eligible snacks. Retailers barely change shelf prices. Losing SNAP's sales-tax exemption raises tax-inclusive prices only 1.9 percent. Spending declines are larger at retailers with greater SNAP use and remain large online, where checkout stigma is less salient. Nationwide adoption would impose 13 percent of the out-of-pocket burden of an equivalent excise tax, while welfare gains reach roughly $1.1 billion annually.

Crossan Cooper
Yale University
crossan.cooper@yale.edu

Jintaek Song
University of Florida
jintaek.song@ufl.edu

Katja Seim
Yale University
School of Management
Department of Economics
and NBER
katja.seim@yale.edu

1 Introduction

The Supplemental Nutrition Assistance Program (SNAP), one of the largest social safety net programs in the United States, provided $95 billion in benefits to roughly 42 million low-income Americans in 2025. Benefits can be used for the purchase of food intended for home consumption, with longstanding exclusions only for alcohol, tobacco, and hot prepared foods. This broad eligibility has raised concern that SNAP supports consumption of products associated with adverse long-run health outcomes: in point-of-sale data from a large national grocery retailer, sugar-sweetened beverages were the second-largest category of SNAP households' grocery spending, accounting for 9.3% (Garasky et al., 2016). And in the broader receipt-based data we rely on, sweetened beverages and candy together account for 11.4% of SNAP households' grocery spending.[1]

Governments increasingly seek to discourage consumption of such products not through the traditional solution of Pigouvian taxes, but by narrowing what SNAP transfers can buy, for example, by excluding sugar-sweetened beverages from eligibility. The effectiveness of this minimally intrusive intervention is unclear. Unlike a tax, an eligibility restriction need not change the price a household faces at the margin: for households whose out-of-pocket food spending exceeds their spending on the restricted category, SNAP benefits are fungible (Southworth, 1945; Moffitt, 1989; Hoynes and Schanzenbach, 2009), and a restriction should merely reshuffle payment methods rather than reduce consumption of the targeted product. Yet households do not appear to treat SNAP benefits as fungible income in practice, with spending and consumption responses to SNAP differing sharply from responses to cash (Hastings and Shapiro, 2018; Cox and Harris-Lagoudakis, 2023; Song, 2026; Chorniy et al., 2025). Whether these frictions are large enough for eligibility restrictions to affect behavior is therefore an empirical question.

We study the SNAP eligibility restrictions on soda – and, in many states, candy, energy drinks, and sports drinks – implemented or federally approved for implementation in 23 U.S. states.[2] These policies marked the first statewide exclusions from SNAP of otherwise eligible foods based on product category or nutritional content (Chrisinger, 2025), allowing us to quantify behavioral responses to such restrictions. Beyond the narrow question of how eligibility restrictions affect households' total purchases of the targeted products, we assess their implications for the nutritional content of the basket of beverage and snack purchases, accounting for substitution toward other products. We then provide evidence on the demand- and supply-side mechanisms behind the observed purchase responses. First,

[1] Sweetened beverages here comprise soda, energy drinks, and sports drinks.

[2] In June 2026, a federal court vacated the federal approvals for five states, while restrictions remain in effect or are moving toward implementation elsewhere.

we examine changes in participants' budget sets, through retailer price adjustments, retailer promotional activity, and the additional sales taxes levied on products that are no longer exempt. Second, we explore mechanisms beyond the budget set, including the informational signal the restriction may send about the healthiness of the excluded products, stigma at the register, and mental accounting. Finally, we consider the policy's normative implications along two dimensions: incidence and efficiency.

To conduct our analysis, we combine a multi-retailer receipt panel from Numerator that records purchases and payment methods across the stores and restaurants visited by participating households with a stacked difference-in-differences research design built on the staggered implementation of the 2026 SNAP eligibility restrictions, comparing implementing to never-implementing states in each month relative to a restriction's introduction. To evaluate retailer price responses, we use an additional, newly collected dataset of soda prices scraped from four large chain retailers operating in Texas, Florida (both treated states), and Georgia (a state without current or planned eligibility restrictions).

Our first result is that the restrictions reduce purchases of the targeted products rather than only changing household payment methods. Among SNAP households, spending on restricted items falls by 11.0%, including declines of 12.0% for soda, 8.8% for candy, 13.4% for energy drinks, and 21.0% for sports drinks. We detect no offsetting aggregate increase in spending on unrestricted grocery items. As a result, total grocery spending falls by roughly the same dollar amount as spending on the restricted items. For soda, we also find no substitution toward private labels or away from small packages with higher per-unit prices. Nor do purchases at dollar, drug, and convenience stores, where such small packages predominate, decline.

The restrictions reduce purchased calories, with little evidence of offsetting substitution toward related products. Following the restrictions on soda, calories from soda purchases fall by 12.3%, total sugar by 12.5%, and sodium by 11.6%. Across all beverage purchases, calories, total sugar, and sodium decline by 7.3%, 8.1%, and 8.9%, respectively.

We find economically modest and frequently insignificant increases in calories, sugar, or sodium for any unrestricted beverage category. The largest increase in calories is among fruit juices. In the five states that do not restrict eligibility of fruit juice products, calories associated with their purchases rise by 5.6% and offset 25.1% of the decline in soda calories.[3] Candy restrictions do not produce a comparable substitution toward similar packaged snacks. Instead, spending on a well-defined basket of still-eligible chips, crackers, ice cream, and

[3]We construct this estimate using data from Kroger and Walmart, where we link receipts to detailed product-level SNAP eligibility information and are able to further decompose the nutrition responses across sweetened and unsweetened juice categories.

frozen novelties falls by 4.2%, resulting in declines in calorie, sugar, and sodium between 5.1% and 7.6% in the packaged-snack basket. Overall, considering the product categories restricted by each state, calories, total sugar, and sodium across the observed beverage and packaged snack basket fall by 5.4%, 7.7%, and 7.5%, respectively.

We next ask why purchases fall. Changes in households' budget sets do not appear to be a primary explanation. Retailers do not systematically adjust shelf prices: no retailer-specific price change estimate exceeds 0.4% in magnitude in our main specification, the signs of the price changes vary, and the observed adjustments do not vary systematically with the size of the local SNAP population. Because purchases made with SNAP are exempt from sales tax, the eligibility restriction can result in purchases households previously made using SNAP becoming newly taxable. If 100% of restricted purchases were made with SNAP pre-policy, this sales tax burden would average 4.9% across households. In practice, only 35–42% of pre-policy spending on restricted products used SNAP. Hence, the additional sales tax amounts to only 1.9%. For most households, pre-policy spending on groceries that would remain eligible was large enough to absorb all observed SNAP-paid grocery spending: they could redirect SNAP spending toward groceries that remained eligible and use the cash freed by that switch to continue purchasing the restricted products.

We then consider nonprice explanations. Related to supply response would be changes in promotional activities. We do not observe a decline in retailer promotions that could drive reduced category spending. If anything, the share of beverage products with promotions rises modestly, with some retailer-level heterogeneity. There are also other demand side mechanisms. Eligibility restrictions may draw attention to the health consequences of the targeted products, but this cannot explain most of the response: soda spending falls by only 3–4% among non-SNAP households, compared with 12% among SNAP households. Similarly, checkout stigma appears unlikely to be central,[4] since the purchase decline is no smaller online (where there is no interaction with a cashier) than in stores. Thus, our preferred interpretation is that households do not treat SNAP and cash as fully interchangeable for psychological reasons. For example, they may associate some shopping trips with SNAP and others with cash. Once a product loses eligibility, they may stop buying it on SNAP-associated trips rather than reorganize payment across the two accounts. Consistent with this interpretation, the decline is largest at retailers where a household used SNAP most frequently before the policy. It is also concentrated at large stores, where SNAP use was more common than at small stores, and operates mainly through fewer shopping trips containing

[4] Barnhill (2011) summarizes the concern. One of USDA's primary objections to any new SNAP restriction is that it could increase the embarrassment and stigma associated with SNAP use and thereby deter participation: a participant unaware of the exclusion would attempt to buy the excluded item and be told at the cash register that it is not allowed.

soda.

We next assess the normative implications of an eligibility restriction along two dimensions. First, we contrast it with an alternative policy of imposing an excise tax on the targeted products that achieves the same purchase reduction. Prior work has noted that Pigouvian taxes reduce lower-income households' purchasing power by raising prices (Allcott et al., 2019) and hence impose a higher financial burden on SNAP participants given that the ban reduces restricted-product purchases with little change in prices. Assuming the excise tax revenues are retained by the government, nationwide adoption of excise taxes on soda, candy, and energy drinks raises the cost of SNAP participants' targeted purchases by an incremental $1.18 billion, relative to the $176 million of new sales tax payments that SNAP households make under an eligibility restriction. The financial burden of the restriction is hence only 13% of the equivalent excise tax burden. Even when the additional tax revenue is returned through equal per-household rebates, the restriction remains less burdensome. Prior estimates suggest that SNAP participants purchase roughly twice as much of these products as nonparticipants,[5] in which case the restriction imposes 26% as much net burden as an equivalent tax. Using the smaller spending differences in our data, the corresponding burden is 67% for soda, candy, and energy drinks and 48% for soda alone.

Second, we use a Marginal Value of Public Funds (MVPF) framework to assess the efficiency gains from adding the three most common types of product restriction to SNAP, relative to otherwise identical SNAP benefits without the restriction. We value the purchase reductions using estimates from Allcott et al. (2019) of health costs that are either not fully internalized by consumers or borne by the health-care system. We apply these estimates to our purchase reductions and extrapolate them to the other restricted products, yielding an estimated $1.01 billion in gross annual corrective benefits. Under a mental accounting interpretation in which SNAP eligibility induces purchases beyond what households would choose under an equivalent cash transfer (Hastings and Shapiro, 2018; Song, 2026), removing eligibility generates an estimated $59 to 155 million in additional annual gains by correcting in-kind distortions. Together, these measured gains amount to approximately $1.07 to $1.16 billion per year, which increases the MVPF of SNAP by 0.2% to 1% depending on assumptions about SNAP's baseline MVPF in the absence of product restrictions.

Contributions and related literature We make four main contributions. First, this paper provides some of the first estimates of the effects of an implemented, at-scale restriction on the purchases made by SNAP recipients. Prior evidence is largely simulated or experi-

[5]This benchmark is based on USDA retailer point-of-sale data in Garasky et al. (2016) and is consistent with the income gradient in Allcott et al. (2019, Figure III).

mental. Reduced-form simulations (Basu et al., 2014) and a microfounded structural model (Cox and Harris-Lagoudakis, 2023) predict that making soda SNAP-ineligible would reduce soda calories or purchases by 15–18%. Harnack et al. (2016) and French et al. (2017) study an experiment among SNAP-eligible nonparticipants – who generally have more resources than SNAP participants – in which soda, candy, and sweet baked goods were excluded from a study-provided food benefit. Their experimental restriction was enforced through participant instructions and receipt auditing rather than at the register, so the friction and stigma emphasized by critics are absent from the experiment.

Two contemporaneous papers study the recent SNAP restrictions. Frisvold and Lozano Rojas (2026), using NIQ Omnishopper data on self-reported SNAP participants, estimate a 13% decline in soda quantity. Allcott et al. (2026), using the same Numerator panel we employ, estimate a 12.4% decline in purchases of excluded drinks. Our estimate of a 12% drop in soda spending is closely in line. Our analysis, developed independently, is complementary in three respects.

First, we directly measure changes in sales tax exposure and transaction prices as individual components of price incidence. Our detailed retailer price data follow the same product in the same store over time. This, together with independently assembled UPC-level information on product volumes, can distinguish supply-side price changes from product substitution in driving volume responses. This distinction is not identified from unit values such as spending per ounce or from prices per ounce predicted using auxiliary data.

Second, we account for spending and nutrition more completely. We cover ready-to-drink beverages and packaged snacks in full rather than a selected set of substitutes to allow us to observe any displaced spending or calories. We link each item to package size and nutrient content so that calories, sugar, and sodium are measured at the item level. We also estimate the effects of candy restrictions, which have been implemented or are scheduled for implementation in fourteen states. This reveals an asymmetry: beverage restrictions induce some substitution toward eligible drinks, whereas candy restrictions reduce purchases of nearby eligible snacks.

Third, we use the multi-retailer coverage of our panel and the payment method on each receipt to study mechanisms. Comparing households' pre-policy SNAP spending with their spending on unrestricted eligible goods, and comparing responses by SNAP status, shopping channel, and retailers' pre-policy SNAP intensity, we test whether the budget set binds, whether the restriction served as a government nutrition signal, whether checkout frictions matter, and what form mental accounting takes.

Finally, this paper contributes to the broader literature on in-kind transfers, which are frequently motivated by the fact that the consumption of goods subsidized by programs

ranging from housing to health care to education may generate social benefits beyond those captured by recipients (Currie and Gahvari, 2008). As in the case of SNAP, recent reforms increasingly seek to steer choices within these broad domains – for example, toward higher-opportunity neighborhoods (Chetty et al., 2016; Bergman et al., 2024) or education and training programs with stronger employment and earnings outcomes (Cellini et al., 2020; U.S. Department of Education, 2026). While fungibility concerns arise most naturally in the context of SNAP, the questions behind these reforms generalize: whether finer targeting within an in-kind program changes behavior and raises welfare depends on how restrictions are implemented, on recipients' ability to adjust along other margins, and on supply-side responses. We provide new evidence on each of these components in the context of the SNAP eligibility restrictions.

The remainder of the paper proceeds as follows. Section 2 describes the state policy reforms and their implementation. Section 3 describes the data, and Section 4 presents the empirical strategy. Section 5 reports the reduced-form effects on purchases and nutrition. Section 6 asks whether changes in the household's budget set can explain the purchase response and then considers other consumer and firm mechanisms. Section 7 compares the restriction with equivalent excise taxes and evaluates its welfare effects. Section 8 concludes.

2 Policy Background

Since its introduction as food stamps in 1964, SNAP has placed few restrictions on which foods recipients may purchase using the benefits, the only exceptions being alcohol, tobacco, hot prepared foods, and non-food items. As a result, the 2026 wave of soda, candy, and energy drink restrictions is the most significant narrowing of the set of SNAP-eligible products in decades. Between April 2025 and March 2026, a total of 23 states requested waivers of standard federal rules governing product eligibility under SNAP. Under the "Make America Healthy Again" (MAHA) agenda, the USDA granted the states' waiver applications to exclude their proposed food categories from SNAP eligibility.

The 2026 restrictions are the culmination of a debate that has extended for more than two decades. As early as 2004, Minnesota petitioned USDA to restrict SNAP recipients from using the benefits to purchase soda and candy. USDA refused, reasoning that such a restriction would "perpetuate the myth that [SNAP] participants do not make wise food purchasing decisions." In 2010, New York City sought a waiver to exclude sugar-sweetened beverages from the list of eligible products, which USDA again denied, citing, among others, the difficulty of defining eligible products and operationalizing the restriction at the register, and an evaluation design it judged inadequate for measuring consumption. Maine's 2015

request met the same fate in the final months of the Obama administration. In each case, the USDA invoked two concerns – administrative burden and the potential stigmatization of recipients – that remain the core objections of critics today. In approving the recent set of waiver applications, the USDA reversed course, framing taxpayer subsidies for sugar-sweetened beverages and candy as a public health and fiscal crisis.

The approved waivers have been implemented in several phases. Our August 2025-June 2026 panel contains three implementation cohorts and ten treated states. Nebraska, Indiana, Iowa, Utah, and West Virginia implement SNAP waivers in January 2026, Idaho, Oklahoma, and Louisiana begin in February, and Texas and Florida begin in April. States with approved implementation dates after June 2026 are excluded from the analysis rather than being used as controls.

States differ in the scope and definition of their product restrictions. All treated states restrict soda, typically defined as non-alcoholic beverages sweetened with sugar or artificial sweeteners, excluding milk-based drinks and beverages that are predominantly juice. Seven of the ten states in our sample restrict candy, eight restrict energy drinks, and five restrict sports drinks, either explicitly or through broader restrictions on soft drinks or sweetened beverages. In four states, these broader definitions also cover sweetened fruit drinks and sweetened tea. Florida is the only state in our sample to restrict prepared desserts. Texas implemented a broader restriction on "sweetened drinks", while the other states' category definitions also differ slightly at the margin. We describe the exact content of the restrictions in detail in Appendix E.7.

For the purposes of this study, the key feature of the waivers is that they operate entirely through the payment method. The restriction prohibits a household from paying for a restricted item with SNAP, but the item remains available for sale and can be bought with cash. As a result, though the policy mechanically removes restricted items from the SNAP-eligible product list, its effects on consumption depend on whether households offset the restrictions by paying for these items out of pocket.

3 Data

Our primary dataset is the Numerator consumer panel, a longitudinal record of household purchases collected by the market research firm Numerator. Rather than scanning barcodes at home, panelists photograph the receipts from every store they shop at (e.g., grocery stores, large retailers, convenience stores, and restaurants) or upload electronic receipts using a mobile application, and Numerator extracts shopping trip and item-level information from the images. Each trip record contains information on the retailer, date, and method of

payment, including whether SNAP was used. Each item record contains information on the price, quantity, brand, and category. We use the 2025-2026 panel of "static panelists," a subset of participants who upload consistently and reliably.

The data have two features essential to our study. First, because SNAP payment information is extracted by Numerator from uploaded receipts, we identify SNAP households directly from payment behavior rather than relying on self-reported SNAP participation, which is substantially underreported in major household surveys (Meyer et al., 2022). We classify a household as a SNAP participant if at least 10% of its in-store grocery trips were paid with SNAP in at least two consecutive months before implementation.[6] Among all unique households in the classification sample, 10.9% meet this definition.[7] For context, USDA administrative records indicate that SNAP served 41.7 million people, or 12.3% of U.S. residents, in an average month of FY2024 (Rabbitt et al., 2025).

Second, the panel records a household's purchases across many retailers, a feature that is particularly important for soda and candy. Households buy these items both at their main grocery store and at convenience, drug, and dollar stores, and these channels differ sharply in how frequently SNAP is used to pay. SNAP use is concentrated at grocery retailers, where most food spending – especially on fresh food – takes place, and is much lower at convenience and dollar stores. Indeed, Castner et al. (2025) report that about 78% of SNAP benefits are redeemed at supermarkets and mass merchandisers versus roughly 5% at convenience stores. For soda, our data show a similar pattern: of pre-policy soda spending recorded on receipts where SNAP was used, 77.5% occurs at large stores (grocery, mass-merchandise, and club stores), while 3.2% occurs at convenience stores. Because the restrictions only affect purchases made using SNAP benefits, a household's response to the waivers can differ across these stores; single-retailer panels from large grocery chains may misestimate the true response to product restrictions. By following each household across stores, our data recover the full effect of the restrictions.

For our primary category-level analysis, we define four product groups directly from Numerator's item-level product information. We define an item as a soda product if it is listed in one of Numerator's fourteen carbonated soft drink categories, as an energy drink if it is listed

[6]We rely on in-store receipts to classify SNAP participation because payment information on online receipts may be less reliable. Following Song (2026), we require SNAP payment on at least 10% of in-store grocery trips to prevent an infrequent miscoded payment method from classifying a non-SNAP household as a SNAP participant.

[7]We calculate the share of SNAP participants relative to all households with an eligible cohort assignment, including those whose SNAP status is ambiguous. At the household-by-implementation-cohort level, we classify units with some observed SNAP use but that do not meet the main SNAP criterion as ambiguous and exclude them from the estimation sample for that cohort. To reduce misclassification in the non-SNAP participant analysis, the non-SNAP sample includes only households with no observed SNAP payment in any channel, in-store or online.

in the energy drink or energy shot category, as a sports drink if it is listed in Numerator's sports drink category, and as candy if it is listed in the candy department or the chocolate and non-chocolate confectionery categories[8]. Some waivers also restrict sweetened products within Numerator's broader Fruit Juice and Ready To Drink Tea categories. Because these categories contain both restricted and eligible products, we analyze them separately rather than treating the entire category as restricted. We describe the full list of product categories in Appendix E.5. Across panelists in treated states, the four restricted categories represent 11.4% of SNAP households' spending on SNAP-eligible groceries: 4.9% from soda, 4.4% from candy, 1.3% from energy drinks, and 0.7% from sports drinks. Non-SNAP households have a similar aggregate expenditure share on these categories (10.6%), but spend around 31% less on soda. Other pre-restriction spending information for SNAP and non-SNAP households is reported in Table 1, along with details on the analysis sample.

Many brands offer multiple package sizes of a given product. While our data record the number of units purchased and the price paid, package size information is not readily available. We therefore rely on product descriptions from the receipts to assign volume information to a particular purchase occasion, where possible. Receipts differ in their information content. Some report a UPC, which we match to package size information scraped from Kroger and Walmart product listings and, for packaged snacks, the USDA's FoodData Central packaged snack records. When no UPC match is available, we extract the size from the product description using string processing (e.g., we assign a volume of 72 fluid ounces to a product described as "12 fl oz, 6 Pack"), which raises coverage from 26% to 49% of beverage spending and from 26% to 45% of packaged snack spending.[9] For soda, which is restricted by all state waivers, we furthermore extract volume information from retailer product listings to increase coverage at sixteen retailers. These steps allow us to assign a package size to 83% of soda spending at sixteen retailers: Walmart, Dollar General, Walmart.com, Kroger, Costco, Sam's Club, Family Dollar, Walgreens, Walmart Neighborhood Market, ShopRite, Target, Meijer, Amazon.com, Albertsons, WinCo Foods and Fred Meyer. These retailers account for 55% of the soda spending of SNAP households. We show in Appendix F that consumer spending patterns across the sixteen retailers are broadly similar to the full panel: large and

[8]The legal definition of candy differs slightly across states – for instance, in whether a flour-based bar such as a KitKat counts as candy. For comparability, we classify products using Numerator's product-category coding rather than state-specific statutory definitions.

[9]We accept a size read from a description only if it is no larger than the largest package a UPC reports for the same category and if the implied price per unit is within an order of magnitude of that category's typical price. A reading that fails either check is treated as missing. The beverage coverage rate excludes coffee, tea, powdered mixes and liquid enhancers, which are sold by weight rather than by volume. Including them, the beverage coverage rate is 41%. Appendix G provides additional details, including how our string processing performs relative to UPC-based sizes for products with UPCs.

small stores account for approximately 61% and 27% of soda spending in this subsample, respectively, compared with 65% and 23% in the full panel.

We then employ the package size information to construct the nutritional content of a purchase. Beyond package size, this requires information on the product's nutrient density per common weight unit. We obtain nutrient density data from USDA FoodData Central product records, which we supplement with Open Food Facts and Kroger product pages. When a receipt reports a UPC that appears in the resulting product reference file, we assign the nutrient information of that exact product. When an exact UPC match is unavailable, we first compare the item's brand and description with reference products of the same product type that use the same sweetener and have the same flavor. If no such match is available, we impute nutrient density from progressively broader groups of reference products: a closely matching brand within the same detailed category and sweetener group, products from the same manufacturer in that category and sweetener group, all products in the detailed category and sweetener group, and finally a broader category or sweetener group average. For several relatively homogeneous categories, such as plain water, tea, and coffee, we move directly from an unsuccessful UPC match to a narrowly defined category average. We apply this procedure separately to each nutrient – sugar content, sodium content, and calories.

Last, we scale nutrient density by package size. Since package size is not available for a share of receipts, we impute total nutrient content for these purchases using a nutrient-per-dollar approach. We first identify purchases for which both size and nutrient density are observed to calculate nutrients per dollar of spending in this matched sample. For products with missing information, we then assign the average nutrients per dollar of spending of products from the same brand, product category, and sweetener group – or broader groupings if necessary – with complete information. We multiply this nutrient-per-dollar rate by the spending for unmatched purchases. We consider the robustness of our results on the nutritional implications of product eligibility restrictions to dropping purchases where we impute nutritional content using the nutrient-per-dollar approach. (Appendix G)

Finally, we aggregate calories, total sugar, sodium, and other nutrients across all purchased restricted items to form household-by-month outcomes. Appendix G provides the complete matching hierarchy and additional validation details. Pre-policy, soda accounts for 56.7% of SNAP households' beverage calories, 60.7% of their beverage sugar and 34.7% of their beverage sodium. Adding energy and sports drinks raises those shares to 64.4%, 68.1% and 59.4%. Candy accounts for 20.9% of SNAP households' packaged-snack calories.

To evaluate firm responses to the restrictions, we focus on beverages and construct a weekly panel of retailer beverage prices by scraping item-level price data from Dollar General, Kroger, Walmart, and Publix at the store-by-item-by-week level from February 9, 2026 to

June 28, 2026. The median number of items per store-week ranges from 213 (Dollar General) to 427 (Walmart) across retailers. To ensure our coverage is representative, we collect data from a large sample of around 300 stores per retailer across the treated states of Florida and Texas and the control state of Georgia.[10] We select these retailers for two reasons. First, they operate across treatment and control states, unlike smaller local retailers or state-level chains. Second, they span several major retail formats: traditional supermarkets (Kroger and Publix), supercenters (Walmart), and dollar stores (Dollar General). We describe this dataset in more detail in Appendix H.

Song (2026) points out that a receipt panel records only a fraction of a household's trips, so that the spending levels it captures will understate true purchases. As a result, we interpret our estimates primarily in percentage terms. So long as households do not differ in their selective reporting before and after the waivers, this under-reporting scales the estimated treatment effect and the pre-waiver averages by the same factor. Whenever a levels interpretation is needed, we follow the same approach as Song (2026) and scale the average SNAP benefit identified in the Numerator receipt data up to the aggregate SNAP benefits reimbursed over the same period, benchmarking the panel's partial coverage against a comprehensive total. As a simple benchmark, the panel records an average of $199.79 in monthly grocery spending paid with SNAP, compared with the FY2024 average monthly household benefit of $341 (Ward and Monkovic, 2026). Thus, the panel appears to capture about 59% of SNAP-funded spending.

4 Empirical Strategy

We estimate the effects of the SNAP restrictions using a stacked difference-in-differences design, as in Cengiz et al. (2019). Because states implement the restrictions in different months, we first group together states with the same implementation month. Let g index these implementation groups and let G_g denote the calendar month in which group g's restriction takes effect.

For each implementation group, we construct a separate comparison dataset containing households in the states implementing in month G_g and households in a common set of control states. This control group, which contains households in any state without either an implemented or scheduled SNAP restriction, consists of 28 states and the District of

[10] Store coverage differs by retailer. Walmart stores were randomly-selected after setting state-specific quotas (100 in Georgia, 80 in Florida, and 120 in Texas). We oversample Dollar General stores located in higher-poverty ZIP codes, and we collect data from the universe of Krogers in Georgia and Texas. We cover around 50% of stores in Publix.

Columbia.[11] We exclude states scheduled to implement SNAP restrictions after June 2026 from the analysis sample rather than using them as not-yet-treated controls.

For each implementation group, we construct a separate comparison sample covering the four months before implementation and up to four months after the SNAP restrictions begin. Because the data end in June 2026, the January and February 2026 groups g are observed through event month +4, but the April 2026 group is observed only through event month +2. Since the same control states are used for each implementation group, a control household may enter the combined data once for each group; household-by-group identifiers distinguish these repeated observations.

We estimate a separate regression for each spending category k. The soda regression includes the ten states that implement an eligibility restriction during the sample period since all ten restrict soda. The candy, energy-drink, and sports-drink regressions include only treated states that restrict the relevant category – seven states for candy, eight for energy drinks, and five for sports drinks – together with the common control group. A state implementing a restriction that does not cover category k is excluded from that category's regression rather than being treated as a control.

For each category k, the static estimating equation is

$$y^k_{isgt} = \beta_k \left(T^k_{sg} \times \text{Post}_{tg}\right) + \alpha_{i,g} + \gamma_{t,g} + \varepsilon^k_{isgt}, \tag{1}$$

where y^k_{isgt} is household i's spending on category k in state s, implementation group g, and calendar month t, T^k_{sg} is an indicator equal to one if state s belongs to implementation group g and restricts category k, Post_{tg} is an indicator equal to one in calendar months at or after that group's implementation month G_g, and the household-by-implementation-group fixed effects $\alpha_{i,g}$ and calendar-month-by-implementation-group fixed effects $\gamma_{t,g}$ absorb time-invariant differences across households within each comparison dataset and shocks common to treated and control households in a given comparison dataset and month. As a result, β_k measures the post-policy change in spending on category k in states that exclude that category from SNAP eligibility, relative to the change over the same period in the common control group. Our primary design pools these comparisons across implementation groups into a single coefficient. Standard errors are clustered by state.

We estimate Equation (1) in levels by OLS, recording household-months without a purchase as zeros. The coefficient β_k is measured in dollars per household-month. Following the

[11] We include Colorado in the control group because the federal approval of its SNAP restrictions was vacated at least 3 months prior to implementation. We do not include Tennessee because federal approval was vacated in the month directly prior to implementation.

logic of Song (2026), we typically report the percentage effect

$$100 \times \frac{\widehat{\beta}_k}{\overline{y}^k_{\text{treated,pre}}},$$

where $\overline{y}^k_{\text{treated,pre}}$ is average pre-policy spending on category k among households in states implementing a SNAP restriction.

In parallel, we estimate the same conditional-mean specification using Poisson pseudo-maximum-likelihood (PPML):

$$E\left[y^k_{isgt} \mid \cdot\right] = \exp\left[\beta^P_k \left(T^k_{sg} \times \text{Post}_{tg}\right) + \alpha_{i,g} + \gamma_{t,g}\right]. \tag{2}$$

PPML similarly retains household-months with zero spending and directly measures the proportional change in average spending, an estimate that is invariant to the units in which spending is measured (Chen and Roth 2024). We report the implied percentage effect as

$$100 \times \left[\exp\left(\widehat{\beta}^P_k\right) - 1\right].$$

Identification requires that, absent the product restrictions, spending in treated and control states would have followed parallel trends within each comparison dataset. To examine this assumption and trace the timing of the response to the SNAP restrictions, we estimate dynamic event study specifications. For this purpose, we convert calendar period t into event time

$$e_{tg} = t - G_g,$$

so that $e_{tg} = 0$ is the implementation month and $e_{tg} = -1$ is the month immediately before implementation. We then estimate

$$y^k_{isgt} = \sum_{\tau \neq -1} \beta_{k\tau} \left(T^k_{sg} \times \mathbf{1}\{e_{tg} = \tau\}\right) + \alpha_{i,g} + \gamma_{t,g} + \varepsilon^k_{isgt}, \tag{3}$$

where each $\beta_{k\tau}$ measures the treated-control difference in the change from the month immediately preceding implementation to event period τ.

5 Reduced-form findings

In this section, we present estimates of the effects of SNAP food restriction waivers on household purchases and nutrition, using the research design described in Section 4.

5.1 Effects on purchases

Table 2 reports the static difference-in-differences OLS estimates of Equation (1). Our first result is that the effect of the restrictions on consumption is large, even though most households' pre-policy spending suggests that they could redirect SNAP payment toward groceries that remained eligible. We find that soda spending falls by 12.0%, candy spending falls by 8.8%, energy drink spending falls by 13.4%, and sports drink spending falls by 21.0%. Spending on the state-specific core restricted basket falls by 11.0%.[12]

One concern is that households might substitute spending on soda in grocery stores for soda purchases at restaurants. In Appendix Figure 8, we show that there is no offsetting increase in food-away-from-home soda spending, including at restaurants where SNAP benefits generally cannot be used.

Next, to assess whether the restrictions differentially affect households with children, we separately estimate Equation (1) in the samples of SNAP households with and without children. Effects are larger among households with children (16.1% versus 8.3% for soda, 11.0% versus 6.0% for candy, and 28.0% versus 11.4% for sports drinks). In both samples, spending on other grocery items is unaffected by the restrictions. The percent spending declines are robust to using a PPML specification (Appendix Table 18).

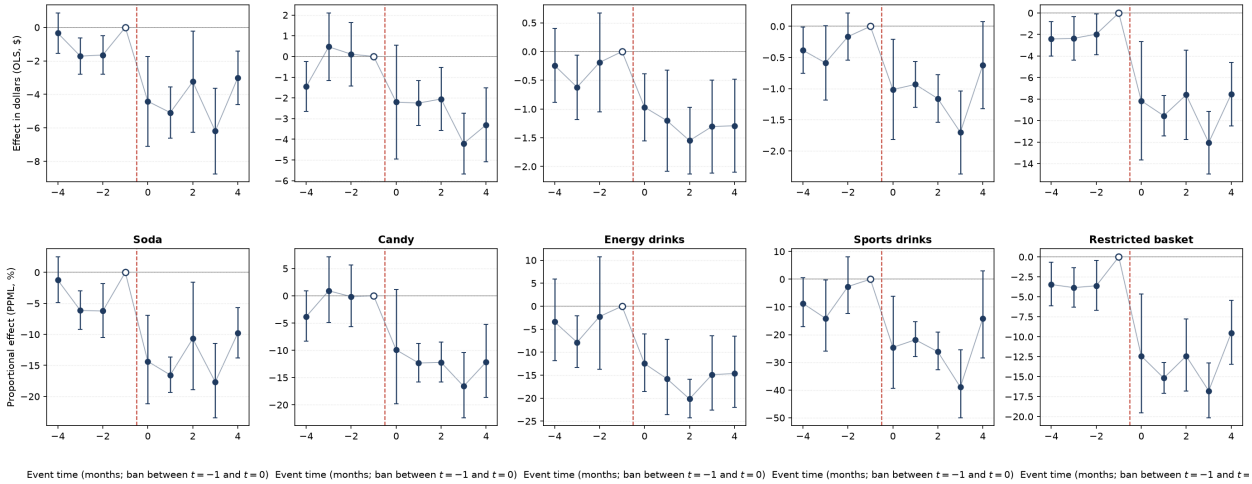
To assess households' dynamic responses, Figure 1 reports the event study estimates of the restriction effects in terms of the dollar response (OLS; top row) and the proportional response (PPML; bottom row). All five outcomes – soda, candy, energy drinks, sports drinks, and each state's overall restricted basket of items – fall at implementation and remain below baseline throughout the post-waiver period, whether measured in dollars or proportional terms. Figure 2 reports the same OLS estimates for other grocery spending. We find no evidence of a dynamic reallocation of spending on products in the restricted categories to other SNAP-eligible food purchases.

5.2 Margin of adjustment

We now consider the underlying sources of the observed spending decline. We focus on soda spending at the 16 retailers where we are able to identify volume information for the

[12]The restricted basket follows the details of the SNAP waivers in each state: soda is restricted in all ten treated states. Energy drinks are restricted in all states but UT and WV. Sports drinks are restricted in IN, IA, ID, OK, and TX. Candy is restricted in those five states plus LA and FL. Packaged cookies and snack cakes (including snack donuts, breakfast cakes, brownies, and bars) are only restricted in FL. Tonic water and flavored powdered drink mixes are only restricted in TX. Numerator's *Fruit Juice* and *Ready To Drink Tea* categories contain both restricted and unrestricted products. Because the category codes do not distinguish them, we exclude both from any state's restricted basket. Appendix E.5 describes the separate item-level classifications of *Fruit Juice* and *Ready To Drink Tea* used in supplemental analyses.

Figure 1: Dynamic effects of SNAP restrictions on purchase of restricted items



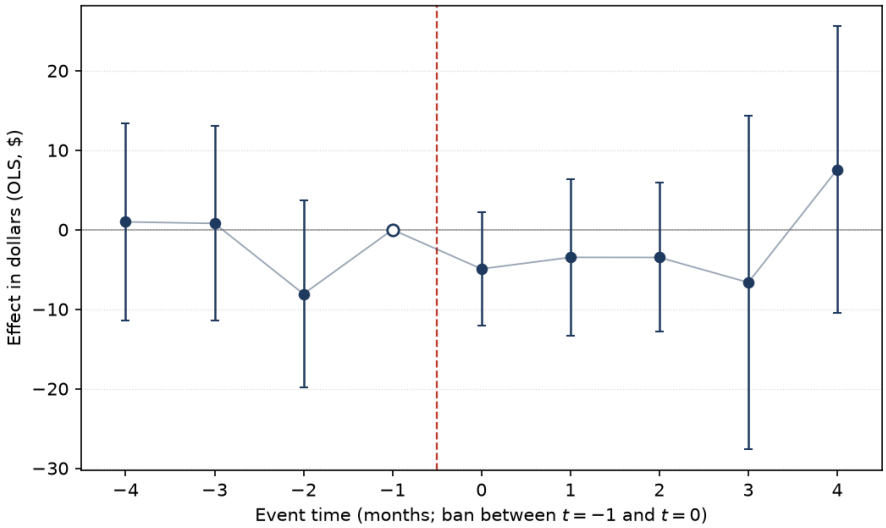
Notes: The top row of Figure 1 reports OLS estimates of effects in dollars per household-month, while the bottom row reports PPML estimates of proportional effects. For each product category, treated states are those that restrict items in the category, and control states are those that have not adopted any SNAP food restriction. Bars represent 95% confidence intervals, where standard errors are clustered by state. The dashed line marks the implementation date (in months), where event time $t = -1$ is omitted. Event times -4 through $+4$ are shown. Estimates for event times $+3$ and $+4$ use only the January and February waiver cohorts.

vast majority of products. For this subsample, the decline in household purchases operates primarily along the extensive margin.[13] Decomposing soda spending into the number of shopping trips that include soda and spending per soda trip, we find that the number of soda trips falls by 11.8%, while spending per soda trip only falls by an insignificant 1.8%. When limiting the sample to soda purchases with available volume information, spending per milliliter shows no statistically significant change. This aggregate lack of a response may mask a change in the composition of products purchased, for example in response to price changes after the policy. We assess supply-side responses in detail in Section 6.1.1. Here, we note that in addition to a lack of a response in spending per milliliter, the private-label share similarly remains unaffected, so households do not appear to downgrade their soda purchases to less expensive unbranded varieties in an effort to reduce spending; instead, spending declines are primarily a result of purchasing the restricted soda products less frequently. The full set of estimates is available in Table 3.

Purchases adjust differently across retailers. Table 4 and Figure 3 show that soda spending at large stores (grocery, mass-merchandise, and club stores) falls by 16.2%. The corresponding estimate at small stores (dollar, drug, and convenience stores) is an imprecise

[13]The same pattern holds when we use purchases at all retailers: soda trips fall by 9.9%, while spending per soda trip falls by a statistically insignificant 1.5%.

Figure 2: Dynamic effects of SNAP restrictions on other grocery purchases

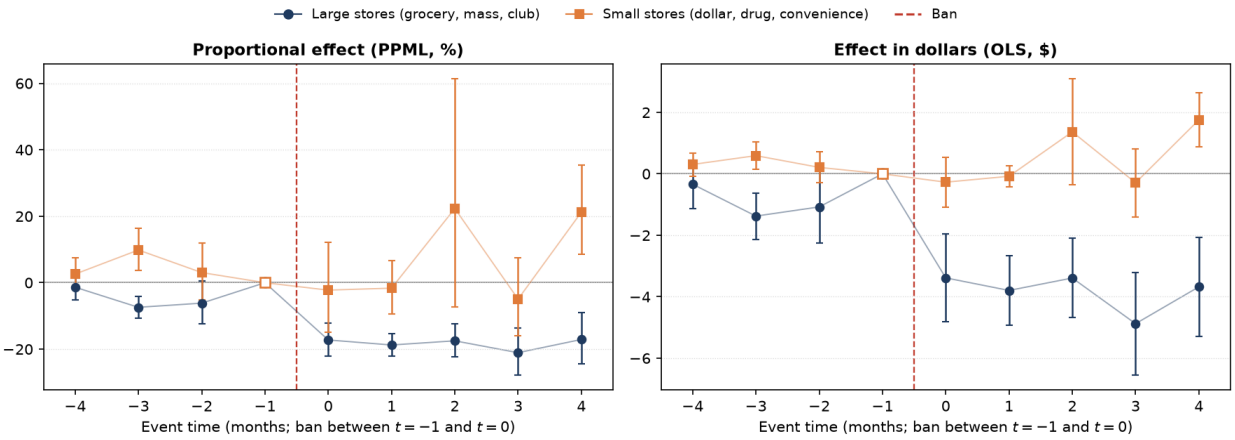


Notes: The outcome in Figure 2 is the change in monthly spending on all grocery items not subject to the state waivers. We report OLS estimates of $\beta_{k\tau}$ from Equation (3) in dollars per household-month. Bars represent 95% confidence intervals, where standard errors are clustered by state. The dashed line marks the implementation date (in months), where event time $t = -1$ is omitted. Event times -4 through $+4$ are shown. Estimates for event times $+3$ and $+4$ use only the January and February 2026 waiver cohorts.

increase of 1.9%. Appendix C reports the heterogeneity analysis for candy, energy drinks, sports drinks, and the full basket of restricted items. Spending reductions are larger at large stores than at small stores in those categories as well, though the precision varies by category. While SNAP use was more pronounced at large stores before the policy, purchases at these smaller stores nevertheless account for 22.5% of SNAP households' soda purchases.

This divergence in spending responses highlights the value of our data's broad coverage of shopping outlets; estimates based on large-retailer scanner data would overstate the overall spending decline of 12.0% by about one-third. The asymmetric response by store format is also relevant given the rapid growth of dollar stores in low-income areas, which Caoui et al. (2026) show displaces independent grocers and reduces fresh produce purchases by low-income households. Our results suggest that eligibility restrictions have limited effectiveness in this growing retail segment.

Figure 3: Dynamic effects of SNAP restrictions on soda spending by store type



Notes: The left panel reports PPML estimates of proportional effects and the right panel reports OLS estimates of effects in dollars per household-month. Event time $t = -1$ is omitted. Bars are 95% confidence intervals with standard errors clustered by state. The corresponding static OLS estimates, reported in Table 4, are a 16.2% decline at large stores and an imprecise 1.9% increase at small stores.

Table 1: Sample composition and summary statistics

Panel A. Sample composition

	Count
Unique SNAP households	23,324
Households in treated states	4,375
Households in control states	18,949
States	39
Treated states	10
Control states (includes DC)	29
Stacked household-month observations	398,081
Unique household-months	209,339

Panel B. Monthly pre-waiver characteristics of households in treated states

	SNAP households		Non-SNAP households	
	Mean	Std. dev.	Mean	Std. dev.
SNAP-eligible grocery spending (all tenders)	\$569.45	\$385.88	\$495.40	\$371.92
Grocery spending paid with SNAP	\$199.79	\$194.13	–	
Soda spending	\$28.11	\$42.09	\$19.39	\$32.67
Candy spending	\$25.30	\$32.34	\$24.03	\$42.88
Energy drink spending	\$7.55	\$24.36	\$6.37	\$21.30
Sports drink spending	\$3.73	\$9.94	\$2.71	\$8.86
Soda, candy, energy drink, and sports drink spending	\$64.68	\$71.16	\$52.51	\$67.24
Soda, candy, energy drink, and sports drink spending paid with SNAP	\$19.87	\$32.16	–	
Grocery trips	20.78	12.68	15.22	10.30
Grocery trips paid with SNAP	5.38	4.92	–	
Share of in-store grocery trips paid with SNAP	0.302	0.223	–	

Notes: Panel A describes the primary estimation sample, covering Sep. 2025–Jun. 2026. Control households may enter once for each implementation cohort, so stacked observations exceed unique household-months. Panel B reports means and standard deviations from the respective main soda regression samples, using 17,403 SNAP and 161,272 non-SNAP household-months in treated states before waiver implementation. Non-SNAP households are defined as having no observed SNAP payments in the receipt panel. Dashes denote SNAP payment measures that are not applicable to that sample. Spending is monthly spending observed in uploaded receipts and includes all payment methods unless identified as "paid with SNAP". The SNAP trip share is defined among household-months with an in-store grocery trip.

Table 2: Effect of SNAP restrictions on SNAP households' spending

	(1) Soda	(2) Candy	(3) Energy drinks	(4) Sports drinks	(5) Restricted basket	(6) Other groceries
Panel A. All households						
Treated × Post	−3.385***	−2.185***	−0.984***	−0.774***	−6.984***	−1.677
	(0.991)	(0.603)	(0.293)	(0.125)	(1.548)	(3.572)
Pre-policy mean ($)	28.109	24.696	7.344	3.694	63.396	506.056
Effect (% of mean)	−12.0	−8.8	−13.4	−21.0	−11.0	−0.3
Observations	398,081	395,426	396,104	382,040	896,955	896,955
Panel B. Households without children						
Treated × Post	−2.103***	−1.256***	−0.794***	−0.312***	−4.549***	−4.318
	(0.485)	(0.428)	(0.216)	(0.094)	(0.708)	(5.630)
Pre-policy mean ($)	25.486	21.040	5.817	2.747	55.342	458.003
Effect (% of mean)	−8.3	−6.0	−13.6	−11.4	−8.2	−0.9
Panel C. Households with children						
Treated × Post	−5.074***	−3.244***	−1.239**	−1.340***	−9.993***	3.124
	(1.717)	(0.742)	(0.569)	(0.170)	(2.684)	(6.902)
Pre-policy mean ($)	31.583	29.580	9.384	4.782	74.062	569.699
Effect (% of mean)	−16.1	−11.0	−13.2	−28.0	−13.5	0.5
Group × cohort FE	Yes	Yes	Yes	Yes	Yes	Yes
Time × cohort FE	Yes	Yes	Yes	Yes	Yes	Yes

Notes: This table reports stacked difference-in-differences estimates of β_k in Equation (1), estimated by OLS, using the household-by-month panel of SNAP participants. The dependent variable is spending in dollars per household-month and equals zero in non-purchasing months. Columns (1)–(4) use a stacked sample of households in "treatment" states with the corresponding SNAP restriction and a sample of households in "control" states with no restrictions. Columns (5)–(6) stack all combinations of restrictions and implementation dates, leading to a larger sample than in Columns (1)–(4) for category-specific restrictions alone. Standard errors are clustered by state and reported in parentheses. Treated states in each column apply that particular product restriction, while control states apply no restriction. Panels B-C split each column using self-reported information on the presence of children. Each reported percentage effect is constructed using that subgroup's own pre-policy average among households in treated states. All specifications include household-by-cohort and calendar-month-by-cohort fixed effects. $^{*}p < 0.10$, $^{**}p < 0.05$, $^{***}p < 0.01$.

Table 3: Heterogeneity in soda effects by margin of adjustment

	(1) Soda trips (extensive)	(2) $ per soda trip (intensive)	(3) Monthly volume (mL)	(4) Log $ per mL	(5) Private label share	(6) Grocery trips (all grocery)
Treated × Post	−0.221***	−0.174	−1,422.7***	−0.00727	0.00714	−0.388***
	(0.061)	(0.141)	(453.2)	(0.0115)	(0.00525)	(0.109)
Pre-policy mean	1.876	9.927	9,993.3	—	—	20.78
Effect (% of mean)	−11.8	−1.8	−14.2	—	—	−1.9
Observations	398,081	185,627	398,081	167,336	185,388	398,081
Group × cohort FE	Yes	Yes	Yes	Yes	Yes	Yes
Time × cohort FE	Yes	Yes	Yes	Yes	Yes	Yes

Notes: Each column reports a stacked difference-in-differences estimate of Equation (1). Standard errors are clustered by state and reported in parentheses. The outcome in Column (1) is the number of distinct grocery baskets containing soda in a household-month, with zero assigned when no soda is purchased. The outcome in Column (2) is household-month soda spending divided by the number of soda-containing baskets and is defined only for household-months with at least one such basket. The outcome in Column (3) is total monthly volume of soda purchased in milliliters, using only items with observable size information. Household-months that do not contain a soda purchase with observed size information are assigned zero. The outcome in Column (4) is the log of total spending on size-observed soda divided by total milliliters purchased. The outcome in Column (5) is private-label soda spending divided by total soda spending. Columns (1)–(5) limit the sample to soda purchases at the sixteen retailers listed in Section 3. Columns (1), (2), and (5) use all soda purchases at these retailers regardless of whether size is observed, whereas Columns (3)–(4) require product-size information. Column (6) is the total number of grocery baskets per household-month. All specifications include household-by-cohort and calendar-month-by-cohort fixed effects. Pre-policy means are calculated for treated households before implementation. $^{*}p < 0.10$, $^{**}p < 0.05$, $^{***}p < 0.01$.

Table 4: Heterogeneity in soda effects by store type

	Large	Small	Online	Other outlets
Panel A. Estimated effects				
Soda-spending change ($)	-2.954^{***}	$+0.121$	-0.626^{***}	$+0.073$
	(0.545)	(0.417)	(0.163)	—
Effect (% of pre-policy mean)	-16.2	$+1.9$	-22.2	—
Panel B. Pre-policy statistics				
Share of soda spending (%)	64.8	22.5	10.0	2.7
Soda dollars per trip ($)	9.21	7.52	12.23	—
SNAP used on observed-tender trips (%)	39.6	25.0	—	—
Soda dollars per SNAP trip ($)	10.09	9.25	—	—
Soda dollars per non-SNAP trip ($)	8.51	6.90	—	—

Notes: Panel A reports difference-in-differences estimates using the stacked sample of SNAP households and the specification in Equation (1). Standard errors are clustered by state and reported in parentheses. The "Other outlets" estimate sums the remaining store-format coefficients. Percentage effects are expressed relative to the store-type-specific pre-policy mean. The corresponding all-store estimate is reported in Table 2. Panel B reports pre-policy statistics for treated households. Share of soda spending is each store type's share of treated households' pre-policy soda spending. Dollars per soda trip is total soda spending divided by the number of distinct baskets containing soda and includes all order channels. The tender rows are limited to in-store soda trips at large and small retailers for which tender is observed. Non-SNAP includes every observed tender other than SNAP, not only cash. $^{*}p < 0.10$, $^{**}p < 0.05$, $^{***}p < 0.01$.

5.3 Robustness checks

We perform two checks to assess the robustness of our main results. First, treated and control states may differ in their responsiveness to the national MAHA agenda, potentially confounding the estimated effects. To address this concern, we estimate Equation (1) using a narrower control group of seven politically conservative states that did not implement a SNAP food restriction during the study period. This gives an estimated effect of −$3.191 per household-month, compared with −$3.385 in our main specification. The original estimate is well within the 95% confidence interval of this new estimate.

Second, differential growth in GLP-1 utilization could confound our estimates if it independently reduced soda purchases in treated states around the time of implementation. To address this concern, we augment the specification in Equation (1) with quarterly state-level measures of GLP-1 prescriptions among Medicaid enrollees.[14] At the time of the analysis, prescription data were available only through 2026Q1. As a result, we end this exercise in March 2026, which excludes the April implementation cohort. To distinguish the effect of adding the GLP-1 controls from the effect of changing the sample, we first estimate the baseline specification on this same shortened sample. The resulting coefficient is −5.143 dollars per household-month. Adding the broad GLP-1 measure as a covariate reduces the estimate to −5.128, while adding the obesity-label measure as a covariate further reduces the estimate to −5.096. Both estimates with the GLP-1 controls are within 1% of the baseline estimate. Table 5 reports these results. Details are provided in Appendix B.

Table 5: Robustness of the soda-spending estimates

Specification	Estimate	Standard error	Observations
Main specification	-3.385^{***}	0.991	398,081
Conservative control states	-3.191^{***}	0.836	85,813
Through March 2026, no GLP-1 control	-5.143^{***}	0.680	202,558
Control for all GLP-1 prescriptions	-5.128^{***}	0.679	202,558
Control for obesity-label products	-5.096^{***}	0.679	202,558

Notes: Estimates are in dollars per household-month. The final three rows use observations only through March 2026 given the 2026Q1 coverage of the CMS data. Standard errors are clustered by state. $^{***}p < 0.01$. Additional details are available in Appendix B.

[14] We construct these measures from the public CMS State Drug Utilization Data. For each state-quarter, we sum reported fee-for-service and managed-care prescriptions and divide by the sum of monthly Medicaid enrollment over the same three months. We multiply the resulting rate by 1,000 and enter its inverse hyperbolic sine in the regression. The broad measure includes all identified GLP-1 products. The narrower obesity-label measure includes Wegovy, Zepbound, and Saxenda.

Table 6: Effect of the restrictions on spending and purchased nutrients

Outcome	Effect	Standard error	Pre-policy mean	Effect (%)
Panel A. Soda ban → beverages				
Beverage spending ($)	−4.900**	1.901	77.368	−6.33
Calories	−808.1**	305.3	11,051.9	−7.31
Total sugar (g)	−215.2***	76.8	2,662.9	−8.08
Sodium (mg)	−716.7***	246.9	8,071.2	−8.88
Panel B. Candy ban → packaged snacks				
Packaged snack spending ($)	−4.291***	0.812	89.057	−4.82
Calories	−1,372.4***	226.6	26,866.2	−5.11
Total sugar (g)	−134.4***	9.4	1,759.2	−7.64
Sodium (mg)	−1,864.5***	466.3	29,247.8	−6.37
Panel C. Beverage + Packaged snack basket				
Combined spending ($)	−8.937***	2.305	167.397	−5.34
Combined calories	−2,061.3***	485.5	38,232.6	−5.39
Combined total sugar (g)	−342.4***	79.3	4,457.1	−7.68
Combined sodium (mg)	−2,830.0***	560.0	37,956.8	−7.46

Notes: Stacked difference-in-differences; standard errors are clustered by state. Panels A and C use the full soda-ban sample ($N = 398{,}081$). Beverage outcomes in Panels A and C exclude dry coffee and tea, powdered drink mixes, liquid enhancers, and hot cocoa. In Panel B, treated states are the seven states that restricted candy, and control states are those that adopted no SNAP food restriction ($N = 395{,}426$). Panel B reports the observed packaged snack basket. Because Florida also restricts packaged cookies and snack cakes, it is a net basket effect rather than a pure substitution estimate. Panel C reports the observed beverage-plus-packaged snack basket rather than the total diet. Percentage effects divide by the matching treated pre-policy mean. $^{*}p < 0.10$, $^{**}p < 0.05$, $^{***}p < 0.01$.

5.4 Nutrition and substitution

The nutritional effects of the SNAP waivers depend both on changes in purchases of restricted items and on substitution toward other items. If households substitute toward unrestricted products with similar nutritional content, even a large decline in targeted purchases may produce little improvement in overall diet quality. In this subsection, we estimate the effects of the restrictions on overall spending and nutrition for the categories most likely to absorb displaced SNAP demand.

The soda ban and beverages. The ban cuts purchased beverage calories by 7.3%, total sugar by 8.1%, and sodium by about 8.9% (Table 6, Panel A). The decline is concentrated in soda. Juice calories rise insignificantly – about 8% of the soda-calorie decline. Energy and sports drink calories also fall, but this is driven by the presence of combined restrictions in some states rather than substitution away from these products.

Table 7: Product substitution: spending and calories

Product bucket	$ DiD (s.e.)	Δ calories (s.e.)	Pre-policy calorie share (%)
Panel A. Soda ban → beverages, by bucket			
Soda	−3.385*** (0.991)	−771.7*** (231.4)	56.7
Energy & sports	−1.207** (0.494)	−95.5** (45.2)	7.8
Water	−0.252 (0.212)	+1.8 (1.6)	0.4
Juice	+0.051 (0.271)	+64.6 (49.9)	19.8
Coffee	−0.139 (0.108)	−13.9 (10.2)	1.6
Tea	−0.148 (0.091)	−21.4 (27.7)	4.5
Tonic	−0.010 (0.007)	+0.9 (0.9)	0.1
Other	+0.190 (0.231)	+27.1 (27.0)	9.2
Total beverages	−4.900** (1.901)	−808.1** (305.3)	100.0
Panel B. Candy ban → packaged snacks			
Candy (restricted)	−2.185*** (0.603)	−590.9*** (122.1)	20.9
Non-candy packaged snacks			
Savory snacks (eligible)	−0.575** (0.273)	−186.4 (150.1)	27.4
Frozen desserts (eligible)	−0.720*** (0.175)	−135.1*** (34.5)	8.3
Other non-candy snacks (residual)	−0.810 (0.499)	−460.0*** (124.0)	43.4
Non-candy total	−2.105*** (0.664)	−781.5*** (168.0)	79.1
Total packaged snacks	−4.291*** (0.812)	−1,372.4*** (226.6)	100.0

Notes: Table 7 reports stacked difference-in-differences estimates with group×cohort and time×cohort fixed effects; standard errors are clustered by state. Panel A uses the full soda-ban sample ($N = 398{,}081$). In Panel B, treated states are the seven states that restricted candy, and control states are those that adopted no SNAP food restriction ($N = 395{,}426$). Beverage buckets partition beverage spending. The three indented rows partition the non-candy total; candy plus the non-candy total partitions total packaged snacks. Savory snacks include plainly savory chips, crackers, pretzels, cheese snacks, meat snacks, and plain popcorn. Frozen desserts combine packaged ice cream and frozen novelties. The other non-candy row is a mixed residual and includes packaged cookies and snack cakes that Florida restricts concurrently. The final column reports each row's share of treated pre-policy calories within its own panel: beverage calories in Panel A and packaged-snack calories in Panel B.
$^{*}p < 0.10$, $^{**}p < 0.05$, $^{***}p < 0.01$.

The candy ban and packaged snacks. Comparing the seven states that restricted candy with states that adopted no SNAP food restriction, total packaged-snack spending falls by 4.8%, calories by 5.1%, total sugar by 7.6%, and sodium by 6.4% (Table 6, Panel B).

Table 8: Fruit-beverage substitution at Kroger and Walmart

Low-juice policy	Product group	Pre-policy share (%)	Calorie effect (%)	p-value	Soda offset (%)
Eligible	Sweetened low-juice drinks	42.0	+4.7	.116	+8.8
	Juice and other retained	34.0	+5.8	.171	+8.8
	Unclassified	24.0	+7.0	.017	+7.5
	Total *Fruit Juice*	100.0	+5.6	.002	+25.1
	Soda, for reference	—	−10.9	$< .001$	—
Restricted	Sweetened low-juice drinks	40.4	−19.2	.040	−19.0
	Juice and other retained	32.8	+37.0	.001	+29.7
	Unclassified	26.9	+17.4	.098	+11.4
	Total *Fruit Juice*	100.0	+9.0	.184	+22.2
	Soda, for reference	—	−9.6	.031	—

Notes: In the five-state group (NE, UT, WV, LA, and FL), the sweetened low-juice products shown here remain SNAP-eligible. In the four-state group (IN, IA, ID, and OK), those products are restricted. Texas is excluded because its sweetener-based rule cannot be mapped to these categories. The analysis uses Kroger and Walmart, where receipt items can be linked to item-level eligibility evidence. Calorie effects are percentages of each row's treated pre-policy mean, and standard errors are clustered by state. Each soda offset divides the row's calorie level effect by the absolute soda-calorie level effect estimated in the same retailer and policy-group sample. The three component offsets sum to the total *Fruit Juice* offset up to rounding; the unclassified component is not allocated.

Combined basket. Combining beverages and packaged snacks, monthly spending falls by 5.3%, purchased calories by 5.4%, total sugar by 7.7%, and sodium by 7.5%. The estimates are reported in Panel C of Table 6. The larger proportional decline in sugar than in calories suggests that the restrictions disproportionately reduce sugar-intensive purchases. Because several states also restrict candy, energy drinks, or prepared desserts, these estimates should be interpreted as capturing the net effect of all waivers, not just the soda restriction. Figure 4 reports the corresponding event studies.

Substitution across categories. We consider spending responses in restricted and related, but unrestricted categories in Table 7. Panel A shows that the full soda-ban sample shows little substitution toward other beverages, with only products in the juice category experiencing a small, but insignificant spending increase. This pooled estimate, however, combines states that restricted fruit drinks and other low-juice products within Numerator's broad *Fruit Juice* category with states that left them eligible, potentially confounding declines in the former states with substitution in the latter.[15] We therefore examine the latter states separately.

Numerator does not report item-level SNAP eligibility, so we limit this exercise to purchases at Kroger and Walmart, where we are able to link receipt items with state-specific information about SNAP eligibility. At these retailers, soda calories fall by 10.9% and fruit-juice calories rise by 5.6% in the five states that left fruit drinks eligible, representing a 25.1% offset of the soda decline (Table 8). Allocating unclassified products in proportion to the classified shares, we estimate that the offset is almost equally divided between low-juice products (12.9%) and juice and other retained products (12.2%). This provides evidence of some degree of product substitution to nearby product categories.

For candy restrictions, spending on candy falls by 8.8%, while spending on savory snacks and frozen desserts falls by 2.6% and 8.5%, respectively (Table 7, Panel B). The latter two categories remain eligible. Thus, nearby eligible snack purchases fall rather than rise.

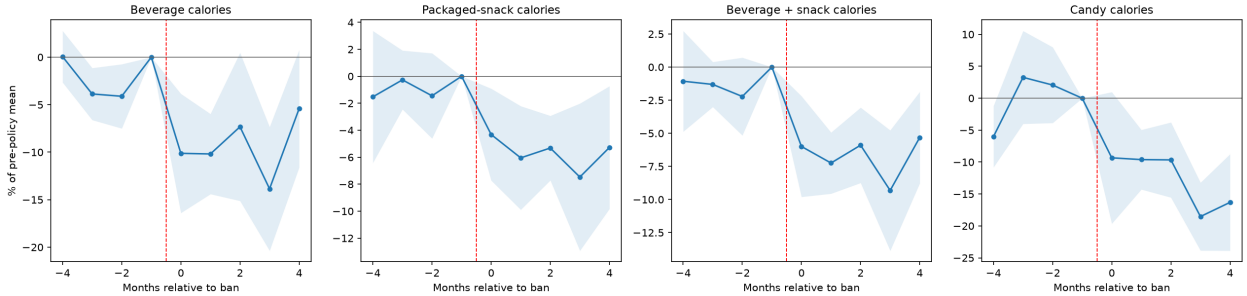
There are two possible explanations for these asymmetric product substitution patterns. First, households may not have fully understood the scope of product restrictions.[16] This uncertainty may have reduced purchases of adjacent eligible snacks. Alternatively, the restrictions may have served as a broader nutrition signal, reducing purchases of foods viewed

[15] Ready-to-drink tea poses the same classification issue as fruit juice: some products are restricted in some states but remain eligible in others. Repeating the same analysis, we find no statistically significant change in tea spending in either policy group (Appendix Table 20).

[16] Louisiana Health Secretary Bruce Greenstein described candy classification as "quite tricky" and noted that a Chips Ahoy cookie remained eligible while a Twix did not (O'Donoghue, 2026). The introduced and Senate committee versions of Texas S.B. 379 included potato or corn chips and packaged cookies, whereas the enrolled law covered only sweetened drinks and candy (Texas Legislature, 2025a,b).

as unhealthy even when they remained eligible. Although we cannot conclusively distinguish these explanations, we find suggestive evidence in support of misinformation about the scope of eligibility restrictions: SNAP households reduce purchases of eligible snacks in stores but increase spending online, where the website displays each product's SNAP eligibility (Appendix Table 19).

Figure 4: Dynamic effects on purchased calories



Notes: From left to right, the first three panels in Figure 4 report effects on beverage, packaged snack, and combined calories under the soda ban using the full sample. In the final (rightmost) panel, treated states are the seven states that restricted candy, and control states are those that adopted no SNAP food restrictions. Each series is scaled by the corresponding treated pre-policy mean. Lines connect stacked-DiD coefficients, with $t = -1$ normalized to zero. Shaded regions are 95% confidence intervals with standard errors clustered by state. The dashed line marks implementation.

6 Understanding consumer responses

In this section, we investigate which mechanisms explain the observed reduction in SNAP households' purchases of restricted items after the implementation of food restriction waivers. We first ask whether the restrictions change households' budget sets and whether those changes can explain the decline. We then consider non-budget explanations, including nutritional signaling, checkout stigma and payment frictions, and psychological departures from fungibility.

6.1 Does a change in the budget set explain the consumption response?

6.1.1 Changes in effective prices

We first ask whether the decline in purchases can be explained by changes in the prices paid by consumers, which could occur for two reasons. First, retailers may raise shelf prices in response to the policy-induced change in demand, which could potentially account for part of the observed decline. Second, when restricted products lose SNAP eligibility, households that continue to purchase them with cash may have to pay sales tax, raising the tax-inclusive price. We examine these two explanations in turn.

Retailer price responses. We use the scraped price dataset to estimate the effects of the SNAP soda restrictions on prices using a difference-in-differences research design. We compare item-level prices in stores in two states implementing the ban (i.e., TX and FL)

to stores in a nearby control state without the ban (GA), separately for each retailer. The analysis uses the full available beverage panel over event weeks −5 through +9. For each implementing state, we separately estimate the linear model

$$\log p_{ijst} = \alpha_{ijs} + \lambda_{ijt} + \beta_j \times \text{Treated}_s \times \text{Post}_t + \epsilon_{ijst}, \tag{4}$$

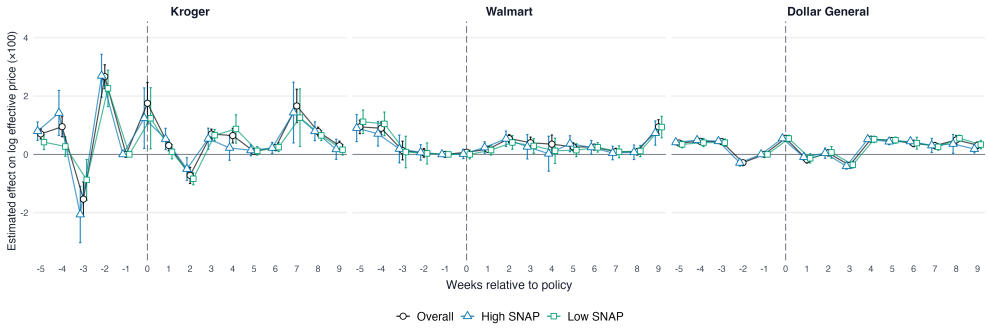
where p_{ijst} is the posted price (inclusive of promotions) of product i at retailer j, store s, and week t. The variables Treated_s and Post_t are indicators equal to one for stores located in treated states and for periods after the SNAP soda ban is implemented, respectively. The store-by-product fixed effects, α_{ijs}, absorb time-invariant differences in price levels across stores for a given product, including persistent differences across states and local markets. The product-by-retailer-by-timing-cell fixed effects, λ_{ijt}, absorb retailer-specific national price changes, promotions, and product-level shocks common to all of a retailer's stores in a given promotional cycle. For Dollar General, Kroger, and Walmart, we use the exact scrape date as the timing cell.[17] Because Publix promotions follow market-specific weekly advertising calendars, we assign each observation to its store's inferred promotional cycle. These cycles generally run Thursday-Wednesday in Florida and Wednesday-Tuesday in Georgia. As a result, β_j is identified from changes in the same item's price at treated and control stores for each retailer in the same promotional cycle. We cluster standard errors at the level of the physical store. Finally, we report a pooled specification obtained by stacking the retailer-level samples and retaining retailer-specific item and timing fixed effects.

To ensure that we do not miss differential responses by retailers in markets with greater exposure to the waivers, we also examine whether prices respond more in markets with higher shares of SNAP participants. In practice, within each treated-control state comparison we sort unique ZIP codes into terciles using the share of households receiving SNAP and retain only stores in the top and bottom terciles of SNAP exposure. We also estimate the price response separately for low- and high-SNAP ZIP codes. Table 9 reports the static estimates of Equation (4), and Figure 5 reports the corresponding event study estimates that allow for dynamic effects.

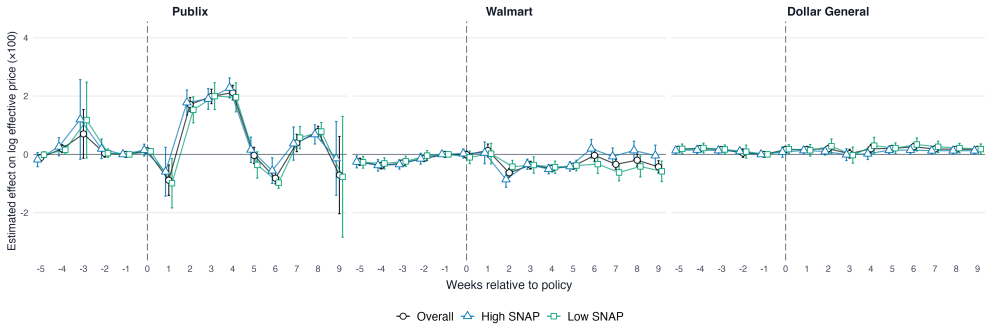
The estimates are small but differ across retailers. In Texas, the retailer-specific estimates imply price changes of 0.06% at Kroger, -0.06% at Walmart, and -0.03% at Dollar General. None of these estimates is statistically significant. In Florida, effective prices rise by 0.33% at Publix and 0.06% at Dollar General and fall by 0.07% at Walmart. The Publix and Dollar General estimates are statistically significant, while the Walmart estimate is not. Every

[17] There is one exception: the Florida–Georgia Dollar General comparison uses an event-week timing cell at event week −4 because the two states' scrape dates do not overlap.

Figure 5: Dynamic retailer price responses to soda ban



(a) Texas vs. Georgia



(b) Florida vs. Georgia

Notes: Figure 5 reports dynamic event study estimates of the effect of the SNAP soda restrictions in Texas (April 1) and Florida (April 20), relative to the control state of Georgia. The black series uses all stores in the treatment and control states. The blue (green) series restricts the sample to stores located in ZIP codes in the top (bottom) tercile of SNAP exposure across the two states. The sample is identical to Table 9. Bars represent 95% confidence intervals. Standard errors are clustered at the store level.

retailer-specific point estimate is smaller than 0.4% in magnitude, suggesting at most very modest responses to the restrictions.

Consistent with this interpretation, the difference-in-differences estimates do not increase systematically with local SNAP exposure. The differences in price responses between stores in the top and bottom terciles of SNAP exposure are small and statistically insignificant in Texas. In Florida, the differences are 0.22 percentage points at Publix and 0.23 percentage points at Walmart. Only the latter is statistically significant, and it points towards larger responses in high-SNAP ZIP codes. The estimated Dollar General difference is -0.09 percentage points and is marginally significant at the 10% level, indicating a smaller (though not economically significant) response in high-SNAP ZIP codes. These opposite-signed differences provide no consistent evidence that retailers raised prices more in areas with greater SNAP exposure.

The event studies in Figure 5 show substantial week-to-week variation around these small point estimates, particularly at Publix. The Publix estimate reflects episodic price movements rather than a smooth, persistent increase in prices after the restriction.[18] In Appendix H, we extend the Publix panel through the week ending August 23, 2026. The extended panel displays the same ongoing promotional cycles, with effective prices in Florida periodically elevated relative to Georgia. Overall, the results are similar if we limit the price sample to a balanced panel of items present in most stores in most weeks, restrict it to specific product categories, or assign the treatment across stores with similar pre-policy pricing behavior, as in Adams and Williams (2019). These results are reported in Appendix H. Altogether, the estimates provide no evidence that retailer pricing responses contribute significantly to the observed decline in consumption.

Finally, retailers do not appear to reduce the frequency of promotions, which could reduce demand through behavioral channels (Sahni et al., 2017). In Appendix Table 28, we report estimates of Equation (4) with measures of promotional intensity or frequency used as the outcome. Pooling across retailers, we estimate a significant 1.21pp increase in promotional frequency in Florida and an insignificant 0.11pp decrease in promotional frequency in Texas, inconsistent with the large observed decline in demand.

Sales tax changes. We next examine how the restriction changes the net price paid by SNAP participants through the loss of SNAP's sales tax exemption. Federal law exempts purchases made with SNAP benefits from state and local sales taxes (U.S. Department of

[18] In practice, the Publix effect is driven by Georgia switching to a buy-two-get-two free promotion for specific Coca-Cola products, while Florida remains on a buy-two-get-one free promotion across ad cycles.

Table 9: Retailer price responses by local SNAP exposure

	Texas vs. Georgia			Florida vs. Georgia		
	(1) Kroger	(2) Walmart	(3) Dollar General	(4) Publix	(5) Walmart	(6) Dollar General
All stores	0.06	−0.06	−0.03	0.33**	−0.07	0.06***
	(0.09)	(0.06)	(0.02)	(0.13)	(0.05)	(0.02)
Low-SNAP ZIPs	0.04	−0.14	−0.02	0.13	−0.20**	0.09**
	(0.11)	(0.10)	(0.03)	(0.21)	(0.08)	(0.04)
High-SNAP ZIPs	−0.06	−0.12	−0.03	0.35**	0.03	−0.01
	(0.19)	(0.10)	(0.03)	(0.17)	(0.08)	(0.02)
High minus low	−0.10	0.03	−0.01	0.22	0.23**	−0.09*
	(0.22)	(0.14)	(0.04)	(0.27)	(0.11)	(0.05)
Observations, all store model	1,421,274	1,369,089	631,161	1,222,906	1,137,851	637,786
Observations, low-SNAP model	523,596	489,408	190,846	438,767	407,641	217,726
Observations, high-SNAP model	421,618	427,753	250,926	393,344	354,689	218,605
Stores, all store model	367	220	200	273	180	200
Stores, low-SNAP model	136	76	60	100	62	68
Stores, high-SNAP model	108	70	80	87	58	69

Notes: Entries in Table 9 are estimates of $100 \times \beta_j$ from Equation (4). Columns (1)-(3) compare Texas stores with Georgia stores around the April 1 restriction. Columns (4)-(6) compare Florida stores with Georgia stores around the April 20 restriction. The "all store" estimates use the full sample. The low- and high-SNAP estimates come from separate DiD regressions on ZIP codes in the bottom and top terciles of SNAP participation. Terciles are formed over unique ZIP codes separately within each policy–retailer comparison. High minus low is the difference between the two independently estimated coefficients. All models include store-by-item and item-by-timing-cell fixed effects. Standard errors, clustered by physical store, are in parentheses. *, **, and *** denote significance at the 10%, 5%, and 1% levels, respectively.

Agriculture, Food and Nutrition Service, 2026). Once a good becomes SNAP-ineligible, however, a participant who continues buying it must pay with cash or a card and may owe sales taxes. This channel applies across much of the country. In 25 states and the District of Columbia, which together account for about 70 percent of national SNAP benefits, at least one restricted category is taxed more heavily than ordinary groceries.

The restriction changes the after-tax price only for purchases that would otherwise have been paid with SNAP. Let φ_k denote the share of pre-restriction spending on category k appearing on receipts on which SNAP was used. Let τ_{sk} denote the sales tax rate on category k in state s and τ_{gs} the rate on ordinary groceries, so the relevant tax differential is $\Delta\tau_{sk} = \tau_{sk} - \tau_{gs}$. Holding purchases fixed, the average effective consumer price increase is

$$\Delta p_{sk}^{\text{eff}} = \varphi_k \Delta\tau_{sk}. \tag{5}$$

These payment patterns sharply limit the price effect of the restriction. In states where a restricted category was taxed more heavily than ordinary groceries, only 40.6 percent of soda spending, 35.4 percent of candy spending, 35.5 percent of energy-drink spending and 42.0 percent of sports-drink spending appeared on receipts on which SNAP was used before the policy. The remaining purchases had no SNAP tax exemption to lose. Across the

30 treated state-category combinations in our study, removing eligibility raises the average tax-inclusive price by only 1.88 percent of the pre-tax price. Among combinations with a positive tax differential, the increase is 2.36 percent. Had households used SNAP for every pre-policy purchase ($\varphi_k = 1$), the corresponding increases would instead have been 4.94 and 6.22 percent.[19]

If this *small* tax-inclusive price change were an important driver of the purchase response, soda spending should fall more in states with a positive soda-versus-grocery sales tax differential. Among the states that started the policy in January, soda spending falls by 13.0 percent in states with no differential and by 13.8 percent in states with a positive differential. The two dollar estimates are not statistically different ($p = 0.379$; Table 10). Together with the retailer price estimates, this comparison indicates that changes in the effective price paid by participants are unlikely to explain the observed decline in purchases of items targeted by the restrictions.

Table 10: Soda effects by sales-tax differential within the January cohort

Tax differential	States	Effect ($)	SE	Pre-policy mean	Effect (% of mean)
Zero	NE, UT	−4.514	0.778	34.842	−13.0
Positive	IN, IA, WV	−5.738	1.155	41.686	−13.8
Zero minus positive	—	1.224	1.373	$p = 0.379$	

Notes: The table applies the paper's stacked difference-in-differences specification to states in the common January implementation cohort. Standard errors are clustered by state. The final row tests the difference between the two dollar coefficients.

6.1.2 Could a binding budget constraint explain the decline?

The restriction could mechanically reduce purchases if households do not buy enough other SNAP-eligible groceries to absorb the benefits used on the restricted products. Otherwise, a household can redirect its SNAP benefits toward groceries that remain eligible and use the cash freed by that switch to keep buying the restricted products. In that case, the household can still afford its pre-policy bundle even though the payment method changes (Hoynes and Schanzenbach, 2009).

This explanation is unlikely, because the ban does not bind the pre-ban choices of the vast majority of SNAP participants. We show this directly using data on households' pre-policy purchases, when they received SNAP without the product restriction. For 95.8% of households, spending on groceries that would remain eligible was at least as large as eligible-

[19]The average across states reflects each state's SNAP population, since we weight by total SNAP benefits. Appendix A describes it in detail.

grocery spending recorded on SNAP-paid receipts. This implies that these households could redirect the corresponding amount of SNAP spending toward groceries that remained eligible and pay cash for the restricted products without changing their consumption bundle.[20]

6.2 Mechanisms beyond the budget set

In this subsection, we investigate mechanisms other than changes in households' budget sets. The restriction and accompanying media coverage may signal that the targeted products are unhealthy and lead households to reassess them. The restrictions may also create stigma or payment frictions at checkout if SNAP participants are forced to switch between payment methods. Finally, we consider behavioral mechanisms, including mental accounting and labeling effects, that may lead households to treat SNAP benefits differently from cash.

6.2.1 Could the policy itself convey a nutritional signal?

The restrictions may provide an additional nutritional signal above and beyond their impact on product eligibility. For example, local news covering the waivers may make a product's nutritional consequences more salient or convey that the state considers it unhealthy. Such a signal could affect non-SNAP households as well as SNAP households. The response of these households helps us gauge how much of the SNAP-household decline could reflect a broader reaction to the policy. This comparison cannot, however, distinguish nutritional signaling from direct spillovers of the decline in SNAP participants' demand onto nonparticipants. As a result, we interpret the results with caution.

Using the same stacked design, we estimate the OLS and PPML event studies among non-SNAP households, with soda spending as the outcome. The estimates imply a persistent decline, equal to 3.4% of the non-SNAP pre-policy mean spending. The proportional PPML specification yields a similar 3–4% decline after implementation, compared with the earlier 12% decline among SNAP households (Figure 9 in Appendix E and Table 11). Thus, some decline in soda demand occurred among non-SNAP participants, consistent with a nutritional signal.

Still, the non-SNAP decline is only about one-quarter to one-third as large as the SNAP-household decline. Given this difference in the relative magnitudes of the estimates, we treat the non-SNAP estimate as evidence that broader forces may contribute to the decline in SNAP purchases, although such forces cannot explain most of the SNAP-household decline.

[20] In this calculation, we average only over months in which the household is classified as using SNAP. Our calculation includes estimated restricted fruit-beverage and ready-to-drink tea spending in Indiana, Iowa, Idaho, and Oklahoma. It excludes Texas because its sweetener-based beverage rule cannot be mapped using the same product classification. See Appendix E.3.

Table 11: Heterogeneity in soda effects for SNAP and non-SNAP households

	SNAP	Non-SNAP
Treated × Post	−3.385***	−0.665***
	(0.991)	(0.118)
Pre-policy mean ($)	28.109	19.372
Effect (% of mean)	−12.0	−3.4
Observations	398,081	2,341,435
Group × cohort FE	Yes	Yes
Time × cohort FE	Yes	Yes

Notes: Stacked difference-in-differences; soda spending in dollars per household-month (non-purchasing months entered as zeros); standard errors clustered by state in parentheses. The percentage row divides the level coefficient by the corresponding treated pre-policy mean. Non-SNAP households have no observed SNAP tender in any channel. This table reports the dollar-level specification; the preferred proportional placebo inference comes from the PPML event study in Figure 9. The store-type decomposition is in Table 4. $^{*}p < 0.10$, $^{**}p < 0.05$, $^{***}p < 0.01$.

6.2.2 Could checkout stigma explain the decline?

Another explanation is that the restriction changes the checkout experience. USDA raised this concern when rejecting earlier restriction proposals, warning that a participant who was unaware of an exclusion might first learn about it after attempting to purchase the item with SNAP and being told at the register that the purchase was not allowed. This interaction could make the participant's SNAP status visible and cause embarrassment (U.S. Department of Agriculture, Food and Nutrition Service, 2007, pp. 4–5). The participant might then leave the item at the register.

Stigma could potentially also affect participants who already knew about the restriction. A participant who used SNAP for eligible groceries and cash for a restricted item might worry that the cashier or other customers would view the split payment as an attempt to circumvent the policy.

We examine these explanations by comparing online and in-store purchases. Online purchases may still require households to select separate payment methods, but they eliminate the interaction with a cashier and make the split payment less publicly visible. If checkout stigma were important, the decline should therefore be larger in stores than online. Instead, the proportional decline in soda spending is no smaller online than in store: it is 19.1% online and 11.2% in store (Appendix Table 21).[21] This pattern suggests that checkout stigma may

[21]These percentages divide the OLS dollar difference-in-differences estimates by each channel's own pre-policy treated mean in the common SNAP household-month stacked sample.

not be the primary explanation for the purchase response.

6.2.3 Psychological departures from full fungibility

Having found little evidence that the mechanisms above explain the decline, our preferred interpretation is a psychological departure from full fungibility: households may not treat SNAP benefits and cash as interchangeable when making grocery purchases. This interpretation is consistent with evidence that food spending responds more strongly to SNAP benefits than to an equivalent amount of cash (Hastings and Shapiro, 2018). Several mechanisms can produce this pattern, and our goal here is not to distinguish among them.

1. *Eligibility may directly raise desired quantity.* A program label can direct spending toward the goods associated with it. For example, recipients spent a labeled cash transfer for winter fuel disproportionately on fuel (Beatty et al., 2014).

2. *Eligibility may make an item seem cheaper.* Households may view SNAP eligibility as an implicit government subsidy attached to eligible goods even though it does not change their shelf prices. Removing eligibility may therefore make a restricted item feel more expensive.

3. *Mental Accounting: Households may keep SNAP and cash accounts separate across trips.* Households may not treat SNAP benefits and cash as a single pooled budget. Instead, they may mentally associate some grocery trips with SNAP and others with cash. Before the restriction, soda could be purchased within the SNAP account on a SNAP-associated trip. After the restriction, buying soda on the same trip requires the household to draw on its cash account and redirect SNAP benefits toward other eligible groceries. If reconsidering and coordinating the two accounts is mentally costly, the household may reduce its soda purchases rather than fully reorganize payment across the two accounts. Purchases can therefore fall even when the household has enough total resources to afford the same bundle.

All three mechanisms can reduce soda purchases even though households can still afford their pre-ban bundle after the restriction. While our goal is not to determine which psychological mechanism drives the response, two patterns reported in Section 5.2 are consistent with the separate accounts story. First, the decline is larger at larger retailers (where SNAP was more commonly used). Second, most of the decline comes from fewer trips containing soda, rather than lower soda spending on trips in which households continue to buy it. One possible interpretation is that, after soda loses eligibility, households stop adding it to some

SNAP-associated trips rather than fully reorganizing spending across their SNAP and cash accounts.

The following analysis of the spending response as a function of a household's SNAP association of a store supports this interpretation. We measure this association using each household's own pre-policy payment history. For each household, we label a retailer SNAP-heavy if the household used SNAP on at least half of its trips there, mixed-use if the share was positive but below one-half, and never-SNAP if it used SNAP on none. The same chain can therefore receive different labels for different households. To avoid comparing different households, we keep only households that shopped at all three retailer types before the policy. Using the 50 largest chains in grocery, drug, mass, convenience, and dollar channels where SNAP acceptance is universal, restricted-product spending fell by 19.8% at SNAP-heavy retailers and 10.6% at mixed-use retailers, while the estimated 3.8% decline at never-SNAP retailers was statistically indistinguishable from zero (Table 12). The Appendix reports similar results using the 100 largest chains and all chains with at least one observed SNAP-paid receipt in a pre-policy period.

Table 12: Restricted-product spending by household-specific SNAP use at retailers

	SNAP-heavy	Mixed-use	Never-SNAP
Effect (% of pre-policy mean)	-19.8^{***}	-10.6^{***}	-3.8
	(5.9)	(2.0)	(5.7)
Observations		316,080	

Notes: The sample contains households that shopped at all three retailer types before implementation. Labels are formed separately for each household using pre-policy in-store grocery trips. All retained retailer–state–cohort cells have observed SNAP use, so never-SNAP describes the focal household's payment behavior, not whether the retailer accepted SNAP. Percentage effects divide each OLS coefficient by its retailer-type-specific treated pre-policy mean. Standard errors in parentheses are scaled by the same means. Models use the main stacked difference-in-differences fixed effects and cluster standard errors by state. Appendix D shows that using the 100 largest chains or all qualifying chains yields the same gradient and also reports the corresponding soda estimates. $^{*}p < 0.10$, $^{**}p < 0.05$, $^{***}p < 0.01$.

6.2.4 Why do firms barely change prices when demand falls?

A final question is why retailers do not adjust prices (see Section 6.1.1) in response to falling demand from SNAP participants. We first note that under imperfect competition, a downward shift in demand does not necessarily change a firm's optimal price. Ultimately, what matters is whether the restriction changes the price elasticity of total market demand. In Appendix I, we calibrate the three psychological models in Section 6.2.3 using elasticity estimates from prior studies and plausible market shares for SNAP participants and nonparticipants. Although the models make different predictions about how the ban changes SNAP participants' price elasticity, they all imply only a small change in the overall market elasticity, leaving retailers with little incentive to adjust prices, particularly if they face adjustment costs. The two psychological eligibility models also allow us to calibrate the households' elasticity of demand in a cash-only world, in the absence of SNAP benefits. We use these calibrated elasticities in the cost-benefit analysis of the eligibility restrictions below.

7 Policy implications

Here, we consider two policy implications of our findings: financial incidence and efficiency. First, we compare the out-of-pocket burden of a restriction that reduces purchases through a nonprice channel with that of an equivalent Pigouvian tax. Second, we assess the efficiency gains from adding the restriction to SNAP relative to otherwise identical benefits without the restriction.

7.1 Financial burden relative to equivalent excise taxes

Pigouvian taxes are a conventional policy tool for reducing consumption of harmful goods, but concerns that they impose financial burdens on low-income households have featured prominently in debates over taxes on sugary drinks, tobacco, and energy (Allcott et al., 2019; Douenne, 2020). The restriction, by contrast, produces sizable purchase reductions while leaving shelf prices largely unchanged. We therefore compare its financial burden on SNAP participants with that of excise taxes calibrated to generate the same purchase reductions among those participants.

Given the lack of retailer price responses we demonstrate above, we focus on sales tax and excise tax implications, assuming full pass-through. The first step is to quantify the additional sales tax payments generated by the restriction under nationwide adoption. Let T_{sk} denote the additional annual sales tax payment by SNAP households in state s for category k. For the analysis, we focus on three items: soda, candy, and energy drinks, the three categories restricted by the most states. For a like-for-like comparison, we count the taxes paid on the targeted goods under each policy but do not assign the spending released by the purchase reduction to another use. It is

$$T_{sk} = B_s \rho_k \left[\varphi_k \Delta\tau_{sk} - c_k \tau_{sk} \right]. \tag{6}$$

Here, B_s is annual SNAP benefits in state s, ρ_k is pre-policy spending on category k per dollar of benefits, φ_k is the share of that spending paid with SNAP, c_k is the estimated proportional decline in category spending, τ_{sk} is the tax rate on the restricted good, and $\Delta\tau_{sk}$ is its tax differential relative to ordinary groceries. The first term captures the loss of the SNAP exemption on purchases previously paid with benefits, net of the tax treatment of the eligible groceries that absorb those benefits. The second subtracts the tax on restricted purchases that are no longer made.[22]

We extrapolate the estimated spending declines – 12.0 percent for soda, 8.8 percent for candy, and 13.4 percent for energy drinks – to all states while retaining each state's own SNAP benefits and sales tax differential. For any set of product categories $\mathcal{K}$, define SNAP participants' total additional sales-tax payments under the restriction as

$$P_B(\mathcal{K}) \equiv \sum_s \sum_{k \in \mathcal{K}} T_{sk}.$$

The resulting payments are $100.7 million for soda alone and $176.4 million for soda, candy,

[22]For comparability, both calculations omit any tax on the subsequent use of the released spending. Appendix A considers ordinary groceries and taxable nonfood as alternative uses.

and energy drinks together. At the FY2024 average household benefit of roughly $340 per month, these amounts correspond to about $4.4 and $7.8 per SNAP household-year, respectively.

We then construct payments under excise taxes calibrated to generate the same purchase reductions among SNAP participants, using the uncompensated sugar-sweetened beverage price elasticity $|\varepsilon| = 1.373$ from Allcott et al. (2019). If spending in category k falls by c_k, the consumer-price increase needed to produce the same decline is approximately

$$\delta_k \simeq \frac{c_k}{|\varepsilon|}.$$

Under full pass-through, δ_k is also the excise-tax rate. Let $P_E(\mathcal{K})$ denote SNAP participants' total additional payments under the equivalent excise tax. Applying this calculation gives $670.8 million for soda alone and $1.356 billion across the three categories, or about $29.5 and $59.7 per SNAP household-year, respectively. If the government retains the revenue, the restriction's financial burden relative to the equivalent excise tax, P_B/P_E, is 13–15 percent: 15 percent for soda alone and 13 percent for the three-category basket.

If instead all additional revenue is returned as an equal lump-sum rebate to each household, the comparison also depends on relative purchases by SNAP and non-SNAP households. Let r denote average pre-policy spending per SNAP household on the restricted goods in $\mathcal{K}$, relative to the corresponding average per non-SNAP household. Assuming that the two groups have the same price elasticity, we approximate the net-burden comparison as

$$\frac{N_B}{N_E} \approx \frac{P_B}{P_E}\frac{r}{r-1}.$$

Table 13 reports both revenue assumptions. The external benchmark $r = 2$, motivated by Garasky et al. (2016) and Allcott et al. (2019, Figure III), implies net-burden ratios of 30 percent for soda and 26 percent for the three-category basket.[23] In our more recent data, $r = 1.45$ for soda and $r = 1.22$ across the three categories. Using the exact category-specific expression gives net-burden ratios of 48 percent for soda and 67 percent for the three-category basket.[24] Equal household rebates therefore narrow, but do not eliminate, the restriction's

[23]The USDA study reports annual soft-drink expenditures of $357.7 million for SNAP households and $1,263.3 million for non-SNAP households, with average monthly counts of 3.2 million and 23.3 million households, respectively. Dividing annual expenditures by 12 times these counts gives approximately $9.3 versus $4.5 per household-month, a ratio of about 2.1.

[24]If the SNAP-to-non-SNAP purchase ratio differs across categories, the exact expression is $N_B/N_E = P_B/\sum_{k\in\mathcal{K}} P_{E,k}(r_k-1)/r_k$, where r_k is the ratio of average pre-policy spending on category k per SNAP household to that per non-SNAP household, $P_{E,k}$ is SNAP households' payment under the corresponding equivalent excise tax, and $\mathcal{K}$ is the set of categories included in the comparison. If $r_k = r$ for every category, this reduces to the displayed expression with equality. For simplicity, the main text uses the aggregate-r

incidence advantage.

Table 13: Restriction burden relative to equivalent excise-tax burden

Product scope	r (our data)	Revenue retained	Equal rebates ($r = 2$)	Equal rebates (our data)
Soda only	1.45	15%	30%	48%
Soda, candy, and energy drinks	1.22	13%	26%	67%[a]

Notes: The last three columns report the restriction's financial burden as a percentage of the burden under excise taxes calibrated to produce the same purchase reductions among SNAP households. "Revenue retained" assumes no rebate. Equal rebates return all additional revenue in equal amounts per household. r is average pre-policy spending per SNAP household relative to a non-SNAP household. [a] For the three-category basket, the final column uses the exact category-specific formula. Using the aggregate $r = 1.22$ approximation instead gives 71 percent. Calculations use unrounded inputs. Additional details are in Appendix A.

This comparison points to a broader tradeoff between incidence and reach. If protecting low-income households' purchasing power is a central concern, a restriction attached to an income-support program can reduce purchases through a nonprice channel and impose less out-of-pocket burden than a tax. The difference would be even larger if purchases made with benefits received no special sales tax treatment, because removing eligibility would then create no additional sales tax payment. Such a restriction, however, directly affects only program participants, whereas a Pigouvian tax applies throughout the market. The preferred policy therefore depends both on protecting participants' purchasing power and on the value of reducing consumption among nonparticipants.

7.2 Cost-benefit analysis

Finally, we use the marginal value of public funds (MVPF) framework of Hendren and Sprung-Keyser (2020) to evaluate the efficiency of adding the restriction to SNAP. The MVPF is the social willingness to pay for a program divided by its net cost to the government:

$$\text{MVPF} = \frac{\text{Willingness to pay}}{\text{Net government cost}}.$$

We ask how the MVPF of SNAP without product restrictions changes when restrictions on soda, candy, and energy drinks are adopted nationwide. Our main calculation considers this full restricted basket, though we also report an incremental MVPF based only on restricting soda.

expression for intuition; applying it to the three-category basket gives 71 percent. Additional details are in Appendix A.

Let B denote baseline SNAP benefit outlays, g_0 the baseline net government cost per dollar of benefits, and m_0 the baseline MVPF of SNAP without the restriction. Then baseline net government costs are $g_0 B$ and baseline willingness to pay is $m_0 g_0 B$. The MVPF after adding the restriction is

$$\text{MVPF}_{\text{SNAP+restriction}} = \frac{m_0 g_0 B + \Delta W}{g_0 B + \Delta G}.$$

Here, ΔW and ΔG are the changes in willingness to pay and net government costs caused by the restriction. We take m_0 and g_0 from published estimates, use administrative SNAP outlays for B, and estimate ΔW and ΔG using our results. Appendix J.3 provides a full derivation.

Willingness to pay for the restriction. Calculating willingness to pay requires a view about which demand curve is welfare relevant. We treat baseline purchases as distorted for three reasons. First, consumers may not fully internalize their own future health costs because of imperfect nutrition information and self-control problems (Allcott et al., 2019). This is the internality. Second, they may not internalize the health-care costs their purchases impose on others. This is the externality. For simplicity, our baseline assumes that these external health-care costs are ultimately borne collectively by SNAP participants through higher insurance premiums.[25] We refer to the gains from correcting these first two distortions as health-related corrective benefits.

We value the estimated reduction in each category using the corrective wedges from Allcott et al. (2019):

$$S^{\text{gross}} = \sum_k (\gamma_k + e_k) \Delta Q_k,$$

where γ_k and e_k denote the internality and health-care externality per unit, and ΔQ_k is our causal estimate of the reduction in category k. This calculation yields approximately \$1.01 billion in annual gross corrective benefits from a nationwide restriction. This estimate considers only changes in the restricted products; accounting for broader changes in beverage and packaged-snack purchases raises it to approximately \$1.20 billion. Appendix J provides the details.

The third distortion that the restriction may correct results from the provision of benefits in kind rather than in cash. This channel exists even if consumers fully internalize their own

[25] We therefore include the full reduction in external health-care spending in willingness to pay. If the government instead bears a share of these costs, we remove that share from willingness to pay and count it as an equal reduction in net government cost. This changes where the benefit enters the MVPF calculation, but not the total surplus generated by the restriction. Appendix J.3 provides the formal accounting.

future health costs ($\gamma = 0$) and impose no health-care costs on others ($e = 0$). Because the benefits are in-kind, the SNAP label or mental accounting may lead households to buy more food, including more soda, than they would under an equivalent cash transfer (Hastings and Shapiro, 2018). If choices under cash are the normative benchmark, removing eligibility reverses some of these excess purchases and generates a private gain: under the cash-demand curve, the eliminated units were worth less to households than they cost.

Figure 6 shows two reasons why the ban can generate welfare gains by correcting these distortions. First, the solid line shows demand under an equivalent cash transfer. Pre-ban demand lies above this line: while the products are SNAP-eligible, households purchase more than they would with cash, consistent with labeling or mental accounting. The ban moves purchases back toward the cash-equivalent choice, generating the yellow gain. Second, even the cash-equivalent choice may be socially excessive because consumers underweight future health costs and some health-care costs fall on others. The lower orange line accounts for these costs. The vertical gap between cash demand and social willingness to pay is the health-cost wedge, $\gamma + e$, and the green area is the corresponding corrective benefit.

Combining the calibrated price elasticities from Section 6.2.4 with our causal estimate of the purchase reduction allows us to quantify the yellow area – the welfare gain from correcting the in-kind-transfer distortion, separate from the internality and externality benefits above. For the full restricted basket, this gain is \$59–155 million annually. Appendix J.2 presents the calculation under the models described in Section 6.2.3.

Sales tax revenue and administrative costs. Additional sales taxes paid by SNAP participants enter both parts of the MVPF calculation equally: they reduce willingness to pay and, as government revenue, reduce net government cost by the same amount. In our baseline, we set other unmeasured costs – including government administrative costs, household hassle costs, and retailer compliance costs – to zero.

Calculations of MVPF We use three sets of published estimates for the baseline MVPF of SNAP. Our preferred estimate is the MVPF of Food Stamps, the precursor to SNAP, estimated by Hendren and Sprung-Keyser (2020) to be $m_0 = 1.04$ (with $g_0 = 1.05$). We also examine how the waivers affect the baseline estimates of Barr and Smith (2023), who obtain an MVPF of 3.86–7.74 after accounting for long-run crime effects (with $g_0 = 0.77$), and Bailey et al. (2024), who obtain an MVPF of 62.25 for families whose children were exposed to Food Stamps between birth and age five (with $g_0 = 1$).

Table 14 reports the main welfare estimates. Under the main benchmark MVPF, the three-category restriction raises SNAP's MVPF from 1.04 to 1.052, an increase of 0.012, or

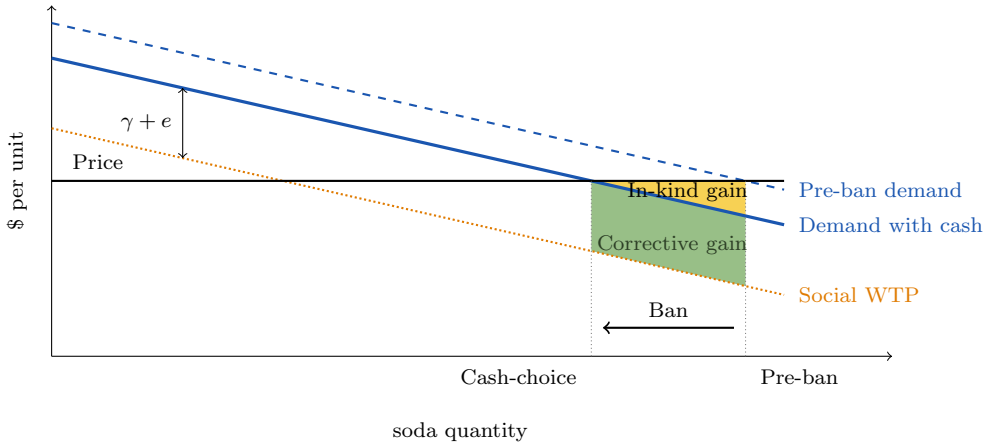


Figure 6: Two welfare gains from the ban. Moving from pre-ban purchases toward the cash-equivalent choice generates the yellow in-kind-transfer gain. The health-cost wedge, $\gamma + e$, between cash demand and social willingness to pay generates the green corrective benefit. The diagram is schematic; Table 32 reports the corresponding estimates.

1.11%. Across the alternative MVPF estimates, the full restrictions increase the MVPF by 0.21–1.11%.

Table 14: Effect of the purchase restrictions on published baseline SNAP MVPFs

Baseline scenario	g_0	Baseline m_0	Post-restriction MVPF	Gain (% of baseline)
Panel A: Soda only				
Hendren–Sprung-Keyser, full population	1.05	1.04	1.05	0.74
Barr–Smith, crime valuation (lower)	0.77	3.86	3.87	0.38
Barr–Smith, crime valuation (upper)	0.77	7.74	7.76	0.26
Bailey et al., exposure at ages 0–5	1.00	62.25	62.32	0.12
Panel B: Restricted basket				
Hendren–Sprung-Keyser, full population	1.05	1.04	1.05	1.11
Barr–Smith, crime valuation (lower)	0.77	3.86	3.88	0.59
Barr–Smith, crime valuation (upper)	0.77	7.74	7.77	0.42
Bailey et al., exposure at ages 0–5	1.00	62.25	62.38	0.21

Notes: m_0 is the published baseline MVPF and g_0 is the associated net government cost per dollar of benefits. The final two columns apply equation (11) using $B = \$92.72$ billion in FY2024 SNAP benefits and the benchmark restriction inputs in Table 34. To collapse the two behavioral models to one private-gain input, we use the midpoint of the lowest and highest estimates across the models and reported sensitivity values: \$54.20 million for soda and \$107.20 million for the restricted basket. The restricted basket contains soda, candy, and energy drinks.

8 Conclusion

This paper provides some of the first evidence from statewide SNAP restrictions on purchases of soda, candy, energy drinks, and sports drinks. Spending on restricted products falls by 11.0%, reflecting lower total purchases rather than a shift across payment methods. Across the combined beverage and packaged-snack basket, calories fall by 5.4% and total sugar by 7.7%, with little evidence of offsetting substitution. Retailers barely change prices, and the associated sales-tax increase is small, indicating that purchases respond primarily through a nonprice channel. The patterns across payment methods, retailers, and shopping channels are consistent with SNAP benefits and cash not being fully fungible. Nationwide, the restriction would impose a substantially smaller direct financial burden on SNAP participants than equivalent excise taxes. A back-of-the-envelope behavioral-welfare calculation incorporating nutrition-related corrective benefits yields annual welfare gains of \$1.07–\$1.16 billion.

An important question for future research is how restrictions on government income support should be combined with conventional Pigouvian taxes. A restriction tied to an income-support program directly affects only program participants, whereas a Pigouvian tax reaches the broader population. When restrictions leave retail prices largely unchanged, however, they may reduce purchases among low-income households while imposing much less direct financial burden than a tax. This creates a tradeoff between reach and incidence. Future research could examine whether the two policies are substitutes or complements and how the scope of restrictions and the level of taxation should be chosen jointly to reduce harmful consumption while limiting the burden on low-income households.

Another future research question is how the rapid diffusion of GLP-1 drugs changes the welfare value of nutrition-oriented purchase restrictions and how their design should adapt in response. Our calculations use corrective-benefit estimates that predate widespread GLP-1 adoption. As use expands, these drugs may change both the health returns to reducing unhealthy consumption and the extent to which mental accounting translates into food overconsumption.

References

Adams, Brian and Kevin R. Williams, "Zone Pricing in Retail Oligopoly," *American Economic Journal: Microeconomics*, 2019, *11* (1), 124–156.

Albertsons Companies, Inc., "Annual Report on Form 10-K for the Fiscal Year Ended February 22, 2025," U.S. Securities and Exchange Commission filing 2025. Filed April 15, 2025.

Allcott, Hunt, Amy Finkelstein, Anna Grummon, and Matthew J. Notowidigdo, "The Effects of SNAP Sugary Drink Restrictions on Consumption and Welfare," Working Paper 35659, National Bureau of Economic Research August 2026.

_ , Benjamin B. Lockwood, and Dmitry Taubinsky, "Regressive Sin Taxes, with an Application to the Optimal Soda Tax," *The Quarterly Journal of Economics*, 2019, *134* (3), 1557–1626.

Andreyeva, Tatiana, Michael W. Long, and Kelly D. Brownell, "The Impact of Food Prices on Consumption: A Systematic Review of Research on the Price Elasticity of Demand for Food," *American Journal of Public Health*, 2010, *100* (2), 216–222.

Bailey, Martha J., Hilary Hoynes, Maya Rossin-Slater, and Reed Walker, "Is the Social Safety Net a Long-Term Investment? Large-Scale Evidence from the Food Stamps Program," *The Review of Economic Studies*, 2024, *91* (3), 1291–1330.

Barnhill, Anne, "Impact and Ethics of Excluding Sweetened Beverages from the SNAP Program," *American Journal of Public Health*, 2011, *101* (11), 2037–2043.

Barr, Andrew and Alexander A. Smith, "Fighting Crime in the Cradle: The Effects of Early Childhood Access to Nutritional Assistance," *Journal of Human Resources*, 2023, *58* (1), 43–73.

Basu, Sanjay, Hilary K. Seligman, Christopher Gardner, and Jay Bhattacharya, "Ending SNAP Subsidies for Sugar-Sweetened Beverages Could Reduce Obesity and Type 2 Diabetes," *Health Affairs*, 2014, *33* (6), 1032–1039.

Beatty, Timothy K. M., Laura Blow, Thomas F. Crossley, and Cormac O'Dea, "Cash by Any Other Name? Evidence on Labeling from the UK Winter Fuel Payment," *Journal of Public Economics*, 2014, *118*, 86–96.

Bergman, Peter, Raj Chetty, Stefanie DeLuca, Nathaniel Hendren, Lawrence F. Katz, and Christopher Palmer, "Creating Moves to Opportunity: Experimental Evidence on Barriers to Neighborhood Choice," *American Economic Review*, 2024, *114* (5), 1281–1337.

Caoui, El Hadi, Brett Hollenbeck, Matthew Osborne, and Elina Page, "The Impact of Dollar Store Expansion on Local Market Structure and Food Access," *Quantitative Marketing and Economics*, 2026, *24* (1), 7.

Castner, Laura, Leah Shiferaw, Andrew Wen, Clare Wolfendale, and Elisabeth Appel, "Benefit Redemption Patterns in the Supplemental Nutrition Assistance Program in Fiscal Year 2022," Research Report, U.S. Department of Agriculture, Food and Nutrition Service, Alexandria, VA 2025.

Cawley, John and David Frisvold, "The Pass-Through of Taxes on Sugar-Sweetened Beverages to Retail Prices: The Case of Berkeley, California," *Journal of Policy Analysis and Management*, 2017, *36* (2), 303–326.

_ , Chelsea Crain, David Frisvold, and David Jones, "The Pass-Through of the Largest Tax on Sugar-Sweetened Beverages: The Case of Boulder, Colorado," Working Paper 25050, National Bureau of Economic Research September 2018.

_ , David Frisvold, Anna Hill, and David Jones, "Oakland's Sugar-Sweetened Beverage Tax: Impacts on Prices, Purchases and Consumption by Adults and Children," *Economics & Human Biology*, 2020, *37*, 100865.

Cellini, Stephanie R., Rajeev Darolia, and Lesley J. Turner, "Where Do Students Go When For-Profit Colleges Lose Federal Aid?," *American Economic Journal: Economic Policy*, 2020, *12* (2), 46–83.

Cengiz, Doruk, Arindrajit Dube, Attila Lindner, and Ben Zipperer, "The Effect of Minimum Wages on Low-Wage Jobs," *The Quarterly Journal of Economics*, 2019, *134* (3), 1405–1454.

Chen, Jiafeng and Jonathan Roth, "Logs with Zeros? Some Problems and Solutions," *The Quarterly Journal of Economics*, 2024, *139* (2), 891–936.

Chetty, Raj, Nathaniel Hendren, and Lawrence F. Katz, "The Effects of Exposure to Better Neighborhoods on Children: New Evidence from the Moving to Opportunity Experiment," *American Economic Review*, 2016, *106* (4), 855–902.

Chorniy, Anna, Amy Finkelstein, and Matthew J. Notowidigdo, "Paternalistic Social Assistance: Evidence and Implications from Cash vs. In-Kind Transfers," Working Paper 34506, National Bureau of Economic Research November 2025.

Chrisinger, Benjamin W., "Characteristics of State Waivers to Establish Nutritional Restrictions in the Supplemental Nutrition Assistance Program," *Health Affairs Scholar*, 2025, *3* (12), qxaf221.

Cox, Christian and Katherine Harris-Lagoudakis, "SNAP Eligible Products and Behavioral Demand," Working Paper 23-04 2023. Later circulated as "In-Kind Benefits and Behavioral Demand".

Currie, Janet and Firouz Gahvari, "Transfers in Cash and In-Kind: Theory Meets the Data," *Journal of Economic Literature*, 2008, *46* (2), 333–383.

Douenne, Thomas, "The Vertical and Horizontal Distributive Effects of Energy Taxes: A Case Study of a French Policy," *The Energy Journal*, 2020, *41* (3), 231–254.

French, Simone A., Sarah A. Rydell, Nathan R. Mitchell, J. Michael Oakes, Brian Elbel, and Lisa Harnack, "Financial Incentives and Purchase Restrictions in a Food Benefit Program Affect the Types of Foods and Beverages Purchased: Results from a Randomized Trial," *International Journal of Behavioral Nutrition and Physical Activity*, 2017, *14*, 127.

Frisvold, David E. and Felipe Lozano Rojas, "Short-Run Impacts of SNAP Restriction Waivers on Beverage Purchases," Working Paper 35613, National Bureau of Economic Research August 2026.

Garasky, Steven, Kassim Mbwana, Andres Romualdo, Alex Tenaglio, and Manan Roy, "Foods Typically Purchased by Supplemental Nutrition Assistance Program (SNAP) Households," Research Report, U.S. Department of Agriculture, Food and Nutrition Service, Alexandria, VA 2016.

Harnack, Lisa, J. Michael Oakes, Brian Elbel, Timothy Beatty, Sarah Rydell, and Simone French, "Effects of Subsidies and Prohibitions on Nutrition in a Food Benefit Program: A Randomized Clinical Trial," *JAMA Internal Medicine*, 2016, *176* (11), 1610–1618.

Hastings, Justine and Jesse M. Shapiro, "How Are SNAP Benefits Spent? Evidence from a Retail Panel," *American Economic Review*, 2018, *108* (12), 3493–3540.

Hendren, Nathaniel and Ben Sprung-Keyser, "A Unified Welfare Analysis of Government Policies," *The Quarterly Journal of Economics*, 2020, *135* (3), 1209–1318.

Hoynes, Hilary W. and Diane Whitmore Schanzenbach, "Consumption Responses to In-Kind Transfers: Evidence from the Introduction of the Food Stamp Program," *American Economic Journal: Applied Economics*, 2009, *1* (4), 109–139.

Mandal, Abir, "State and Local Sales Tax Rates, 2026," Tax Foundation January 2026.

Meyer, Bruce D., Nikolas Mittag, and Robert M. Goerge, "Errors in Survey Reporting and Imputation and Their Effects on Estimates of Food Stamp Program Participation," *Journal of Human Resources*, 2022, *57* (5), 1605–1644.

Moffitt, Robert, "Estimating the Value of an In-Kind Transfer: The Case of Food Stamps," *Econometrica*, 1989, *57* (2), 385–409.

O'Donoghue, Julie, "Q&A: SNAP Benefit Restrictions with LDH Secretary Bruce Greenstein," Louisiana Illuminator March 2026.

Powell, Lisa M., Jamie F. Chriqui, Tamkeen Khan, Roy Wada, and Frank J. Chaloupka, "Assessing the Potential Effectiveness of Food and Beverage Taxes and Subsidies for Improving Public Health: A Systematic Review of Prices, Demand and Body Weight Outcomes," *Obesity Reviews*, 2013, *14* (2), 110–128.

Rabbitt, Matthew P., Madeline Reed-Jones, Laura J. Hales, Shellye Suttles, and Michael P. Burke, "Household Food Security in the United States in 2024," Economic Research Report ERR-358, U.S. Department of Agriculture, Economic Research Service December 2025.

Sahni, Navdeep S., Dan Zou, and Pradeep K. Chintagunta, "Do Targeted Discount Offers Serve as Advertising? Evidence from 70 Field Experiments," *Management Science*, 2017, *63* (8), 2688–2705.

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

Song, Jintaek, "Propensity to Consume Food Out of SNAP and Its Welfare Implications," Working Paper 2026.

Southworth, Herman M., "The Economics of Public Measures to Subsidize Food Consumption," *Journal of Farm Economics*, 1945, *27* (1), 38–66.

Tax Foundation, "Facts & Figures 2023: How Does Your State Compare?," Technical Report, Tax Foundation, Washington, DC 2023. Table: "Sales Tax Treatment of Groceries, Candy, & Soda".

Texas Legislature, "S.B. 379: Purchase of Certain Food and Drink Items Prohibited Under the Supplemental Nutrition Assistance Program," 89th Texas Legislature, Regular Session, Senate committee report version 2025.

_ , "S.B. 379: Purchase of Sweetened Drinks and Candy Prohibited Under the Supplemental Nutrition Assistance Program," 89th Texas Legislature, Regular Session 2025.

The Kroger Co., "Annual Report on Form 10-K for the Fiscal Year Ended February 1, 2025," U.S. Securities and Exchange Commission filing 2025. Filed March 25, 2025.

U.S. Bureau of Labor Statistics, "Historical Consumer Price Index for All Urban Consumers (CPI-U): U.S. City Average, All Items," Consumer Price Index supplemental table December 2024.

U.S. Department of Agriculture, Food and Nutrition Administration, "SNAP Food Restriction Waivers," Food and Nutrition Administration website 2026.

U.S. Department of Agriculture, Food and Nutrition Service, "Implications of Restricting the Use of Food Stamp Benefits," Technical Report, U.S. Department of Agriculture, Food and Nutrition Service, Washington, DC March 2007.

_ , "No Sales Taxes on SNAP Purchases," 7 C.F.R. § 272.1(b) 2026.

U.S. Department of Education, "Accountability in Higher Education and Access Through Demand-Driven Workforce Pell: Pell Grant Exclusion Relating to Other Grant Aid," *Federal Register*, May 2026, *91* (96), 29254–29338.

Ward, Ben and Mia Monkovic, "Characteristics of Supplemental Nutrition Assistance Program Households: Fiscal Year 2024," Nutrition Assistance Program Report SNAP-24-CHAR, U.S. Department of Agriculture, Food and Nutrition Service May 2026.

A Sales tax incidence and financial burden

The restriction can change participants' tax-inclusive prices even when retailers do not adjust shelf prices. We first quantify this mechanical price effect and the resulting sales tax payments, and then compare those payments with excise taxes calibrated to produce the same purchase reductions.

A.1 Effective consumer prices

Purchases paid with SNAP are exempt from sales tax. The loss of this exemption changes the price only for purchases that would otherwise have been paid with SNAP. Thus, holding purchases fixed, the average effective price increase for category k in state s is

$$\Delta p_{sk}^{\text{eff}} = \varphi_k \Delta\tau_{sk},$$

where φ_k is the pre-policy SNAP-paid share of category spending and $\Delta\tau_{sk}$ is the difference between the tax rate on the restricted product and the rate on ordinary groceries. We estimate φ_k using pre-policy in-store purchases by households in the main regression sample during months in which they satisfy our robust SNAP-participation measure, restricting to jurisdictions with a positive category-specific tax differential. The resulting payment shares are

$$\varphi_{\text{soda}} = 0.4058, \qquad \varphi_{\text{candy}} = 0.3544, \qquad \varphi_{\text{energy}} = 0.3555, \qquad \varphi_{\text{sports}} = 0.4204.$$

Across the 30 treated state–category combinations, the average increase is 1.88 percent of the pre-tax price, or 2.36 percent among combinations with a positive tax differential. If every pre-policy purchase had been paid with SNAP, the corresponding increases would be 4.94 and 6.22 percent.

These are weighted averages. Writing T for the set of 30 treated state–category combinations,

$$\overline{\Delta p^{\text{eff}}} = \frac{\sum_{(s,k)\in T} B_s \rho_k \, \varphi_k \Delta\tau_{sk}}{\sum_{(s,k)\in T} B_s \rho_k},$$

where B_s is annual SNAP benefits in state s and ρ_k is pre-policy spending on category k per dollar of benefits, so that B_s weights across states and ρ_k across categories. Setting $\varphi_k = 1$ gives the full-use figures, and restricting T to combinations with $\Delta\tau_{sk} > 0$ gives the positive-differential ones. Sixteen of the thirty combinations carry a positive differential, but they hold 79 percent of the weight, since Texas and Florida make up about two-thirds of it

and tax every category they restrict.

Table 15 reports the second input to the effective-price calculation, the statewide category-specific tax differential $\Delta\tau_{sk}$. A zero denotes equal tax treatment of the restricted product and ordinary groceries, rather than necessarily indicating that both are tax exempt. The tax-base classification begins with the Tax Foundation's 2023 grocery, candy, and soda table and is updated using state statutes. The rate grid uses 2026 statewide rates (Tax Foundation, 2023; Mandal, 2026); we do not add county or city sales taxes.

Table 15: Statewide restricted-good sales-tax differentials used in the calculation

State	Δ soda	Δ candy	Δ energy	Δ sports	State	Δ soda	Δ candy	Δ energy	Δ sports
AK	0	0	0	0	MT	0	0	0	0
AL	0	0	0	0	NC	4.75	4.75	4.75	4.75
AR	6.5	6.5	6.5	6.5	ND	5	5	5	5
AZ	0	0	0	0	NE	0	0	0	0
CA	7.25	0	7.25	7.25	NH	0	0	0	0
CO	2.9	2.9	2.9	2.9	NJ	6.625	6.625	6.625	6.625
CT	6.35	6.35	6.35	6.35	NM	0	0	0	0
DC	8	6	8	8	NV	0	0	0	0
DE	0	0	0	0	NY	4	4	4	4
FL	6	6	6	6	OH	5.75	0	5.75	5.75
GA	0	0	0	0	OK	0	0	0	0
HI	0	0	0	0	OR	0	0	0	0
IA	6	6	6	6	PA	6	0	6	6
ID	0	0	0	0	RI	7	7	7	7
IL	6.25	6.25	6.25	6.25	SC	0	0	0	0
IN	7	7	7	7	SD	0	0	0	0
KS	0	0	0	0	TN	0	3	0	0
KY	6	6	6	6	TX	6.25	6.25	6.25	6.25
LA	0	0	0	0	UT	0	0	0	0
MA	0	0	0	0	VA	0	0	0	0
MD	6	6	6	6	VT	6	0	6	6
ME	5.5	5.5	5.5	5.5	WA	6.5	0	6.5	6.5
MI	0	0	0	0	WI	5	5	5	5
MN	6.875	6.875	6.875	6.875	WV	6	0	6	6
MO	0	0	0	0	WY	0	0	0	0
MS	0	0	0	0					

Notes: Entries are percentage points, with $\Delta k = \tau_{sk} - \tau_{gs}$ equal to the statewide rate on category k minus the statewide rate on ordinary groceries. A zero denotes equal tax treatment, not necessarily tax exemption. The rates are statewide; county and city sales taxes are not added. The energy-drink column includes products with Nutrition Facts labels in the soft-drink category. The sports-drink column uses the same statewide rate convention as energy and is shown for completeness; it does not enter the three-category incidence calculation reported in this section. Sources are described in the text.

A.2 Sales tax payments under nationwide adoption

Let $B = \$92.7205$ billion denote total SNAP benefits paid in FY2024. The approximate \$340 monthly benefit is used only to translate national amounts into household equivalents; the FY2024 SNAP Quality Control mean is \$341 (Ward and Monkovic, 2026).

Let B_s be SNAP benefits in state s. For category k, let φ_k be the pre-policy SNAP-paid share of category spending, ρ_k be all-tender category spending relative to SNAP-funded grocery spending, and c_k be the all-channel proportional spending reduction. Let τ_{sk} be the restricted-good tax rate, τ_{gs} the ordinary-grocery rate, $\Delta\tau_{sk} = \tau_{sk} - \tau_{gs}$, and τ_s^r the tax rate on any purchase that replaces the spending reduction. The change in sales taxes is

$$\frac{T_{sk}(\tau_s^r)}{B_s} = \rho_k \left[\varphi_k \Delta\tau_{sk} - c_k(\tau_{sk} - \tau_s^r)\right] = \rho_k \left[\varphi_k \Delta\tau_{sk} - c_k\tau_{sk} + c_k\tau_s^r\right].$$

Given that our detail price data do not show a price response, we focus on sales tax and excise tax implications under the assumption of full pass through. The first term applies the restricted-good–grocery tax differential to the pre-policy SNAP-paid share. The next term subtracts the tax on restricted purchases that no longer occur. The last term adds the tax on whatever replaces that spending. For the comparison with the equivalent excise tax, we set $\tau_s^r = 0$, so that both calculations exclude downstream tax consequences of the spending released when purchases fall. This is an accounting convention for making the two policies comparable, not an assumption that households literally discard the money, and it gives equation (6). We also report the cases in which the released spending buys ordinary groceries ($\tau_s^r = \tau_{gs}$) or taxable nonfood, for which we use $\tau_s^r = \tau_{sk}$ as a proxy for the general sales rate.

We measure φ_k and ρ_k on pre-policy in-store purchases, where the payment method is observed reliably, and use these measures as proxies for their all-channel counterparts. We use the paper's all-channel response for c_k. The category inputs are

$$\begin{aligned}(\varphi_s, \rho_s, c_s) &= (0.4058, 0.0938, 0.1204),\\ (\varphi_c, \rho_c, c_c) &= (0.3544, 0.0896, 0.0885),\\ (\varphi_e, \rho_e, c_e) &= (0.3555, 0.0252, 0.1341).\end{aligned}$$

Table 16: National sales tax transfer under alternative uses of reduced spending

Category	Positive-differential jurisdictions	SNAP-benefit share (percent)	No purchase or untaxed use	Ordinary groceries	Taxable other goods
			Net transfer ($ millions per year)		
Soda	25	67.8	100.7	100.7	143.3
Candy	20	45.5	54.7	55.2	73.6
Soda and candy	26	69.5	155.4	156.0	216.8
Energy drinks	25	67.8	21.0	21.0	33.7
Three-category total	26	69.5	176.4	176.9	250.5

Notes: The no-purchase-or-untaxed-use column is the measure used for the comparison with an equivalent excise tax and applies equation (6). It excludes downstream taxes on the use of the released spending under both policies. The ordinary-grocery and taxable-other-goods columns show alternative uses of that spending; the final column uses the restricted-good rate as a proxy for the general sales rate on taxable nonfood. Across the three categories, 25 states and the District of Columbia have a positive differential. The energy-drink row includes products with Nutrition Facts labels in the soft-drink category; the result therefore depends on that classification. Equal-treatment cells in the rate grid are set to zero, so the table measures the policy-induced differential rather than total sales tax receipts in equal-treatment jurisdictions.

A.3 Comparison with equivalent excise taxes

We calculate each state–category cell using its statewide tax rates and then sum the cells; we do not use a national average tax rate. For the policy comparison, restriction payments are $\rho_k[\varphi_k \Delta\tau_{sk} - c_k\tau_{sk}]$, while excise-tax payments are $\rho_k\delta_k(1-c_k)$. In both calculations we leave the subsequent use of the released spending outside the comparison.[26] The sensitivity table above shows how the restriction estimate changes when that spending instead buys ordinary groceries or taxable nonfood.

Allcott et al. (2019) estimate an uncompensated SSB price elasticity of 1.373. Under isoelastic demand, the exact price increase associated with a proportional quantity decline c_k satisfies

$$1 - c_k = (1+\delta_k)^{-|\varepsilon_k|}, \qquad \delta_k = (1-c_k)^{-1/|\varepsilon_k|} - 1.$$

The paper uses the local approximation $\delta_k \simeq c_k/|\varepsilon_k|$, consistent with treating the estimated spending response as a local price-equivalent response. For soda, this gives $\delta_s = 0.120430/1.373 = 0.08771$. Consumer excise-tax payments relative to SNAP benefits are

$$\frac{R_s^{\text{excise}}}{B} \simeq \rho_s\delta_s(1-c_s) = 0.093770(0.08771)(1-0.120430) = 0.007234.$$

This is the excise-tax payment per dollar of SNAP benefits. Multiplying it by the $92.7205 billion in FY2024 SNAP benefits gives $670.8 million in annual excise-tax payments for soda.

[26] Under the excise tax, ordinary sales-tax payments on the targeted goods also change. We omit this second-order term because the price increase offsets most of the quantity decline; for soda, it is about $7 million nationwide.

Table 16 shows that the restriction instead generates $100.7 million in additional soda sales-tax payments. The restriction therefore imposes about 15 percent as much financial burden for soda.

We repeat the calculation for candy and energy drinks. The equivalent excise tax rates are 6.445 and 9.764 percent and generate $488.2 million and $197.3 million in payments, respectively. Including soda, the equivalent excise-tax payments for the full restricted basket total $1.356 billion. The corresponding restriction payments in Table 16 are $54.7 million for candy and $21.0 million for energy drinks; including soda, they total $176.4 million. Thus, for the full basket, the restriction imposes 13.0 percent as much financial burden as the equivalent excise taxes. The SSB elasticity applies directly to soda; its application to candy and energy drinks is an extrapolation. Dividing FY2024 SNAP benefits by $340 per month times 12 months gives 22.73 million household-equivalent years. The two national totals are therefore $7.76 and $59.67 per household-equivalent year, which give the rounded $8 and $60 figures in the text.

Equal per-household revenue recycling. Let p be the share of households receiving SNAP. For each category $k \in \mathcal{K}$, let $P_{B,k}$ and $P_{E,k}$ be SNAP households' payments under the restriction and the equivalent excise tax, respectively, and let r_k be average pre-policy spending per SNAP household relative to a non-SNAP household. The restriction raises revenue only from SNAP households. Equal per-household recycling therefore leaves them with net payments

$$N_B = (1-p)\sum_{k\in\mathcal{K}} P_{B,k}.$$

Under full pass-through, non-SNAP households pay $P_{E,k}(1-p)/(r_k p)$ under the excise tax in category k. Because SNAP households receive the fraction p of the combined revenue as rebates, their net category-k payment is

$$P_{E,k} - p\left[P_{E,k} + \frac{P_{E,k}(1-p)}{r_k p}\right] = (1-p)P_{E,k}\frac{r_k - 1}{r_k}.$$

Summing across categories and cancelling $1-p$ gives the exact comparison

$$\frac{N_B}{N_E} = \frac{\sum_{k\in\mathcal{K}} P_{B,k}}{\sum_{k\in\mathcal{K}} P_{E,k}(r_k-1)/r_k}.$$

Let $P_B = \sum_k P_{B,k}$ and $P_E = \sum_k P_{E,k}$. If $r_k = r$ for every category (or the restriction covers a single category), this expression reduces to

$$\frac{N_B}{N_E} = \frac{P_B}{P_E}\frac{r}{r-1}.$$

Thus, under the external benchmark $r_k = 2$ for every category, the equal-rebate ratio is twice the no-rebate ratio: 0.30 for soda and 0.26 for the three-category basket.

For soda alone, Table 1 gives $r_{\text{soda}} = 28.1089/19.3918 = 1.4495$, so the equal-rebate ratio is

$$\frac{100.7}{670.8(1 - 1/1.4495)} = 0.48.$$

Across soda, candy, and energy drinks, the aggregate ratio is $r = 60.95/49.80 = 1.22$. Combining it with the full-basket no-rebate ratio gives the 0.71 approximation reported in the text. The category-specific ratios for candy and energy drinks are $25.2955/24.0346 = 1.0525$ and $7.5459/6.3704 = 1.1845$. The exact calculation is therefore

$$\frac{176.4}{670.8(1 - 1/1.4495) + 488.2(1 - 1/1.0525) + 197.3(1 - 1/1.1845)} = 0.67.$$

These calculations use unrounded inputs and assume that each excise tax is fully passed through, so the consumer price increase and tax revenue coincide.

B Robustness checks

We first address the possibility that the estimates depend on comparisons between the implementing states and a politically or economically different set of controls. We re-estimate the soda specification using only Alabama, Arizona, Georgia, Kentucky, Mississippi, North Carolina, and South Dakota as controls. These are the untreated conservative states used in our matched-control exercise. We exclude Montana from this set because its SNAP food-restriction waiver was approved for implementation in September 2026. Column (2) of Table 17 shows that the soda coefficient remains negative, statistically significant, and similar in magnitude to the full-sample estimate in column (1).

We next address the rapid growth of GLP-1 use, which could independently reduce demand for soda and could differ across states during the implementation period. We construct a state-by-quarter measure from the public CMS State Drug Utilization Data. The numerator is the sum of reported fee-for-service and managed-care GLP-1 prescriptions; the denominator is the sum of monthly total Medicaid enrollment over the same quarter. We express the resulting measure as prescriptions per 1,000 Medicaid enrollee-months and apply the inverse hyperbolic sine transformation, which retains zeros while reducing the influence of states with very high prescription rates. The broad measure includes all identified GLP-1 products, while the obesity-label measure includes Wegovy, Zepbound, and Saxenda. The prescription data measure claims rather than unique patients and do not report diagnoses, so they cannot identify off-label use. CMS also suppresses small cells; we do not recode those cells as zero, and the sum of reported prescriptions is therefore a lower bound in affected state-quarters.[27]

CMS data for 2026Q2 were not yet available, so this check ends in March 2026 and omits the April implementation cohort, which would have no post-policy observation in that window. Columns (3)–(5) use exactly the same January and February cohort sample through March 2026. Column (3) estimates the specification without a GLP-1 control on this shorter sample; columns (4)–(5) add the broad and obesity-label GLP-1 rates. The treatment coefficient changes by only 0.015 and 0.047, respectively, and neither GLP-1 control is statistically significant.

[27] The public sources are the CMS State Drug Utilization Data and the monthly State Medicaid and CHIP Applications, Eligibility Determinations, and Enrollment Data.

Table 17: Robustness to control-state composition and Medicaid GLP-1 utilization

	(1) Full sample	(2) Conservative controls	(3) Sample through March 2026 Baseline	(4) Broad GLP-1	(5) Obesity-label GLP-1
Treated × Post	−3.385***	−3.191***	−5.143***	−5.128***	−5.096***
	(0.991)	(0.836)	(0.680)	(0.679)	(0.679)
asinh(broad GLP-1 rate)	—	—	—	−0.0833	—
				(0.356)	
asinh(obesity-label GLP-1 rate)	—	—	—	—	−0.0716
					(0.133)
Observations	398,081	85,813	202,558	202,558	202,558
Treatment cohorts	Jan./Feb./Apr.	Jan./Feb./Apr.	Jan./Feb.	Jan./Feb.	Jan./Feb.

Notes: The dependent variable is soda spending in dollars per household-month, with non-purchasing months entered as zeros. All columns use the stacked difference-in-differences specification with household-by-cohort and event-time-by-cohort fixed effects; standard errors clustered by state are in parentheses. Column (1) is the headline sample. Column (2) limits controls to Alabama, Arizona, Georgia, Kentucky, Mississippi, North Carolina, and South Dakota. Columns (3)–(5) use an identical sample ending in March 2026; the April cohort is omitted because it has no post-policy month in that window. GLP-1 rates are reported prescriptions per 1,000 Medicaid enrollee-months in the contemporaneous quarter. $^{*}p < 0.10$, $^{**}p < 0.05$, $^{***}p < 0.01$.

Table 18: PPML estimates of SNAP restrictions on SNAP households' spending

	(1) Soda	(2) Candy	(3) Energy drinks	(4) Sports drinks	(5) Restricted basket	(6) Other groceries
Panel A. All households						
Treated × Post	−0.115***	−0.117***	−0.139***	−0.216***	−0.115***	−0.006
	(0.029)	(0.026)	(0.035)	(0.033)	(0.021)	(0.008)
Effect (%)	−10.8	−11.1	−13.0	−19.5	−10.9	−0.6
95% confidence interval	[−15.7, −5.7]	[−15.4, −6.5]	[−18.7, −6.9]	[−24.5, −14.1]	[−14.5, −7.1]	[−2.2, 1.0]
Observations	383,832	392,628	267,516	270,622	888,130	896,955
Separated observations (%)	3.6	0.7	32.5	29.2	1.0	0.0
Panel B. Households without children						
Treated × Post	−0.074***	−0.089***	−0.143***	−0.105***	−0.084***	−0.013
	(0.017)	(0.024)	(0.029)	(0.029)	(0.013)	(0.011)
Effect (%)	−7.2	−8.5	−13.4	−10.0	−8.1	−1.3
95% confidence interval	[−10.2, −4.1]	[−12.7, −4.1]	[−18.1, −8.3]	[−15.0, −4.6]	[−10.4, −5.7]	[−3.5, 0.8]
Observations	232,310	239,584	152,827	153,582	542,894	549,900
Separated observations (%)	4.5	1.0	36.9	34.3	1.3	0.0
Panel C. Households with children						
Treated × Post	−0.159***	−0.146***	−0.140**	−0.298***	−0.146***	0.002
	(0.051)	(0.030)	(0.058)	(0.041)	(0.038)	(0.015)
Effect (%)	−14.7	−13.5	−13.0	−25.8	−13.6	0.2
95% confidence interval	[−22.9, −5.7]	[−18.5, −8.3]	[−22.3, −2.6]	[−31.5, −19.6]	[−19.7, −7.0]	[−2.7, 3.1]
Observations	151,522	153,044	114,689	117,040	345,236	347,055
Separated observations (%)	2.1	0.3	25.4	21.1	0.5	0.0
Group × cohort FE	Yes	Yes	Yes	Yes	Yes	Yes
Time × cohort FE	Yes	Yes	Yes	Yes	Yes	Yes

Notes: This table reports PPML counterparts to the estimates in Table 2. Effects are calculated as $100[\exp(\widehat{\beta}) - 1]$; brackets report 95% confidence intervals. Standard errors are clustered by state. Non-purchasing household-months enter as zeros. PPML retains zeros except when an absorbed fixed-effect group is separated; `ppmlhdfe` drops those observations because they provide no finite identifying variation. The observation rows report the resulting PPML estimation samples, and the separation rows report excluded observations as a percentage of the corresponding input sample. Category-specific treated and control states, fixed effects, and child-status panels otherwise match Table 2. $^{*}p < 0.10$, $^{**}p < 0.05$, $^{***}p < 0.01$.

Table 19: Candy-ban snack spending by channel

Channel	Product	Pre-policy spending ($)	Change ($)	Effect (%)
In store	Candy *(restricted)*	22.28	-1.875^{***}	−8.4
			(0.545)	
	Savory snacks *(eligible)*	19.99	-0.752^{***}	−3.8
			(0.234)	
	Frozen desserts *(eligible)*	7.63	-0.856^{***}	−11.2
			(0.137)	
	Other non-candy snacks *(mixed)*	29.19	-1.312^{***}	−4.5
			(0.378)	
Online	Candy *(restricted)*	2.41	-0.311^{**}	−12.9
			(0.119)	
	Savory snacks *(eligible)*	2.54	$+0.176^{**}$	+6.9
			(0.066)	
	Frozen desserts *(eligible)*	0.85	$+0.136^{**}$	+16.0
			(0.061)	
	Other non-candy snacks *(mixed)*	4.16	$+0.502^{***}$	+12.1
			(0.175)	

Notes: This table splits the SNAP-household spending estimates in Table 7, Panel B, by purchase channel. The specification and candy-ban sample otherwise match that table. Savory snacks include non-candy chips, crackers, pretzels, cheese and puffed snacks, meat snacks, and plain popped popcorn. Frozen desserts combine packaged ice cream and frozen novelties, which remain eligible in all states. Other non-candy snacks are the residual packaged-snack products after removing savory snacks and frozen desserts. This is a mixed category rather than an eligible-product category because it includes cookies and snack cakes restricted in Florida. Across the four product rows, the in-store and online dollar coefficients sum to the total packaged-snack estimate in Table 7, Panel B, up to rounding. Effects are percentages of the treated pre-policy mean in each channel. Standard errors, clustered by state, are in parentheses. $^{*}p < 0.10$, $^{**}p < 0.05$, $^{***}p < 0.01$.

Table 20: Ready-to-drink tea spending by policy group

	Tea products remain eligible NE, UT, WV, LA, FL	Tea products restricted IN, IA, ID, OK
Treated × Post	0.022	−0.257
	(0.114)	(0.180)
Effect (%)	0.8	−9.0
Pre-policy mean	2.770	2.855
Observations	381,120	272,882

Notes: The dependent variable is all-retailer ready-to-drink tea spending in dollars per household-month, with non-purchasing months entered as zeros. The two columns use the same five-state and four-state policy groups as the fruit-beverage analysis. Both estimates use the stacked difference-in-differences specification with household-by-cohort and event-time-by-cohort fixed effects. Standard errors clustered by state are in parentheses.

Table 21: Soda spending responses by order channel

Channel	DiD estimate ($)	Standard error	Pre-policy mean ($)	Effect (%)	Observations
In store	-2.831^{***}	0.877	25.199	-11.2	398,081
Online	-0.555^{***}	0.173	2.910	-19.1	398,081

Notes: Each row reports the stacked difference-in-differences specification used for the main soda estimate, with standard errors clustered by state. Percentage effects divide the dollar estimate by the channel's treated pre-policy mean. The two channel outcomes add to total soda spending. Numerator's order-method variable indicates whether a purchase was made in store or online. It therefore differs from the "Online" column of Table 4, which denotes retailers classified as eCommerce: an online order from a brick-and-mortar retailer enters the online row here but not the eCommerce retailer category there. $^{***}p < 0.01$.

C Heterogeneity by store type across categories

Section 5 reports that the soda reduction is concentrated in large stores and is essentially absent at the small stores. Table 22 extends the same store-type decomposition to candy, energy drinks, and the restricted basket, and Figure 7 plots the corresponding event studies. The point estimates generally show larger proportional reductions at large stores and eCommerce than at small stores, although the precision varies by category.

The pre-policy trip-level pattern behind the gradient is as follows. Soda-carrying trips averaged \$9.21 in soda spending at large stores against \$7.52 at small stores, and SNAP was used on 39.6% of in-store soda trips with observed tender at large stores versus 25.0% at small stores. SNAP-used trips carried more soda than other trips in both formats (\$10.09 versus \$8.51 at large stores, \$9.25 versus \$6.90 at small stores; Table 4). The gap in SNAP use is not specific to soda: 38.8% versus 20.9% for candy and 37.6% versus 22.0% for energy drinks, with SNAP-used trips again carrying more category spending, although the premium is negligible for energy drinks at small stores. This is a descriptive pre-policy pattern, not a causal decomposition of the store-format effects.

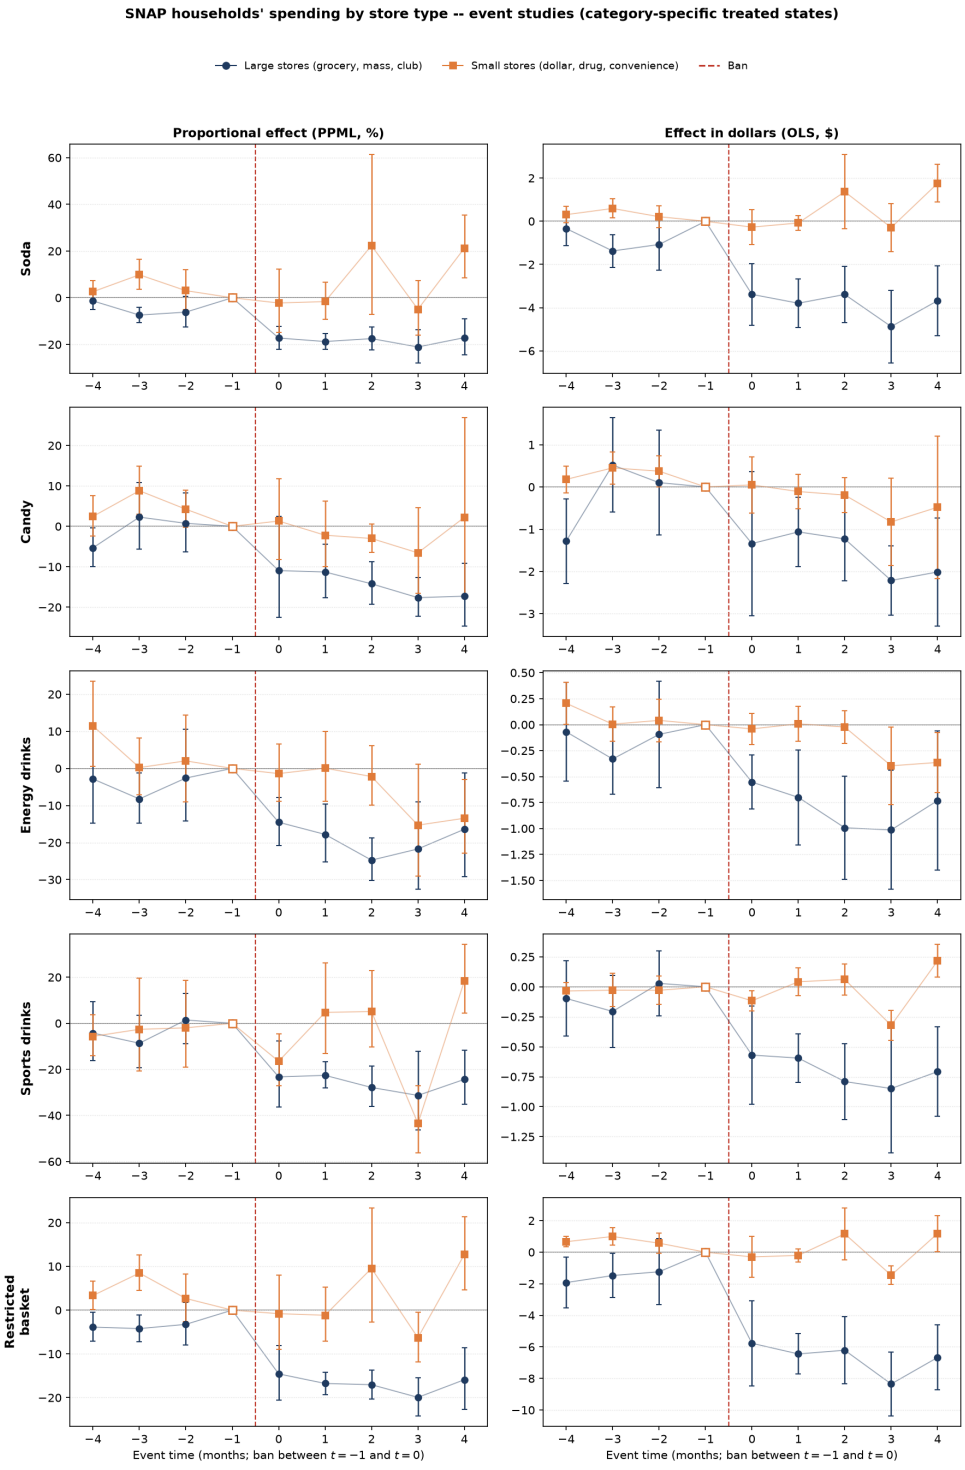


Figure 7: Event studies of SNAP households' spending at large and small stores for soda, candy, energy drinks, sports drinks, and the restricted basket. The left column reports proportional PPML effects and the right column OLS effects in dollars per household-month. Candy, energy, and sports use their category-specific ban-state samples; the restricted basket uses each state's own restricted scope. Event time $t = -1$ is omitted; bars are 95% confidence intervals with standard errors clustered by state, and the dashed line marks the restriction.

Table 22: Where the reduction lands, by store type – all restricted categories

Store type	Share	Reduction ($)	(SE)	% of mean
Soda (all treated states)				
Large (grocery, mass, club)	64.8%	−2.954	(0.545)	−16.2
Small (dollar, drug, conv.)	22.5%	+0.121	(0.417)	+1.9
eCommerce	10.0%	−0.626	(0.163)	−22.2
Candy (candy-ban states)				
Large	58.3%	−1.180	(0.481)	−8.2
Small	27.3%	−0.419	(0.218)	−6.2
eCommerce	9.6%	−0.327	(0.114)	−13.8
Energy drinks (energy-ban states)				
Large	53.4%	−0.643	(0.197)	−16.4
Small	26.5%	−0.137	(0.071)	−7.0
eCommerce	14.8%	−0.246	(0.091)	−22.6
Restricted basket (each state's own banned items)				
Large	60.8%	−4.911	(0.677)	−12.9
Small	23.6%	−0.383	(0.537)	−2.6
eCommerce	10.1%	−1.172	(0.271)	−18.5

Notes: SNAP households; non-purchasing months zero-filled; stacked difference-in-differences. For each category and store type, *Share* is that store type's share of the category's pre-policy spending (the three channels shown cover ≈91–97%; the rest is restaurants, specialty, and other formats); *Reduction* is the effect in dollars per household-month, with its standard error (clustered by state) in parentheses; and *% of mean* is that effect relative to the store type's own pre-policy treated mean. Category totals are in Table 2. Soda is estimated on all treated states; candy, energy drinks, and the restricted basket on the states that restrict *that* category. Eight in-window states restrict energy drinks. Event studies by store type are in Figure 7.

D Household-specific SNAP use across retailers

To test the trip-level mental-accounting interpretation without treating a retailer as uniformly SNAP-oriented for every customer, we define retailer type at the household–retailer–state–implementation-cohort level. Thus, the same chain can be classified as SNAP-heavy for one household and never-SNAP for another. We rank FMCG chains by pre-policy grocery spending and use the 50 largest chains in the main specification. Within this set, we retain only retailer-by-state-by-cohort cells in which at least one SNAP-paid trip is observed before implementation. Thus, a never-SNAP retailer in Table 12 is a retailer that demonstrably accepted SNAP in that state and cohort but at which the focal household did not use SNAP.

Classification uses event months -4 through -1 and is fixed before treatment. A household–retailer pair is SNAP-heavy if SNAP was recorded on at least 50% of its pre-policy in-store grocery trips, mixed-use if the share was positive but below 50%, and never-SNAP if it was zero; classification requires at least two such trips. Outcomes include all tenders so that purchases remain observed if a household switches from SNAP to cash or card after implementation.

Table 23 shows that the preferred same-household gradient is essentially unchanged when coverage expands from the 50 largest chains to the 100 largest chains or to all qualifying chains.

Table 23: Robustness of the same-household retailer gradient to chain coverage

	SNAP-heavy	Mixed-use	Never-SNAP	Observations
Panel A. Restricted products				
Top 50 chains (main)	-19.8^{***}	-10.6^{***}	-3.8	316,080
	(5.9)	(2.0)	(5.7)	
Top 100 chains	-18.8^{***}	-8.6^{***}	$+1.3$	389,408
	(6.2)	(2.0)	(4.3)	
All qualifying chains	-18.8^{***}	-9.4^{***}	$+0.0$	463,725
	(5.3)	(1.8)	(2.5)	
Panel B. Soda				
Top 50 chains (main)	-24.8^{***}	-12.0^{***}	-6.1	139,560
	(7.3)	(3.2)	(10.2)	
Top 100 chains	-23.9^{***}	-10.7^{***}	-5.1	171,478
	(7.6)	(2.7)	(9.3)	
All qualifying chains	-24.7^{***}	-11.9^{***}	-4.2	205,408
	(6.8)	(2.6)	(3.3)	

Notes: Entries are percentage effects, calculated by dividing each OLS coefficient by the corresponding retailer-type-specific treated pre-policy mean; standard errors in parentheses are scaled by the same mean. Each row retains households observed at all three retailer types before implementation. "All qualifying chains" removes the rank cutoff but retains the FMCG and revealed-SNAP-acceptance requirements. Models use the main stacked difference-in-differences fixed effects and cluster standard errors by state. $^{*}p < 0.10$, $^{**}p < 0.05$, $^{***}p < 0.01$.

E Data and sample construction

E.1 The Numerator panel

Our purchase data come from the Numerator consumer panel. Panelists photograph the receipt from every shopping trip – at grocery stores, supermarkets, club and dollar stores, convenience stores, and restaurants – using the Receipt Hog mobile application, and Numerator extracts trip- and item-level records from these images. We assemble three raw Numerator tables: an item-level transaction table (household, retailer, date, item, price, quantity, and item spending), a basket-level summary table (order method, delivery method, and *payment method*, including whether SNAP/EBT was used), and an item catalog recording each product's description, brand, and classification at four nested levels: sector, department, major category, and category. Our analysis uses a panel of "static panelists" – consistent uploaders whose demographic composition is broadly consistent with the national population – from August 2025 through June 2026.[28] The spending and nutrition outcomes use all observed purchase channels. Because SNAP tender is most reliably observed at the physical register, the pre-policy tender shares used in the sales tax calculation are measured in store. The item-level detail serves two purposes: separating SNAP-eligible grocery spending from non-grocery spending within a receipt, and linking each purchase to nutrition (Appendix G). The data have two measurement limitations we address in turn – panelists do not upload every receipt, and the receipts carry no nutritional content – and their central advantage is *multi-retailer* coverage: because we observe a household's receipts across a variety of stores, we can see soda purchases with cash and SNAP across retailers (something missed by single-chain scanner panels).

E.2 Identifying SNAP households

Because the Numerator dataset records the payment method used on each basket, we identify SNAP participation from observed payments rather than self-reports. We define a shopping trip as SNAP-paid when the basket's payment field equals `SNAP`. For each household-month we compute the share of that month's *in-store grocery* trips paid with SNAP, and call the household SNAP-active that month when the share is at least 10%. To avoid labeling one-off EBT-like transactions as participation, our primary flag `robust_snap` requires the household to be SNAP-active in the month *and* in an adjacent month ($t-1$ or $t+1$). The estimation sample keeps households that are `robust_snap` in at least two of the three pre-policy months.

[28] We define static panelists by retaining households whose latest `END_DATE` across Numerator static-panel membership spells is on or after June 1, 2025.

For our results on non-SNAP households, we define households as non-SNAP only when they have *zero* observed SNAP-paid trips in any channel, in store or online. Because the category-specific restrictions retain different treated states, the stacked SNAP samples contain 398,081 household-months for soda, 395,426 for candy, 396,104 for energy drinks, and 896,955 for the restricted basket.

E.3 Pre-policy payment-routing capacity

We calculate payment-routing capacity on the treated household–subcohort units that enter the paper's SNAP sample. For each unit, we require at least two pre-policy months $t = -4, \ldots, -1$ with positive spending on groceries that were SNAP-eligible before the restriction and explicitly recorded by Numerator as purchased in store. Let R denote spending on the products that the state will restrict, G spending on other SNAP-eligible groceries, and $F = R + G$. Let S denote eligible-grocery spending on receipts with an observed SNAP tender. We average dollar amounts within household before making any comparison.

State(s)	Products included in R
NE	Soda; energy drinks
IN, IA, ID, OK	Soda; energy drinks; sports drinks; candy; estimated restricted portions of Numerator's *Fruit Juice* and *Ready To Drink Tea* categories
UT, WV	Soda
LA	Soda; energy drinks; candy
FL	Soda; energy drinks; candy; packaged cookies; snack cakes

In Indiana, Iowa, Idaho, and Oklahoma, Numerator's *Fruit Juice* and *Ready To Drink Tea* categories combine products that are covered by the state rules with products that remain eligible. We use Kroger and Walmart only to estimate the restricted-product share within each category; the routing-capacity calculation itself uses spending at all retailers. To estimate the shares, we pool positive pre-policy Kroger and Walmart spending across the four states, regardless of tender, and classify products using the item-level product list developed for Indiana. We can classify 71.2% of this spending in *Fruit Juice* and 72.5% in *Ready To Drink Tea*. Among the classified spending, the restricted shares are 54.0% and 89.8%, respectively.

We apply these two shares to each household-month's all-retailer spending in the corresponding category that remains in G. This proportional allocation assigns unclassified products the same restricted shares and holds the shares fixed across the four states, retailers, households, and months. Products already included in R are not counted again; F and

S are unchanged. Texas is excluded because its sweetener-based beverage rule cannot be mapped using the same classification. Thus, $G = F - R$ is the observed grocery spending that remains eligible under the preferred coding.

The observed routing condition is

$$\bar{G}_i \geq \bar{S}_i.$$

It asks whether the household could assign all of its recorded SNAP-paid eligible-grocery spending to products that remain eligible and pay cash for the restricted products while holding its observed bundle fixed. It does not ask what the household would have bought before receiving SNAP. Among months in which households are classified as using SNAP, the condition holds for 2,818 of 2,942 household–subcohort units (95.8%). Using all valid pre-policy months, it holds for 2,832 units (96.3%).

E.4 Geography, weighting, and panel rotation

Each household is assigned a state from the ZIP code of its most recent address spell, mapped through a 2020 Census ZCTA crosswalk. We both assign treatment and cluster all standard errors at the state-level. Numerator's static panel rotates: each household holds one or more twelve-month membership spells. The headline sample is an *unbalanced* panel – a household-month appears only when the household actually shopped. Numerator supplies national-projection weights, but our primary estimates are *unweighted*.

E.5 Classifying products

Every purchased item carries Numerator's taxonomy, and we define the ban categories from it at the item level. *Soda* is any of fourteen carbonated-soft-drink categories (cola, citrus/berry, pepper-style, root beer/cream, ginger ale, specialty/natural soda, and their foodservice analogues); *energy drinks* are the energy-drink and energy-shot categories; *sports drinks* are Numerator's sports-drink category; *candy* is the candy department together with the chocolate and non-chocolate confectionery categories. Numerator's *Fruit Juice* and *Ready To Drink Tea* categories are broader than the corresponding state restrictions: each contains both restricted sweetened products and products that remain eligible. We therefore do not classify either category as wholly restricted. When this distinction is needed, we use item descriptions and Kroger and Walmart eligibility information to classify products within these categories. Table 8 and Appendix E.3 describe the treatment of products that cannot be classified. Grocery status is the Grocery sector; SNAP-eligibility additionally excludes alco-

hol; a trip is in-store when the order method is "In Store." For the substitution analysis we also flag packaged ice cream and frozen novelties (the Frozen-Foods "Ice Cream & Novelties" major category), which remain SNAP-eligible and serve as a never-banned control snack.

E.6 Spending and trip variables

All spending variables are item spending summed with the relevant indicators and collapsed to the household-month. Soda spending is item spending on grocery soda over all tenders; its *SNAP-paid* counterpart additionally requires SNAP tender; the in-store and online components are additive to the total. Total SNAP grocery spending (SNAP-tendered, SNAP-eligible grocery) is the denominator for the budget shares in the incidence calculation. Soda *trips* are counted once per basket. The all-tender versus SNAP-paid distinction is what separates a mechanical loss of the SNAP-payment channel from a true behavioral reduction; the in-store versus online split is what tests whether the ban merely relocates soda to a lower-stigma channel.

Soda spending outside grocery stores We separately test whether households shift soda purchases outside Numerator's Grocery sector. The outcome sums all-tender spending on the paper's enumerated soda categories in nonblank sectors other than Grocery. The sector restriction is applied before the product-category restriction: an item categorized as Cola in the Grocery sector is excluded, while the same category label is counted only when it appears in a non-Grocery sector. In constructing non-grocery soda spending, we also include purchases assigned to four categories in Numerator's current limited-service-restaurant beverage taxonomy: bottled soda, fountain drinks, dirty soda, and specialty soda. We exclude energy drinks, hot beverages, beverage additives, and frozen drinks, adding frozen drinks only in a sensitivity check.

Using the same stacked soda specification, spending outside Grocery falls by 3.3% and is not statistically distinguishable from zero ($p = 0.19$). Figure 8 likewise provides no evidence of an offsetting increase: every post-policy confidence interval includes zero, and the pre-policy leads are jointly insignificant ($p = 0.16$). Restricting the outcome to the four explicit limited-service-restaurant soda categories gives the same estimate.

E.7 SNAP food restriction waivers by state

The restrictions we study are USDA-approved state waivers that "restrict the purchase of non-nutritious items like soda and candy" using SNAP. The waivers are issued in waves throughout 2025 under the administration's "Make America Healthy Again" initiative and

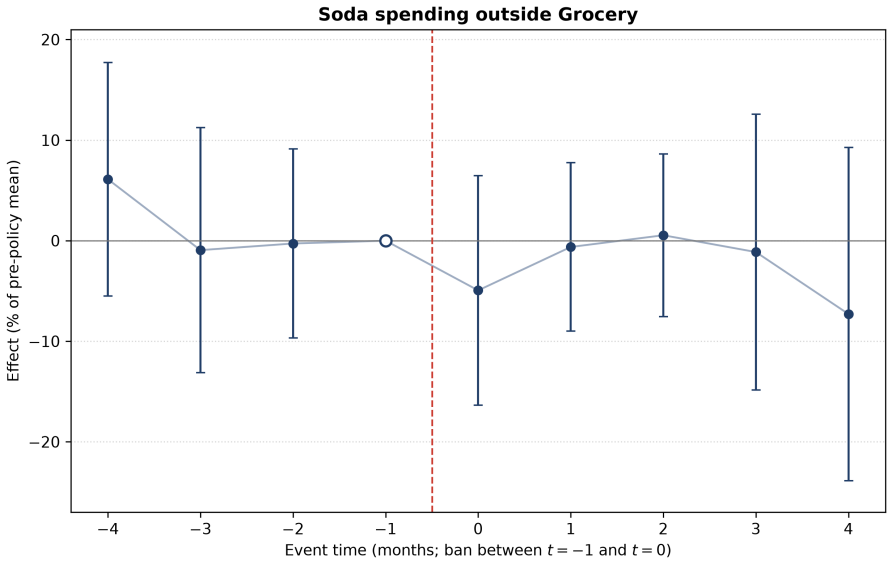


Figure 8: Event study of SNAP households' soda spending outside Grocery. The outcome is all-tender spending on enumerated soda categories in nonblank sectors other than Grocery. Points are OLS coefficients divided by the pre-policy treated mean; bars are 95% confidence intervals with standard errors clustered by state. The specification and stacked sample match the headline soda-spending estimates. Event time $t = -1$ is omitted, and the dashed line marks the restriction. The pre-policy coefficients are jointly insignificant ($p = 0.16$).

take effect on staggered dates in 2026-2027. Table 24 lists the approved waivers, their effective dates, and the categories of items restricted by each waiver, as recorded on the USDA Food and Nutrition Service SNAP waiver homepage. Ten state waivers take effect within our August 2025–June 2026 window and form the three treated cohorts. The remaining waivers take effect after the study period and are *excluded* from the control group rather than being used as not-yet-treated controls, so the comparison group remains untreated. The scope of the restrictions varies across states – most states place restrictions on soft drinks, several add candy, and some add energy drinks, sports drinks, or prepared desserts.

E.8 Non-SNAP households

Figure 9 plots the event study of soda spending for strictly non-SNAP households using the same stacked research design as in Section 4. The preferred proportional PPML specification shows a post-policy decline of roughly 3–4%, far smaller than the SNAP-household response but not zero. The OLS level specification yields a similar −3.4% ratio.

Table 24: SNAP food-restriction waivers by state

State	Effective date	Restricted categories
Panel A. In-sample treated cohorts		
Nebraska	Jan 1, 2026	Soft drinks, energy drinks[a]
Indiana	Jan 1, 2026	Soft drinks, candy
Iowa	Jan 1, 2026	Taxable foods[b]
Utah	Jan 1, 2026	Soft drinks
West Virginia	Jan 1, 2026	Soft drinks
Idaho	Feb 15, 2026	Soft drinks, candy
Oklahoma	Feb 15, 2026	Soft drinks, candy
Louisiana	Feb 18, 2026	Soft drinks, energy drinks, candy
Texas	Apr 1, 2026	Sweetened drinks, candy[c]
Florida	Apr 20, 2026	Soda, energy drinks, candy, prepared desserts
Panel B. Future waves excluded from controls		
Arkansas	Jul 1, 2026	Soda, low/no-calorie soda, <50%-juice drinks, energy drinks, other unhealthy drinks, candy
South Carolina	Aug 31, 2026	Soft/sweetened drinks, energy drinks, candy
North Dakota	Sep 1, 2026	Sweetened drinks, energy drinks, candy
Montana	Sep 30, 2026	High-sugar drinks, energy drinks, candy, prepared desserts
Ohio	Oct 1, 2026	Sugar-sweetened beverages, fountain drinks
Virginia	Oct 1, 2026	Sweetened beverages
Hawaii	Apr 1, 2027	Soft drinks
Kansas	Feb 15, 2027	Soft drinks, candy
Missouri	Feb 15, 2027	Candy, prepared desserts, certain drinks
Wyoming	Feb 1, 2027	Sweetened carbonated beverages
Nevada	Feb 1, 2028	Sugar-sweetened beverages, candy
Panel C. Approval vacated before implementation		
Colorado	Vacated Jun 22, 2026	Soft drinks
Tennessee	Vacated Jun 22, 2026	Processed foods and beverages such as soda, energy drinks, and candy

Source: USDA Food and Nutrition Service waiver status as of July 7, 2026 (U.S. Department of Agriculture, Food and Nutrition Administration, 2026). The ten in-sample states form the January, February, and April cohorts. [a] Nebraska's later candy expansion never entered the analysis window. [b] Iowa's rule follows the state sales tax definition and is broader than a conventional list of soda and candy. It explicitly allows SNAP recipients to purchase "food producing plants and seeds". [c] Texas defines a sweetened drink as a water-based nonalcoholic beverage that meets the statute's caloric-sugar threshold or contains any artificial sweetener. Energy drinks that are otherwise SNAP-eligible and meet this definition are restricted as sweetened drinks (Texas Legislature, 2025b). On June 22, 2026, a federal court vacated the approvals for Colorado, Iowa, Nebraska, Tennessee, and West Virginia. Iowa, Nebraska, and West Virginia had operated from January 1; their June observations are event time +5 and fall outside the paper's $[-4, +4]$ estimation window. Colorado and Tennessee never implement a restriction.

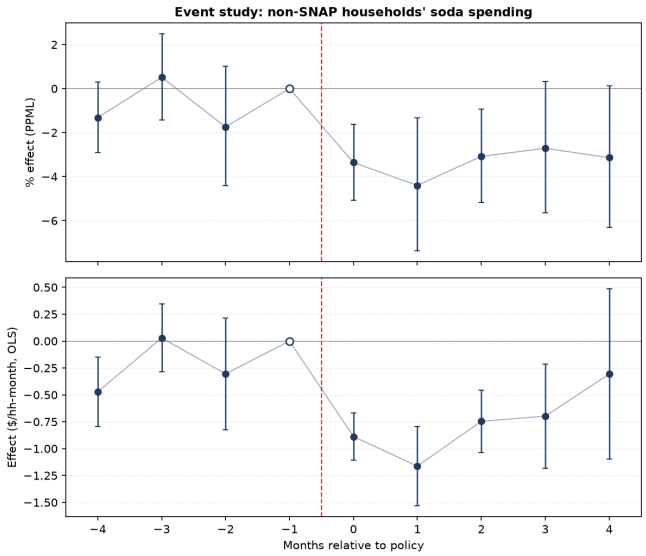


Figure 9: Event study of non-SNAP households' soda spending. The top panel reports proportional PPML effects and the bottom panel OLS effects in dollars per household-month. Event time $t = -1$ is omitted; bars are 95% confidence intervals with standard errors clustered by state, and the dashed line marks the restriction.

F Package size for the soda volume decomposition

The decomposition needs a size (volume) of each purchase. We therefore run it on the sixteen retailers where one of two routes covers the volume of most of a retailer's soda spending: a barcode on the receipt that an external catalog also carries, or a receipt description written in a form we can read. Those sixteen are 55.2 percent of the panel's soda spending, and we size 83.0 percent of it. Table 26 reports each retailer on its own.

Table 25 shows that 16 retailers are representative of all retailers of our sample in composition of store type–

Table 25: Store-type composition of the soda volume sample

	Sixteen retailers	Panel
	Share of soda spending (%)	
Large format	61.1	64.8
Small format	26.8	22.5
Online	12.2	10.0
Other outlets	—	2.7
Total	100.0	100.0

Notes: Both columns are soda spending by treated households before the restrictions, the basis Table 4 uses, and the panel column reproduces the store-type shares reported there. Table 26 instead covers every household over the whole sample period, so the retailer shares it reports do not aggregate to the first column here. The sixteen retailers include no outlet of the fourth type.

At eight of them the barcode route is enough on its own. Walmart, Dollar General, Walmart.com, Family Dollar, Walmart Neighborhood Market, Walgreens, Meijer and Fred Meyer each have size information for 81–87% of their soda spending. To raise coverage at the Walmart banners we collected size information for Walmart private label products, which the Kroger catalog does not carry, and linked it by barcode.

At four the barcodes leave a third of the shelf unsized and the receipt description closes it. Target, WinCo Foods, ShopRite and Albertsons write package configuration in a consistent form – `12Z/355Ml Cn 1X20P` is a case of twenty twelve-ounce cans – so we read the per-container size and the pack count and convert them to milliliters. Barcodes alone size between 59 and 69 percent of their soda spending; with the description they reach between 76 and 97 percent. We accept a reading only if it falls below the size at the 99.5th percentile of the packages that barcode-matched purchases report in its category, with five percent of slack, and only if the price per milliliter it implies lies within a factor of ten of what that category charges. A reading that fails either test is treated exactly as a missing one.

At the remaining four neither route reaches far enough, so we looked the items up ourselves. Costco, Sam's Club and Amazon.com are almost absent from the catalogs, at 0.0, 1.5 and 5.6 percent sized, and Kroger prints a barcode on approximately 96 percent of its soda spending while the catalogs return a size for 38.7 percent. Within each of the four we rank the unresolved items by spending, identify the product from its brand and receipt description, and confirm its package configuration against the retailer's own listing. We continue down the ranked list until roughly 90% of the retailer's previously unsized spending has been reviewed. Items that cannot be identified confidently remain unresolved. Where a hand-checked size and a parsed reading cover the same item, the hand-checked value takes precedence. A size from either of these two sources must also fall between 100 mL and 50 L.

Table 26: Package-size coverage in the soda volume sample

Retailer	Share of soda spending (%)	Barcode catalogs only (%)	All sources (%)	Price per 12 oz ($)	Principal new source
Walmart	30.0	70.3	81.0	0.54	Additional UPC catalogs
Dollar General	16.6	82.2	86.5	0.66	Additional UPC catalogs
Walmart.com	12.0	70.4	82.4	0.50	Additional UPC catalogs
Kroger	8.3	38.7	75.1	0.52	Manual review
Costco	5.7	0.0	86.6	0.54	Manual review
Sam's Club	4.8	1.5	86.8	0.53	Manual review
Family Dollar	3.0	79.5	82.1	0.62	Additional UPC catalogs
Walgreens	2.5	83.9	85.9	0.72	Additional UPC catalogs
Walmart Neighborhood Market	2.5	77.4	86.1	0.50	Additional UPC catalogs
ShopRite	2.4	63.4	85.8	0.55	Receipt-text parser
Target	2.4	58.7	93.0	0.65	Receipt-text parser
Meijer	2.4	74.8	82.5	0.52	Additional UPC catalogs
Amazon.com	2.3	5.6	84.9	0.61	Manual review
Albertsons	2.1	58.9	75.9	0.56	Receipt-text parser
WinCo Foods	1.7	68.6	96.6	0.48	Receipt-text parser
Fred Meyer	1.6	77.1	82.2	0.55	Additional UPC catalogs
All sixteen retailers	100.0	—	83.0		0.56

Notes: Shares are of soda spending within the sixteen-retailer sample and may not sum to 100 because of rounding. The two coverage columns give the share of a retailer's soda spending that carries a size when only the barcode catalogs are used, and when every source described above is used. The price column is spend-weighted over the sample period and covers the sized spending only. The sixteen retailers are 55.2% of soda spending by SNAP households in the panel, so 83% coverage within them supplies a size for approximately 46% of panel soda spending.

Sizing the spending also lets us price it, and what a household pays for soda turns on where it shops. The last column of Table 26 gives the spend-weighted price of a twelve-ounce equivalent at each retailer. It runs from $0.48 at WinCo Foods to $0.72 at Walgreens.

G Nutrition linkage

Numerator does not report the nutritional content of the items households buy. We therefore link each purchased item to an external reference product, read a package size, and multiply

the two. This appendix describes the reference data, the products each basket contains, how items are matched, where package size comes from, how a purchase becomes a nutrient total, and how those totals are summed to the household-month.

Our procedure differs from the approach in Song (2026), which links products to nutrition information through a NielsenIQ Homescan category crosswalk and assigns each crosswalk category a fixed nutrient-per-dollar value. In that approach, nutritional content varies across categories but not across products within a category. Our primary measure instead allows nutrient density and package size to vary across products whenever item-level linkage is possible. We use the spending-based rate only as a fallback when package size information is unavailable.

G.1 Reference nutrition data

Our primary nutrition source is the December 2025 release of the USDA FoodData Central Branded Foods database, which compiles nutrition information reported on branded food labels. This release contains 1,993,975 product records. Because USDA creates a new record when information about an existing product changes, this raw row count is not a count of distinct products. For descriptive purposes, we standardize product identifiers and retain the most recently published record for each product. This results in 441,915 unique products (normalized UPCs), of which 440,808 are identified as U.S. products. For beverages, we supplement USDA records with nutrition information compiled from U.S. Open Food Facts.

We construct separate reference tables for beverages and solid foods. The beverage table contains 44,483 rows representing 42,007 unique products. Of these rows, 41,821 come from USDA and 2,662 come from the supplemental U.S. Open Food Facts database. The solid food table contains 677,067 USDA records representing 161,334 unique products.

A third source enters through the barcode rather than through these tables. The Kroger product catalog behind the beverage package sizes also reports nutrition per 100 mL, and we read it where a purchase carries a barcode for which the reference tables hold no value. It covers the same nutrients on the same basis.

For each record we retain, when reported, calories, total sugar, sodium, carbohydrates, fat, protein, and serving size information. Beverage nutrient densities are measured per 100 mL, while solid food nutrient densities are measured per 100 g. We combine these densities with package size information to construct package-level nutrients. During matching, we assign Numerator items and reference records to one of three sugar classes – zero, reduced, or regular – using product descriptions and, for USDA records, reported sugar content.

G.2 What each basket contains

The beverage basket covers the drinks a household buys as a ready-to-drink liquid. We exclude ten categories sold as a dry good – leaf and convenience tea, single-serve pods, dry coffee, powdered drink mixes, liquid enhancers, and hot cocoa. A household that stops buying soda is unlikely to replace it with loose tea or a canister of drink mix, and their spending does not change once the restrictions take effect: the estimated change is +0.7 percent, with a ninety-five percent interval running from −3.0 to +4.4 percent. They account for approximately 16 percent of beverage spending.

The packaged snack basket covers candy, chips and other snacks, crackers, cookies and snack cakes, sweet bakery products, and retail ice cream. We use it to test whether reductions in soda or candy purchases are offset by substitution toward other energy-dense foods, so it is intended to capture plausible margins of substitution rather than the household's entire diet. For reporting we group these categories into the broad policy-relevant buckets shown in Panel B of Table 7.

G.3 Package size

A beverage purchase is sized from a barcode match against the Kroger product catalog, and otherwise from a size read out of the product description. Appendix F describes the description reader and the limits it must satisfy.

For solid foods, package weight comes from the *package weight* field in the USDA FoodData Central database and is linked by UPC. We parse weights into total package grams. USDA creates a new record whenever a product's information changes, so one barcode can carry several package weights that disagree. Where no weight has a clear majority, where the majority implies a density twenty times that of water, or where it conflicts with the product's own history – which usually means it is the weight of one unit rather than of the pack – we drop the barcode rather than choose among them. The Kroger catalog behind the beverage sizes reports a single size per barcode, so beverages need no equivalent step. When the field reports a frozen novelty or ice cream package in fluid ounces, it measures volume rather than weight. We convert that volume to grams using a product-specific density recovered from USDA serving information. The median estimated ice cream density is approximately 0.6 g/mL. These conversions reproduce package weights from an independent Kroger ice cream scrape to within ±10% for 97% of products.

Where no barcode supplies a size, we read one from the product description and accept it only if it passes two checks. The first compares the reading with the largest package that

barcode-matched purchases report in the same category: we accept it only if it falls below the 99.5th percentile of that distribution, with five percent of slack. The second compares the price the reading implies with what the category ordinarily charges: for beverages we require the implied price per milliliter to fall within a factor of ten of the category norm, and for packaged snacks the implied price per gram to fall within a factor of five.[29] A reading that fails either check is discarded and the purchase is treated exactly as though no size had been found. This matters for the two sections that follow: a discarded reading can neither become a package nutrient nor influence the rate we use in its place.

Both checks can be graded, because a purchase the barcode sized never had its description read. A reading agrees with the barcode to within five percent for 86.9 percent of text-read spending across packaged snacks and 87.7 percent across beverages. The checks discard 4.4 and 6.7 percent of that spending respectively, and what survives is wrong by ten times or more for 0.7 and 0.1 percent of it.

Reading the description raises the share of spending that carries a validated size from approximately 26 to 49 percent across the beverages the totals retain and from approximately 26 to 45 percent across packaged snacks.

G.4 Nutrient density

The second ingredient is how much of each nutrient a product contains per unit of size – calories per 100 mL for a beverage, per 100 g for a solid. That figure is printed on the product's own label and is what the USDA reference records. Reading it is not the hard part. Deciding which reference product a purchase corresponds to is.

Numerator records a store-specific description and brand information alongside a barcode, but barcodes are missing or truncated often enough that most purchases have to be placed some other way. We therefore match hierarchically, taking the most granular match a purchase supports and falling back to progressively broader ones. Spending that carries both a validated size and a density is 49.0 percent of the basket, and reading down the hierarchy, as a share of that spending:

1. **A barcode match to the USDA reference:** 47.8 **percent.** The purchase takes that product's own label. Because USDA writes a new record whenever a product changes, one barcode can carry several; we use it only where those records agree on

[29] The two factors differ because the categories differ in how tightly unit price is distributed; each is set from the category's own barcode-matched purchases. For beverages we also refuse description-derived sizes outright in the dry categories above, where a gram figure would be read as a volume.

sugar class and lie within five calories and two grams of sugar of one another, and send the rest down to the later steps rather than choosing among them.

2. **A barcode match to the Kroger catalog:** 10.2 **percent.** The catalog behind the beverage package sizes also reports nutrition per 100 mL, and we read it where the reference tables hold no value for that barcode.

3. **The commodity average:** 13.7 **percent.** Water, sparkling water, plain coffee, cocoa and tea take an average by design rather than by failure, because within those categories the products are close enough in composition that an average built from plain reference products is the better answer.

4. **A brand match that also agrees on sugar class, form and flavor:** 12.4 **percent.** We match Numerator's brand and parent-brand fields against the reference's brand: directly where the two spell it the same way, and through a hand-checked crosswalk where they do not.[30] A brand is not enough on its own, because the step asserts that one product *is* another. Product form and flavor or product line must agree as well, which is what stops a powder being read as a ready-to-drink beverage, and so must sugar class, which is what stops a diet product taking a regular-sugar reference.[31]

5. **A wider group average:** 13.0 **percent.** No product-level match is made, so the purchase takes the average of a set of reference products. The sets widen in turn:

   - brands whose names share enough tokens with the purchase's, within its category and sugar class — 8.3;
   - the manufacturer's other brands, within category and sugar class — 1.1;
   - the category and sugar class alone — 2.5, and 0.1 more for the fountain categories;
   - the major category — 0.2;
   - sugar class by itself, the set that always exists — 1.0.

6. **A named exception:** 2.2 **percent.** A check that rejects a correct match sends a product we could have named to the average of its category.[32] For twenty-one product

[30] Automatically generated crosswalk entries are kept only above a similarity threshold, without which the automatic pass matches a brand on a fragment of an unrelated description – finding a brand named Ice inside a description of iced tea.

[31] Sugar class is a join condition where the brand strings match directly, so a diet product cannot take a regular-sugar reference there. At the crosswalk and parent-brand stages, together 4.9 percent of this spending, it only ranks the candidates, so a diet product with no same-class candidate can take an opposite-class one.

[32] A receipt reading `MTN DEW BAJA BLAST` carries no barcode and matches the Mountain Dew brand, but fails the flavor check against the record it reaches and falls to the average of its category – even though the reference does carry Baja Blast.

lines where this happened often enough to matter, a written rule restores the named product. None of them overrides a barcode.

7. **A label reporting the nutrient as zero:** 0.6 **percent.**

Nearly three-fifths of this spending is settled by a barcode and a further 13.7 percent by the commodity categories, leaving 13.0 percent that reads a group average because no product-level match could be made. A purchase that reaches the end of the list without even a group is carried as missing rather than assigned a nearest guess.

Packaged snacks and desserts are matched to USDA reference foods by the same descending-precision cascade, run separately with solid categories and per-gram nutrient amounts.

G.5 From size and density to household nutrients

When a validated package size is available, we multiply nutrient density by package size to obtain the nutrient content of one package and then multiply by the number of packages recorded in the Numerator receipt data to obtain the total nutritional content of a household purchase.

Where no size survives we multiply spending instead by a nutrient-per-dollar rate, estimated from the purchases for which both size and density are observed. The rate is taken from the narrowest cell that has one. Of beverage spending in the basket:

- 49.0 percent carries a size and needs no rate
- 46.0 percent takes the rate of its own brand, category and sugar class
- 4.9 percent has no cell that narrow and takes the rate of its category and sugar class
- under 0.1 percent falls further, to the rate of its category alone
- nothing is left over. Every beverage dollar the totals retain carries a nutrient value.[33]

Household-month nutrient totals remain missing whenever a selected purchase lacks the corresponding dense nutrient value; those observations are omitted only from regressions for that nutrient.

Finally, explicitly reported zero nutrient values are retained in our analysis sample, but an unreported value is not treated as zero. The dense nutrition outcomes used in the beverage

[33]This is a property of the basket rather than of the procedure. Across all beverage categories 3.8 percent of spending resolves to no nutrient value at all, and every dollar of it sits in the dry categories, where roughly a quarter of spending reaches the end of both sequences with nothing to assign.

and combined-basket regressions are calories, total sugar, sodium, carbohydrates, fat, and protein.

G.6 Estimates without the imputed purchases

A purchase whose package size we could not read takes a nutrient-per-dollar rate in place of its own nutrients. To see whether the estimates rely too heavily on this step, we rebuild the household-month totals from the purchases whose size we did read, drop the rest, and re-estimate. This exercise drops about half of "matched" calories – we capture 54 percent of pre-policy beverage calories and 52 percent of pre-policy packaged-snack calories. Expressed as percentage changes, the effect estimates are similar using this restricted subsample of items.

Table 27: Effects with and without the imputed purchases

	All purchases	Size observed
	Percent of the pre-policy mean	
Beverage calories	−7.3	−6.8
Beverage total sugar	−8.1	−7.5
Beverage sodium	−8.9	−6.5
Packaged snack calories	−5.1	−4.0
Packaged snack total sugar	−7.6	−6.5
Packaged snack sodium	−6.4	−5.3

Notes: The two columns use the same specification and the same households and differ only in which purchases contribute nutrients. Each divides by its own pre-policy mean, which is why the columns are comparable in percent rather than in levels. Beverage sodium moves most, and that difference sits in energy and sports drinks rather than in the restricted products: soda sodium falls by 11.6 percent on all purchases and by 11.3 percent on those we could size. The three packaged-snack rows all move by about a point.

G.7 Aggregation to the household-month panel

We sum package-level nutrients within each household and calendar month, separately for beverages and packaged snacks and desserts. We merge these totals onto the full set of active household-months, defined as months in which the household records at least one purchase in the Numerator data. An active household-month with no purchase in a given nutrition category is assigned zero for that category. As a result, our estimand captures the average effect over all active household-months rather than an effect conditional on purchasing soda, candy, or another restricted category.

We construct beverage-plus-snack outcomes by summing the two category-specific measures. Because the component outcomes are constructed over the same active household-month sample, category-specific effects are additive and sum to the combined effect. Table 6, Panel C reports estimates for the full sample.

As in any receipt panel, some purchases may go unreported. Following Song (2026), we emphasize percentage effects. If receipt reporting is approximately proportional and stable across the treatment and control groups, it scales both the estimated level effect and the pre-policy mean and has little effect on their ratio.

H Retailer price responses

H.1 Data construction

We scrape weekly soda price data from four large retail chains: Kroger, Dollar General, Walmart, and Publix. We select these retailers for two reasons. First, they operate across treatment and control states, unlike smaller local retailers or state-level chains. Second, they span several major retail formats: traditional supermarkets (Kroger and Publix), supercenters (Walmart), and dollar stores (Dollar General). For computational reasons, we limit our data collection to three states: Texas, Georgia, and Florida. We choose these states because Texas and Florida implemented SNAP restrictions covering soda in April 2026 (during our data collection period), while Georgia provides a regional comparison state that did not impose a restriction. Walmart and Dollar General data are collected from Texas, Georgia, and Florida, while Publix data are collected from Florida and Georgia and Kroger data are collected from Texas and Georgia. Because collecting prices from every store is intractable, we use a fixed subset of stores across each retailer (i.e., we scrape prices from the same set of stores in each week).

Our sample contains weekly price observations from 300 Dollar General stores, 273 Publix stores, 367 Kroger stores, and 300 Walmart stores. The median number of unique products per store-week ranges from 213 (Dollar General) to 427 (Walmart) across retailers.

Products are linked over time using the retailer SKU or item identifier when available and the product URL otherwise. Stores are linked using a combination of store name and address. Our processed files retain product name and URL, store name and address, scrape dates, shelf or listed price, promotion and coupon text, and a constructed effective price. The effective price is the primary outcome variable in Section 6.1.1. To construct the effective price, we apply deterministic rules using the text of promotions and coupons to construct the effective shelf price, including buy-x-get-y-free, multiple-units-for-a-fixed-dollar-amount, dollar-off, and percentage-off promotions, and otherwise retain the best available posted price as the shelf price.

H.2 Additional results

H.2.1 Promotions

The main analysis uses effective price, defined as the product price inclusive of promotions. To distinguish changes in the frequency of promotions from changes in discount size, we re-estimate Equation (4) for two additional outcomes. First, we construct an indicator equal

to one whenever an item's promotional price falls below its listed price. Second, we construct a measure of the promotion depth, or the discount per item.

Table 28 reports promotion exposure effects in percentage points and promotion depth effects in cents. Pooling retailers, promotion exposure falls by 0.11 percentage points in Texas but rises by 1.21 percentage points in Florida. Promotion depth falls in both comparisons, by 0.64 cents in Texas and 1.07 cents in Florida. The underlying retailer estimates are heterogeneous. Most notably, Publix combines a 2.69 percentage point increase in promotion incidence with a 1.76 cent reduction in average promotion depth. Thus, the Florida results reflect more frequent but shallower discounts rather than a uniform expansion in promotional intensity.

Table 28: Effects on promotions

	Texas vs. Georgia		Florida vs. Georgia	
	(1) Pr(promotion), pp	(2) Depth, cents	(3) Pr(promotion), pp	(4) Depth, cents
All retailers	−0.11	−0.64**	1.21***	−1.07***
	(0.10)	(0.26)	(0.11)	(0.36)
Kroger	1.11***	−2.56***	—	—
	(0.14)	(0.51)	—	—
Walmart	−1.21***	0.64	0.36***	−0.68**
	(0.17)	(0.41)	(0.14)	(0.29)
Dollar General	−0.35***	0.47***	−0.12***	−0.42***
	(0.08)	(0.08)	(0.04)	(0.14)
Publix	—	—	2.69***	−1.76**
	—	—	(0.19)	(0.82)
Num. observations	3,421,524	3,421,524	2,998,543	2,998,543
Num. stores	787	787	653	653
Num. items	2,401	2,401	3,360	3,360

Notes: The sample and fixed effects in Table 28 match the specification in Table 9. Promotion is an indicator equal to one when promotional price is below the listed price. Promotion depth is $\max\{P^{list} - P^{promo}, 0\}$. All models include store-by-item and item-by-timing-cell fixed effects. Standard errors, clustered by physical store, are in parentheses. *, **, and *** denote significance at the 10%, 5%, and 1% levels, respectively.

H.2.2 Stable panel

The full sample is unbalanced because observed products change frequently across stores and weeks. To assess whether this changing composition drives the price estimates, we repeat the analysis using a more stable, though not completely balanced, basket of products. Within each retailer-policy comparison, we retain an item if it appears in at least 95% of stores in at least 80% of observed weeks. This sample restriction dramatically reduces the

number of items in the panel. Aggregating across the three retailer comparisons, the Texas-Georgia sample retains 336 of 2,401 retailer-item pairs (14.0%) and 1,345,207 of 3,421,524 observations (39.3%). The Florida-Georgia sample retains 281 of 3,360 retailer-item pairs (8.4%) and 859,703 of 2,998,543 observations (28.7%).

Table 29 reports the stable panel estimates. The pooled Texas coefficient increases from -0.01% in the full panel to 0.31% in the balanced panel, while the pooled Florida coefficient remains at 0.12%. Retailer-specific coefficients typically have the same sign but differ modestly in magnitude, ranging from −0.34% at Florida Walmart to 0.73% at Publix. These results qualitatively preserve the retailer patterns in the larger sample, though this small subsample generally yields somewhat larger estimates in magnitude.

Table 29: Retailer price responses in the balanced item panel

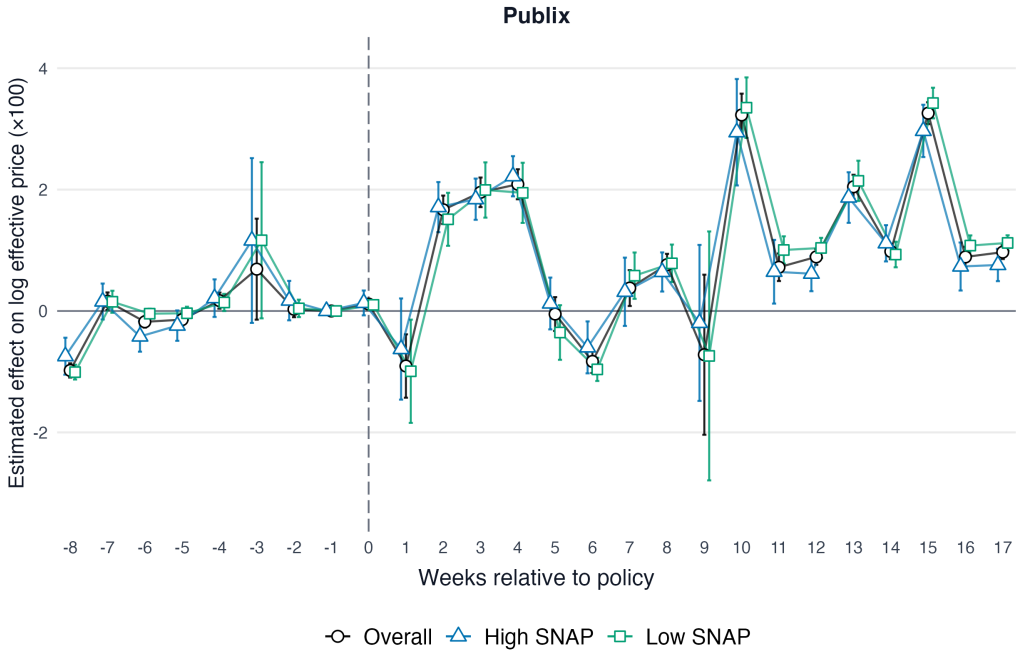
	Texas vs. Georgia (1)	Florida vs. Georgia (2)
All retailers	0.31***	0.12**
	(0.05)	(0.05)
Kroger	0.61***	—
	(0.10)	—
Walmart	0.03	−0.34***
	(0.04)	(0.04)
Dollar General	0.03	0.03
	(0.02)	(0.02)
Publix	—	0.73***
	—	(0.14)
Num. observations	1,345,207	859,703
Num. stores	787	653
Num. items	336	281

Notes: Entries are estimated coefficients on log price multiplied by 100. Within each retailer-policy comparison, an item must appear in at least 95% of stores in at least 80% of observed weeks to be retained in the sample. All models include store-by-item and item-by-timing-cell fixed effects. Standard errors, clustered by physical store, are in parentheses. *, **, and *** denote significance at the 10%, 5%, and 1% levels, respectively.

H.2.3 Extended Publix results

Figure 10 extends the Publix event study through the week ending August 23, 2026, seventeen Publix ad cycles (weeks) after the implementation of Florida's SNAP restrictions on April 20. We re-estimate the primary specification in Equation (4) on the full Publix panel, with physical-store-by-item and item-by-Publix-ad-cycle fixed effects and standard errors clustered by physical store.

Figure 10: Extended Publix price event study



Notes: The figure reports event study coefficients for log effective price at Publix. Effective price equals the promotional price when one is observed and the listed shelf price otherwise. The black series uses all eligible stores, while the blue and green series use stores in ZIP codes in the top and bottom terciles, respectively, of SNAP participation in the Florida-Georgia Publix sample. Event time $t = -1$ is omitted, the dashed line marks the implementation week, and bars report 95% confidence intervals based on standard errors clustered by store.

In the extended sample, the overall coefficient is −0.91% in week +1, rises to between 1.67% and 2.09% in weeks +2 through +4, and returns to approximately zero or below in weeks +5 and +6. Later peaks recur in weeks +10 and +15 (about 3.2%), while the other estimates from weeks +10 through +17 range from 0.72% to 2.05%. Using the extended event window, the average treatment effect is 1.10% (standard error 0.08), compared with 0.33% (standard error 0.13) in the main −5 through +9 window. However, the small difference between areas with different SNAP exposure and recurring peaks in across-state price differences are more consistent with state-level promotion cycle differences than with a one-time, persistent price response to the restrictions.

I Theory: mental accounting and retail pricing

This appendix asks why retailers have little incentive to change prices when the restriction reduces demand from SNAP households. We first relate a change in market elasticity to the retailer's optimal price. We then calibrate three ways in which the psychological mechanisms discussed in the main text can change SNAP households' price elasticity.

I.1 Retailer pricing and the optimal price

Consider a retailer that charges the same price to SNAP participants and nonparticipants. Let π_S be SNAP households' share of soda volume, and let ε_S and ε_N be the positive magnitudes of the two groups' price elasticities. The elasticity of total market demand is

$$\bar{\varepsilon} = (1 - \pi_S)\varepsilon_N + \pi_S\varepsilon_S.$$

With marginal cost c, optimal price p^*, and conduct parameter $\theta \in [0, 1]$, the retailer's first-order condition is

$$\frac{p^* - c}{p^*} = \frac{\theta}{\bar{\varepsilon}} \equiv L,$$

where L is the retailer's markup as a share of price.

The restriction does not change marginal cost. It affects the target price only by changing the elasticity of market demand. Linearizing the first-order condition around the pre-restriction optimum gives

$$\Delta \ln p^* \simeq -\rho L \frac{\Delta\bar{\varepsilon}}{\bar{\varepsilon}}, \tag{7}$$

where $\rho = \partial p^*/\partial c$ is local cost pass-through. Thus, a restriction that makes market demand less elastic raises the retailer's target price. The response is small when the change in market elasticity is small.

I.2 Calibration inputs

The pricing calculation requires SNAP households' share of soda volume rather than their share of customers. If R is their customer share and they buy r times as much soda per household as nonparticipants, then

$$\pi_S = \frac{rR}{rR + (1 - R)}.$$

At the national participation rate, $R = 0.123$, the benchmark $r = 2$ implies $\pi_S = 0.219$.

The calibration proceeds in two stages. The purchase decline d and the model-specific inputs $(s_g, \varepsilon_{gs}, \varepsilon_S, \varphi)$ determine the change in SNAP households' elasticity. The inputs (R, r, ε_N) then map that change into market elasticity, and ρL maps the market elasticity change into the retailer's target-price response.

Table 30: Inputs to the retailer-pricing calibration

Input	Value	Source and use
SNAP soda decline, d	0.12	Estimate from this paper
SNAP customer share, R	0.10–0.15; 0.30 sensitivity	USDA FNS participation counts; national rate 0.123
Relative soda purchases, r	2.0	USDA retailer point-of-sale data in Garasky et al. (2016) and the income gradient in Allcott et al. (2019, Figure III)
SNAP-associated soda share, φ	0.40	Pre-policy share of in-store soda spending on receipts using SNAP; used only in Model (3)
Pre-restriction SNAP elasticity, ε_S	1.20	Mean SSB elasticity in Powell et al. (2013); 1.373 from Allcott et al. (2019) used as a sensitivity
Non-SNAP elasticity, ε_N	0.8–1.4	Range spanning Andreyeva et al. (2010), Powell et al. (2013), and Allcott et al. (2019)
Other eligible-grocery budget share, s_g	0.10–0.20	Cash-equivalent share excluding soda; used only in Model (1)
Other-grocery cross-price elasticity, ε_{gs}	0.20 upper bound	Conservative sensitivity; used only in Model (1)
Pricing multiplier, ρL	0.18; $[0.095, 0.28]$	Beverage-tax pass-through and grocery-company filings

The benchmark $r = 2$ comes from per-household soft-drink spending in the USDA point-of-sale study and is consistent with the income gradient in Allcott et al. (2019). For the pricing multiplier, the central values $\rho = 0.70$ and $L = 0.25$ give $\rho L \simeq 0.18$. The range combines pass-through estimates of 0.43–1.00 from beverage-tax studies (Cawley and Frisvold, 2017; Cawley et al., 2018, 2020; Seiler et al., 2021) with fiscal-2024 company-wide gross margins of 22.3 percent at Kroger and 27.7 percent at Albertsons (The Kroger Co., 2025; Albertsons Companies, Inc., 2025). We use these accounting margins only to anchor the range for L; they are not soda-specific Lerner indices.

Why tax pass-through is informative. The same local curvature of demand that governs tax pass-through also governs the response to a change in elasticity. Differentiating the retailer's pricing condition with respect to a cost shock and to an elasticity shock gives the common multiplier ρL in Equation (7). We use the tax evidence to discipline this multiplier, not to assume that the restriction itself raises marginal cost.

I.3 From SNAP demand to market elasticity

The restriction can change market elasticity in two ways. It may change SNAP households' own price elasticity, denoted by $\Delta\varepsilon_S$, and it reduces their weight in market demand. If their soda purchases fall by a proportion d, their post-restriction volume share is

$$\pi'_S = \frac{\pi_S(1-d)}{1-d\pi_S}.$$

The change in market elasticity is therefore

$$\Delta\bar{\varepsilon} = \pi_S'\Delta\varepsilon_S + (\pi_S - \pi_S')(\varepsilon_N - \varepsilon_S). \tag{8}$$

The first term is the change in SNAP households' price sensitivity. The second is a composition effect: after the decline, nonparticipants receive more weight in market demand. The three models below supply $\Delta\varepsilon_S$.

I.4 Three behavioral models

We begin with the cash-equivalent benchmark. Let s denote soda, g other eligible groceries, n all other goods, Y cash resources, B SNAP benefits, and p_s the price of soda. If B were paid as cash and fully pooled with Y, the household would choose

$$\big(s^U(p_s), g^U(p_s), n^U(p_s)\big) \in \arg\max_{s,g,n\geq 0} U(s,g,n) \quad \text{s.t.} \quad p_s s + g + n \leq Y + B,$$

where the prices of g and n are normalized to one. The superscript U denotes demand under this equivalent unrestricted cash transfer. Each model changes the household's decision rule while holding this benchmark fixed.

I.4.1 Model (1): eligibility raises desired quantity

Model (1) represents eligibility as a multiplicative quantity shifter rather than a perceived price wedge. Relative to cash-equivalent demand, labeled goods are scaled by $\ell \geq 1$, while the remaining goods are scaled to satisfy the budget constraint. Before the restriction, both soda and other eligible groceries are labeled; afterward, only the other eligible groceries remain labeled:

$$\begin{aligned} \text{before:} \quad & (s^{\text{pre}}, g^{\text{pre}}, n^{\text{pre}}) = (\ell s^U, \ell g^U, \kappa_0 n^U), \\ \text{after:} \quad & (s^{\text{post}}, g^{\text{post}}, n^{\text{post}}) = (\kappa_1 s^U, \ell g^U, \kappa_1 n^U). \end{aligned}$$

where κ_0 and κ_1 are budget-balancing scalars. Matching the observed proportional decline d requires $\kappa_1/\ell = 1 - d$.

Let

$$s_g \equiv \frac{g^U(p_s)}{Y + B},$$

denote the cash-equivalent budget share of other eligible groceries. Because these groceries remain labeled after the restriction, post-restriction budget balance requires

$$\ell s_g + \kappa_1(1 - s_g) = 1.$$

Combining the two restrictions gives

$$\kappa_1 = \frac{1 - \ell s_g}{1 - s_g}, \qquad \ell = \frac{1}{1 - d(1 - s_g)}.$$

Using our estimate $d = 0.12$ and the calibration range $s_g \in [0.10, 0.20]$ in Table 30 implies $\ell \simeq 1.11$–1.12 and $\kappa_1 \simeq 0.97$–0.99.

Because ℓ is price invariant, pre-restriction soda demand has the same elasticity as s^U. The post-restriction scalar κ_1 varies with p_s only through the response of g^U. Writing $\varepsilon_{gs} = \partial \ln g^U / \partial \ln p_s$ and reporting own-price elasticities as positive magnitudes,

$$\Delta\varepsilon_S = -\frac{\partial \ln \kappa_1}{\partial \ln p_s} = (\ell - 1)\varepsilon_{gs}\frac{s_g}{(1 - \ell s_g)(1 - s_g)}.$$

This expression requires no parametric functional form for s^U or g^U; it follows from the multiplicative allocation rule, a binding budget constraint, a price-invariant ℓ, and differentiable cash-equivalent demand. At the conservative upper bound $\varepsilon_{gs} = 0.20$ reported in the same table, $\Delta\varepsilon_S < 0.007$. We therefore treat Model (1) as leaving SNAP households' price elasticity essentially unchanged.

I.4.2 Model (2): eligibility acts like a price subsidy

Model (2) represents eligibility as a perceived specific subsidy λ. The household pays the shelf price p_s but chooses soda as if its price were $p_s - \lambda$; removing eligibility eliminates this perceived subsidy:

$$s^{\text{pre}}(p_s) = s^U(p_s - \lambda), \qquad s^{\text{post}}(p_s) = s^U(p_s).$$

Under the isoelastic specification $s^U(p) = Ap^{-\varepsilon_S^0}$, the pre-restriction elasticity with respect to the shelf price is

$$-\frac{\partial \ln s^{\text{pre}}(p_s)}{\partial \ln p_s} = \varepsilon_S^0 \frac{p_s}{p_s - \lambda}.$$

Normalizing the calibration price to one therefore gives

$$\varepsilon_S^{\text{pre}} = \frac{\varepsilon_S^0}{1 - \lambda}, \qquad \varepsilon_S^{\text{post}} = \varepsilon_S^0.$$

Unlike Model (1), this elasticity mapping is functional-form dependent: it uses isoelastic demand and a price-invariant specific subsidy. The corresponding quantity ratio is

$$\frac{s^{\text{post}}(1)}{s^{\text{pre}}(1)} = (1 - \lambda)^{\varepsilon_S^0} = 1 - d.$$

Thus, the observed pre-restriction elasticity and purchase decline jointly identify ε_S^0 and λ:

$$\varepsilon_S^0 = \varepsilon_S(1-\lambda), \qquad (1-\lambda)^{\varepsilon_S^0} = 1-d.$$

Using our estimate $d = 0.12$ and the mean SSB elasticity $\varepsilon_S = 1.20$ from Powell et al. (2013), as summarized in Table 30, the economically relevant low-subsidy solution is

$$\lambda = 0.113, \qquad \varepsilon_S^{\text{post}} = 1.064, \qquad \Delta\varepsilon_S = -0.136.$$

The equations also yield a second solution, $\lambda \simeq 0.97$. We do not use it because it implies a cash-demand elasticity of about 0.04, far below estimates from soda-tax studies. Thus the perceived subsidy is 11.3 percent of the normalized shelf price, and its removal lowers measured elasticity from 1.20 to 1.064. This rotation is a calibrated implication of the isoelastic cash-demand benchmark, not a separately estimated elasticity change.

Both models can reproduce the same purchase decline. Their difference is what happens to price sensitivity: Model (1) mainly shifts the level of demand, whereas Model (2) changes the mapping from the shelf price to the price the household perceives.

I.4.3 Model (3): payment rerouting on SNAP-associated trips

Model (3) treats the need to coordinate SNAP and cash as a money-metric planning cost on SNAP-associated trips. Let φ be the share of pre-restriction soda purchases made on these trips. On cash-only trips, the household chooses

$$q_C(p) \in \arg\max_{q\geq 0}\{u(q) - pq\}$$

both before and after the restriction. On a SNAP-associated trip, the choices before and after the restriction are

$$q_A^{\text{pre}}(p) \in \arg\max_{q\geq 0}\{u(q) - pq\},$$
$$q_A^{\text{post}}(p) \in \arg\max_{q\geq 0}\{u(q) - pq - \Psi(q)\}.$$

The restriction changes the choice problem only on SNAP-associated trips by adding the planning cost $\Psi(q)$. We specify isoelastic marginal utility and a quadratic planning cost,

$$u'(q) = q^{-1/\varepsilon_S}, \qquad \Psi(q) = \frac{\kappa}{2}q^2.$$

The restriction therefore leaves the willingness-to-pay schedule unchanged but raises the marginal cost of purchasing soda on an affected trip from p to $p+\kappa q$. Before the restriction, both trip types have elasticity ε_S, and their aggregate elasticity is also ε_S:

$$\varepsilon_S^{\text{pre}} = (1-\varphi)\varepsilon_S + \varphi\varepsilon_S = \varepsilon_S.$$

Normalize the calibration price to $p = 1$, so the pre-restriction quantity on each trip is one, and let $x = q_A^{\text{post}}(1)$. The observed aggregate decline identifies x from

$$1 - d = (1 - \varphi) + \varphi x, \qquad x = 1 - \frac{d}{\varphi}.$$

The post-restriction first-order condition is

$$q^{-1/\varepsilon_S} = p + \kappa q.$$

Evaluating it at $(p, q) = (1, x)$ identifies the cost parameter:

$$\kappa = \frac{x^{-1/\varepsilon_S} - 1}{x}.$$

Implicit differentiation of the first-order condition gives the affected-trip elasticity

$$\varepsilon_{S,A}^{\text{post}}(p) = -\frac{p}{q_A^{\text{post}}(p)} \frac{\partial q_A^{\text{post}}(p)}{\partial p} = \frac{p}{q^{-1/\varepsilon_S}/\varepsilon_S + \kappa q}.$$

At the normalized price, $p = 1$ and $q = x$. Aggregating affected and unaffected trips gives

$$\varepsilon_S^{\text{post}} = \frac{(1 - \varphi)\varepsilon_S + \varphi x \varepsilon_{S,A}^{\text{post}}}{1 - d}.$$

The first term is the unchanged cash-only component; the second is the post-restriction SNAP-associated component. The calibration uses $d = 0.12$, $\varepsilon_S = 1.20$, and $\varphi = 0.40$, where φ is the pre-policy share of in-store soda spending on receipts using SNAP, which we use as a proxy for the corresponding all-channel share (Table 30). These inputs give

$$x = 0.70, \qquad \kappa = 0.4945, \qquad \varepsilon_S^{\text{post}} = 1.0349, \qquad \Delta\varepsilon_S = -0.1651.$$

I.5 Implied price responses

For each model, we insert $\Delta\varepsilon_S$ into Equation (8) and then apply Equation (7).

Table 31: Calibrated elasticity and target-price responses

	Post-ban ε_S	$\Delta\varepsilon_S$	Target-price response
Model (1): quantity license	$\simeq 1.20$	< 0.007	-0.06% to $+0.19\%$
Model (2): per-unit premium	1.064	-0.136	$+0.25\%$ to $+0.83\%$
Model (3): trip rerouting cost	1.035	-0.165	$+0.31\%$ to $+0.97\%$

Notes: Price-response ranges use $R = 0.10$–0.15, $\varepsilon_N = 0.8$–1.4, $r = 2$, and the central pricing multiplier $\rho L = 0.18$.

The models differ materially in their prediction for SNAP households' own elasticity, but SNAP

households account for a limited share of total market volume. Consequently, each model implies a target-price response below one percent over the benchmark national range. Even when SNAP participants make up 30% of a market's customers, the largest response is 1.53% under the central pricing multiplier and 2.38% under the upper multiplier. Menu costs and uniform chain pricing can further reduce the observed adjustment. The calibration therefore rationalizes the small retail-price responses in Section 6.1.1.

J Welfare estimates

The calculations below ask what the annual effects would be if the estimated responses applied to SNAP spending nationwide. Let $B = \$92.7205$ billion denote the total SNAP benefits paid in FY2024. The approximate \$340 monthly benefit is used only when translating a national amount into a household equivalent; the FY2024 SNAP Quality Control mean is \$341 (Ward and Monkovic, 2026). Because the corrective benchmark below was reported for sugar-sweetened beverages, the soda calculation is our primary estimate. We extrapolate that benchmark to candy and energy drinks using their nutrient equivalents, report those extrapolations separately, and add them to soda for the restricted-basket total.

J.1 Corrective-benefit mapping

Step 1: estimate the calorie reduction per SNAP dollar. We build calories and total sugar from package sizes and nutrient densities, as described in Appendix G, and aggregate them to the household-month. For each nutrient, we use the same stacked difference-in-differences specification as for spending. The outcomes measure purchased nutrients.

Let $\widehat{\beta}_K$ denote the estimated change in soda calories per household-month and let $\bar{G}_{\text{obs}}$ denote average pre-policy SNAP-funded grocery spending among treated households in the same regression sample.[34] We calculate

$$\theta_K \equiv \frac{|\widehat{\beta}_K|}{\bar{G}_{\text{obs}}}.$$

[34] If any purchased item line in a category–household–month lacks complete nutrition information, we exclude that household-month from both reported nutrient regressions rather than count the missing nutrients as zero. We do not recode its spending as missing in the underlying data. For this welfare calculation, however, we average grocery spending only over the household-months retained in the nutrient regression, so the calorie estimate and its spending denominator use the same sample. Across the monthly category audits, the maximum incomplete shares are 2.04 percent of item lines, 0.05 percent of spending, and 6.46 percent of category-household-months. Because Table 7 uses the broader category sample, its coefficient need not equal the complete-nutrition estimate used here.

We use this ratio because $\widehat{\beta}_K$ need not equal the actual reduction in household calorie purchases when some receipts are missing. Dividing by $\bar{G}_{\text{obs}}$ corrects for this problem under proportional underreporting because the numerator and denominator are measured from the same receipt sample. If receipts capture the same fraction p of purchased calories and SNAP-funded grocery spending, both $|\widehat{\beta}_K|$ and $\bar{G}_{\text{obs}}$ are scaled by p. The common reporting rate therefore cancels in θ_K. We then multiply this calorie reduction per SNAP dollar by total SNAP benefits from administrative data to obtain the national calorie reduction. The current estimates are $\widehat{\beta}_K = -732.200$ calories per household-month (SE = 222.992) and $\bar{G}_{\text{obs}} = \$196.211$, giving

$$\theta_K = 3.731692 \text{ calories per SNAP dollar.}$$

Step 2: scale nationally and express the reduction in equivalent ounces. Multiplying θ_K by B, the administrative total for all FY2024 SNAP benefits paid, applies the estimated calorie reduction per SNAP dollar to SNAP spending nationwide. Because the corrective values used below are measured per ounce, we then divide by 140 calories per 12 fluid ounces from the Coca-Cola Original nutrition label and obtain

$$\Delta Q_s = \frac{3.731692}{140/12} B = 29.6575 \text{ billion ounces per year.}$$

This is a calorie-equivalent conversion, not a claim that every affected beverage has the same calorie density. In particular, diet soda carries essentially no calories, so its lost purchases add nothing to ΔQ_s and receive no corrective credit, although they remain in the spending estimates.

Step 3: apply the Allcott et al. corrective benchmark. Allcott et al. (2019) use Nielsen Homescan household purchases as their empirical quantity measure. They report a response-weighted marginal internality of 0.93 cents per SSB ounce and a health-system externality of 0.85 cents per ounce. Following their purchase-based convention, we apply the sum, 1.78 cents per ounce, to the estimated reduction in regular-cola-equivalent purchased ounces. We treat the reported values as 2016 dollars, the final year of their 2006–2016 estimation sample, and do not add the redistributive adjustment that enters their optimal-tax formula. Converting both components to 2024 dollars using annual CPI-U (U.S. Bureau of Labor Statistics, 2024),

$$\gamma_{2024} = 0.0093(313.689/240.007), \qquad e_{2024} = 0.0085(313.689/240.007),$$
$$I_s = \gamma_{2024}\Delta Q_s = \$360.490 \text{ million}, \qquad E_s = e_{2024}\Delta Q_s = \$329.480 \text{ million}.$$

These are gross soda corrective benefits, before subtracting any private costs. They do not subtract health effects from substitution toward other goods. The national calculation assumes that the calorie reduction per SNAP dollar estimated in our sample applies to SNAP spending nationwide and that the Allcott et al. per-ounce values apply to SNAP households. Because Allcott et al. (2019) estimate a larger bias for low-income households (1.10 rather than the population-average 0.93 cents per ounce), using the population-average value for SNAP households is conservative. Using total sugar rather than calories to construct soda-equivalent ounces gives a similar gross corrective benefit of approximately $645 million.

Step 4: extrapolate the SSB benchmark to candy and energy drinks. Allcott et al. (2019) do not estimate category-specific corrective values for either product, so we report both calculations as extrapolations. Using the same reduction-per-SNAP-dollar normalization as for soda, the candy regression gives 0.427813 grams of total sugar per SNAP dollar and the energy-drink regression gives 0.193757 calories per SNAP dollar. Under the assumptions that each gram of sugar content (for candy) or calorie (for energy drinks) is nutritionally equivalent across these product categories, we convert candy purchases into soda ounces using 39 grams of total sugar per 12 ounces of Coca-Cola and energy drink purchases into soda ounces using 140 calories per 12 ounces:

$$\Delta Q_c = \frac{0.427813}{39/12} B = 12.2052 \text{ billion equivalent ounces per year,}$$
$$\Delta Q_e = \frac{0.193757}{140/12} B = 1.5399 \text{ billion equivalent ounces per year.}$$

We apply the same CPI-adjusted γ_{2024} and e_{2024} from Step 3, yielding $148.4 million of internality benefits and $135.6 million of externality benefits for candy, and $18.7 million and $17.1 million, respectively, for energy drinks. The restricted basket total adds these two extrapolations to the direct soda calculation.

Table 32: Components of the welfare calculation

Category	I	E	$I+E$	Model (1) gain	Model (2) gain
	Millions of 2024 dollars per year				
Soda (primary)	360.5	329.5	690.0	37.4–50.3	71.0
Candy (extrapolated)	148.4	135.6	283.9	—	—
Energy drinks (extrapolated)	18.7	17.1	35.8	—	—
Restricted-basket total	527.6	482.2	1,009.7	59.5–108.8	109.9–154.9
Effect on broad-basket considering substitution	627.8	573.8	1,201.6	59.5–108.8	109.9–154.9

Notes: I and E use the CPI-adjusted SSB values from Allcott et al. (2019), as described above. The candy and energy-drink rows are extrapolations. Model-based gains are calculated only for soda and the combined restricted basket, which is why separate gains are not reported for candy and energy drinks. Soda uses $|\varepsilon_S| = 1.20$; the basket range uses $|\varepsilon| \in [1.0, 1.373]$. Totals are calculated before rounding. The broad-basket analysis changes only the corrective-benefit calculation; model-based gains are held at the restricted-basket values.

Accounting for substitution across products. The calculation above counts nutritional changes in soda, candy, and energy drinks. To account for changes in nearby products, we make two adjustments. We replace the soda- and energy-drink calorie effects with the effect on calories from all beverages, and replace the effect on candy sugar with the effect on sugar from all packaged snacks. The first adjustment includes substitution toward juice and other beverages; the second includes changes in savory snacks, frozen desserts, and other non-candy snacks. Using the same national scaling and nutrient-equivalent conversions, the beverage and packaged-snack estimates imply corrective benefits of $747.8 million and $453.7 million, respectively, or $1.20 billion in total. Adding the private gains estimated below gives total gains of $1.26–$1.36 billion per year.

J.2 Model-based private gains

Appendix I defines the behavioral models (Models (1) and (2)) and shows how each model maps the observed purchase decline and pre-ban elasticity into a cash-equivalent demand curve. Here, we use those mappings to calculate the corresponding private gains. Model (3) does not by itself specify demand under an equivalent cash transfer. Using it for welfare would require an additional estimate of how much households would buy under cash. One could, for example, use the difference between spending responses to SNAP and cash – estimated for food by Hastings and Shapiro (2018) and for product categories by Cox and Harris-Lagoudakis (2023) – to infer the gap between pre-ban SNAP purchases and cash-equivalent purchases. We do not add that calibration here.

Let p_0 denote the actual pre-ban price. We normalize prices by p_0, so $p = 1$, represent

pre-ban spending $S_0 = p_0 Q_0$ by $q_0 = S_0$, and set $q_1 = (1-d)q_0$. Given cash-equivalent elasticity η and cash optimum q^U, inverse cash-equivalent demand is

$$P^U(q) = \left(\frac{q}{q^U}\right)^{-1/\eta}, \qquad P^U(q^U) = 1.$$

Model (2). Appendix I.4.2 implies

$$q^U = q_1, \qquad (1-\lambda)^{\varepsilon(1-\lambda)} = 1-d, \qquad \eta = \varepsilon(1-\lambda).$$

Because post-ban purchases equal the cash optimum, we approximate the private gain by the area between price and the inverse cash-equivalent demand curve. This consumer-surplus measure is exact under quasilinear preferences and provides a local approximation when income effects are small, which is plausible given the small share of income spent on restricted goods:

$$M_2 = \int_{q_1}^{q_0} \left[1 - P^U(q)\right] \, dq.$$

Model (1). Appendix I.4.1 implies

$$\eta = \varepsilon, \qquad \ell = \frac{1}{1-d(1-s_g)}, \qquad q^U = \frac{q_0}{\ell}, \qquad q^U - q_1 = d\, s_g q_0.$$

The restriction therefore moves purchases slightly below the cash optimum. Using the same consumer-surplus approximation, we calculate

$$M_1 = \int_{q^U}^{q_0} \left[1 - P^U(q)\right] \, dq - \int_{q_1}^{q^U} \left[P^U(q) - 1\right] \, dq.$$

To express these areas in annual dollars, we set $S_0 = B\rho_{\text{all}}$. This gives $S_0 = \$9.967$ billion for soda and $S_0 = \$22.383$ billion for the restricted basket. The corresponding declines are $d = 12\%$ and $d = 10.84\%$. We use $\varepsilon = 1.20$ for soda and report $\varepsilon \in \{1.0, 1.20, 1.373\}$ for the restricted basket. Table 33 reports the resulting gains.

J.3 Updating SNAP's MVPF

This subsection first constructs the welfare and fiscal inputs generated by the purchase restrictions and then combines them with estimate of SNAP's baseline MVPF from the literature. Let $S^{\text{gross}} = I + X$ be the gross corrective benefit, consisting of the internality correction I and the health-system externality reduction X. Let M denote the private gain from reversing label-induced excess purchases, R the additional sales tax revenue, and

Table 33: Model-based private gains on a common elasticity grid

	$\varepsilon = 1.00$	$\varepsilon = 1.20$	$\varepsilon = 1.373$
Restricted basket, model (2)	154.93	127.01	109.92
Restricted basket, model (1)	80.92–108.79	67.81–91.18	59.47–79.98
Soda, model (2)	—	71.04	—
Soda, model (1)	—	37.36–50.29	—

Notes: Millions of 2024 dollars per year. Model-(1) ranges vary $s_g \in [0.10, 0.20]$; higher shares give lower gains. For a given elasticity, model (2)'s gain is larger because it treats the post-restriction quantity as the cash optimum; model (1) subtracts the loss from moving slightly below that optimum. Soda is reported at the paper's pre-ban elasticity input of 1.20. The basket columns are elasticity sensitivities. The MVPF calculation collapses the two-model results to a single private-gain input by averaging the lowest and highest estimates for each policy scope: \$54.20 million for soda and \$107.20 million for the restricted basket. All entries are reported to two decimal places.

$X_G \in [0, X]$ the share of medical savings accruing to the government. Finally, let ΔE denote implementation, administration, and enforcement cost, H unmeasured household hassle or stigma costs, and C_R retailer compliance costs.

Sales tax revenue is paid by households and received by the government, so it reduces willingness to pay and net government cost by the same amount. Public medical savings are treated analogously. The restriction's contribution to aggregate willingness to pay and net government cost is therefore

$$\begin{aligned} \Delta W &= S^{\text{gross}} + M - X_G - R - H - C_R, \\ \Delta E - \Delta C &= \Delta E - R - X_G. \end{aligned} \tag{9}$$

Here $\Delta C = R + X_G$ denotes government savings and additional revenue. The accounting identity

$$\Delta W - (\Delta E - \Delta C) = S^{\text{gross}} + M - H - C_R - \Delta E \tag{10}$$

shows that sales tax revenue and the public-private allocation of medical savings cancel from unweighted, consolidated surplus.

Table 34 reports the benchmark inputs. We include the observed sales tax transfer and set $\Delta E = X_G = H = C_R = 0$ because these terms are not measured in our data. These normalizations provide a transparent benchmark rather than an assertion that the omitted costs or public medical savings are zero.

Table 34: Inputs to the update of SNAP's MVPF

Policy scope	S^{gross}	M	R	ΔW	$\Delta E - \Delta C$
	Millions of 2024 dollars per year				
Soda only	689.97	54.20	100.75	643.42	−100.75
Restricted basket	1,009.74	107.20	176.42	940.52	−176.42

Notes: The restricted basket contains soda, candy, and energy drinks. $S^{\text{gross}} = I + X$ is the gross corrective benefit, M is the model-based private gain, and R is additional sales tax revenue. The benchmark sets $\Delta E = X_G = H = C_R = 0$, so $\Delta W = S^{\text{gross}} + M - R$ and $\Delta E - \Delta C = -R$. To collapse the two behavioral models to one private-gain input, M is the midpoint of the lowest and highest estimates across the models and reported sensitivity values: \$54.20 million for soda and \$107.20 million for the restricted basket. Calculations use the exact unrounded components from Tables 32 and 33; all entries are reported to two decimal places.

To update the program-level MVPF, let $B = \$92.7205$ billion be FY2024 SNAP benefit spending, m_0 a literature estimate of conventional SNAP's MVPF, and g_0 the associated net government cost per dollar of benefits. Baseline net government cost is $g_0 B$, and baseline willingness to pay is $m_0 g_0 B$. After adding the restriction, SNAP's MVPF is

$$m_{\text{post}} = \frac{m_0 g_0 B + \Delta W}{g_0 B + (\Delta E - \Delta C)}. \tag{11}$$

The incremental change relative to the literature baseline is

$$m_{\text{post}} - m_0 = \frac{\Delta W - m_0(\Delta E - \Delta C)}{g_0 B + (\Delta E - \Delta C)}. \tag{12}$$

Table 14 applies these equations to the literature benchmarks described in the main text. In the benchmark calculation, $\Delta W > 0$ and $\Delta E - \Delta C < 0$, so both terms in the numerator of equation (12) raise SNAP's MVPF. Alternative assumptions about administration, compliance, hassle, or medical savings enter through equations (9) and (12).