Promotions measured with error: Implications for health policy and firm behavior*

Emilio Gutierrez† Adrian Rubli‡ Jose Tudon§

October 18, 2023

Abstract

We estimate price elasticities and promotion effects across many consumer goods, leveraging nationwide retail data with observable promotional discounts from purchase receipts. We contrast price elasticities with and without promotions, and document demand shifts due to promotions' non-price attributes. We repeat this exercise using a traditional, algorithm-driven approach that imputes promotions whenever prices decrease, leading to underestimated promotion effects and overestimated price effects. We emphasize the importance of these distinctions through several policy-relevant findings: reduced frequency of promotions following "sin food" taxes, similar promotion elasticities for healthy and unhealthy products, and higher sensitivity to promotions among high-BMI consumers.

Keywords: price promotions, retail industry, unhealthy demand products, high-BMI consumer demand.

JEL: D22, L11, D12, I18.

*David Ontaneda and David Sosa provided excellent research assistance. The authors acknowledge support from the Asociación Mexicana de Cultura, AC. Data from Kantar Worldpanel made available via a prior data-sharing contract with ITAM. All remaining errors are our own.

†Instituto Tecnológico Autónomo de México (ITAM), Department of Economics. Avenida Camino a Santa Teresa 930. Heroes de Padierna, Mexico City CDMX 10700, Mexico. Email: emilio.gutierrez@itam.mx.

‡ITAM, Department of Business Administration. Email: adrian.rubli@itam.mx.

§ITAM, Department of Business Administration. Email: jtudon@itam.mx.

1 Introduction

This paper studies price and promotion effects in the retail grocery industry and analyzes how promotion effects relate to the healthiness of products, demographics, and how firms' use of promotions reacts to product taxation. For the analysis, we use purchase receipts data from the Kantar Worldpanel consumer panel for Mexico, which gives us a comprehensive picture of the national retail industry. Our data has the advantage of providing direct observations of promotional discounts, which is lacking in similar data widely used by industry practitioners and by researchers in economics and marketing (e.g., Nielsen RMS scanner data). We provide insights on how firms use promotions and how consumers react heterogeneously to them, and we caution against the usage of promotions measured with error, which would be the case of data with unobservable promotions. Because we estimate price and promotion elasticities across a wide range of consumer products, our findings are generalizable, and not specific to particular sets of products, stores, or retailers, as in most of the literature (Hitsch, Hortaçsu and Lin, 2021).

In the first part of the paper, we estimate the demands of products from 800 different brands, sold in 46 chain stores, in 57 cities. Throughout, we use industry-standard empirical strategies to identify price and promotion elasticities (Hitsch, Hortaçsu and Lin, 2021); namely, we leverage our detailed data and employ a rich set of fixed effects. Overall, we find elastic demand curves and significant promotion effects, even after accounting for price effects. On average, promotions increase demand by 21 to 27%.

In the second part of the paper, we repeat our estimations as if promotions data were unobservable: we use a state-of-the-art algorithm developed by Hitsch, Hortaçsu and Lin (2021) to identify promotional discounts from dips in prices. We find similar price elasticities of demands, but, due to (non-classical) measurement error in promotions, we find null effects of the non-price attributes of promotions when these are not directly observed. In light of these results, we argue that direct data on promotions are crucial if we are interested in analyzing promotion elasticities.[1]

Finally, in the third part of the paper, we analyze promotion effects through different lenses, we discuss policy implications, and note that the analysis is only possible when measuring promotions with (relative) precision. We first present a

[1] For a review of measurement error in binary variables, see Aigner (1973), who also shows we have attenuation bias.

breakdown of price and promotion elasticities by the "healthiness" of food products (Hut and Oster, 2022), in order to provide insights into the heterogeneous effects of prices and promotions. We find unhealthy products have a slightly higher price elasticity of demand than healthy ones, and both types of products have similar and large promotion elasticities.

We then study how firms change their promotions and price strategies after the nationwide introduction of a tax on sugary drinks and high-calorie foods. Using cleaning products as controls, an event study approach reveals that this "sin food" tax not only increased prices, but also decreased the frequency of promotions for affected products.

Moreover, we analyze the relationship between the household's body-mass index and price and promotion effects. We find that price elasticity is constant with respect to body-mass index. However, high body-mass index households exhibit much higher promotion elasticities.

These findings have practical policy implications. Policymakers who are concerned about public health might take note that promotion regulation is an important tool at their disposal, because demands of unhealthy products are indeed affected by promotions, firms do adjust promotion strategies when facing taxation, and people with high body-mass index, who are the arguably policy-relevant target, are more susceptible to promotions.

This paper is related to a literature that analyzes price and promotion elasticities across different markets and products. Bijmolt, Van Heerde and Pieters (2005) and Tellis (1988) offer a meta-analysis. We follow closely the empirical strategy of Hitsch, Hortaçsu and Lin (2021), but we augment it by allowing for promotions to shift demand slopes and not only intercepts. Importantly, we also use data where we observe promotions directly from purchase receipts.

This paper also relates to a literature that ties price elasticities and promotion effects to market characteristics (Bolton, 1989), consumer demographics and competition (Boatwright, Dhar and Rossi, 2004; Hoch et al., 1995), and product categories Narasimhan, Neslin and Sen (1996). While most of the literature studies price effects, we focus on how promotion effects relate to product characteristics, demographics, and how promotions are affected by taxation.

Finally, this paper relates to a literature that studies various ways in which retailers and policymakers may affect consumption and health outcomes like obesity. From the perspective of retailers, many potential levers have been analyzed that may help reduce obesity, such as product offerings (Allcott et al., 2019), placement of products (Vogel et al., 2021; Shaw et al., 2020; Brimblecombe et al., 2020), promotional materials, and prices and discounts (Alsubhi et al., 2023; Ni Mhurchu et al., 2010; Glanz, Bader and Iyer, 2012). In terms of public policies, there is a growing literature devoted to understanding the appropriateness, optimal design, and impact of different levers aimed at improving consumers' diets, such as taxes on sugar-sweetened beverages (Cawley, 2015; Cawley and Frisvold, 2023; Allcott, Lockwood and Taubinsky, 2019*b*,*a*), front-package labeling (Barahona, Otero and Otero, 2023; Alé-Chilet and Moshary, 2022; Araya et al., 2022), mandatory calorie posting (Bollinger, Leslie and Sorensen, 2011), advertising bans (Dubois, Griffith and O'Connell, 2018), and information campaigns (Ippolito and Mathios, 1990, 1995).

2 Data description

We use data from the Kantar Worldpanel (KW) consumer panel for Mexico during the years 2010 to 2015. KW is a firm specializing in data collection for marketing purposes. This dataset records all retail purchases of selected households over a period of time. The data are not necessarily representative of all household purchases or of the average consumption of a Mexican household. However, the KW consumer panel does seem to line up with self-reported information in national surveys (Aguilar, Gutierrez and Seira, 2021).

Participating households are visited by surveyors on a weekly basis, who record information on purchases directly from receipts. For each transaction, we observe the product name (including brand, size, and category), the amount purchased, the price paid, whether there was a promotion and of what type, the city where the transaction was made, and the name of the retailer. Due to this structure, we can only observe the price of a product in a given city-retailer-week if at least one household made a purchase.

We do not observe the precise location of each retailer, and hence refer to them as chain stores instead of individual locations. For instance, although there are five Wal-Mart stores in the Tampico metropolitan area, the data do not specify at which of these five stores a household made purchases. In the full dataset, we observe 73 cities and 83 different retail chain stores.

2.1 Restricting the sample

We follow stardard practices (see, e.g., Hitsch, Hortaçsu and Lin (2021)). We aggregate products at the brand level. For each description, we group products that differ only in their presentation size. Thus, for example, a two liter bottle of Coca-Cola and a six-pack of Coca-Cola cans would be classified as the same brand, while Coca-Cola Light would be a different brand.

We calculate equivalent units —such as grams or liters— to standardize quantities and calculate the appropriate weighted price aggregation. For each transaction in a city-week, we calculate equivalent units as $q \times \text{unit volume}/N_{HH}$, where q is the purchased quantity and N_{HH} is the number of households that made at least one purchase in that city-week. We then aggregate to the brand level, which effectively corresponds to an inverse probability weighting aggregation.

Therefore, all quantities are measured in equivalent units, and price differences across brands reflect actual price-per-unit differences and not the distribution of purchases across presentation sizes. From the original 70,142 products (UPCs) in the full data, this exercise results in 8,700 different brands.

We further make sample restrictions to avoid estimating demand models on data with many missing values, because prices are only observed conditional on a purchase. First, we identify the top 2,000 brands in the data based on their total revenue at the national level, which allows us to screen out uncommon brands. Second, for our main specification, we exclude the city-chain store combinations for which we observe less than 156 weeks with non-zero purchases, corresponding to half of our total time frame of 312 weeks. As a robustness test, we show results for all of the top 2,000 brands in the data, regardless of how many non-zero purchases weeks they have.

Table 1 shows a comparison between the full data and the estimating sample. Our sample restrictions leave us with 800 different brands, 57 cities, and 46 retail chains. Because the panel is not balanced, it corresponds to 4,703 brand-city pairs and a total of 8,583 brand-city-chain store combinations. Our demand models will be estimated for each of these brand-city pairs. As a robustness check, we also present results on the full dataset without these sample restrictions. Table 2 further shows descriptive statistics and a difference in means test comparing the full and restricted datasets.

TABLE 1: Comparison between full data and estimating sample

	Full data	Estimating sample
Brands	8,700	800
Cities	73	57
Chain stores	83	46
Brand-city pairs	128,067	4,703
Brand-city-chain combinations	450,535	8,583
Non-zero purchase weeks	6,880,710	1,913,867
Share non-zero purchase weeks	0.05	0.71

Notes: This table shows descriptives of the full KW dataset from 2010-2015 and the estimating sample for our empirical exercise.

TABLE 2: Descriptive statistics

	Full data				Estimating sample				Diff. in means
	Mean	p25	p50	p75	Mean	p25	p50	p75	
Price	0.81	0.02	0.04	0.09	0.41	0.02	0.03	0.07	0.56***
Promotion (in KW data)	0.18	0.00	0.00	0.00	0.23	0.00	0.00	0.00	-0.07***
Promotion (algorithm)	0.25	0.00	0.00	1.00	0.45	0.00	0.00	1.00	-0.27***
Quantity	25.48	1.49	6.06	20.00	30.23	2.25	8.70	27.65	-6.59***
Households	189.89	41.00	96.00	237.00	242.69	59.00	147.00	463.00	-73.14***
Chain stores	3.52	1.00	2.00	4.00	15.17	11.00	15.00	19.00	-11.97***
Non-zero purchase weeks	105.19	26.00	78.00	168.00	233.37	190.00	229.00	278.00	-177.57***
Observations	6,880,710				1,913,867				

Notes: This table shows descriptive statistics in the full data and the estimating sample. We show the mean, median, and 25th and 75th percentiles. The price is the per unit price in Mexican pesos. Promotions are measured both directly in the KW data and using the algorithm developed in Hitsch, Hortaçsu and Lin (2021). The quantity is measured in equivalent units. All price, promotion and quantity statistics are calculated at the brand-city-chain store level. We also show the number of households making purchases per city, the number of retailers (chain stores) per brand-city pairs, and the number of non-zero purchase weeks (out of a total of 312 weeks) at the brand-city-chain store level. The last column shows a difference-in-means test. *** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$.

2.2 Identifying promotions

We identify promotions in two different ways. First, we take promotions as recorded directly in the KW panel. Although all types of promotions are recorded, we restrict our attention to those coded as price discounts, which are the most common promotion type by far. Because promotion information is taken directly from purchase receipts, we are certain we have no false positives: every promotion in the data indeed corresponds to a purchase with a price discount.[2] The only potential

[2]For the Nielsen Homescan data in the US, it has been shown there are discrepancies in how prices are recorded (Einav, Leibtag and Nevo, 2010), mainly because this variable is taken directly from store-level data and not from individual purchases. Because our KW panel identifies promotions from information provided in receipts, this type of discrepancy is not an issue for us.

measurement error might be the case of a retailer that did not print out promotion information on the receipt. To minimize this potential measurement error, we only consider a promotion as being present when at least 20% of purchases record it.[3]

The second way in which we record promotions follows the literature in identifying price discounts from large dips in a product's price in the data (Warner and Barsky, 1995; Hendel and Nevo, 2013). Specifically, we replicate the (state-of-the-art) algorithm developed in Hitsch, Hortaçsu and Lin (2021): we identify a base price for each brand-city-chain store combination and classify prices as discounted if they are at least 5% lower than their base price.[4]

Table 2 shows summary statistics for prices, promotions as observed, algorithmically predicted promotions, and quantities, in both the full data and the estimating sample. As detailed below, the algorithm predicts a larger number of promotion events than were actually recorded in the data.

2.3 Predictive power of algorithmic identification of promotions

In this empirical context, the algorithmic identification of prices results in significant measurement error. Overall, the results show the algorithm correctly identifies the promotional status of 69% of prices. Looking specifically at promotional prices recorded directly in the data, this algorithm correctly identifies 33% of them. Additionally, only 23% of the promotions identified by the algorithm correspond to a recorded price promotion. Table 3 shows additional performance statistics, which are common in the machine learning literature. In sum, the algorithm tends to over-predict promotions (many false negatives and few true positives).

Figure 1 shows the relationship between the probability of a product being on promotion and prices and quantities. The algorithm consistently predicts a larger share of products on promotion than the actual incidence in the data. As a function of

[3]In the appendix, we present results when defining a promotion as being present if at least one purchase records it. Both definitions give very similar results.

[4]As Hitsch, Hortaçsu and Lin (2021) explain: "The algorithm distinguishes between regular and promoted prices based on the frequently observed saw-tooth pattern in a store-level time series of prices whereby prices alternate between periods with (almost) constant regular price levels and shorter periods with temporarily reduced price levels. We perform this classification separately for each store-UPC pair. We assume that weeks without sales are non-promoted weeks, which is justified by the large sales spikes that are frequently observed in promoted weeks. Hence, we impute the missing prices using the predictions of the current regular (non-promoted) price levels based on the price classification algorithm" (page 294).

TABLE 3: Predictive performance of algorithmic identification of promotions

Diagnostic statistic	Definition	Value
Accuracy	(TN+TP)/Total	0.69
Sensitivity	TP/(TP+FN)	0.33
Specificity	FN/(FN+FP)	0.76
Precision	TP/(TP+FP)	0.23
Negative predictive value	TN/(TN+FN)	0.84
F Score	2*TP/(2*TP+FP+FN)	0.27
Mean prediction error	1-(TN+TP)/Total	0.31
Pseudo R2	(TN+TP-M)/(Total-M)	0.17

Notes: TN=True Negatives, TP=True Positives, FN=False Negatives, FP=False Positives, Total=Total number of observations, and M=Most frequent outcome number, which in this case is 0.

prices, the likelihood of a product being on promotion is actually, roughly U-shaped. However, the algorithm predicts that the likelihood of promotions is decreasing in prices. On the other hand, as a function of quantities, the likelihood of promotions is roughly increasing. However, the algorithm finds a similar relationship, which, if anything, turns decreasing at higher levels of quantity.

FIGURE 1: Likelihood of promotions with respect to (log) price and quantity

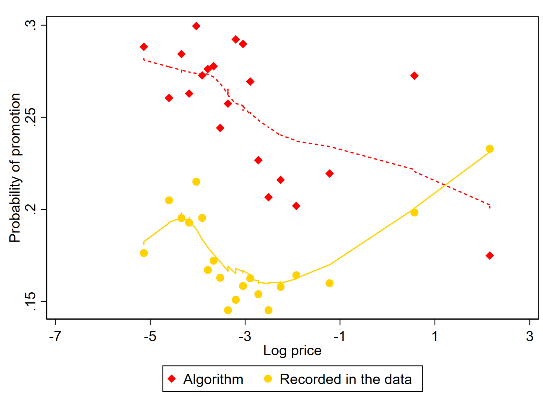

(A) Price

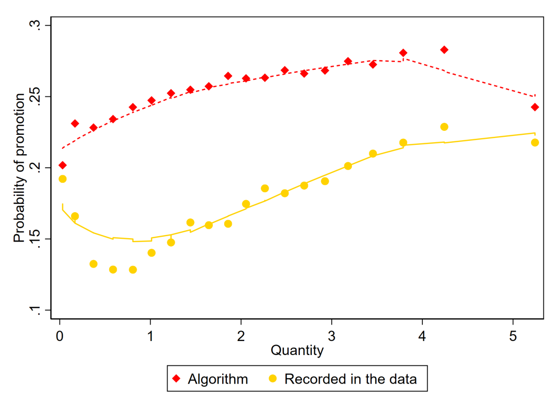

(B) Quantity

Notes: Lowess and binscatter plots of observed likelihood of promotions.

3 Empirical strategy

In this section, we detail our empirical strategy to estimate the demand for each product at the store level. First, we use the Kantar Worldpanel data to estimate price and promotion elasticities. Second, we estimate demands using the algorithmic identification of promotions, and we compare results. We find price elasticities are comparable between both methods, but the algorithmic identification results in much smaller estimates for promotions elasticities, due to measurement error.

3.1 Demand specification

We consider a log-linear demand model for brand j, sold in store s, at city c, in week t:

$$\begin{aligned}\log(1+q_{jsct}) &= \alpha_{jsc} + \tau_{jct} + \beta_{jc}\log(p_{jsct}) + \gamma_{jc}D_{jsct} + \delta_{jc}(\log(p_{jsct}) \times D_{jsct}) \\ &\quad + \sum_{k\in\mathcal{K}_{jsct}} \eta_{jck}\log(p_{ksct}) + \sum_{k\in\mathcal{K}_{jsct}} \theta_{jck}D_{ksct} + \varepsilon_{jsct}, \end{aligned} \tag{1}$$

where q_{jsct} are sold quantities, p_{jsct} are prices, and D_{jsct} indicates if there was a price promotion. We closely follow the specification used by Hitsch, Hortaçsu and Lin (2021), including the use of $1+q_{jsct}$ to accommodate zero sales in the data, but we augment their model by allowing for interactions between prices and promotions. Thus, we not only allow for promotions to change levels of demand curves, but also their slopes. As a robustness check, we consider a Poisson regression as an alternative to $\log(1+q)$ and we find very similar results, which are available upon request.

The model also includes chain store fixed effects α_{jsc} and time period effects τ_{jct}. We further control for the prices and for promotion indicators of at most five competitors based on revenue and identified at the product-chain-city level; we denote by $\mathcal{K}_{jsct}$ the set of competitors.[5] Finally, ε_{jsct} is the error term.

The model implies that the predicted price elasticity of demand is approximately β_{jc} when no promotions are present, because

$$\left.\frac{\partial q_{jsct}}{\partial p_{jsct}}\frac{p_{jsct}}{q_{jsct}}\right|_{D_{jsct}=0} = \beta_{jc}\frac{q_{jsct}+1}{q_{jsct}} \approx \beta_{jc},$$

[5] Products are classified into 84 broad categories made up by different brands (e.g., flavored carbonated beverages). Competitors may differ across locations (for instance, the top five competitors

and similarly for the discrete effect of the promotion, γ_{jc}. Throughout we simply use the term elasticities, and, in the case of promotions, we refer to the discrete change of promotions and its effect on quantities.

The coefficients in equation (1) represent elasticities as long as we can give a causal interpretation to these estimates. Prices and promotions are endogenous if retailers use unobservable information (to the researcher) to determine them. A common solution is to use instrumental variables (Berry, 1994), however, finding valid and strong instruments is specially challenging in our context where we have a multiplicity of markets (Rossi, 2014). Indeed, we are interested in the price and promotion effects which are plausibly available to sophisticated firms, and these effects are largely estimated without instrumental variable strategies (Hitsch, Hortaçsu and Lin, 2021).

Our approach for dealing with endogeneity relies on a rich set of fixed effects. Given our estimation at the brand-city level, chain store fixed effects imply that we rely only on price and promotion variation within each retailer. If chain stores differ in their responsiveness to unobservable shocks or in how they collect and use demand information for a given brand in a given city, these fixed effects would capture those differences so long as they are constant over time. Likewise, our time effects account for demand shocks and consumer trends at a very narrowly defined level (month-year). Therefore, the only remaining threat to a causal interpretation would be unobservable factors related to demand that retailers respond to and that vary at the brand-city-store-month level.[6]

In practice, we estimate the regressions separately for each brand-city pair. We select the 2,000 brands with the highest revenue nationwide and the corresponding city-store pairs for which we observe prices (i.e., non-zero purchases) for at least half of the weeks in our data. Hence, our main estimates below are the results of 4,703 separate regressions, corresponding to 800 different brands across (at most) 57 different cities.

of Coca-Cola in Wal-mart stores in Guadalajara may differ from those of Coca-Cola in Soriana stores in Guadalajara).

[6]Unobservable factors that are related to supply (e.g., costs and markups) would not bias our interpretation of elasticities if these factors do not affect demand.

FIGURE 2: Estimated price and promotion elasticities

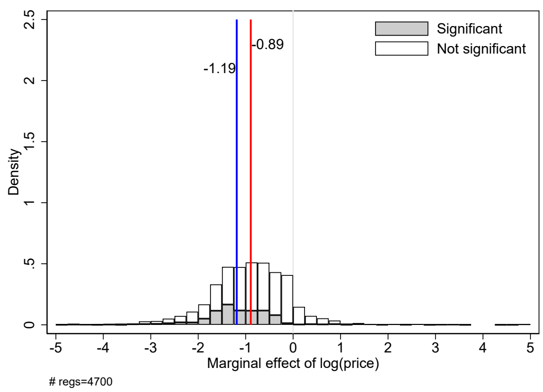

(A) Price elasticity

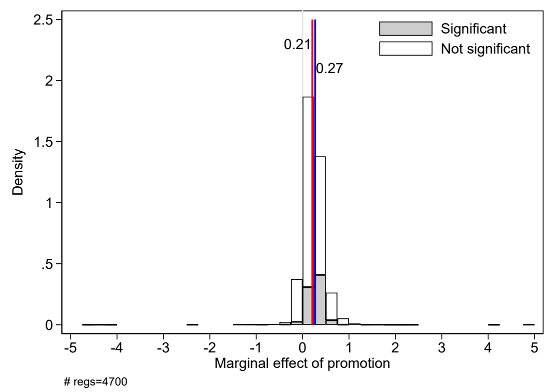

(B) Elasticity of promotions

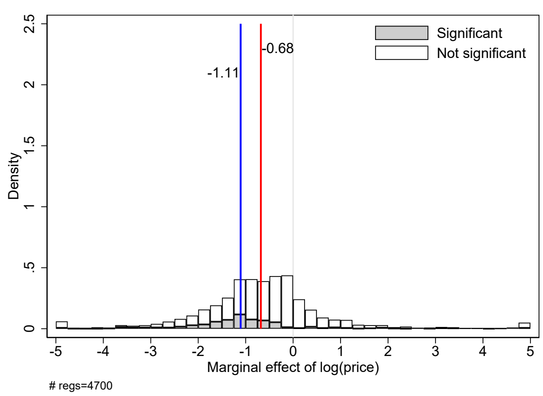

(C) Price elasticity with promotion

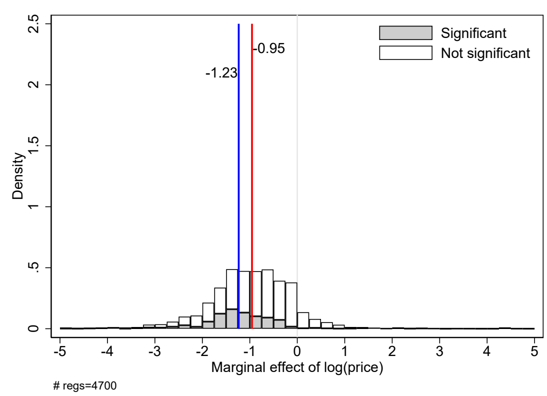

(D) Price elasticity without promotion

Notes: This figure shows price and promotion elasticities from 4,703 demand models as specified in equation (1) for our restricted sample. We identify promotions as those that are recorded from purchase receipts directly in the data. We distinguish between estimates that are significant at the 95% and those that are not. The red vertical line depicts the median of the estimates, regardless of significance; the blue line depicts the median of the estimates significant at the 95%. We winsorize the estimates at -5 and 5 for clarity. The top two plots show price elasticities and price promotion elasticities. The bottom plots show price elasticities conditional on having or not having a price promotion.

3.2 Results

We start by showing our estimates using information on promotions recorded directly in the data from purchase receipts. Figure 2 plots the distribution of the estimates for own-price and promotion elasticities. Estimates are weighted by brand revenue. We distinguish between estimates that are significant at the 95% confidence level (shaded area) and those that are not (unshaded area). Both top panels in Figure 2 evaluate elasticities at means.

The median price elasticity is -0.89, while the median promotion elasticity is 0.21, as depicted by the red vertical lines. If we focus solely on the significant coefficients, the median price elasticity is -1.19 and the median promotion elasticity is 0.27, as depicted by the blue lines. In words, demand increases roughly by 27% in the presence of promotions, all else equal, on average. Note that some price elasticities (16.7% of our estimates) are non-negative, although the majority are insignificant. Likewise, 5.9% of our estimated promotion elasticities are non-positive but mostly insignificant.

The bottom two plots in Figure 2 display estimated own-price elasticities conditional on a promotion and conditional on no promotion, respectively. The median price elasticity with a promotion is -0.68, while that without a promotion is -0.95. These estimates suggest that consumers are more sensitive to prices when there are no promotions, which might happen if, for instance, promotional discounts change the composition of consumers.

As a robustness check, we estimate a similar specification to equation (1) with the addition of a control for the base price at the brand-city-chain store level. That is, we control for discount depth. The distribution of the estimates are shown in Figure 7 in the appendix. The median across the distributions is very similar to our previous estimates, and we also find consumers are more elastic in the presence of promotions. Figure 8 in the appendix shows an additional exercise where we exclude the interaction term between the log price and the promotion indicator in our demand models, and we obtain similar results.

In sum, we find relatively large effects of promotions, even accounting for price effects, which are in line with the literature (Anderson and Fox, 2019). Promotions increase demand by around 21 to 27% and demands are overall elastic.

3.3 Comparison with algorithmic identification of promotions

We now turn to using the traditional, algorithm-driven approach for identifying promotions and comparing with the results presented in Section 3.2. Figure 3 plots the corresponding distribution of the estimates for own-price and promotion elasticities, weighted again by brand revenue and distinguishing between estimates by significance at the 95% confidence level. The median own-price elasticity is -0.85, similar to what was shown when using promotions directly from the data in Figure 2.

However, the median elasticity of promotions is now an order of magnitude smaller at 0.03 when identifying promotions from the algorithm. If we focus on the significant coefficients, the median price and promotion elasticities are -1.13 and 0.08. The bottom two plots in Figure 3 show the estimated price elasticities, conditional on the existence or absence of an algorithmic-identified promotion. Here, we also find a higher sensitivity to prices when there are no promotions than when there are price discounts (median of -1.36 versus -0.34).

In Figure 9 in the appendix, we present results from a robustness check where we add the base price at the brand-city-chain store level as a control in equation (1). This exercise results in very similar estimates across the board. We include an additional robustness check in Figure 10 in the appendix, where we exclude the interaction term between the log price and the promotion indicator in our demand models, obtaining again similar results.

For better assessing the differences in our estimates across methods, we first calculate the difference at the brand-city level between the estimated price elasticities with and without a promotion, and then plot these differences for either approach for identifying promotions. Figure 11 in the appendix shows the distribution of these differences. The plot on the left relies on estimates that identify promotions directly from the data, and the one on the right from the price algorithm. The red vertical line depicts the median of all estimates, while the blue line depicts the median of significant estimates at 95% level.

The distribution of the difference in price elasticities with and without promotions identified directly in the data is fairly symmetric and almost centered at zero, with a median of -0.10 (Figure 11, left panel). In contrast, the distribution of the corresponding difference for promotions identified via the traditional algorithmic approach has a majority of its mass to the left of zero, with a median difference of -0.86 (Figure 11, right panel). Restricting to significant estimates yields a similarly

FIGURE 3: Estimated price and promotion elasticities, using promotions identified with the price algorithm

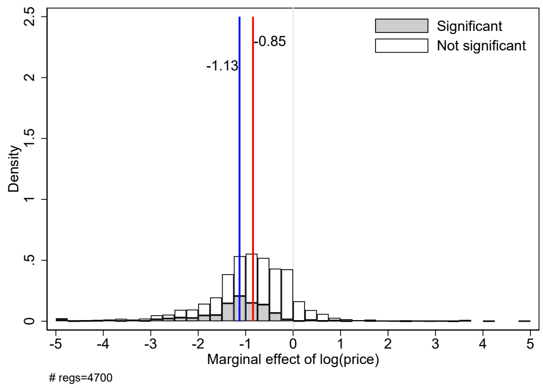

(A) Price elasticity

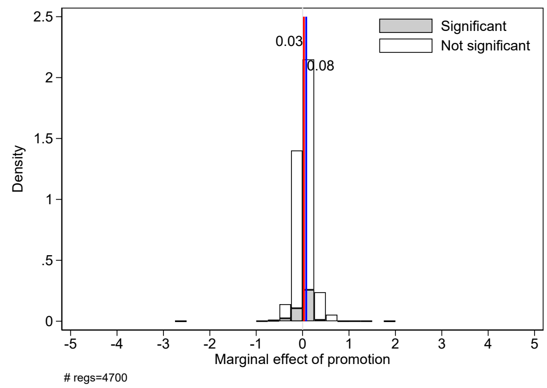

(B) Elasticity of promotions

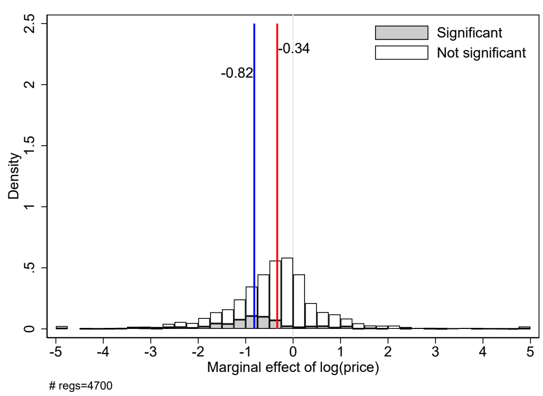

(C) Price elasticity with promotion

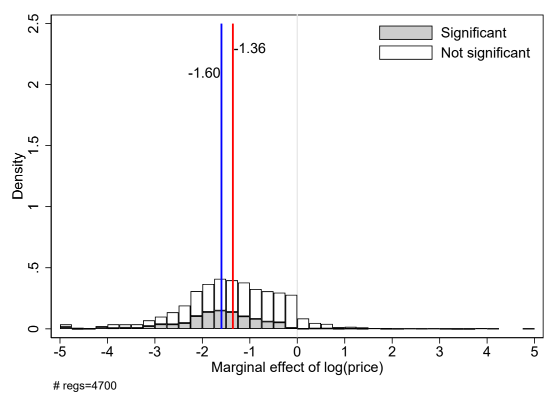

(D) Price elasticity without promotion

Notes: This figure shows price and promotion elasticities from 4,703 demand models as specified in equation (1) for our restricted sample. We identify promotions from prices that are at least 5% below the base price, as determined by the algorithm in Hitsch, Hortaçsu and Lin (2021). We distinguish between estimates that are significant at the 95% and those that are not. The red vertical line depicts the median of the estimates, regardless of significance; the blue line depicts the median of the estimates significant at the 95%. We winsorize the estimates at -5 and 5 for clarity. The top two plots show price elasticities and price promotion elasticities at means. The bottom plots show price elasticities conditional on having or not having a price promotion.

centered distribution for promotions identified directly in the data, with a median difference of -0.20. However, using promotions identified with the algorithm, we find a median difference of -0.63 and a large mass to the left of zero. In sum, the algorithmic identification of promotions consistently biases the price sensitivity of consumers in the presence of promotions.

Figure 12 in the appendix shows the empirical cumulative distribution functions for the estimated price elasticities with and without promotions. The plot in the top left panel considers estimates from promotions recorded directly in the data, while the plot in the top right uses promotions identified via the algorithm. Various non-parametric tests of equality of continuous distributions reject that the distribution with and without promotions are equal in both plots. However, the distribution of elasticities conditional on no promotion only stochastically dominates the distribution of elasticities with promotions for the algorithm-identified promotions. The bottom plots in Figure 12 are restricted to significant estimates and are qualitatively similar.

In summary, we find that estimated own-price elasticities conditional on a promotion are fairly similar to those conditional on no promotion when using directly observed promotions data. However, once we implement a traditional algorithmic approach, we find significant differences, with larger price elasticities when conditioning on no promotion than those that are obtained conditional on having a promotion. Because the algorithm tends to predict promotions when there are none, price elasticity estimates must be higher as they account for consumers not buying more in the presence of a false positive. In other words, measurement error of promotions yields both lower promotion effects and larger baseline price elasticities.

4 Implications

The richness of the data allow us to analyze price and promotion effects under different lenses. In particular, analyzing promotion effects would yield biased results if done under measurement error. In particular, policy might be misguided by the erroneous finding that price elasticities are all that matter, because promotion effects are null. For instance, the current discussion and design of tax policies aimed at discouraging the consumption of sugary drinks may benefit from considering regulating promotions, as they have a large effect by themselves on consumption.

In this section, we present a breakdown of price and promotion elasticities by the "healthiness" of the products. We also analyze how promotions and prices change after the introduction of a tax on sugary drinks and high-calorie products. Finally, we break down price and promotion elasticities by the body-mass index of the household head. Throughout we note how the analysis is only appropriate when promotions are measured without error.

4.1 Comparison between healthy and unhealthy products

We analyze the price and promotion elasticities of healthy and unhealthy products. For this purpose, we use the classification of Hut and Oster (2022) where they measured the healthfulness of groups of products based on a survey implemented to 17 primary care doctors in the United States. We thus classify products dichotomously into being healthy or unhealthy according to whether they are above or below the median healthfulness index in Hut and Oster (2022). This analysis provides insight into the heterogeneity of price and promotion effects as a function of healthfulness.

Figure 4 plots the distribution for price and promotion elasticities. On the top left panel we show the distribution for price elasticities of unhealthy products. The median price elasticity is -1.15. The top right panel then shows the distribution for price elasticities of healthy products. The median price elasticity is -1.06. The bottom two panels of Figure 4 show the promotion elasticity for unhealthy (left panel) and healthy (right panel) products. Their median elasticities are 0.20 and 0.19, respectively.

Figure 13 in the appendix shows the distribution for price and promotion elasticities using the promotions identified through the algorithm developed by Hitsch, Hortaçsu and Lin (2021). The median price elasticity for unhealthy products is -0.99, and for healthy products is -0.99. Additionally, the promotion elasticities for unhealthy products is 0.03 and for healthy products is 0.03.

These plots show that healthy products' demands have a lower price elasticity, reflecting their lower sensitivity to price changes. On the other hand, promotion elasticities seem very similar across both categories. Finally, it's worth noting that the algorithm seems to capture correctly the price elasticities, but it greatly underestimates the promotion elasticities due to measurement error.

We confirm that a "sin food" tax would have large effects on unhealthy products' demands. But policymakers might also find interesting that promotions are another tool at their disposal to tackle public health concerns. Indeed, taxing or regulating

FIGURE 4: Estimated price and promotion elasticities for healthy and unhealthy products

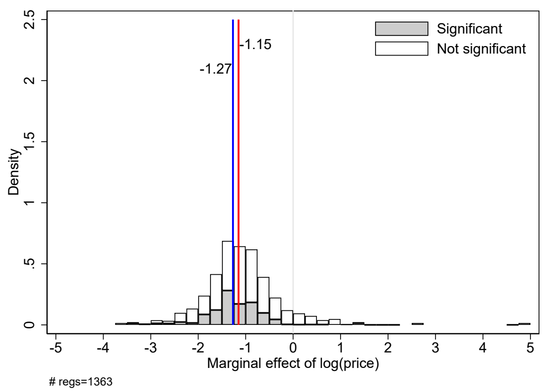

(A) Price elasticity of unhealthy products

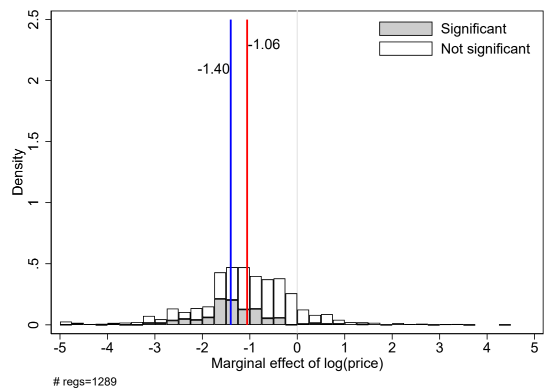

(B) Price elasticity of healthy products

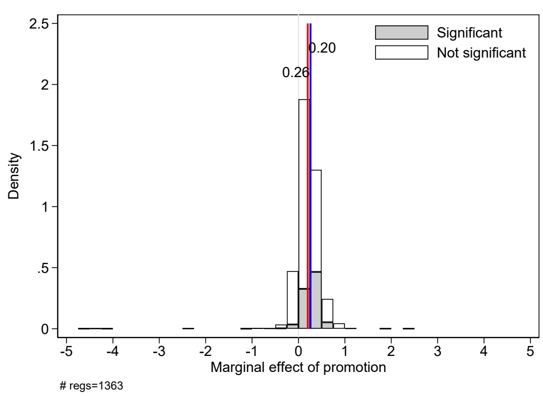

(C) Promotion effects of unhealthy products

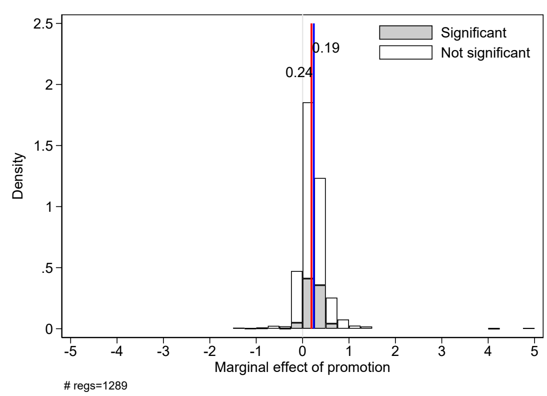

(D) Promotion effects of healthy products

Notes: This figure shows price and promotion elasticities at means from 4,703 demand models as specified in equation (1) for our restricted sample while separating healthy and unhealthy products. We distinguish between estimates that are significant at the 95% and those that are not. The red vertical line depicts the median of the estimates, regardless of significance; the blue line depicts the median of the estimates significant at the 95%. We winsorize the estimates at -5 and 5 for clarity.

promotions, on top of prices, might be effective on the margin.

4.2 The effect of a sin food tax on promotions

In this section we focus on the effect of a nationwide tax introduced on January 1, 2014 for drinks with added sugar and for high-calorie foods. We are interested in observing if there is an effect of this tax on the periodicity of promotions. This is relevant for public policy, because if promotions respond to taxes, these effects should also be considered in the cost-benefit analysis of these kind of policies.

To analyze the impact of the taxes on promotions, we generate a Laspeyres index where we fix the average consumption basket for each city-seller in 2013. Following Aguilar, Gutierrez and Seira (2021) we construct this index to avoid confounding the observed changes in the presence of promotions for a given set of products with the shift in consumption towards products more frequently on promotion (regardless of the tax). For each city-seller we are interested in three categories: taxed food and beverages, untaxed food and beverages, and the control group (we decided a suitable control would be cleaning products). For each category we calculate the share of sales that came from products in promotion.

Using an event study framework we analyze the evolution of prices and promotions in Figure 5. Collapsing to quarterly data, we regress our measures of promotion frequency and price on leads and lags of the introduction of the tax (first quarter of 2014), with brand-city-store fixed effects. We weight our estimates by revenue and cluster the standard errors by brand-city-store. Coefficients on the leads and lags are interpreted relative to the excluded category (the quarter prior to the tax). We estimate regressions separately for taxed and untaxed products, using cleaning products as a control. The plots in Figure 5 show estimates for a three-year window centered around the tax.

The introduction of the tax resulted in a sharp reduction of promotions for taxed products, but not for untaxed products. The share of taxed products bought with a promotion declined around 3% in the first quarter of 2014 and it declined up to 10% after 6 quarters of the tax introduction (although estimates after a year of the tax introduction are noisy). Meanwhile, untaxed products did not observe any significant change in the share of products bought with a promotion for the year after the tax introduction (a slight reduction appears to happen after a year, but estimates are also noisy).

The bottom two panels of Figure 5 show prices of taxed products had a sharp increase at the beginning of 2014, while the prices of untaxed products remained unchanged. These price changes are exactly what we would expect, and, reassuringly, are similar to what Aguilar, Gutierrez and Seira (2021) find.

These results show that instruments of public policy like the sin food tax introduced in Mexico in 2014 not only affect prices, but also the frequency of promotions. While analyzing the reasoning behind the decline in promotions is beyond the scope of this paper, we believe the results can inform policy design as they show that there may be some unintended consequences of these public policy actions. Governments may then want to also consider these secondary effects when designing them.

4.3 Heterogenous consumption by BMI level

Thanks to our rich dataset, we can also observe the body-mass index (BMI) of the household head.[7] In our sample the median BMI is 27 and 72% of the households' heads have a BMI greater or equal to 25, which is the overweight definition by the World Health Organization. In addition, 25% of households' heads have a BMI greater or equal to 30, which is the threshold for obesity.

Figure 6 shows the relationship between estimated price and promotion effects and the average BMI at each city-seller pair. We find that the price elasticity is constant with respect to the BMI level. However, the promotion elasticity follows a U-shape where the elasticity is much higher for very high levels of BMI.

These findings have practical policy implications. If people who are overweight are more susceptible to promotions, then a policymaker who is concerned about public health issues, might find it worthwhile to regulate promotions on high-calorie products.

5 Concluding remarks

We estimate demands for a wide range of consumer products using nationwide data from Mexico. We find demands elastic to prices, and large promotion effects. We note how these large promotion effects disappear if promotions are not directly observed

[7] The body-mass index of a person is their body mass, measured in kilograms, divided by the square of their body height, measured in meters: kg/m^2.

FIGURE 5: Prices and promotions evolution after "sin food" tax

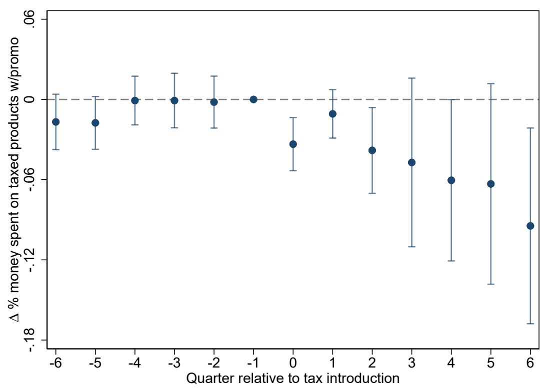

(A) Promotions of taxed products

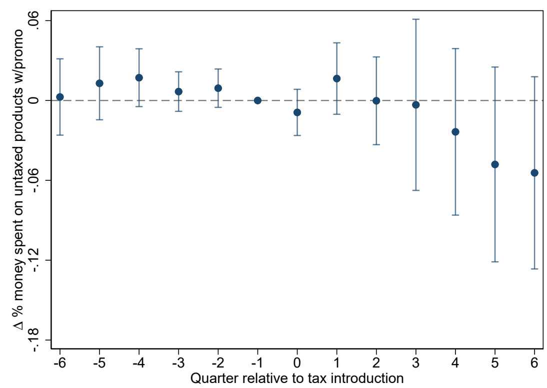

(B) Promotions of untaxed products

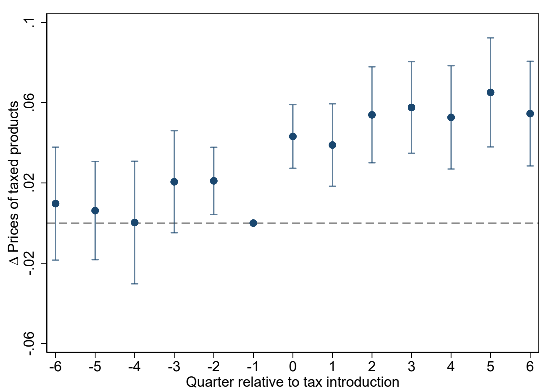

(C) Prices of taxed products

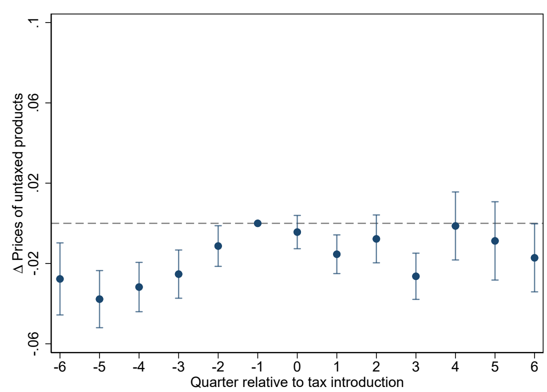

(D) Prices of untaxed products

Notes: This figure shows the evolution of taxed and untaxed products around a tax to sugary drinks and high-calorie foods was introduced in January 2014. We identify promotions as those that are recorded from purchase receipts directly in the data. Coefficients from an event study approach shown, with 95% confidence intervals constructed from standard errors clustered at the brand-city-store level.

FIGURE 6: Estimated price and promotion elasticities by BMI

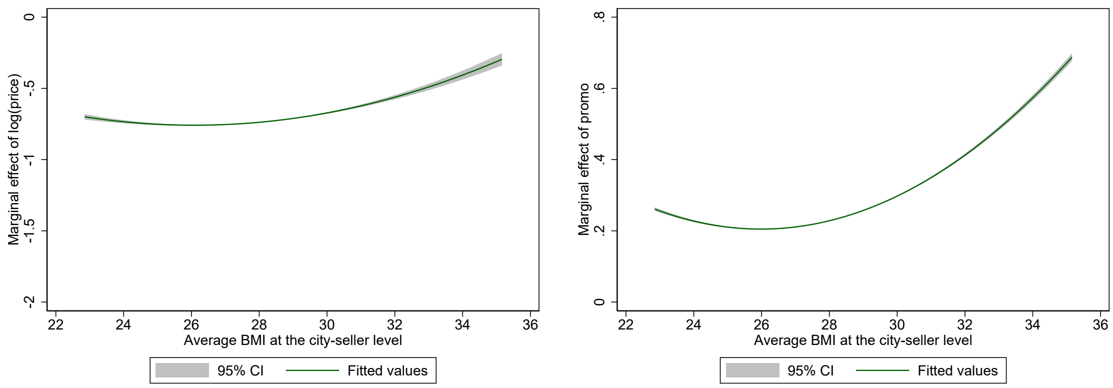

(A) Price elasticity

(B) Elasticity of promotions

Notes: This figure shows price and promotion elasticities at means from 4,703 demand models as specified in equation (1) for our restricted sample. We identify promotions as those that are recorded from purchase receipts directly in the data.

in the data, and are instead predicted using algorithmic strategies.

With estimated promotion effects, we break them down by product healthfulness, by household BMI, and we analyze how firms react to "sin food" taxes by changing the frequency of promotions. We do not find significant heterogeneity in the promotion effects of healthy and unhealthy products. We do find firms use promotions less frequently after sugary drinks and high-calorie products are taxed. Finally, we find that high-BMI consumers are more sensitive to promotions.

From a policy perspective, the results suggest that policymakers could also regulate promotions, on top of prices, if their aim is to reduce high-calorie foods consumption by using economic tools. While "sin food" taxes already decrease the usage of promotions, their disproportionate effect on high-BMI people might warrant a direct regulation.

References

Aguilar, Arturo, Emilio Gutierrez, and Enrique Seira. 2021. "The effectiveness of sin food taxes: Evidence from Mexico." *Journal of Health Economics*, 102455.

Aigner, Dennis J. 1973. "Regression with a binary independent variable subject to errors of observation." *Journal of Econometrics*, 1(1): 49–59.

Alé-Chilet, Jorge, and Sarah Moshary. 2022. "Beyond consumer switching: Supply responses to food packaging and advertising regulations." *Marketing Science*, 41(2): 243–270.

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

Allcott, Hunt, Benjamin B Lockwood, and Dmitry Taubinsky. 2019*b*. "Should we tax sugar-sweetened beverages? An overview of theory and evidence." *Journal of Economic Perspectives*, 33(3): 202–227.

Allcott, Hunt, Rebecca Diamond, Jean-Pierre Dubé, Jessie Handbury, Ilya Rahkovsky, and Molly Schnell. 2019. "Food deserts and the causes of nutritional inequality." *The Quarterly Journal of Economics*, 134(4): 1793–1844.

Alsubhi, Moosa, Miranda Blake, Tan Nguyen, Ishani Majmudar, Marj Moodie, and Jaithri Ananthapavan. 2023. "Consumer willingness to pay for healthier food products: A systematic review." *Obesity Reviews*, 24(1): e13525.

Anderson, Eric T., and Edward J. Fox. 2019. "How price promotions work: A review of practice and theory." In *Handbook of the Economics of Marketing*. Vol. 1, 497–552. Elsevier B.V.

Araya, Sebastian, Andres Elberg, Carlos Noton, and Daniel Schwartz. 2022. "Identifying food labeling effects on consumer behavior." *Marketing Science*, 41(5): 982–1003.

Barahona, Nano, Cristóbal Otero, and Sebastián Otero. 2023. "Equilibrium effects of food labeling policies." *Econometrica*, 91(3): 839–868.

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

Bijmolt, Tammo HA, Harald J Van Heerde, and Rik GM Pieters. 2005. "New empirical generalizations on the determinants of price elasticity." *Journal of marketing research*, 42(2): 141–156.

Boatwright, Peter, Sanjay Dhar, and Peter E. Rossi. 2004. "The Role of Retail Competition, Demographics and Account Retail Strategy as Drivers of Promotional Sensitivity." *Quantitative Marketing and Economics*, 2(2): 169–190.

Bollinger, Bryan, Phillip Leslie, and Alan Sorensen. 2011. "Calorie posting in chain restaurants." *American Economic Journal: Economic Policy*, 3(1): 91–128.

Bolton, Ruth N. 1989. "The relationship between market characteristics and promotional price elasticities." *Marketing Science*, 8(2): 153–169.

Brimblecombe, Julie, Emma McMahon, Megan Ferguson, Khia De Silva, Anna Peeters, Edward Miles, Thomas Wycherley, Leia Minaker, Luke Greenacre, Anthony Gunther, et al. 2020. "Effect of restricted retail merchandising of discretionary food and beverages on population diet: a pragmatic randomised controlled trial." *The Lancet Planetary Health*, 4(10): e463–e473.

Cawley, John. 2015. "An economy of scales: A selective review of obesity's economic causes, consequences, and solutions." *Journal of health economics*, 43: 244–268.

Cawley, John, and David Frisvold. 2023. "Taxes on sugar-sweetened beverages: Political economy, and effects on prices, purchases, and consumption." *Food Policy*, 117: 102441.

Dubois, Pierre, Rachel Griffith, and Martin O'Connell. 2018. "The effects of banning advertising in junk food markets." *The Review of Economic Studies*, 85(1): 396–436.

Einav, Liran, Ephraim Leibtag, and Aviv Nevo. 2010. "Recording discrepancies in Nielsen Homescan data: Are they present and do they matter?" *QME*, 8(2): 207–239.

Glanz, Karen, Michael DM Bader, and Shally Iyer. 2012. "Retail grocery store marketing strategies and obesity: an integrative review." *American journal of preventive medicine*, 42(5): 503–512.

Hendel, Igal, and Aviv Nevo. 2013. "Intertemporal price discrimination in storable goods markets." *American Economic Review*, 103(7): 2722–51.

Hitsch, Günter J, Ali Hortaçsu, and Xiliang Lin. 2021. "Prices and promotions in US retail markets." *Quantitative Marketing and Economics*, 19(3): 289–368.

Hoch, Stephen J, Byung-Do Kim, Alan L Montgomery, and Peter E Rossi. 1995. "Determinants of store-level price elasticity." *Journal of marketing Research*, 32(1): 17–29.

Hut, Stefan, and Emily Oster. 2022. "Changes in household diet: determinants and predictability." *Journal of Public Economics*, 208: 104620.

Ippolito, Pauline M, and Alan D Mathios. 1990. "Information, advertising, and health choices: a study of the cereal market." *Rand J. Econ*, 21: 211–246.

Ippolito, Pauline M, and Alan D Mathios. 1995. "Information and advertising: The case of fat consumption in the United States." *The American Economic Review*, 85(2): 91–95.

Narasimhan, Chakravarthi, Scott A. Neslin, and Subrata K. Sen. 1996. "Promotional Elasticities and Category Characteristics." *Journal of Marketing*, 60(2): 17–30.

Ni Mhurchu, Cliona, Tony Blakely, Yannan Jiang, Helen C Eyles, and Anthony Rodgers. 2010. "Effects of price discounts and tailored nutrition education on supermarket purchases: a randomized controlled trial." *The American journal of clinical nutrition*, 91(3): 736–747.

Rossi, Peter E. 2014. "Invited Paper —Even the Rich Can Make Themselves Poor: A Critical Examination of IV Methods in Marketing Applications." *Marketing Science*, 33(5): 655–672.

Shaw, Sarah C, Georgia Ntani, Janis Baird, and Christina A Vogel. 2020. "A systematic review of the influences of food store product placement on dietary-related outcomes." *Nutrition reviews*, 78(12): 1030–1045.

Tellis, Gerard J. 1988. "The price elasticity of selective demand: A meta-analysis of econometric models of sales." *Journal of marketing research*, 25(4): 331–341.

Vogel, Christina, Sarah Crozier, Daniel Penn-Newman, Kylie Ball, Graham Moon, Joanne Lord, Cyrus Cooper, and Janis Baird. 2021. "Altering product placement to create a healthier layout in supermarkets: Outcomes on store sales, customer purchasing, and diet in a prospective matched controlled cluster study." *PLoS Medicine*, 18(9): e1003729.

Warner, Elizabeth J, and Robert B Barsky. 1995. "The timing and magnitude of retail store markdowns: Evidence from weekends and holidays." *The Quarterly Journal of Economics*, 110(2): 321–352.

A Additional figures

FIGURE 7: Estimated price and promotion elasticities; robustness to controlling for base price

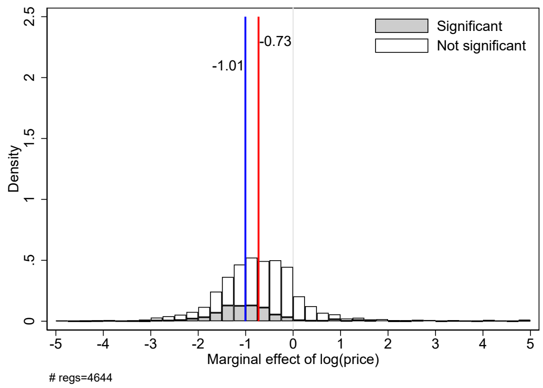

(A) Price elasticity

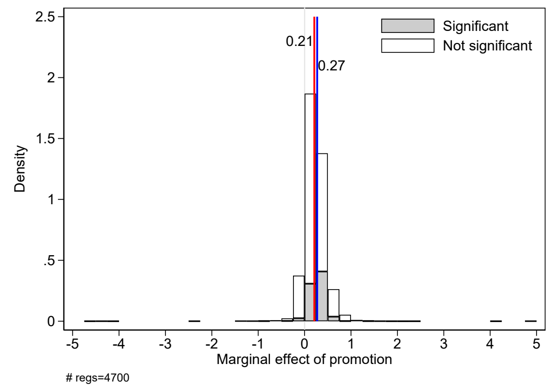

(B) Elasticity of promotions

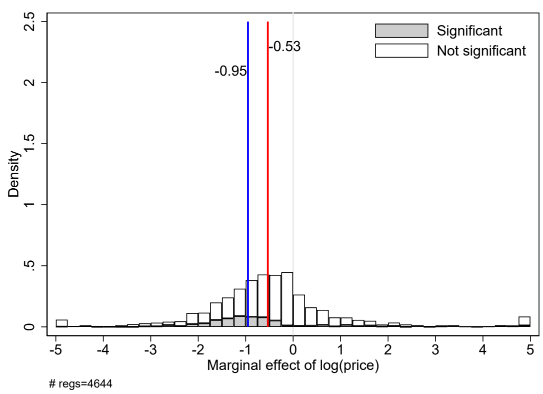

(C) Price elasticity with promotion

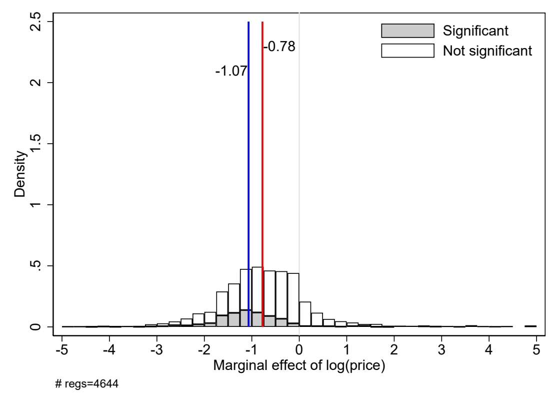

(D) Price elasticity without promotion

Notes: This figure shows price and promotion elasticities from 4,703 demand models as specified in equation 1 for our restricted sample with the addition of a control for the base price at the brand-city-chain store level. We identify promotions as those that are recorded from purchase receipts directly in the data. We distinguish between estimates that are significant at the 95% and those that are not. The red vertical line depicts the median of the estimates, regardless of significance; the blue line depicts the median of the estimates significant at the 95%. We winsorize the estimates at -5 and 5 for clarity. The top two plots show price elasticities and price promotion elasticities at means. The bottom plots show price elasticities conditional on having or not having a price promotion.

FIGURE 8: Estimated price and promotion elasticities; without price and promotion interaction

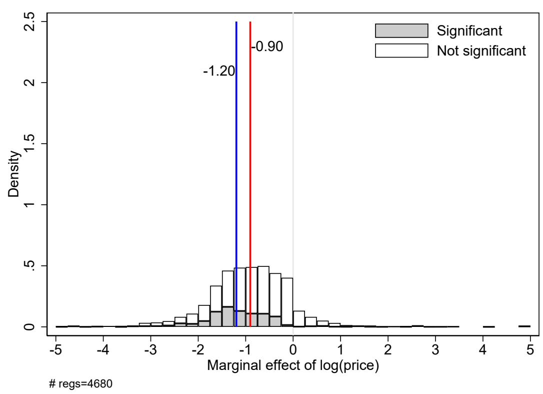

(A) Price elasticity

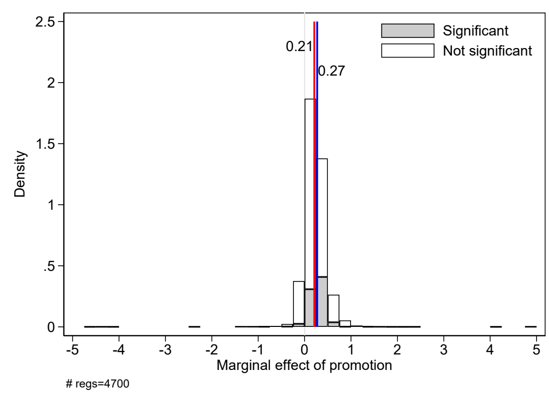

(B) Elasticity of promotions

Notes: This figure shows price and promotion elasticities from 4,703 demand models as specified in equation (1) for our restricted sample excluding the interaction term $\log(p_{jsct}) \times D_{jsct}$. We identify promotions as those that are recorded from purchase receipts directly in the data. We distinguish between estimates that are significant at the 95% and those that are not. The red vertical line depicts the median of the estimates. The blue line depicts the median of significant estimates at 95% level. We winsorize the estimates at -5 and 5 for clarity.

FIGURE 9: Estimated price and promotion elasticities, using promotions identified with the price algorithm; robustness to controlling for base price

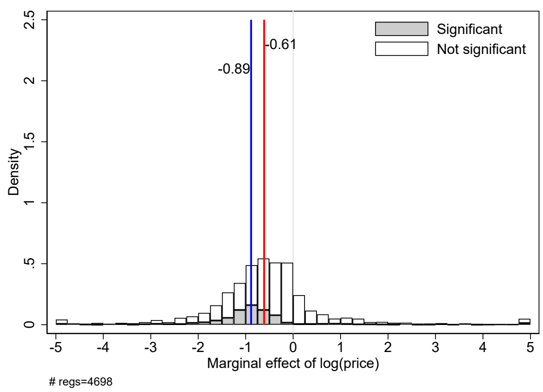

(A) Price elasticity

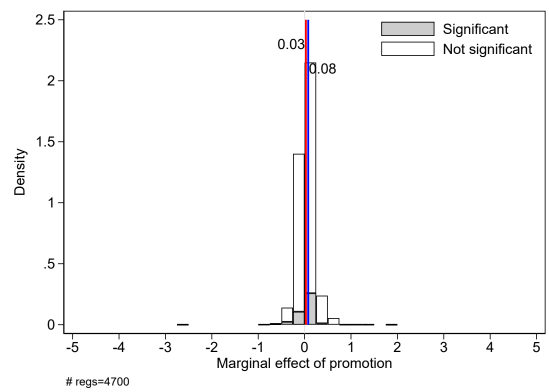

(B) Elasticity of promotions

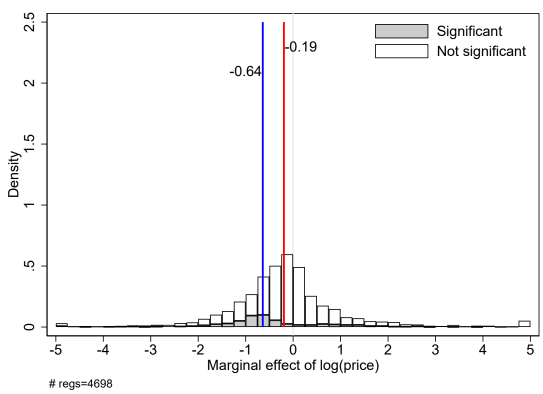

(C) Price elasticity with promotion

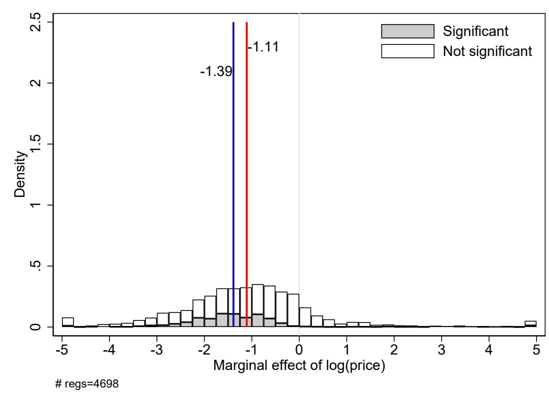

(D) Price elasticity without promotion

Notes: This figure shows price and promotion elasticities from 4,703 demand models as specified in equation 1 for our restricted sample with the addition of a control for the base price at the brand-city-chain store level. We identify promotions from prices that are at least 5% below the base price, as determined by the algorithm in Hitsch, Hortaçsu and Lin (2021). We distinguish between estimates that are significant at the 95% and those that are not. The red vertical line depicts the median of the estimates, regardless of significance; the blue line depicts the median of the estimates significant at the 95%. We winsorize the estimates at -5 and 5 for clarity. The top two plots show price elasticities and price promotion elasticities at means. The bottom plots show price elasticities conditional on having or not having a price promotion.

FIGURE 10: Estimated price and promotion elasticities, using promotions identified with the price algorithm; without price and promotion interaction

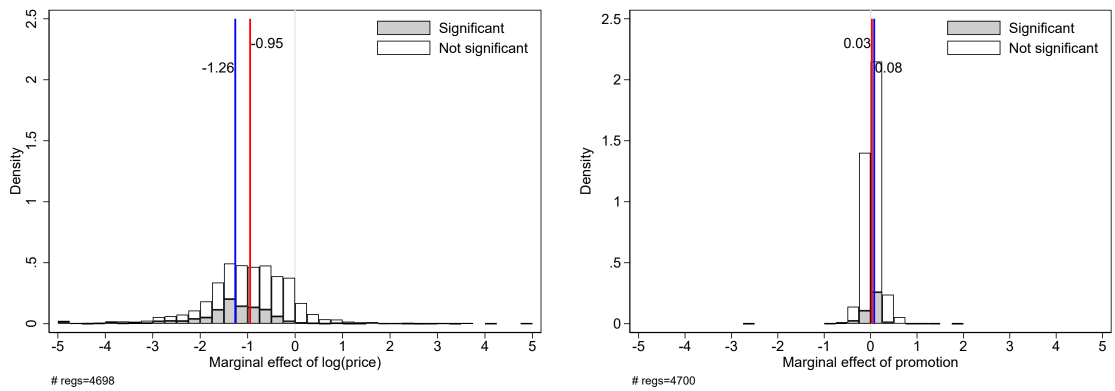

(A) Price elasticity (B) Elasticity of promotions

Notes: This figure shows price and promotion elasticities from 4,703 demand models as specified in equation 1 for our restricted sample excluding the interaction term $\log(p_{jsct}) \times D_{jsct}$. The plots identify promotions from prices that are at least 5% below the base price, as determined by the algorithm in Hitsch, Hortaçsu and Lin (2021). We distinguish between estimates that are significant at the 95% and those that are not. The red vertical line depicts the median of the estimates. The blue line depicts the median of significant estimates at 95% level. We winsorize the estimates at -5 and 5 for clarity.

FIGURE 11: Differences in estimated price elasticities with and without promotions

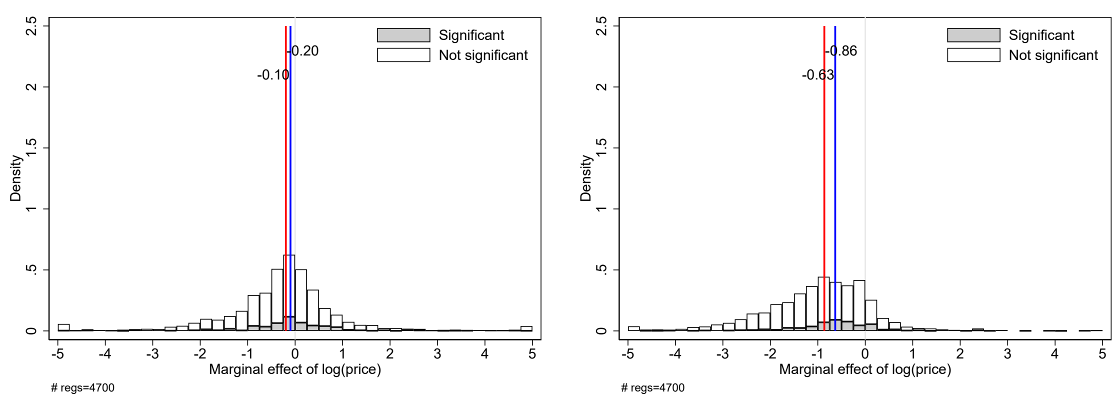

(A) Promotions recorded directly in data (B) Promotions using the price algorithm

Notes: This figure shows the distribution of the differences in price elasticities with and without promotions from 4,703 demand models as specified in equation 1 for our restricted sample. The plots on the left identify promotions as those that are recorded from purchase receipts directly in the data. The plots on the right identify promotions from prices that are at least 5% below the base price, as determined by the algorithm in Hitsch et al. (2019). The top two plots consider all estimates, while those in the bottom restrict to significant estimates (at the 95% level). The red vertical line depicts the median of the estimates, regardless of significance; the blue line depicts the median of the estimates significant at the 95%. We winsorize the estimates at -5 and 5 for clarity.

FIGURE 12: Cumulative distributions of price elasticities with and without promotions

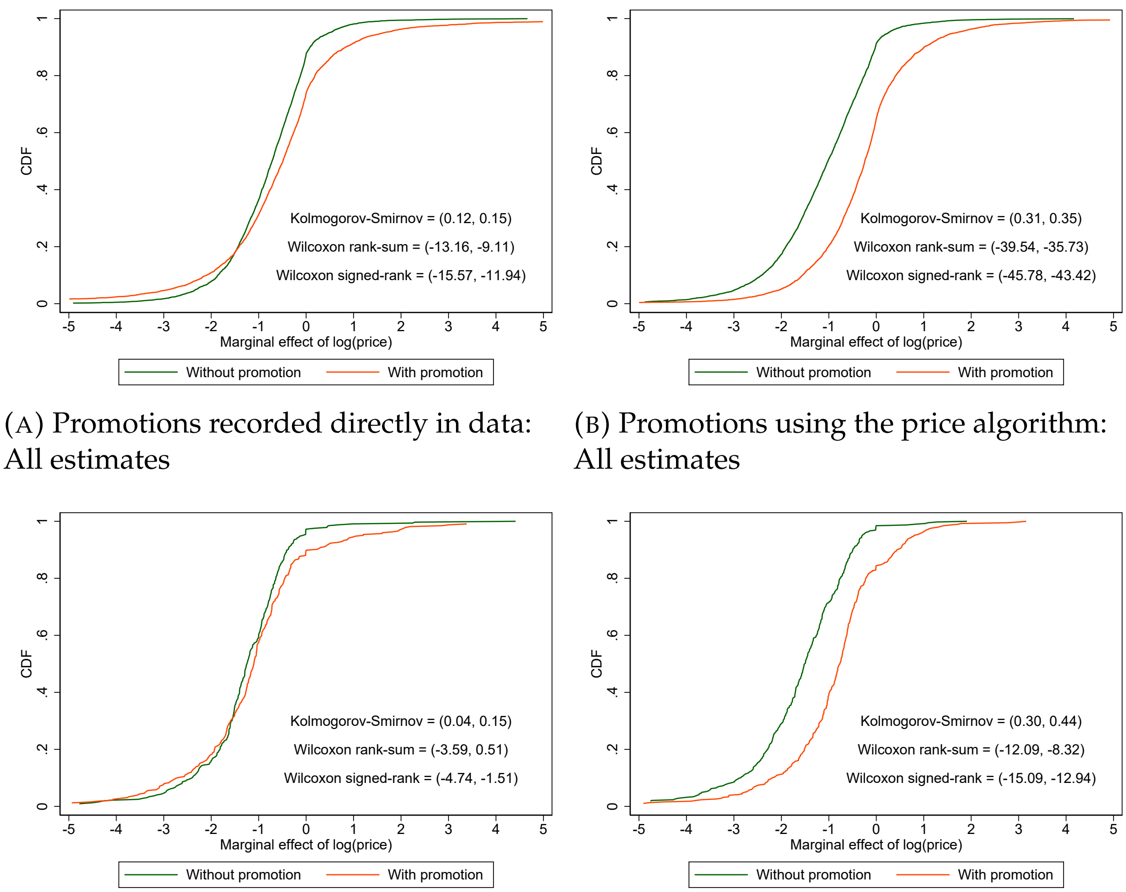

(A) Promotions recorded directly in data: All estimates

(B) Promotions using the price algorithm: All estimates

(C) Promotions recorded directly in data: Significant estimates

(D) Promotions using the price algorithm: Significant estimates

Notes: This figure shows empirical cumulative distribution functions of price elasticities with and without promotions from 4,703 demand models as specified in equation 1 for our restricted sample. The plots on the left identify promotions as those that are recorded from purchase receipts directly in the data. The plots on the right identify promotions from prices that are at least 5% below the base price, as determined by the algorithm in Hitsch, Hortaçsu and Lin (2021). The top two plots consider all estimates, while those in the bottom restrict to significant estimates (at the 95% level). We winsorize the estimates at -5 and 5 for clarity. We show three non-parametric tests of equality of continuous distributions.

FIGURE 13: Estimated price and promotion elasticities for healthy and unhealthy products, using promotions identified with the price algorithm

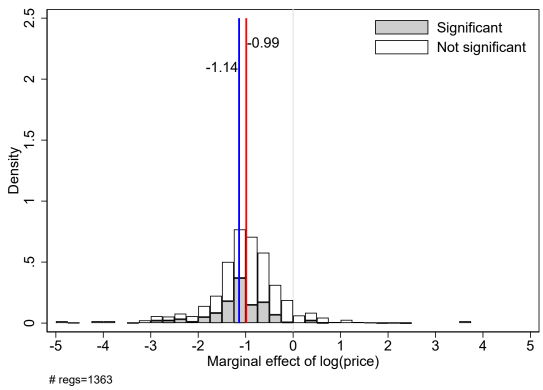

(A) Price elasticity of unhealthy products

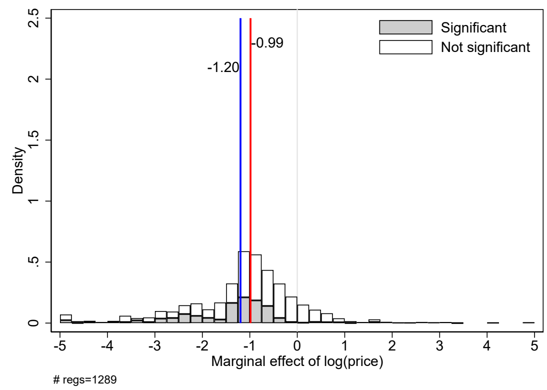

(B) Price elasticity of healthy products

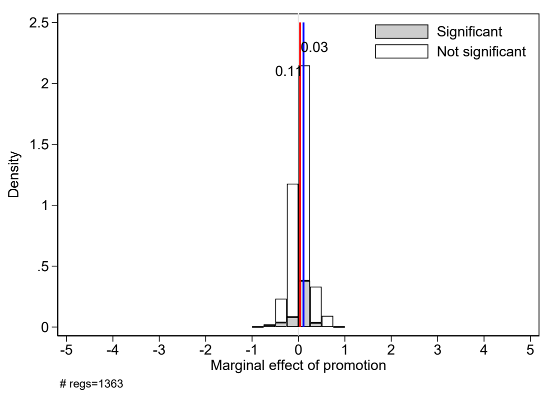

(C) Promotion effects of unhealthy products

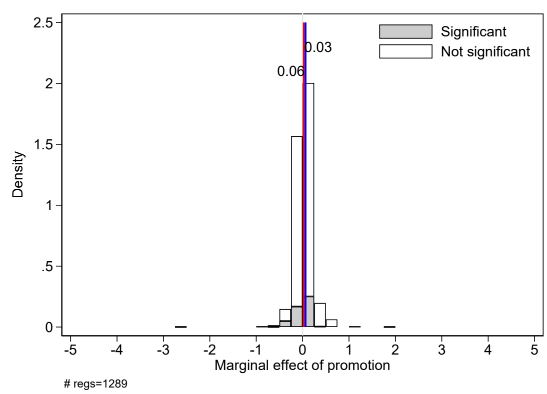

(D) Promotion effects of healthy products

Notes: This figure shows price and promotion elasticities at means from 4,703 demand models as specified in equation 1 for our restricted sample while separating healthy and unhealthy products. We identify promotions from prices that are at least 5% below the base price, as determined by the algorithm in Hitsch, Hortaçsu and Lin (2021). We distinguish between estimates that are significant at the 95% and those that are not. The red vertical line depicts the median of the estimates, regardless of significance; the blue line depicts the median of the estimates significant at the 95%. We winsorize the estimates at -5 and 5 for clarity.