Consumer Inattention and Quantity Surcharging

PRELIMINARY: Please do not circulate

Giovanni Compiani, Chicago Booth
Jean-Pierre Dubé, Chicago Booth and NBER
Joonhwi Joo, UT Dallas Naveen Jindal School of Management*

November 6, 2026

Abstract

We exploit the widespread incidence of *quantity surcharging* in consumer packaged goods to study consumer price inattention. QS generates independent variation in price and size that deviates from the standard quantity discounts. Suggestive descriptive analysis shows less quantity surcharing in states with unit-pricing mandates, and a lower consumer propensity to choose QS'd products when prices are salient due to promotions. We then use the QS variation to identify a nonparametric test for price inattention that is robust to a broad class of search models, finding strong evidence of inattention. Through counterfactual analyses, we quantify the role of price inattention and the extent to which it arises from QS specifically, along with the impact on consumer welfare and retailer revenues.

keywords: quantity surcharging, consumer inattention, search, choice model, nonparametric identification

*Authors contributed equally with names listed in alphabetical order. Researchers' own analyses calculated (or derived) based in part on data from Nielsen Consumer LLC and marketing databases provided through the NielsenIQ Datasets at the Kilts Center for Marketing Data Center at The University of Chicago Booth School of Business. The conclusions drawn from the NielsenIQ data are those of the researchers and do not reflect the views of NielsenIQ. NielsenIQ is not responsible for, had no role in, and was not involved in analyzing and preparing the results reported herein. We are grateful for comments from seminar participants at the 2026 Marketing Science Conference. Compiani and Dubé acknowledge the Kilts Center for Marketing. Dubé also acknowledges the Charles E. Merrill Faculty fund. Contact information: Giovanni Compiani (gio1compiani@gmail.com), Jean-Pierre Dubé (jdube@chicagobooth.edu), Joonhwi Joo (joonhwi.joo@gmail.com).

1 Introduction

An extant literature using primary data collected in situ has found that many consumers do not check grocery shelf prices (Wells and LoSciuto, 1966; Harrell et al., 1976; Dickson and Sawyer, 1990), but has not measured its economic consequences.We exploit the widespread practice in consumer packaged goods (CPG) categories of *quantity surcharging* (QS) as an instrument. QS arises when larger packsizes of a product cost more per unit than smaller packsizes, a price premium customers may be paying due to inattention.[1] We first document the incidence of QS across a large number of CPG categories and provide suggestive evidence that QS is at least partly driven by consumers' price inattention. We then leverage the independent variation in price and size generated by QS to implement a nonparametric test of price inattention, identify the drivers of QS purchases, and quantify the broader impact of inattention on consumer welfare and retailer revenues.

The incidence of QS is economically interesting in its own right since the conventional wisdom from the non-linear pricing literature would predict that the unit price should decline with the quantity purchased due to diminishing marginal utility (e.g., Armstrong, 2016). For decades, industry observers have noted the incidence of QS raises concerns that inattentive consumers may purchase the larger pack size under the incorrect belief it is cheaper: the "larger economy size belief" a phenomenon documented in lab studies (e.g., Granger and Bilson, 1972; Widrick, 1979; Manning et al., 1998). Concerned by the potentially deceptive nature of QS if consumers are indeed inattentive to prices, the European Union mandated unit pricing in member countries in 1979 (European Union, 2026) and, in spite of a lack of a federal mandate in the U.S., nine states have mandatory unit pricing regulations NIST (2015).[2]

In our CPG data, we find evidence of QS on the shelves 15% of store-weeks. Approximately 12% of retailers' CPG sales revenues are generated by QS'd products, with 3% of sales stemming from the potentially avoidable surcharges.[3] Consistent with the "larger economy size belief," we find that QS is empirically rare for a specific brand and that, under rational expectations, a consumer would logically prioritize search for larger packsizes of their preferred brands if their primary goal is to save money. In this regard, QS offers a

[1] According to the National Institute of Standards and Technology (NIST): "Unit pricing is the cost per unit of measure of goods, such as the cost per milliliter or per ounce."(NIST, 2015, p. 2).

[2] Mandated unit pricing regulation has also been implemented in Canada, UK, Australia, Argentina, Chile and Singapore.

[3] While it is not our goal herein to test theories for why retailers use QS pricing, the extant literature has hypothesized it could be due to price discrimination on consumers with incorrect beliefs (Agrawal et al., 1993; Gabaix and Laibson, 2006) or due to retailer price image management (Sprott et al., 2003). The fact that we observe systematically different rates of QS at the shelf across categories does suggest the practice is not merely due to random human error in setting the prices.

useful natural experiment for studying consumer price inattention. The fact that consumers pay a price premium for surcharged large packsizes serves as prima facie evidence of price inattention since these large packs are dominated by cheaper bundles of smaller packsizes of the same product. Moreover, QS creates useful relative price deviations that we can exploit to implement our tests.

We conduct two descriptive analyses to provide suggestive evidence that price inattention may play a role in the incidence of QS and the payment of the QS price premium. First, we find that, in aggregate, retailers are less likely to charge QS at the shelf and consumers are less likely to pay the QS price premium in states that have mandatory unit pricing regulations. Second, turning to the demand side, we find that a consumer's propensity to choose the surcharged packsize, conditional on choosing a brand with at least some QS'd sizes, decreases when both the surcharged size and the smaller size are promoted by price advertising and/or in-store displays that make the prices salient.

The occurrence of QS in the data does not unequivocally imply that consumers are inattentive to prices. For example, consumers may prefer larger QS'd packs over multiple smaller packs due to convenience (e.g., Gerstner and Hess, 1987), an alternative explanation past work on QS has not been able to rule out (e.g., Clerides and Courty, 2017). We build on recent work by Abaluck et al. (2026) (ACZ hereafter) and develop a test for price inattention that accommodates controls for preferences for large sizes. Our approach comprises three key features: (i) it is robust to a wide range of search conducts that consumers may follow when discovering prices, including those commonly used in the literature, (ii) it does not rely on strong parametric assumptions on utility, and (iii) it does not require data on consumer search or consideration sets.

The test involves derivatives of the choice probabilities with respect to the pack size and price of appropriately chosen products. Estimating these derivatives requires sufficient independent variation in price and pack size conditional on other product attributes, and QS induces exactly this type of variation. Thus, in addition to the suggestive evidence of inattention described above, QS also provides the variation needed to implement our structural test.

We find statistically and economically large degrees of attenuation in our CPG case studies. We also find that the degree of attenuation increases for the subset of QS trips, during which a consumer would likely be "surprised" to learn that purchasing the large pack size is a "mistake" on the basis of the unit price since they could buy multiple units of a smaller pack to obtain approximately the same quantity more cheaply. These results are robust to controls for heterogeneous preferences for larger packsizes (e.g., larger households may consider large packsizes more convenient).

We then extend ACZ's approach to analyze the economic implications of price inattention and QS. Moments based on the second-order properties of choice behavior semi-parametrically identify consumers' preferences up to a known distribution of random utilities, even under incomplete price information. We assume nested logit random utility to allow for correlation in tastes between products that differ only in pack size. We use the utility parameter estimates to conduct a counterfactual analysis of demand under complete price information. Interestingly, we find that many households have a relative preference for larger packsizes, consistent with a taste for convenience. Nevertheless, a decomposition indicates that inattention accounts for as much as 80% of the observed QS purchases (Laundry Detergent); although we see considerable heterogeneity including beer where 73% is due to convenience and only about 1% is due to inattention. Under full information, we find that consumer surplus increases as much as 20% of the average price (detergent); although again we see heterogeneity. On the supply side, full information would reduce retailer revenues by as much as 25% (Detergent) and 16% (Toilet Paper), confirming that retailers monetize the price inattention in some categories, at the expense of consumer surplus. These results should be of interest to advocates for unit pricing regulation, assuming mandatory disclosures of unit pricing are themselves salient to consumers.

Our work is highly related to the extant literature seeking to quantify limited attention through the size of consumers' self-reported consideration sets (e.g., (Roberts and Lattin, 1991; Hauser and Wernerfelt, 1991; Honka, 2014; Hauser, 2014)). Our test only requires choice data and can establish limited consideration even when consideration sets and search behavior are unobserved under our assumption that price is the underlying source of uncertainty. Our work also contributes more broadly to the literature studying the impact of incomplete information and information-gathering frictions on consumer demand Stigler (1961); Howard and Sheth (1969).

Our findings also add to the growing literature on consumer price inattention (Wells and LoSciuto, 1966; Harrell et al., 1976; Dickson and Sawyer, 1990), where inattention to the objective all-in price-per-equivalent-unit price leads to the appearance of consumers making choice errors. Our work also relates to the literature on left-digit bias (Anderson and Simester, 2003; Strulov-Shlain, 2023; List et al., 2023) and on price partitioning (Chetty et al., 2009; Finkelstein, 2009; Grubb and Osborne, 2015; Greenleaf et al., 2016), where consumers are inattentive to components of the total price. Our nonparametric tests should be applicable to these other forms of price inattention. Our work is also related to the estimation of models of rational inattention (e.g., Matejka and McKay (2015); Joo (2023); Natan (2025); Brown and Jeon (2024); although our framework does not explicitly nest these parametric choice models. Our findings are also related to the broader literature on

consumers appearing to make choice errors when products are sold using nonlinear pricing (DellaVigna and Malmendier, 2006; Lambrecht and Skiera, 2006).

Our work is also closely related to recent literature studying "shrinkflation" where firms use deceptive, hard-to-notice packsize reductions with disproportionately small price reductions to increase the effective price Kim (2024); Lee (2024); Wasserman and Pires (2025); Janssen and Kasinger (2025); Grigolon et al. (2026). This literature focuses on inattention to packsizes (as opposed to prices) since it focuses on case studies of consumer response to small and hard-to-detect changes in the packsize. In our case studies, a product's packsize assortment is mostly time-invariant in contrast with prices which vary on a regular basis.

Finally, our findings contribute to the growing body of research using price search models to capture consumers' incomplete price information (e.g., de los Santos et al. (2012); Honka (2014); Bronnenberg et al. (forthcoming); Seiler (2013)) and limited consideration sets (e.g., Roberts and Lattin (1991); Hauser (2014)). Our test for inattention is consistent with many of the stylized models of search used in the literature. However, it is also robust to a broader set of conducts leading to incomplete price information. In this regard, we also contribute to an emerging literature using a structural approach to study behavioral economics, as advocated by DellaVigna (2018).

2 Data

2.1 NielsenIQ Homescan Panel (HMS)

We obtained the HMS data through the relationship between NielsenIQ (hereafter NIQ) and the Kilts Center for Marketing at the University of Chicago Booth School of Business. The HMS database consists of a rolling panel of about 60,000 unique households per year and their product purchases in the CPG industry between 2004 and 2023. We use the sample period from 2006 to 2020 to conform with the RMS sample, which starts in 2006, and to avoid the major structural changes to the NIQ data after 2021. For each product purchased, we observe the Universal Product Code (UPC), the price paid and the quantity purchased.

2.2 Retail Measurement System (RMS) Databases

We also obtained the RMS data through the relationship between NIQ and the Kilts Center for Marketing at the University of Chicago Booth School of Business. The RMS database contains point-of-sale data from 2006 to 2023. We again focus on the sample period between 2006 and 2020 due to the aforementioned changes in the data structure.

The 2006-2020 RMS data contain the weekly sales data for a non-random sample of 61,174 unique stores, spanning 268 chains. For each store-week, we observe the total unit volume and average price for each product with positive sales, designated by its UPC. For a subset of 8.1% of stores, we also observe promotional information for each UPC, with indicators for whether a UPC was displayed in the store and/or featured in the chain's newspaper circular that week.

2.3 Classifying the UPCs into Unique Product Items

We classify UPCs into "items" that are identical except for packsize and conduct our subsequent QS analysis at the item level. Hereafter, we refer to a unique UPC as a product and a set of comparable UPCs as an item.

First, we use NIQ's UPC-level product attribute file. While NIQ reports some attributes consistently across product modules, like brand name and packsize, others are module-specific, such as format (e.g., pill vs. gel tablet and sachet vs. bottle). In some cases, NIQ's brand name field contains partial information about flavor and other functional attributes.

For a more rigorous analysis, we supplement the attribute file with additional product attributes for six product categories: beer, cereal, cracker, chewing gum, laundry detergent, and toilet paper, designated by the NIQ product group codes 5001, 1005, 1506, 0505, 4501 and 4507, respectively.[4] We use web searches and external UPC database searches in barcodelookup.com, barcodespider.com, and upcdatabase.org.

Table 1 reports each category's total number of unique UPCs matched and successfully assigned a unique item identifier using the procedure described above. The coverage in terms of RMS dollar sales share ranges from 88.2% for cereal to 100% for beer and toilet paper. Some categories exhibit a relatively low fraction of matched UPCs in terms of the UPC count, due primarily to NIQ's policy of masking store-brand UPCs. We treat each unidentified UPC as unique item during our analysis. In our analysis below, we use the coarse item classification across all categories to assess the overall pervasiveness of QS. For our subsequent analysis of price inattention, we focus our analysis on the more granular item classification for the 6 case studies.

[4]For product group 4501, we exclude dish soap from our analysis.

Table 1: Identified and Matched UPCs

Category	Total # of UPC	Matched # of UPC	Matched UPC dollar sales share
beer	42,624	42,510	1.000
cereal	32,147	14,864	0.882
cracker	17,191	10,316	0.918
gum	8,610	7,371	0.996
laundry detergent	12,603	7,885	0.974
toilet paper	18,956	11,163	1.000

Notes. This table summarizes the total number of unique UPCs matched and assigned a unique item identifier for each category. Dollar sales share is calculated using the entire 2006-2020 RMS data for the respective category.

2.4 Choice Panel Construction

We construct the choice panel for each of the six categories by merging the HMS and RMS datasets by store and week. RMS data provide the complete point-of-sale information about the available products during a trip. For each category, we retain only those matched trips during which the consumer purchases a single UPC, accounting for 65.7%-97.7% of the trips, depending on the category. We also retain only those trips to stores for which NIQ tracks the promotional activity. The resulting samples used for our descriptive analysis contain 115,826-771,225 trips with 34.4-251.7 unique UPCs available on average. See Table 14 in Appendix A.2 for a complete summary of the filtering process and retained sample sizes used below for analysis.

For our structural analysis, we drop "fringe" items that rank below the 30% most frequently chosen items in each category to reduce the number of the alternatives in the choice set. See Table 2 for descriptive statistics of our choice panels.

Table 2: Summary Statistics of the Choice Panels

Category	Trimming	# trips	Sample size	Avg. choice set size	Unit price ($)	Feature or display	$\mathcal{QS}$ propensity
beer	X	169,115	38,305,639	226.51	0.097	0.132	0.129
					(0.053)	(0.339)	(0.335)
cereal	X	771,225	194,124,860	251.71	0.269	0.158	0.013
					(0.140)	(0.365)	(0.114)
cracker	X	438,486	74,192,321	169.20	0.361	0.234	0.021
					(0.214)	(0.423)	(0.144)
gum	X	115,826	11,459,396	98.94	0.093	0.062	0.032
					(0.288)	(0.241)	(0.175)
laundry detergent	X	154,322	15,217,792	98.61	0.131	0.168	0.048
					(0.107)	(0.374)	(0.213)
toilet paper	X	205,924	7,085,496	34.41	0.003	0.140	0.061
					(0.001)	(0.347)	(0.239)
beer	O	46,741	1,678,707	35.92	0.075	0.205	0.323
					(0.049)	(0.404)	(0.468)
cereal	O	221,848	6,852,621	30.89	0.258	0.234	0.057
					(0.148)	(0.423)	(0.231)
cracker	O	130,900	3,065,498	23.42	0.300	0.274	0.130
					(0.133)	(0.446)	(0.337)
gum	O	33,377	702,549	21.05	0.063	0.073	0.056
					(0.017)	(0.261)	(0.229)
laundry detergent	O	45,283	1,121,986	24.78	0.113	0.217	0.101
					(0.053)	(0.412)	(0.301)
toilet paper	O	54,751	483,106	8.82	0.002	0.136	0.072
					(0.001)	(0.342)	(0.259)

Notes. This table reports summary statistics for our choice panels drawn from the HMS data. # trips and Sample size columns report the number of choice incidence and total number of choice alternatives in the data, respectively. Unit Price ($), Feature or display, and $\mathcal{QS}$ propensity columns report the mean of the variables, with standard deviations in parentheses. "X" indicates the raw choice panel without trimming. "O" indicates the choice panel after trimming the fringe items, used for the structural analysis.

3 Descriptive Analysis: Stylized Facts About Quantity Surcharging

We follow the literature's convention defining *quantity surcharging* (QS) as the incidence of a large packsize of an item with a higher unit price than at least one available smaller packsize of the same item Widrick (1979) that can be bundled to approximately the same total quantity. To address the indivisibility of packsizes, we define the variable $\mathcal{QS}_j$ to

indicate when an item j sells a given packsize at a price such that approximately the same quantity (up to a margin of 15% volume difference) could be obtained more cheaply by buying multiple units of a smaller packsize of the same item.

3.1 The Incidence of Quantity Surcharging

To capture the incidence of QS and its economic relevance to the CPG sector, we use the 2018 RMS store-level data. For each of the $\tilde{1}$,100 product modules, we compute: total dollar sales across all UPCs, total dollar sales generated by QS'd items according to our $\mathcal{QS}$ definition, the share of total category sales generated by QS'd items, and the cross-store-week average probability that at least one QS'd item was available.

Table 3 reports these descriptive statistics across all NIQ modules. For the average module, 5% of sales are generated by products satisfying $\mathcal{QS}$ and 15% of store-weeks have at least one QS'd item available. These results are not weighted by the modules' sales volumes. In addition, these results should be interpreted as an upper bound on potential QS since the NIQ attribute file may omit some of the differentiating attributes across packsizes of the same brand. Table 21 in Appendix B.2 reports total sales and QS sales rolled up to the department level for each of the 10 departments and overall; although the Deli category should be read with caution due to the random-weight nature of the category. We find that, overall, 12% of total RMS sales are generated by products satisfying $\mathcal{QS}$ and that 3% of total sales ($10.8 Billion) represent the surcharge—opportunity cost—from consumers purchasing a large packsize of a QS'd item instead of purchasing approximately the same quantity through a cheaper smaller packsize of the same item. A back-of-the-envelope calculation puts total U.S. CPG sales in 2018 at $530 Billion[5] so that, extrapolating, consumers paid a total surcharge of approximately $17.5 Billion in 2018 by purchasing a QS'd large packsize.[6] That said, for our six categories with more granular data, the median ratio of the corrected opportunity cost to the RMS opportunity cost is 0.36. Applying that correction to the full data, we end up with a more conservative $5-$6 Billion estimate of total 2018 surcharges.

[5] Online CPG sales in 2018 were $58.6B Ma (2025) and represented 11% of total US CPG sales Melton (2019).

[6] $10.853 \times (530/327.937) = 17.5$

Table 3: $\mathcal{QS}$ Incidence Across All RMS Product Modules, 2018

Variable	Mean	Min	Max	Median	SD
sales ($)	$ 293,324,672	$ 1	$ 12,596,736,605	$ 64,527,780	$ 735,794,531
$\mathcal{QS}$ sales ($)	$ 36,168,646	$ 0	$ 2,124,345,533	$ 1,104,244	$ 129,594,264
$\mathcal{QS}$ sales share	0.050	0	0.499	0.019	0.071
prob. at least one $\mathcal{QS}$ can be found	0.150	0	0.963	0.052	0.195

Notes. This table summarizes $\mathcal{QS}$ incidence across product modules in the 2018 RMS data using the NIQ attributes for item definitions.

Tables 4 and 5 provide a more conservative assessment of QS by focusing on the granular item classifications in our six case studies. Using $\mathcal{QS}$, we find that the incidence of at least one QS'd item ranges from as low as 31.4% of store-weeks (cereal) to as high as 86.1% (beer). The cross-store-week average proportion of QS'd items ranges from as low as 1.8% (cereal) to 11.2% (beer) (Col (4) of Table 4) While QS is somewhat exceptional on any given item in a given store-week, it is not unusual for consumers to find at least one QS offer across brands on the shelf in these categories. Comparing Columns (4) and (7) of Table 4 confirms that QS is not an anomaly for obscure or niche products. Rather, the QS'd items are concentrated primarily among the top-selling items in the category.

Turning to Table 5, products satisfying $\mathcal{QS}$ represent a non-trivial share of total sales, ranging from as high as 14.7% (beer) to as low as 1.9% (cereal). As an upper bound on the potential cost of inattention, we also look at the opportunity cost of QS by comparing total sales of surcharged large packsizes of a QS'd item to the alternative sales had a comparably quantity been purchased instead with cheaper, smaller packsizes of the same item. In beer, 1.7% of total sales are due to the opportunity cost of a consumer paying a surcharge for a large packsize when a smaller packsize of the same item could have been purchased instead. For the QS'd items, the opportunity cost represents about 11.5% (toilet paper) to 21.6% (beer) of the QS'ed dollar sales (Col. (3) ÷ Col. (2) in Table 5). Appendix B.1 reports analogous results for three alternative $\mathcal{QS}$ measures (see Tables 15 to 20). We find qualitatively similar evidence for $\mathcal{QS}$, with numbers varying on the stringency with which small packs can be bundled to match the size of the large pack. Appendix Table 22 compares QS under these granular item definitions to the definitions using only NIQ attributes.

Table 4: $\mathcal{QS}$ Incidence Under Granular Item Definitions

category	(1) at least one $\mathcal{QS}$	(2) $\mathcal{QS}$ no. items	(3) total no. items	(4) $\mathcal{QS}$ share (2÷3)	(5) $\mathcal{QS}$ no. top 5% items	(6) total no. top 5% items	(7) $\mathcal{QS}$ share of top 5% items (5÷3)
beer	0.861	7.18	64.38	0.112	7.12	53.49	0.111
cereal	0.314	0.98	83.11	0.012	0.96	47.23	0.012
cracker	0.451	0.84	47.06	0.018	0.82	29.40	0.018
gum	0.685	2.06	40.35	0.051	2.00	26.95	0.050
laundry detergent	0.477	1.34	43.93	0.031	1.21	23.59	0.028
toilet paper	0.446	0.87	17.32	0.050	0.68	7.06	0.039

Notes. This table summarizes $\mathcal{QS}$ by product module in the 2018 RMS data using our granular item definitions. Col. (1) reports the observed proportion of store-weeks during which at least one $\mathcal{QS}$ item is available. Cols. (2) to (4) report the average number of items on the shelf in a store-week, how many of them have at least one surcharged packsize, and the $\mathcal{QS}$ item share. Cols. (5) to (7) report the analogous statistics for the subset of items comprising the top 5% on the basis of annual category-year dollar sales. All statistics are calculated using the entire sample of RMS stores in 2018.

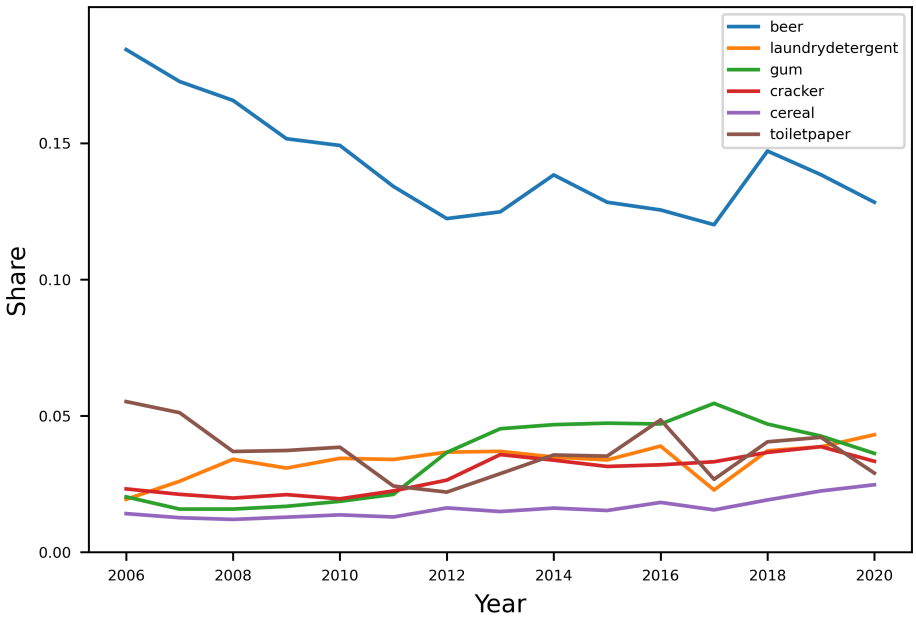
Table 5: $\mathcal{QS}$ Sales by Category Under Granular Item Definitions, 2018

category	(1) \$ sales	(2) $\mathcal{QS}$ sales (\$)	(3) $\mathcal{QS}$ opp. cost (\$)	(4) $\mathcal{QS}$ sales share (2÷1)	(5) $\mathcal{QS}$ opp. cost share (3÷1)
beer	\$ 9,200.6M	\$ 1,353.9M	\$ 157.2M	0.147	0.017
cereal	\$ 4,235.1M	\$ 81.1M	\$ 14.2M	0.019	0.003
cracker	\$ 2,450.8M	\$ 89.7M	\$ 19.4M	0.037	0.008
gum	\$ 976.0M	\$ 45.9M	\$ 8.4M	0.047	0.009
laundry detergent	\$ 2,856.6M	\$ 106.2M	\$ 16.7M	0.037	0.006
toilet paper	\$ 3,271.6M	\$ 132.4M	\$ 15.2M	0.040	0.005

Notes. This table reports the dollar-amount sales subject to $\mathcal{QS}$ in the 2018 RMS data using our granular item definitions. Dollar figures are expressed in millions of dollars (\$ M). Shares are $\mathcal{QS}$ dollar sales divided by total dollar sales.

Lastly, Figure 1 plots the time-series of the dollar-sales share of $\mathcal{QS}$ during our sample period. With the exception of beer, where QS is the most prominent even though it has declined slightly from 18.4% in 2006 to 12.8% in 2020, the QS share has been quite stable over time in the other categories.

Figure 1: Time Trend of QS Dollar Sales Shares



Notes. This figure plots the time trends [($ sales subject to QS)/(total category $ sales)] in the RMS data during 2006-2020.

3.2 Testing the Rarity of *Quantity Surcharging* and Pervasiveness of *Quantity Discounting*

We now validate the "larger economy size belief" (e.g., Granger and Bilson, 1972; Widrick, 1979; Manning et al., 1998) in the HMS choice panel data: consumers with rational price expectations and search frictions would logically prioritize larger packsizes in their search to save money (abstracting from storage costs). We apply a probability integral transform (PIT) test to each store-week-item's price ranks across packsizes, under the null that packsizes are uninformative about prices (*exchangeability*). For QS, we compare the rank-order of the unit price for a given item's packsizes during a store-week to the lowest unit price across its cheaper comparable bundles of smaller packsizes. Intuitively, we locate the break-even point in the distribution and assess whether surcharging is a rare inversion of the typical ordering. For Quantity Discounting (QD), we use the entire set of available packsizes for a given store-week-item to show the pervasiveness of monotonicity of unit prices in size. For technical details, see Appendix C.1.

Quantity Surcharging Table 6 summarizes our evidence on the rarity of QS for an item. Column (3) confirms that for each category, we strongly reject the null that packsize is uninformative for pricing when comparing a given packsize to the cheapest comparable smaller

Table 6: PIT Test for Quantity Surcharging

Category	(1) Break-even pctile	(2) Null rate (%)	(3) Uniformity (AD)
beer	72.1	59.0	reject***
	[69.0, 76.2]		
cereal	88.4	53.3	reject***
	[86.9, 89.9]		
cracker	76.4	58.3	reject***
	[69.9, 87.8]		
gum	88.8	51.4	reject***
	[87.3, 90.3]		
laundrydetergent	76.1	51.7	reject***
	[74.4, 77.7]		
toiletpaper	84.9	56.2	reject***
	[82.0, 87.0]		

Notes. Anderson-Darling test of H_0: PIT values $\sim$ Uniform$(0, 1)$; critical values 2.49 (5%), 3.88 (1%). *** rejects at $p < 0.01$; all categories reject at $p < 10^{-16}$.

Table 7: PIT Test for Quantity Discounting

Category	(1) Discount-ordered (%)	(2) Monotone (%)	(3) Null (%)	(4) Uniformity (AD)
beer	85.8	60.9	50.0	reject***
	[83.7, 87.6]			
cereal	81.2	69.5	50.0	reject***
	[78.5, 83.6]			
cracker	78.8	56.6	50.0	reject***
	[71.6, 84.4]			
gum	90.8	85.2	50.0	reject***
	[89.5, 92.0]			
laundrydetergent	73.5	62.2	50.0	reject***
	[71.7, 75.2]			
toiletpaper	85.0	61.6	50.0	reject***
	[81.0, 88.5]			

Notes. Cell-level Kendall τ_b between pack size and unit price, per store-week-item ($K \geq 2$ sizes). "Discount-ordered" = $\Pr(\tau_b < 0)$; "Monotone" = $\Pr(\tau_b = -1)$; "Null" is the exchangeability benchmark. CI (bracketed) for the discount-ordered share, clustered by store-week-item. AD tests uniformity of the PIT; all categories reject at $p < 10^{-16}$.

packsize avaialable (exchangeability). Turning to columns (1)-(2), comparing against the null rate (column (2)), we see that a positively-valued quantity surcharge is indeed a tail event ranging from the 72nd to 89th percentile across our categories. Substantively, a consumer with *rational expectations* and facing search costs who typically purchases large packs to save money would not likely prioritize searching smaller packsizes.

Quantity Discounting Table 7 summarizes our evidence for the pervasiveness of QD. Once again, column (4) confirms that for each category, we strongly reject the null that packsize is uninformative for pricing (exchangeability); this time at the store-week-item level. Turning to columns (1)-(2), comparing against the null rate (column (3)), column (1) shows a big deviation, with 74% to 91% of the size ladders in the data discount-ordered by comparing the large-pack unit price to the cheapest unit price across smaller packs—the largest packsize is a better deal. Column (2) is more stringent and indicates that, with the exception of crackers, over 60% of size-ladders across store-week-items have zero unit-price inversions whereby a bundle of smaller packs is cheaper than a large pack: every larger packsize is cheaper and the entire size schedule is a volume-discount curve. Substantively, a consumer with rational expectations would anticipate a volume discount from larger packsizes and, accordingly, would logically prioritize large packsizes when search is costly to save money.

Tables 24-25 in Appendix C.2 also show that QS incidence is not a persistent phenomenon concentrated among a small number of items. The 1.62% of the UPCs that are always QS'd within a store are niche items with very small purchase shares. So consumers are unlikely to be able to learn which products are likely QS'd on a given trip.

3.3 Price Salience and Suggestive Evidence of Price Inattention

We now turn to *suggestive* descriptive evidence of price inattention by showing that price salience is associated with a lower propensity for consumers to pay the QS price premium.

3.3.1 Mandatory Unit Pricing Regulations and QS Purchases

We use cross-state variation in mandatory unit-pricing laws as a preliminary test for inattention. While the U.S. does not have a federal unit-pricing mandate, nine states had unilaterally implemented mandatory unit-pricing regulations by the end of 2017[7]: Connecticut, Maryland, Massachusetts, New Hampshire, New Jersey, New York, Oregon, Rhode Island, and Vermont. The District of Columbia also had a mandatory regime (National Institute

[7]This classification is corroborated by a contemporaneous NIST publication reporting in March 2017 that unit pricing remained voluntary in the other 41 states (Sefcik, 2017).

of Standards and Technology, 2015).[8]. In principle, unit-pricing disclosures make the price and QS more salient, ostensibly reducing or eliminating the search cost.

Table 8 compares the $\mathcal{QS}$ dollar-sales share and the item count share across the unit-pricing mandate status for our six focal categories using the RMS data. Standard errors are clustered by state using wild bootstrap, accounting for our small number of treated states.[9] Panel-A column (3) shows that consumers' $\mathcal{QS}$ purchases are 0.58-5.76 percentage points lower across categories in states with unit-pricing mandates. Panel-B, column (3) shows that retailers are also less likely to post prices exhibiting $\mathcal{QS}$ in these states, proxied by the number of QS'ed items divided by the total number of items on the shelf; although our results are less precise. These findings are suggestive that mandating unit-pricing not only reduces the incidence of quantity surcharges on the supply side, it also reduces the propensity for consumers to pay a surcharge on the demand side. The fact that mandated unit pricing does not fully eliminate the use of QS may stem from the fact that some consumers may not pay attention to *unit pricing* disclosures when making purchase decisions (Harrell et al., 1976).

[8]The principal authorities in force by the end of 2017 were Conn. Gen. Stat. §§21a-73-77; Md. Code, Commercial Law §§14-101-105; Mass. Gen. Laws ch. 6, §115A and 202 CMR 5.00; N.H. Rev. Stat. Ann. §438:26-a; N.J. Stat. Ann. §56:8-25 et seq. and N.J.A.C. 13:45A-14.1-15; N.Y. Agriculture and Markets Law §214-h; Or. Rev. Stat. §§616.850-890; R.I. Gen. Laws ch. 6-31; and 6 V.S.A. §§681-687. The corresponding District of Columbia provisions are D.C. Code §§28-5201-28-5210.

[9]Our key results are substantively unchanged if we instead use the jackknife clustered standard errors.

Table 8: Mandatory Unit-Price Display and Quantity-Surcharged Purchases, 2018

category	(1) no mandatory display (r_0)	(2) mandatory display (r_1)	(3) difference ($r_0 - r_1$)	(4) bootstrap SE	(5) one-sided p-value
Panel A: $\mathcal{QS}$ dollar-sales share					
beer	0.1533	0.0958	0.0575	0.03065	0.027
cereal	0.0203	0.0146	0.0058	0.00266	0.002
cracker	0.0387	0.0289	0.0098	0.00420	< 0.001
gum	0.0483	0.0411	0.0072	0.00395	0.040
laundry detergent	0.0403	0.0248	0.0154	0.00856	0.023
toilet paper	0.0428	0.0316	0.0112	0.00633	0.027
Panel B: $\mathcal{QS}$ incidence (store-week-UPC observations)					
beer	0.1266	0.0851	0.0414	0.02148	0.032
cereal	0.0105	0.0093	0.0012	0.00071	0.028
cracker	0.0221	0.0181	0.0040	0.00160	< 0.001
gum	0.0455	0.0406	0.0050	0.00312	0.073
laundry detergent	0.0291	0.0249	0.0042	0.00344	0.116
toilet paper	0.0398	0.0370	0.0028	0.00275	0.171

Notes. This table reports weighted store-week regressions with week fixed effects using the 2018 RMS data. Panel A reports the share of dollar sales attributable to $\mathcal{QS}$, weighting each observation by its total dollar sales. Panel B reports the share of store-week-UPC observations exhibiting $\mathcal{QS}$, weighting each observation by its UPC count. Columns (1) and (2) report predicted shares without and with a statewide mandatory unit-price-display regime, respectively. Col. (3) is Col. (1) minus Col. (2), equal to the no-mandate coefficient. Col. (4) reports the bootstrap standard error, calculated as the standard deviation of the no-mandate coefficient estimates across the null-imposed wild cluster bootstrap replications. Col. (5) reports the one-sided wild cluster bootstrap-t p-value for $H_0 : r_0 - r_1 \leq 0$ against $H_1 : r_0 - r_1 > 0$, using 9,999 replications, state-level Webb six-point multipliers, and the null imposed at $r_0 - r_1 = 0$. Week fixed effects are retained in the bootstrap calculation. Values below 0.001 are reported as $p < 0.001$.

3.3.2 Promotions and QS Purchases in the Choice Panel

We now exploit the HMS choice panel data's variation in promotional activity, feature advertising and in-store displays designed to make a product's shelf price more salient. As explained in Blattberg and Neslin (1990, p. 247), "If the consumer is not informed that a price discount is offered, the price elasticity is likely to be small." Features and displays communicate the price of the promoted item to a consumer by making the price (and the product itself) salient. Assuming that promotions increase the salience of the price of the promoted product, we test whether promotions moderate households' propensity to buy a surcharged larger packsize.

We take a reduced-form approach that follows the convention in the search literature that

consumers must search the price of a product before they can buy it. Therefore, we expect that promoting the dominant smaller packsize should reduce the propensity to purchase a surcharged product for a given item as any consumer considering the surcharged larger packsize would behave as fully informed about prices. The effect of promoting the surcharged product would likely increase the propensity of an error since it would not make the dominant price salient. The effect of combining a promotion for both the surcharged and dominant smaller size could be ambiguous if promotions have persuasive effects in addition to making prices salient.

We retain all trips during which a QS'd *item* is chosen and only one packsize of that item is surcharged.[10] Let $y_{ht} \in \{0, 1\}$ indicate whether a consumer chooses a product satisfying $\mathcal{QS}$, meaning they pay a surcharge by choosing a packsize of that item even though they could purchase approximately the same quantity at a lower total price by choosing (possibly multiple units of) a smaller packsize of the same item. We run the following regression:

$$y_{ht} = \lambda_1 + \lambda_2 \mathbb{I}_{ht}^{promo,surcharged} + \lambda_3 \mathbb{I}_{ht}^{promo,dominant} + \lambda_4 \mathbb{I}_{ht}^{promo,surcharged} \times \mathbb{I}_{ht}^{promo,dominant} + \epsilon_{hjt} \quad (1)$$

where the indicator variables $\mathbb{I}_{ht}^{promo,size}$ indicate each of the price salience states for consumer h on trip t: (1) the surcharged large packsize is promoted, and (2) at least one of the dominant smaller packsizes of the chosen item with a lower unit price and comparable bundled quantity is promoted.

Table 9 reports results for OLS, including with household fixed-effects. The fixed-effects control for the possibility that some households exhibit a persistent taste for larger packsizes and for specific items, which could spuriously create the illusion of systematic choice errors. As expected, the propensity to choose a surcharged item is always smaller when the dominant product is promoted, conditional on a trip where a consumer purchases a QS'd item; although the point estimate is insignificant for gum. As expected, the effect of only promoting the surcharged item exacerbates consumers purchasing it, an effect that is large and significant in each category. For five of the 6 categories, the combined effect of promoting both is directionally negative but insignificant in beer and gum. Crackers exhibits a significant positive effect from promoting both, possibly due to an overwhelming persuasive effect.

These results are only suggestive since promotions could operate purely through a persuasive effect and may be selected on systematically different types of consumers whose preferences correlate with the incidence of promotions. However, if anything, past research finds the opposite with deal-prone consumers tending to prefer larger pack sizes (Ghosh, 2017; Catena et al., 2021), working in favor of our evidence being suggestive of inattention.

[10] This represents over 75% of the trips during which a QS'd item is chosen, where the remainder have two or more surcharged packsizes.

Table 9: Effect of Promotions on Choosing the Surcharged Packsize Conditional on Purchasing a QS'd Item

Beer

	OLS	Fixed Effects
Intercept	0.217***	
	[0.003]	
Surcharge promo	0.151***	0.085***
	[0.011]	[0.010]
Dominant promo	−0.109***	−0.037***
	[0.009]	[0.008]
Both promo	−0.067***	−0.009
	[0.020]	[0.018]
Num.Obs.	24 349	21 418
R2	0.016	0.566

Cereal

	OLS	Fixed Effects
Intercept	0.143***	
	[0.002]	
Surcharge promo	0.097***	0.125***
	[0.009]	[0.010]
Dominant promo	−0.068***	−0.042***
	[0.003]	[0.003]
Both promo	0.003	−0.030***
	[0.010]	[0.011]
Num.Obs.	62 901	55 234
R2	0.017	0.358

Crackers

	OLS	Fixed Effects
Intercept	0.068***	
	[0.002]	
Surcharge promo	0.087***	0.065***
	[0.007]	[0.008]
Dominant promo	−0.027***	−0.017***
	[0.003]	[0.003]
Both promo	0.029***	0.023**
	[0.009]	[0.010]
Num.Obs.	33 857	27 205
R2	0.019	0.391

Gum

	OLS	Fixed Effects
Intercept	0.206***	
	[0.004]	
Surcharge promo	0.054**	0.067***
	[0.023]	[0.026]
Dominant promo	−0.093***	−0.010
	[0.011]	[0.013]
Both promo	0.019	−0.012
	[0.043]	[0.047]
Num.Obs.	9718	6692
R2	0.009	0.591

Laundry Detergent

	OLS	Fixed Effects
Intercept	0.132***	
	[0.004]	
Surcharge promo	0.118***	0.071***
	[0.014]	[0.015]
Dominant promo	−0.059***	−0.033***
	[0.005]	[0.006]
Both promo	−0.065***	−0.046***
	[0.018]	[0.019]
Num.Obs.	15 750	11 330
R2	0.016	0.550

Toilet Paper

	OLS	Fixed Effects
Intercept	0.144***	
	[0.002]	
Surcharge promo	0.135***	0.116***
	[0.007]	[0.009]
Dominant promo	−0.078***	−0.053***
	[0.004]	[0.005]
Both promo	−0.065***	−0.059***
	[0.011]	[0.013]
Num.Obs.	32 830	26 396
R2	0.028	0.407

Notes. This table summarizes the QS purchase behavior under different price salience conditions, proxied by promotions, using the HMS panel data. OLS is a linear probability model. HH FE is a linear probability model with household-specific fixed effects. Standard errors are reported in parentheses. Promo both is the coefficient on the interaction term. Rows not applicable to a selected specification are left blank. For fixed-effect specifications, observations can differ because the singleton fixed-effect groups are dropped. $^{*}p < 0.1$; $^{**}p < 0.05$; $^{***}p < 0.01$.

4 Model

We build on ACZ and specify a choice model that nests a range of different price search conducts without imposing strong functional form restrictions on utility. For now, we focus on an individual consumer, but we will allow for persistent taste heterogeneity in our empirical application.

For a given trip, consumer i faces the choice set $\mathcal{J} \equiv \{1, ..., J\}$. The consumer observes each product alternative's attributes $\mathbf{x}_j, \forall j \in \mathcal{J}$, including the packsize. However, the consumer must search a product j to observe its price, p_j. In Appendix B.3, we provide supporting evidence for this assumption in our CPG case studies by documenting considerably more variation in an item's prices over time (in a store) than in the item's packsize assortment.

The consumer obtains the following choice-specific consumption utilities

$$U_{ij} = u(\mathbf{x}_j) + \beta p_j + \epsilon_{ij} \tag{2}$$

where ϵ_{ij} is an additive, continuously distributed idiosyncratic random utility shock observed by consumers; but not by the researcher. We do not impose any parametric structure on the deterministic component of utility, $u(\mathbf{x}_j)$, other than twice-differentiability. While not strictly necessary, we have implicitly assumed quasi-linearity in a numeraire good to obtain the additively separable "price effect" so that, as is customary in the literature, there are no income effects and our choice probabilities do not depend on the consumer's budget.[11]

Prior to search, the consumer forms price expectations conditional on $\mathbf{x}_j$:

$$\mathbb{E}(p_j|\mathbf{x}_j) = f(\mathbf{x}_j). \tag{3}$$

So the expected utility from good j prior to search takes the form:

$$EU_{ij} = u(\mathbf{x}_j) + \beta f(\mathbf{x}_j) + \epsilon_{ij}. \tag{4}$$

We denote the *surprise price* revealed by searching product j by:

$$\tilde{p}_j \equiv \begin{cases} p_j - f(\mathbf{x}_j) & \text{if } j \text{ is not promoted} \\ 0 & \text{if } j \text{ is promoted.} \end{cases} \tag{5}$$

$\tilde{p}_j$ measures the incremental price of a product relative to the consumer's expectation of the

[11] Formally, we assume utility is quasi-linear in a numeraire good that enforces the "adding-up" condition in random utility models. β then measures the marginal utility of income.

price before searching its price. It also embeds the assumption that promotions eliminate all uncertainty about a product's price (e.g., Mehta, 2007), which captures the salience effects of promotions documented in section 3.3.2.

We write the expected choice probability for product j as

$$\sigma_j = \int \sum_{\mathcal{S}: j \in \mathcal{S}} \mathbb{I}\left\{U_{ij} \geq U_{ik}, \forall j, k \in \mathcal{S}\right\} dG, \tag{6}$$

where $\mathcal{S} \subseteq \mathcal{J}$ denotes the consumer's search set and G denotes the joint distribution of the vector of utility shocks $\boldsymbol{\epsilon}_i$ and any unobserved determinants of search, such as search costs.

Following ACZ, we make the following assumptions about the consumer's search conduct.

Assumption 1. *(Search Conduct Assumptions)*

(i) If searching sequentially, a consumer searches goods in decreasing order of EU_{ij}. If searching simultaneously, a consumer searches the top K goods in terms of EU_{ij} for some non-negative integer K.

(ii) If searching sequentially, a consumer searches product j if and only if $g(\mathbf{x}_j, \epsilon_{ij}, \bar{u}) \geq 0$ where g is a decreasing function of the maximum utility uncovered so far in the search process $\bar{u}$. If searching simultaneously, a consumer searches product j if and only if $g(\mathbf{x}_j, \epsilon_{ij}) \geq 0$ for some function g.

(iii) Consumers choose the good which maximizes utility among the searched goods.

(iv) Only the value of $\tilde{p}_j$ is unknown to consumers prior to search, and search fully reveals $\tilde{p}_j$.

Assumptions 1(i) and 1(iii) are self-explanatory. Assumption 1(ii) ensures that consumers use search rules, $g\left(\mathbf{x}_j, \epsilon_{ij}, \bar{u}\right)$, that depend on the attributes visible to them prior to search and that they prioritize the most ex-ante appealing products. Assumption 1(iv) ensures that price is the only source of incomplete information. In our CPG application, this assumption is relatively uncontroversial since the key drivers of choice in most CPG categories are brand[12] and packsize,[13] both of which are highly salient to the consumer due to the packaging.[14]

ACZ show these assumptions are robust to a broad class of consumer search models and classic models of brand consideration sets (e.g., Roberts and Lattin (1991)). For instance, they are consistent with the conventional sequential search conduct of Weitzman (1979) provided beliefs over the unobserved prices $\tilde{p}$ are distributed i.i.d. across products

[12] See Bronnenberg et al. (2019) for a survey of the importance of brands.

[13] See for instance, Guadagni and Little (1983) and Singh et al. (2005) for the importance of packsize for brand choice.

[14] As mentioned earlier, Appendix B.3 provides supporting evidence for this assumption in our CPG case studies.

(homoskedastic) and the search cost is the same for all goods.[15] The literature typically maintains this homoskedasticity assumption (e.g., Hortaçsu and Syverson (2004), Kim et al. (2010), Honka et al. (2017), Ursu (2018), Morozov et al. (2021), Jiang et al. (2021) and Chung et al. (2025)). The assumptions are also consistent with heuristic search rules, such as "search until a minimum utility threshold is achieved" (Simon, 1955; Stüttgen et al., 2012), "simultaneously search all goods that have expected utility above a threshold" (e.g., Bronnenberg and Vanhonacker, 1996), or the conduct of the directed cognition model of Gabaix et al. (2006). Finally, the assumptions also nest the limiting full-information case where prices are known to consumers for all products prior to search.

However, Assumption 1 is not entirely without loss of generality. For example, Assumption 1(*ii*) could be violated if consumers search sequentially and stop after discovering a good with small $\tilde{p}$ regardless of utility in hand. Assumption 1(*ii*) could also be violated if consumers believe that $\tilde{p}_j$ has sufficiently higher variance than $\tilde{p}_k$ to justify searching product j before k even if the latter has a higher expected utility (e.g., Kim et al., 2017). To relax this condition, Proposition 1 below shows that our key test of inattention is robust to heteroskedastic beliefs under the commonly-estimated Weitzman (1979) sequential search conduct. Finally, Assumption 1(*iv*) rules out the gradual acquisition of information about utilities, as under Bayesian learning (e.g., Erdem and Keane (1996), Ackerberg (2003) and Crawford and Shum (2005)). It also implicitly assumes that consumers are aware of the existence of all products. Proposition 2 below extends our key test to the case where consumers do not initially observe the entire set of products and, instead, sequentially choose which products to discover and, among those discovered, which ones to search (Greminger, 2022).

To apply this model to the QS context, the choice set will include bundles of the cheaper small pack that deliver the same total quantity as the large pack but more cheaply. Under inattention, a consumer purchasing a surcharged large pack would have failed to search the cheaper bundle, even if they searched the smaller packsize. Under inattention, a consumer who has searched the small pack does not automatically update their beliefs about the bundle price, $f(x_{bundle})$, due to the cognitive friction associated with the mental arithmetic of unit pricing. We assume consumers form beliefs about an n-unit bundle as they would about a single pack of the same total size: $\mathbb{E}(p_b \mid x_b) = f(x_b)$, where $size_b = n \cdot size_s$. Consumers who have searched the small pack do not use p_s to update these beliefs, reflecting the cognitive cost of unit-price arithmetic. Bundles may still rank lower in expected utility than a single pack of equal size because of the disutility of buying multiple packs, which we estimate in Secton 6.3.

[15] For bundles, independence refers to consumers' subjective beliefs.

5 A Nonparametric Test for Inattention

In this section, we derive our nonparametric test for inattention based on the model in section 4. We show that ratios of derivatives of the choice propensities, which can be computed using nonparametric estimates of the choice propensity functions, are sufficient to identify preferences and a test for inattention under our conduct assumptions.

We partition the product attributes $\mathbf{x}_j = (size_j, \mathbf{w}_j)$ since we focus on packsize, with the remaining attributes in $\mathbf{w}$ serving as controls. We then write the choice probability for product j as:

$$\sigma_j(\mathbf{size}, \mathbf{p}, \mathbf{w}), \forall j \in \mathcal{J}, \tag{7}$$

where $\mathbf{size} \equiv (size_1, \ldots, size_J)$ denotes the vector of packsizes, $\mathbf{p}$ denotes the corresponding vector of shelf prices, and $\mathbf{w} \equiv (\mathbf{w}_1, \ldots, \mathbf{w}_J)$ contains additional controls.

For each trip, we define good 1 as the product with the lowest surprise price $\tilde{p}$. In practice, good 1 denotes the UPC with the lowest deviation between the actual shelf price and the consumer's expected price conditional on the UPC's size, item association, and promotion status.[16] Consumers do not observe the identity of good 1 prior to search. However, the researcher can identify good 1 by comparing the price realizations in each choice set with the predictions given by the prices beliefs distribution.

We start by stating our key Lemma, which is the nonparametric analog of Lemma 4 in ACZ.

Lemma 1. *Assume that consumers have utility functions of the form (2), and observe $\mathbf{x}_j \equiv (size_j, \mathbf{w}_j)$ and ϵ_{ij} for all j. Let f and EU_{ij} be defined as in (3) and (4), respectively.*

(a) If consumers are fully informed about prices, then, for all $j \neq 1$,

$$\left| \frac{\partial^2 \sigma_1}{\partial p_1 \partial p_j} \Big/ \frac{\partial^2 \sigma_1}{\partial p_1 \partial size_j} \right| = \left| \frac{\partial \sigma_j}{\partial p_j} \Big/ \frac{\partial \sigma_j}{\partial size_j} \right| \tag{8}$$

holds at every point of the support of σ_1 and σ_j where the denominators are nonzero.

(b) Suppose instead that consumers must search good j to observe p_j, with search conduct

[16] Technically, we assume consumers base their expectations on the item and size of each available UPC since, as explained above, we assume promotions eliminate all the price uncertainty for the promoted UPC.

satisfying Assumption 1 with corresponding g function. Further, assume that:

$$\frac{\partial u}{\partial size_j} + \beta \frac{\partial f}{\partial size_j} \equiv \frac{\partial EU_{ij}}{\partial size_j} \geq 0 \tag{9}$$

$$\frac{\partial g}{\partial size_j} \geq 0, \tag{10}$$

at every point in the support of the functions u, f and g. Then, for all $j \neq 1$

$$\left| \frac{\partial^2 \sigma_1}{\partial p_1 \partial p_j} \Big/ \frac{\partial^2 \sigma_1}{\partial p_1 \partial size_j} \right| \geq \left| \frac{\partial \sigma_j}{\partial p_j} \Big/ \frac{\partial \sigma_j}{\partial size_j} \right| \tag{11}$$

at every point of support of σ_1 and σ_j where the denominators are nonzero.

Proof. See Appendix D. □

Lemma 1.(a) follows from the fact that, under full information, choice probabilities only depend on prices and sizes through utilities, so that by the chain rule, both the ratio of first derivatives and the ratio of second derivatives identify the marginal rate of substitution (MRS) between the two attributes:[17]

$$\left| \frac{\partial^2 \sigma_1}{\partial p_1 \partial p_j} \Big/ \frac{\partial^2 \sigma_1}{\partial p_1 \partial size_j} \right| = \left| \frac{\partial \sigma_j}{\partial p_j} \Big/ \frac{\partial \sigma_j}{\partial size_j} \right| = \frac{|\beta|}{\partial u(\mathbf{x}_j) / \partial size_j} \equiv MRS. \tag{12}$$

In contrast, Lemma 1.(b) shows that the ratio of *first* derivatives is attenuated relative to the ratio of *second* derivatives when consumers have incomplete price information, under the conditions in (9) and (10). This attenuation follows from two results, which we discuss in turn: (i) Assumption 1 ensures that the ratio of second derivatives continues to point identify the MRS, even when consumers are only partially informed about prices; and (ii) the ratio of first derivatives underestimates the MRS under conditions (9) and (10).

We now discuss the underlying economics of the result in (i). First, consumers searching good 1 always purchase their utility-maximizing choice, even when they are not fully informed.[18] Consider a consumer i with utility-maximizing choice j^*: $U_{ij^*} \geq U_{ij}$ for all j. Good j^* also has higher expected utility than 1 ($EU_{ij^*} \geq EU_{i1}$), since good 1 dominates in terms of $\tilde{p}$ ($\tilde{p}_1 \leq \tilde{p}_j, \forall j$) and $U_{ij} = EU_{ij} + \beta \tilde{p}_j$ for all j. So if the consumer searches good 1, they must also search j^* under Assumption 1(i), which ensures consumers search products in decreasing order of their expected utilities. Second, the ratio of second derivatives takes the set of consumers who searched good 1 (corresponding to differentiating with respect to

[17] Equivalently, $\frac{1}{MRS}$ represents the pseudo-willingness to pay for one additional equivalent unit of the product conditional on the current size. The identification result follows from an argument similar to that in Lemma 2 in ACZ, but relaxing the assumption of linear utility.

[18] See Lemma 1 of ACZ.

p_1) and measures their sensitivity to the price of good j relative to their sensitivity to the size of good j. The argument in the first step ensures that these consumers behave as if they have full information and, in turn, their sensitivity to price relative to size reveals the true (normative) MRS.

To understand result in (ii), naively measuring the sensitivity to price relative to size across all consumers using $\frac{\partial \sigma_j}{\partial p_j} / \frac{\partial \sigma_j}{\partial size_j}$ (as one would in a standard full-information model) conflates preferences and information, attenuating the estimated MRS. In other words, the first-derivative ratio tends to underestimate the extent to which consumers care about price relative to size because it does not account for their information (or lack thereof): some consumers are unresponsive to price not because they don't care about it, but because they're not informed about it. The attenuation relies on the conditions (9) and (10). Intuitively, the condition (9) requires consumers to value larger packsizes sufficiently to offset the fact that they expect prices to be higher for those larger packs. Similarly, (10) requires consumers to be more likely to search larger packsizes, all else equal.

Lemma 1 yields a natural nonparametric test for incomplete information based on comparing the ratio of first derivatives relative to the ratio of second derivatives. Formally, we consider the hypothesis

$$\begin{aligned} &\mathbb{H}_0 : \text{Consumers have full information about prices} \\ &\mathbb{H}_1 : \text{Consumers need to search to reveal prices.} \end{aligned} \tag{13}$$

Given flexible estimates of $\hat{\sigma}_j$, Lemma 1 supports the following two-tailed test: "reject the null of full information whenever the ratio of first derivatives $\left| \frac{\partial \hat{\sigma}_j}{\partial p_j} / \frac{\partial \hat{\sigma}_j}{\partial size_j} \right|$ is significantly different from the ratio of second derivatives $\left| \frac{\partial^2 \hat{\sigma}_1}{\partial p_1 \partial p_j} / \frac{\partial^2 \hat{\sigma}_1}{\partial p_1 \partial size_j} \right|$." When conditions (9) and (10) hold, Lemma 1 generates a higher-powered, one-tailed test that exploits the attenuation result: "reject the null of full information whenever $\left| \frac{\partial \hat{\sigma}_j}{\partial p_j} / \frac{\partial \hat{\sigma}_j}{\partial size_j} \right| < \left| \frac{\partial^2 \hat{\sigma}_1}{\partial p_1 \partial p_j} / \frac{\partial^2 \hat{\sigma}_1}{\partial p_1 \partial size_j} \right|$ holds statistically."

The test is nonparametric in the sense that it does not require parametric assumptions on the distribution of the unobservables or the functional form of utility, other than the separability of price and $\mathbf{x}$ which is standard in the discrete-choice literature, and the additivity of the random-utility shocks. The test does require the set of sufficient conduct restrictions on the search process in Assumption 1; but these assumptions are weaker than those maintained in fully-specified search models.[19] The condition (9) for our one-tailed test is testable using the choice panel data (see Section 6.2 below). In contrast, the shape condition (10) is not testable without additional data on the search process itself. However, even when (10)

[19]Recall that the quasi-linearity in price is not strictly necessary and that we impose this condition ex ante for our applications since we do not expect income effects in our case studies.

cannot be tested, it holds automatically for all leading search models discussed in Section 4.[20]

When these two conditions fail to hold, we can still test for inattention nonparametrically using the lower-power, two-sided version of the test (13). In our empirical application below, we will first test condition (9) on the data. Then, we will implement the one-tailed test if the condition is satisfied and the two-tailed test if the condition is not satisfied.

We now extend the test in Lemma 1 to accommodate two violations of the baseline search conduct restrictions in Assumption 1 that are likely to matter in many empirical settings.

Robustness to Weitzman (1979) sequential search conduct with heteroskedastic beliefs:

We now show how to relax Assumption 1(i), extending our test to the case of sequential search (e.g., Weitzman, 1979) with heteroskedastic consumer price beliefs. Under heteroskedastic beliefs, a consumer might optimally search a high-variance product before a lower-variance product even if the latter has a higher expected utility. The following proposition gives sufficient conditions for the second-derivative ratio to identify the MRS under partial information when the variance of price beliefs depends on a product's packsize.

To relax Assumption 1(i), we modify our beliefs with a packsize-dependent scale term:

$$p_k \mid x_k \overset{d}{=} f(x_k) + \tau(size_k) \cdot Z_k, Z_k \sim i.i.d. F_Z, E[Z_k] = 0, Var(Z_k) = 1. \tag{14}$$

The following proposition extends Lemma 1 to Weitzman (1979) search when price beliefs take the location-scale form in (14), so that their variance, $\tau(size_k)^2$, depends on packsize while their mean remains $f(\mathbf{x}_k)$. It gives sufficient conditions under which the second-derivative ratio still identifies the MRS under incomplete price information (part i) and the first-derivative ratio is attenuated relative to it (part ii).

Proposition 1 (Test validity under heteroskedastic beliefs)**.** *Assume that consumers have utility functions of the form* (2) *and price expectations of the form* (14)*, and follow Weitzman (1979) sequential search conduct.*

Define the reservation threshold in price-surprise units, $m_i(\tau_k)$, by

$$\frac{c_i}{|\beta|} = \mathbb{E}[(m_i(\tau_k) - \tau_k Z_k)^+],$$

so that the overall reservation utility is $EU_{ik} + \beta m_i(\tau_k)$.

Fix $j \neq 1$. At points where the relevant denominators are nonzero:

[20] More precisely, the condition holds for the search models we discussed under the natural assumption that consumers prefer larger pack sizes all else equal (including pack price).

(i) *If, locally for almost every consumer,*

$$\tilde{p}_1 > m_i(\tau_1) \implies \tilde{p}_j > m_i(\tau_j), \tag{15}$$

then

$$\left| \frac{\partial^2 \sigma_1}{\partial p_1 \partial p_j} / \frac{\partial^2 \sigma_1}{\partial p_1 \partial size_j} \right| = |MRS|.$$

Given our choice of good 1, a sufficient condition for (15) *is* $\tau_1 \leq \tau_j$.

(ii) *If, in addition,* $\partial u / \partial size_j > 0$ *and*

$$\frac{\partial EU_{ij}}{\partial size_j} = \frac{\partial u(\mathbf{x}_j)}{\partial size_j} + \beta \frac{\partial f(\mathbf{x}_j)}{\partial size_j} \geq 0, \qquad \tau'(size_j) \geq 0, \tag{16}$$

then

$$\left| \frac{\partial \sigma_j}{\partial p_j} / \frac{\partial \sigma_j}{\partial size_j} \right| \leq \left| \frac{\partial^2 \sigma_1}{\partial p_1 \partial p_j} / \frac{\partial^2 \sigma_1}{\partial p_1 \partial size_j} \right|.$$

Proof. See Appendix E.1. □

In practice, we cannot test condition (15) because it depends on consumers' reservation values, which are unobserved. However, we can test the sufficient condition $\tau_1 \leq \tau_j$ using the distribution of price surprises $\tilde{p}_k$ in the data. In our empirical analysis, we assume rational expectations and estimate beliefs as the empirical distribution of prices conditional on attributes. We can then compare the estimated standard deviations of beliefs across products. Condition (16) retains the expected-utility restriction in (9) from the baseline case and adds $\tau'(size_j) \geq 0$ so that the variance of price beliefs weakly increases with packsize. Intuitively, since larger packs have higher prices, similar uncertainty about prices in percentage terms would imply larger price surprises in dollar terms. Once again, we can test this restriction using the dispersion of price surprises across packsizes. For bundles, τ_b is evaluated at $size_b$ using the dispersion of single-pack price surprises, consistent with the bundle-belief assumption in Section 4.

Robustness to Greminger (2022) search and discovery (SD) conduct:

The following proposition relaxes Assumption 1(*iv*), extending our test to the case of search and discovery in Greminger (2022), whereby consumers have to "discover" a product before they can search it and learn its price. However, we need to maintain the homoskedastic beliefs assumption.

Proposition 2 (Test restrictions under SD conduct with common beliefs)**.** *Assume that consumers have utility functions and price expectations of the form* (2) *and* (3)*, respectively,*

and maintain Assumptions SD1–SD3 in Appendix E.2. Fix $j \neq 1$. At points where the relevant denominators are nonzero:

(i)
$$\left| \frac{\partial^2 \sigma_1 / \partial p_1 \partial p_j}{\partial^2 \sigma_1 / \partial p_1 \partial size_j} \right| = MRS.$$

(ii) *Under the additional Assumption SD4 and condition (9), namely $\partial EU_{ij}/\partial size_j \geq 0$, we have*
$$\left| \frac{\partial \sigma_j}{\partial p_j} / \frac{\partial \sigma_j}{\partial size_j} \right| \leq \left| \frac{\partial^2 \sigma_1}{\partial p_1 \partial p_j} / \frac{\partial^2 \sigma_1}{\partial p_1 \partial size_j} \right|.$$

Proof. See Appendix E.2. □

6 Semi-parametric Structural Evidence for Inattention

We now implement a semi-parametric analog of the nonparametric test for price inattention derived in section 5. While the population-level arguments in Section 5 do not rely on any parametric restrictions, taking the approach to the data requires some additional structure to avoid the curse of dimensionality. Specifically, we restrict (i) how the choice propensities depend on the controls $\mathbf{w}$, and (ii) packsize to affect utility and price expectations log-linearly. The constraints in (ii) allow for diminishing marginal utility from bigger packs, in line with economic intuition, and for consumers to expect quantity discounts, consistent with the patterns in the data (Section 3.2). We embed these restrictions in our model by defining
$$\widetilde{size}_j \equiv \log(1 + size_j)$$
and using $\widetilde{size}_j$ instead of $size_j$ throughout. Importantly, we do not impose any parametric restrictions on the distribution of the unobservable utility shocks or search costs.

We conduct the test in three stages. First, we estimate price beliefs to identify good 1, the good delivering the lowest surprise price for each trip. Second, we verify assumption (9) to assess whether we can use the higher-powered 1-tailed, as opposed to 2-tailed test for inattention. Third, we estimate the choice propensities which will be used to construct our moments based on first and second derivatives for the test statistic. In those cases where Assumption (9) holds, we constrain the nonparametric estimator of the choice propensities to satisfy the attenuation of the first-derivative ratio relative to the second-derivative ratio in (11). The attenuation restriction helps discipline the nonparametric estimator in a principled theory-based way whenever the data supports it. Finally, we compute the ratios of first and second derivatives of the estimated choice propensity function that will comprise the test

statistic for the hypotheses (13). We report results for six CPG categories: beer, cereal, crackers, gum, laundry detergent and toilet paper.

Throughout this section, we use variation across trips t to identify the primitives of interest, but we omit the t subscripts throughout unless needed for clarity.

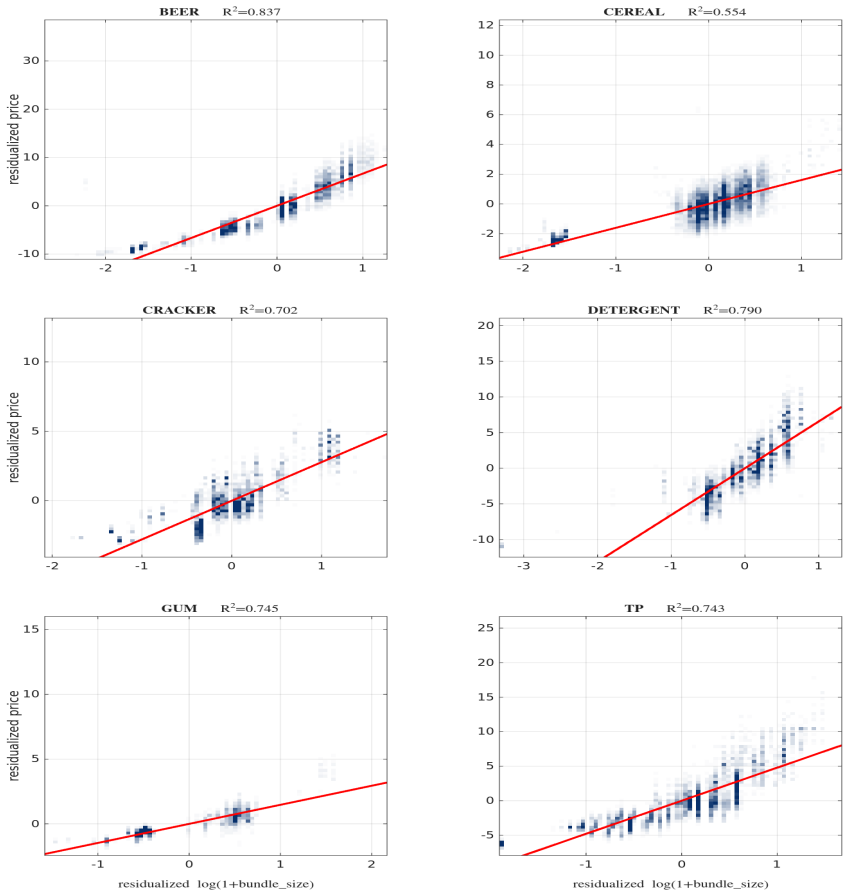
6.1 Step 1: Price Beliefs

We assume price expectations are linear in log size and item dummies:

$$\mathbb{E}(p_j|\mathbf{x}) = \gamma_0 + \gamma_{\widetilde{size}}\widetilde{size}_j + \gamma_{item}\text{item}_j. \tag{17}$$

We specify expected price as a concave function of size to capture the fact that, within an item, larger packs tend to have lower per-unit prices. We estimate (17) on single packs only, so that f reflects the prices consumers observe on the shelf and bundle beliefs are $f(x_b)$. Rational expectations thus refers to single packs.

Figure 2: Relationship Between Log Size and Price



Notes. For each category, this figure plots the log packsize on the x-axis and pack price on the y-axis, along with the best linear fit. Both variables are residualized against item dummies.

For each of the six categories, Figure 2 shows the relationship between price and log size after partialling out the item dummies from each variable. We find that prices tend to increase linearly in the log of packsize ($\gamma_{\widetilde{size}} > 0$) so that, under rational expectations, consumers should expect a quantity discount (i.e., prices to increase disproportionately slower than size). Further, the low-probability of QS incidence comes from large deviations above the red line for relatively large packsizes and/or large deviations below the red line for relatively small packsizes.

This analysis also demonstrates how QS provides variation for the estimation of the first and second derivative ratios that drive our key test statistics. Conditional on all other item attributes, QS generates deviations from the expected relationship between price and size. If instead price and size always co-moved for a given item, it would be impossible to identify the relevant derivatives.

6.2 Step 2: Verifying the Assumption in (9)

We now verify condition (9), using our beliefs estimates from the previous step to measure price surprise. The test follows from an argument in ACZ. When all products generate the same surprise price $\tilde{p}_j$, Assumption 1(i) implies that consumers always choose the product with the highest expected utility. The result follows from the fact that: (i) when all goods deliver the same surprise, the product with the highest expected utility also has the highest overall utility; and (ii) when consumers search as in Assumption 1(i), they never fail to search the product with the highest expected utility. Consequently, when $\tilde{p}_j = \tilde{p}_k$ for all j, k, we can write the choice probabilities as

$$\sigma_j = P\left(EU_{ij} \geq EU_{ik} \quad \forall k\right),$$

which depend only on expected utility and not on the search process. Thus, differentiating these probabilities with respect to pre-search attributes yields valid estimates of the impact of those attributes on expected utility.

To operationalize this argument, we condition on trips with limited variation in $\tilde{p}_j$ across products[21] and estimate choice probabilities flexibly using an approach similar to that described in step 3 below. We can then estimate $\frac{\partial EU_{ij}}{\partial \widetilde{size}_j}$ by differentiating the choice probabilities with respect to the product's own log size, enabling us to verify assumption (9).

We account for persistent heterogeneity in preferences using an approach inspired by Bonhomme and Manresa (2015). We use K-means to cluster households into three groups[22] based on demographics (race, income, household size) and shopping patterns (frequency of purchases of promoted UPCs and frequency of purchases of the various items in each category). We then conduct the verification of the assumption and the analysis of choice behavior in step 3 separately for each group and report the estimated population mean values, computed as weighted averages of the group-level outcomes.[23]

6.3 Step 3: Flexible Estimation of the Choice Propensity Function

We now discuss the estimation of the choice propensity σ_1 in equation (7). Following common practice in the literature (e.g., Hendel and Nevo (2006), Gowrisankaran and Rysman (2012),

[21] The identification argument conditions on trips where $\tilde{p}_j = \tilde{p}_k$ for all j, k. Given the relatively infrequent occurrence of such trips, we instead condition on trips where the standard deviation of $\tilde{p}$ across goods is in the bottom 20%.

[22] Eventually, we will use a data-based method to select the number of clusters.

[23] Motivated by the asymptotic equivalence results of Bonhomme and Manresa (2015), we assume that first-step clustering error has no first-order effect on the asymptotic distribution of our second-step estimators. We therefore conduct inference treating the estimated household group memberships as known.

Ryan (2012)), we add simplifying structure to circumvent the curse of dimensionality.

For each trip, we define the choice set as single-unit and two-unit commodity bundles for each of the available UPCs.[24] We then estimate the choice propensity σ_1 as a function of $\widetilde{size}_1$ and p_1 (where option 1 has the highest surprise utility or lowest $\tilde{p}$ during trip t), and $\widetilde{size}_2$ and p_2 (where option 2 has denotes the closest packsize to option 1 available for the same item)),[25] and $\widetilde{size}_{other}$ and p_{other} for a composite product capturing the other options available during the trip.[26] Since $\mathbf{w}$ serves only as a control, we assume it enters the model linearly. The controls in $\mathbf{w}$ include item fixed effects, promotion status (for option 1 and the average for the other options in that trip), a dummy for whether option 1 corresponds to buying a single pack or multiple packs of a product, and the fraction of other UPCs from the same item. The multi-pack dummy plays an economic role since, under our bundle beliefs, it explains why bundles may be ranked lower in search due to their lower expected utility (Assumption (1(*iii*))).

We allow the choice propensities to depend on promotion in two ways. First, consistent with our descriptive analysis, we allow promotions to eliminate consumer uncertainty about a product's price (see equation (5)). Second, we also allow promotions to have a direct effect on utility.[27] For trips with a QS'd item, the choice set $\mathcal{J}$ includes the "dominating" option (absent an intrinsic preference for larger packsizes) corresponding to buying multiple units of a smaller packsize of the QS'd item at a lower total price relative to the surcharged larger packsize.

For each cluster of households, we estimate the choice propensities, σ_1, using a sieve-GMM estimator as in Compiani (2022). See Appendix F for estimation details. We constrain the polynomial approximation to satisfy the attenuation restriction (that the first-derivative ratio is less than or equal in magnitude to the second-derivative ratio) whenever assumption (9) is verified in step 2. For instance, if we find that the assumption is satisfied for cluster 1 but not for clusters 2 and 3 within a category, we only impose the attenuation constraint for cluster 1.

We also restrict the ratio of second derivatives to be constant across trips, which holds

[24] 82% of the trips in our data involve purchases of a single unit, and 78% of the multi-unit trips involve two units of the same UPC. This formulation also allows us to capture the potential for choice errors if someone chooses a single large packsize of an item even though two units of a smaller packsize might have generated approximately the same quantity more cheaply.

[25] If option 1 is the only packsize available for that item, we instead define option 2 as the product closest to good 1 in terms of size across the entire choice set.

[26] We assign the composite product the average size and average price across the remaining options.

[27] This approach accounts for the effect of promotions on choices. However, since the model is agnostic about price search conduct, it cannot be used to predict the effect of counterfactual changes in promotions except in the extreme case where all products are promoted and, effectively, consumers have complete price information.

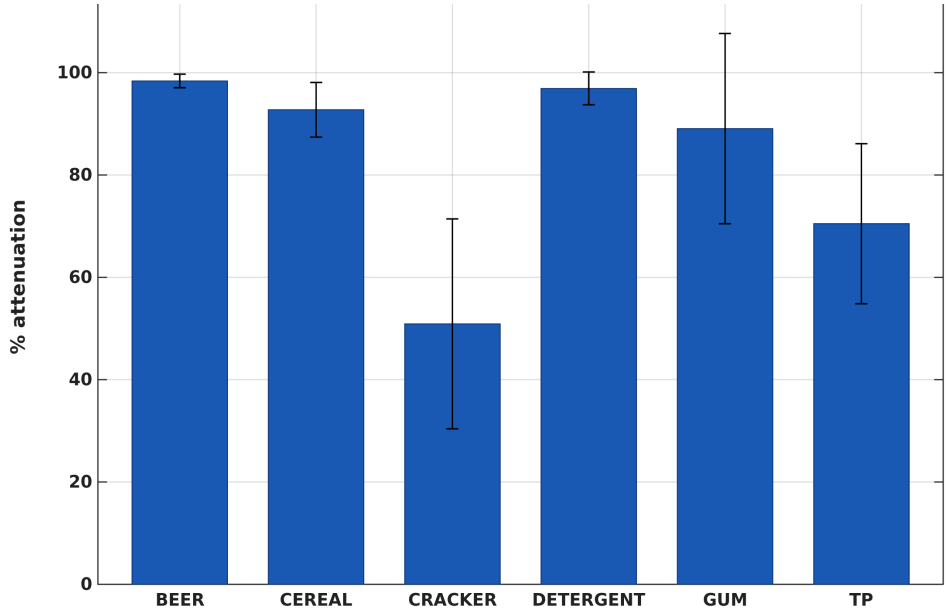
automatically when utility is linear in p and $\widetilde{size}$. Appendix F describes how we leverage partial differential equations results to impose this constraint.

6.4 Test Results

To simplify our notation, we denote the ratios of first and second derivatives as $R_1 \equiv \left| \frac{\partial \sigma_j}{\partial p_j} / \frac{\partial \sigma_j}{\partial \widetilde{size}_j} \right|$ and $R_2 \equiv \left| \frac{\partial^2 \sigma_1}{\partial p_1 \partial p_j} / \frac{\partial^2 \sigma_1}{\partial p_1 \partial \widetilde{size}_j} \right|$, respectively, and let $\hat{R}_1, \hat{R}_2$ denote the corresponding estimates. Recall that R_1 attenuates the MRS under the sign condition (9), which we verified in Section 6.2, and the monotonicity restriction in (10). We now test $\mathbb{H}_0 \colon R_1 = R_2$ against the alternative $\mathbb{H}_1 \colon R_1 \neq R_2$. When we impose the (weak) attenuation constraint, the alternative hypothesis effectively becomes one-tailed, $\mathbb{H}_1 \colon R_1 < R_2$. Imposing the attenuation constraint disciplines the estimates of the relevant derivatives, improving the power of our test.

Figure 3 reports the estimated population means of the test statistics for each of the six categories.[28] The difference between the population mean of R_2 and of R_1 is estimated to be significantly different from zero for all categories. Interestingly, even though we only impose the attenuation constraint when the data supports assumption (9), we nevertheless find that, on average, R_1 tends to be attenuated relative to R_2 for all categories.

Figure 3: Results for Price Inattention Test

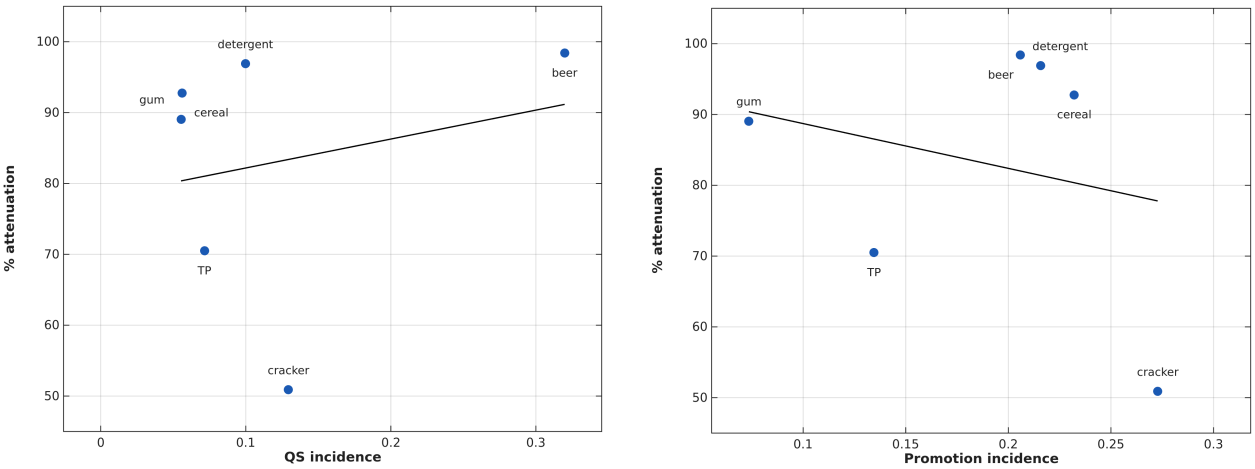


Notes. For each category, this figure plots the estimate of $\frac{R_2 - R_1}{R_2}$ integrated over the distribution of consumer heterogeneity, along with 95% confidence intervals. Positive values indicate attenuation of R_1 relative to R_2.

[28] The population means are obtained by integrating over the population distribution using the cluster approximation of the distribution of tastes.

Next, we explore how the degree of attenuation in R_1 relative to R_2 varies with the incidence of QS and promotions. Figure 4 shows that categories with a higher QS incidence tend to exhibit a higher degree of attenuation, stronger evidence of price inattention. This result is consistent with our finding in section 3.2 that QS is a "surprise" to consumers with rational expectations, making it more likely for price inattention to impact choices during QS trips. By contrast, categories where promotions are more prevalent tend to exhibit lower degrees of attenuation. This results is also intuitive and consistent with our descriptive findings in section 3.3.2, suggesting that promotions reduce price inattention by making prices more salient.

Figure 4: Cross-Category Heterogeneity



Notes. The left plot shows how the attenuation measure $\frac{R_2-R_1}{R_2}$ varies with the incidence of QS, measured as the percentage of all trip/UPC combinations that are QS'd in the choice panel. The right plot shows how the attenuation measure $\frac{R_2-R_1}{R_2}$ varies with the incidence of promotions (measured analogously).

Finally, we investigate how the degree of attenuation, the evidence of price inattention, varies between states with and without unit-price disclosure mandates. Table 10 shows that trips in states with a unit-price mandate tend to exhibit less attenuation. This result is consistent with consumers having to exert effort to search prices: increasing the salience of unit prices through disclosure mandates reduces price inattention. The fact that mandated unit pricing does not fully eliminate the use of QS may stem from the fact that some consumers may not pay attention to *unit pricing* disclosures when making purchase decisions (Harrell et al., 1976).

The results in this section provide strong evidence that consumers are inattentive to prices. We quantify the economic impacts of price inattention in the next subsection.

Table 10: Relationship Between Unit-Price Mandate and Attenuation

	(1)	(2)
Unit-price mandate	-0.134^{***}	-0.070^{***}
	(0.035)	(0.026)
Category fixed effects	No	Yes
Observations	498,583	498,583

Notes. This table shows the results of a regression of the trip-level attenuation estimates on a unit-price mandate dummy (plus a constant). While our R_2 estimates are constrained to be fixed across trips (within a consumer cluster), the attenuation estimates vary across trips since the estimates of R_1 vary. We winsorize the trip-level attenuation measure at the 10%-90% quantiles to reduce the impact of outliers from the nonparametric estimation of the derivative ratios.

7 Parametric Utility Estimates and the Full-Information Counterfactual

Given our evidence of consumer price inattention, we now explore how consumer choices would change if they were fully informed about prices along with the corresponding impact on surplus and retail revenues. To conduct the counterfactual, we impose additional parametric structure on the utilities in (2). First, we assume a nested logit structure on the idiosyncratic shocks ϵ to allow for correlation in unobserved tastes between products that differ only on the basis of packsize. Second, we assume that utility is linear in the product attributes. As shown in ACZ, these restrictions are not crucial for the identification of the utility primitives, but are useful to obtain more precise estimates. We then devise a 2-step estimator of the structural taste coefficients.

7.1 Utility Specification and Parameter Estimates

Consistent with most empirical applications of discrete-choice models, we model utility as a linear function of price and non-price product attributes:

$$U_{ij} = \boldsymbol{\alpha}'\mathbf{x}_j + \beta p_j + \epsilon_{ij}, \tag{18}$$

where $\mathbf{x}_j$ contains $\widetilde{size}_j$, $promo_j$, the number of packs corresponding to option j, and item fixed effects.

Since our items group products that are identical except for packsize, we relax the usual i.i.d. ϵ_{ij} assumption and allow for within-item correlation with a nested logit. We assign each

item to a unique nest containing all the item's available packsizes. Formally, we define the random utilities for a product j that is a specific packsize of item g as: $\epsilon_j = \zeta_g + (1-\lambda)\nu_j$.[29]

To estimate the utility parameters, we follow ACZ and combine the estimates of R_2 from section 6.4 with the price-belief estimates from section 6.1 and estimates of the expected utility parameters. Under the linear utility assumption in (18), the R_2 ratio from the test in section 6.4 identifies $MRS \equiv \frac{\beta}{\alpha_{\widetilde{size}}}$. To obtain estimates of β and $\alpha_{\widetilde{size}}$, we use the fact that in our class of search conducts, consumers always choose the option that maximizes expected utility when $\tilde{p}_j = \tilde{p}_k$ for all j, k (as argued in section 6.2). We can therefore estimate the expected utility coefficients and λ by fitting a (full-information) nested logit model on trips with choice sets exhibiting low variation in $\tilde{p}$ across products. In particular, we can estimate how size affects expected utility, which reflects both the direct impact of size on utility and its impact via price expectations. Given the pice beliefs estimates from section 6.1, we can then disentangle the two channels and separately identify $\alpha_{\widetilde{size}}$ and β. Appendix G provides more details and Figure 5 summarizes the steps of the testing and estimation procedure.

Table 11 reports the point estimates and standard errors for the population mean values of each of the individual utility coefficients across the six categories. Appendix H provides more details on the estimated joint distribution of the coefficients. On average, consumers prefer larger versus smaller pack sizes, conditional on price: $\mathbb{E}(\alpha_{\widetilde{size}}) > 0$. Perhaps most striking is that, across all categories, consumers on average dislike buying multiple smaller packs of an item, as opposed to a single larger pack of the same item: $\mathbb{E}(\alpha_{\#packs}) < 0$. For instance, larger packsizes may be considered more convenient, offering an alternative mechanism to explain the incidence of purchasing QS'd packsizes. This effect is especially strong for beer, which is consistent with the hassle cost of carrying several individual bottles or cans compared to the convenience of the packaging itself for a single commodity bundle or "case". Importantly, our evidence for price inattention is robust to these preferences. In the next section, we will explore the relative economic importance of the convenience mechanism and inattention through a decomposition exercise.

[29] Cardell (1997) shows the conditions under which ϵ is an extreme value random variable when ν is an extreme value random variable. The nest parameter $\lambda \in (0, 1]$ determines the degree of correlation in the random utilities within each nest, where $\lambda = 1$ corresponds to standard logit and where the within-nest correlation increases as λ decreases towards 0.

Figure 5: Testing and Estimation Procedure

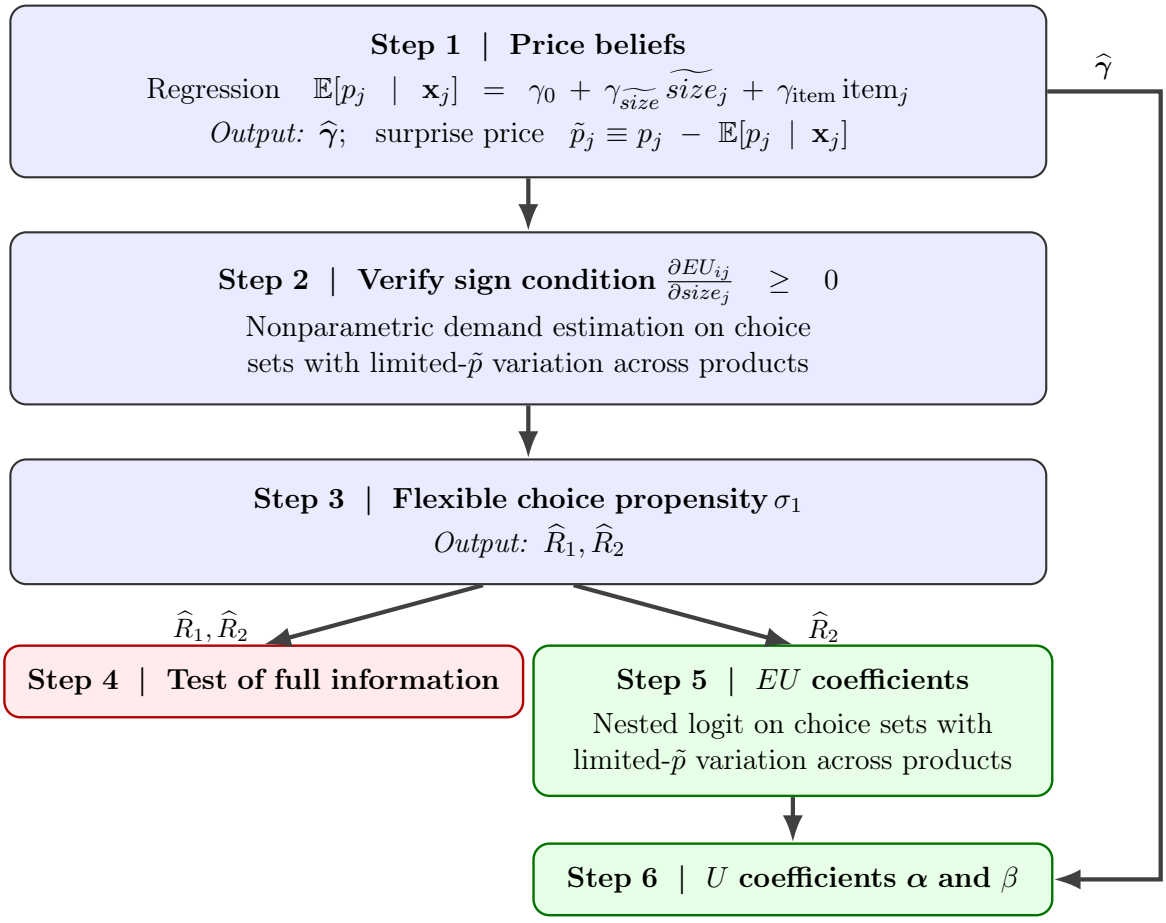


Table 11: Average Utility Parameters

	beer	cereal	cracker	detergent	gum	TP
$\mathbb{E}(\alpha_{\widetilde{size}})$	3.75	3.68	0.48	0.01	0.01	0.07
	(0.85, 7.13)	(3.22, 4.14)	(0.08, 0.93)	$(4.49 \times 10^{-3}, 0.11)$	$(3.17 \times 10^{-4}, 1.18)$	(0.01, 0.29)
$\mathbb{E}(\beta)$	-0.46	-2.06	-0.16	-0.10	-0.11	-0.08
	$(-0.94, -0.06)$	$(-2.33, -1.79)$	$(-0.30, -0.03)$	$(-0.12, -0.09)$	$(-0.91, -0.06)$	$(-0.13, -0.05)$
$\mathbb{E}(\alpha_{promo})$	0.76	0.45	0.58	0.93	0.43	1.18
	(0.31, 0.93)	(0.42, 0.48)	(0.55, 0.61)	(0.86, 1.00)	(0.28, 0.53)	(1.06, 1.26)
$\mathbb{E}(\alpha_{\#\text{ of packs}})$	-2.07	-0.35	-0.70	-0.71	-0.89	-1.04
	$(-2.34, -1.00)$	$(-0.39, -0.30)$	$(-0.75, -0.64)$	$(-0.78, -0.61)$	$(-1.07, -0.68)$	$(-1.17, -0.87)$
$\mathbb{E}(\lambda)$	0.84	0.65	0.68	0.99	0.85	0.98
	(0.41, 0.91)	(0.62, 0.68)	(0.65, 0.72)	(0.95, 1.00)	(0.72, 0.94)	(0.89, 1.00)

Notes. This table reports the point estimates and 95% percentile bootstrap confidence intervals for the utility parameters, averaged across consumer clusters. See Appendix G for implementation details.

7.2 The Drivers of QS Purchases

We use the preference parameters from Table 11 to perform a decomposition of QS purchases. The goal is to quantify the extent to which consumers purchase QS'd products due to price inattention as opposed to other reasons (inconvenience from buying multiple smaller packs, promotions, mismatch between the size of the QS'd product and that of the dominating bundle).

To implement the decomposition, we take every trip with at least one QS'd UPC and compute the probability of choosing a QS'd option under different scenarios (see Appendix I for more details):

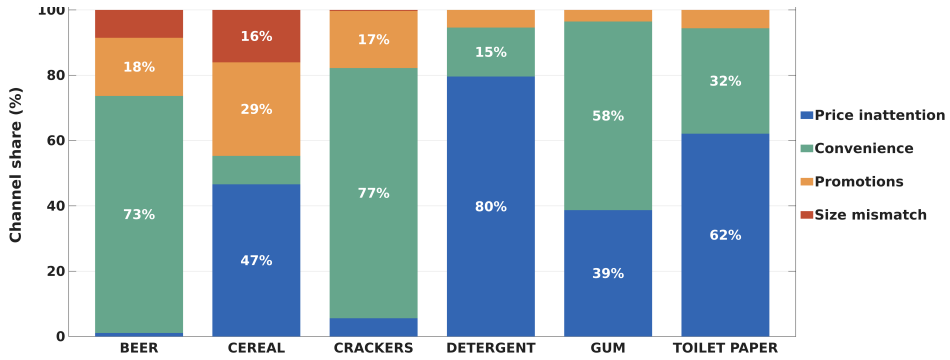
1. Full price information
2. #1 + no disutility from multiple packs
3. #1 + #2 + no promotions for QS'd options
4. #1 + #2 + #3 + no size mismatch between a QS'd option and the cheaper bundles.

Scenarios 1-4 progressively add channels that work against QS purchases. Thus, comparing the probability of a QS purchase across these scenarios allows us to quantify the importance of price inattention relative to competing channels.

Figure 6 shows the results of the decomposition. We find substantial heterogeneity across categories. For beer and crackers, convenience is the main driver of QS purchases. For beer, this result is consistent with our earlier finding that consumers exhibit a strong disutility

from purchasing multiple smaller packs (Table 11). Thus, convenience appears to be the main driver of QS purchases in beer in spite of the evidence of price inattention (Figure 3). In the case of crackers, the result is consistent with our earlier finding of weaker evidence for price inattention relative to other categories (Figure 3). However, for all other categories, the decomposition shows that price inattention seems to be the dominant mechanism, explaining as much as 80% of QS purchases in the case of laundry detergent.[30] Table 12 reports the point estimates and confidence intervals for the percentage of QS purchases attributed to price inattention.

Figure 6: Decomposition of QS Purchases



Notes. For each category, this figure shows a decomposition of the probability of a QS purchases into various channels. See Appendix I for implementation details.

Table 12: Percentage of QS Purchases Attributed to Price Inattention

beer	cereal	cracker	detergent	gum	TP
1.1	46.6	5.6	79.6	38.7	62.1
(0.0, 54.0)	(22.8, 50.6)	(0.7, 21.5)	(74.4, 83.5)	(16.0, 57.6)	(43.9, 72.1)

Notes. This table reports the point estimates and 95% percentile bootstrap confidence intervals for the percentage of QS purchases attributed to price inattention.

7.3 Full Information Counterfactual

Finally, we use the preference parameters from Table 11 to simulate consumers' counterfactual choices under complete and incomplete (status quo) price information, respectively.[31]

[30] In Figure 6, the price inattention channel includes both people who switch to a different packsize of the same item and people who switch to a different item altogether under full information. When we repeat the same exercise conditioning on people who only switch within the same item, we obtain similar results, as shown in Appendix J.

[31] To simulate full-information choices, we plug the preference parameters from Table 11 into the utility specification (2). To simulate status quo choices, we use the same parameters except that we take the price

Table 13: Full-information counterfactual

	Δ revenues ($)		Δ welfare ($)	
Category	Δ	%	Δ	%
beer	-0.05	-0.42%	$+5.91 \cdot 10^{-4}$	+0.00%
cereal	-0.07	-2.44%	+0.01	+0.37%
cracker	-0.03	-0.69%	+0.01	+0.37%
detergent	-2.98	-25.67%	+1.65	+20.74%
gum	-0.19	-8.40%	+0.09	+4.78%
TP	-1.47	-16.52%	+0.64	+9.00%

Notes. This table reports for each category the average per-trip change in surplus and revenues under a full-information counterfactual relative to the status quo. Percentages are computed relative to the status quo value for that outcome, except for welfare, which is expressed as a percentage of the average pack price for the category.

This exercise differs from the QS decomposition in the previous section in two key ways: (i) we consider *all* trips in our data, including those that do not exhibit QS, and (ii) we measure how price inattention affects not just QS purchases, but market outcomes more broadly. For example, this exercises captures the fact that a consumer may be choosing the wrong brand due to price inattention, as opposed to just the wrong size of a given item.

Table 13 summarizes the changes in consumer welfare and retailer revenues under the full-information counterfactual. Full price information boosts consumer welfare, with effects that are again heterogeneous across categories. While for beer, cereal and crackers welfare does not improve materially, the effects are larger for the other categories, reaching over 20% of the average price for laundry detergent. Retailers, in turn, appear to monetize this inattention through revenues that would not survive under full information, falling by as much as 25.7% in the detergent category. Given the relatively low opportunity costs of QS documented in section 3.1, we see that price inattention in these categories is not merely due to inattention to size-related surcharges. For instance, full information reduces detergent revenues by 25%, even though the opportunity cost of QS was 0.3% of sales.

8 Conclusions

We exploit the quantity surcharging in U.S. CPG markets to study consumer price inattention. We find that QS is economically significant and pervasive: it is present in roughly 15% of store-weeks and accounts for as much as 15% of category sales in the categories we

coefficient implied by the naive MRS estimate based on the ratio of *first* derivatives. This approach reflects the lower price sensitivity that consumers exhibit when they do not have full price information.

examine, with U.S. consumers paying an estimated $15.8 billion in unnecessary surcharges in 2018 alone. A novel nonparametric test that is robust to a wide array of search and consideration-set-formation conducts decisively rejects the hypothesis that consumers have full price information. Further, we find that price inattention is the main driver of QS purchases across several categories. Our counterfactual analysis finds that if consumers were fully informed about prices, revenues would decrease on average (by as much as 26% for detergent) and consumer surplus would increase by as much as 21% of the price paid.

Methodologically, our nonparametric test could be applied to any setting in which choices reveal partial information about prices, such as shrinkflation, left-digit bias, price partitioning, drip pricing, shrouded fees, and other forms of inattention that have attracted recent attention in the behavioral and IO literatures.

Our analysis does not account for the underlying mechanisms leading retailers to set QS pricing. Testing the supply-side mechanisms for this practice and whether it merely discriminates deliberately on consumer inattention should be of interest to the academics and regulators. Indeed, many countries have already implemented mandatory *unit pricing* laws such as the EU, UK, and Australia. However, there is no such federal regulation in the U.S., where only a handful of states require unit pricing. On a related note, our counterfactual analysis takes prices and assortments as fixed, abstracting from retailers' equilibrium responses to changes in consumer attentiveness. Future work might explore the equilibrium response on the supply side to policies that offset consumer price inattention.

References

Jason Abaluck, Giovanni Compiani, and Fan Zhang. A method to estimate discrete choice models that is robust to consumer search. *Journal of Political Economy*, 134(7), 2026.

Daniel A Ackerberg. Advertising, learning, and consumer choice in experience good markets: an empirical examination. *International Economic Review*, 44(3):1007–1040, 2003.

Jagdish Agrawal, Pamela E. Grimm, and Narasimhan Srinivasan. Quantity surcharges on groceries. *The Journal of Consumer Affairs*, 27:335–356, 1993.

Eric T. Anderson and Duncan I. Simester. Effects of $9 price endings on retail sales: Evidence from field experiments. *Quantitative Marketing and Economics*, 1(1):93–110, 2003.

T. W. Anderson and D. A. Darling. A test of goodness of fit. *Journal of the American Statistical Association*, 49(268):765–769, 1954.

Mark Armstrong. Nonlinear pricing. *Annual Review of Economics*, 8:583–614, 2016.

Robert C. Blattberg and Scott A. Neslin. *Sales Promotion: Concepts, Methods, and Strategies*. Prentice Hall, Englewood Cliffs, NJ, 1990.

Stéphane Bonhomme and Elena Manresa. Grouped patterns of heterogeneity in panel data. *Econometrica*, 83(3):1147–1184, 2015.

A. E. Brockwell. Universal residuals: A multivariate transformation. *Statistics & Probability Letters*, 77(14):1473–1478, 2007.

Bart J Bronnenberg and Wilfried R Vanhonacker. Limited choice sets, local price response, and implied measures of price competition. *Journal of Marketing Research*, 33(2):163–173, 1996.

Bart J. Bronnenberg, Jean-Pierre Dubé, and Sridhar Moorthy. The economics of brands and branding. In *Handbook of the Economics of Marketing*, volume 1, chapter 6, pages 291–358. Elsevier, 2019. doi: 10.1016/bs.hem.2019.04.003.

Bart J. Bronnenberg, Jun B. Kim, and Carl F. Mela. Zooming in on choice: How do consumers search for cameras online? *Marketing Science*, forthcoming.

Zach Y. Brown and Jihye Jeon. Endogenous information and simplifying insurance choice. *Econometrica*, 92(3):881–911, May 2024.

N. Scott Cardell. Variance components structures for the extreme-value and logistic distributions with applications to models of heterogeneity. *Econometric Theory*, 13(2):185–213, 1997.

Marco Catena, Simon Land, Pieter Reynders, and Martino Zizioli. How precision revenue growth management transforms cpg promotions. *McKinsey Quarterly*, July:1–5, 2021.

Raj Chetty, Adam Looney, and Kory Kroft. Salience and taxation: Theory and evidence. *American Economic Review*, 99(4):1145–1177, September 2009.

Michael Choi, Anovia Yifan Dai, and Kyungmin Kim. Consumer search and price competition. *Econometrica*, 86(4):1257–1281, 2018. doi: 10.3982/ECTA14837.

Jae Hyen Chung, Pradeep Chintagunta, and Sanjog Misra. Simulated maximum likelihood estimation of the sequential search model. *Quantitative Marketing and Economics*, 23(1): 105–164, 2025.

Sofronis Clerides and Pascal Courty. Sales, quantity surcharge and consumer inattention. *The Review of Economics and Statistics*, 99(2):357–370, 2017.

Giovanni Compiani. Market counterfactuals and the specification of multiproduct demand: A nonparametric approach. *Quantitative Economics*, 13(2):545–591, 2022.

Gregory S Crawford and Matthew Shum. Uncertainty and learning in pharmaceutical demand. *econometrica*, 73(4):1137–1173, 2005.

Claudia Czado, Tilmann Gneiting, and Leonhard Held. Predictive model assessment for count data. *Biometrics*, 65(4):1254–1261, 2009.

Babur de los Santos, Ali Hortacsu, and Matthijs R. Wildenbeest. Testing models of consumer search using data on web browsing and purchasing behavior. *American Economic Review*, 102(6):2955–2980, 2012.

Stefano DellaVigna. *Handbook of Behavioral Economics*, volume 1, chapter Structural Behavioral Economics, pages 613–723. North-Holland, 2018.

Stefano DellaVigna and Ulrike Malmendier. Paying not to go to the gym. *American Economic Review*, 96(3):694–719, 2006. doi: 10.1257/aer.96.3.694. URL https://www.aeaweb.org/articles?id=10.1257/aer.96.3.694.

Peter R. Dickson and Alan G. Sawyer. The price knowledge and search of supermarket shoppers. *Journal of Marketing*, 54:42–53, 1990.

Tülin Erdem and Michael P Keane. Decision-making under uncertainty: Capturing dynamic brand choice processes in turbulent consumer goods markets. *Marketing science*, 15(1): 1–20, 1996.

European Union. Price indications on consumer products. https://europa.eu, 2026. Accessed: 2026-07-01.

Lawrence C Evans. *Partial differential equations*, volume 19. American mathematical society, 2022.

Amy Finkelstein. E-ZTax: Tax salience and tax rates. *The Quarterly Journal of Economics*, 124(3):969–1010, August 2009.

Xavier Gabaix and David Laibson. Shrouded attributes, consumer myopia, and information suppression in competitive markets. *The Quarterly Journal of Economics*, 121(2):505–540, May 2006. doi: 10.1162/qjec.2006.121.2.505.

Xavier Gabaix, David Laibson, Guillermo Moloche, and Stephen Weinberg. Costly information acquisition: Experimental analysis of a boundedly rational model. *American Economic Review*, 96(4):1043–1068, 2006.

Eitan Gerstner and James D. Hess. Why do hot dogs come in packs of 10 and buns in 8s or 12s? a demand-side investigation. *The Journal of Business*, 60(4):491–517, 1987. doi: 10.1086/296410.

Amit K. Ghosh. The influence of price and promotion on package size propensity. *American Journal of Business*, 32(2):93–103, 2017.

Gautam Gowrisankaran and Marc Rysman. Dynamics of consumer demand for new durable goods. *Journal of political Economy*, 120(6):1173–1219, 2012.

C.W. Granger and A. Bilson. Consumers' attitudes toward package size and price. *The Journal of Marketing Research*, 9(8):239–248, 1972.

Eric A. Greenleaf, Eric J. Johnson, Vicki G. Morwitz, and Edith Shalev. Nonlinear pricing. *Journal of Consumer Psychology*, 26(1):105–124, 2016.

Rafael P. Greminger. Optimal search and discovery. *Management Science*, 68(5):3904–3924, 2022. doi: 10.1287/mnsc.2021.4085.

Laura Grigolon, Ian Meeker, and Jakob von Ditfurth. Does peter piper pick pepper inattentively? consumer inattention to package content. 2026. Working Paper.

Michael D. Grubb and Matthew Osborne. Cellular service demand: Biased beliefs, learning, and bill shock. *American Economic Review*, 105(1):234–271, 2015.

Peter M. Guadagni and John D.C. Little. A logit model of brand choice calibrated on scanner data. *Marketing Science*, 2:203–238, 1983.

Gilbert D. Harrell, Michael D. Hutt, and John W. Allen. *Universal Product Code: Price Removal and Consumer Behavior in Supermarkets*. MSU Business Studies. Division of Research, Graduate School of Business Administration, Michigan State University, East Lansing, MI, 1976. ISBN 9780877441410.

John R. Hauser. Consideration-set heuristics. *Journal of Business Research*, 67(8):1688–1699, 2014.

John R. Hauser and Birger Wernerfelt. The Competitive Implications of Relevant-Set/Response Analysis. *Journal of Marketing Research*, 28(4):391–407, 1991. doi: 10.1177/002224379102800401. URL https://sagepub.com.

Igal Hendel and Aviv Nevo. Measuring the implications of sales and consumer inventory behavior. *Econometrica*, 74(6):1637–1673, 2006.

Elisabeth Honka. Quantifying search and switching costs in the u.s. auto insurance industry. *Rand Journal of Economics*, 45:847–884, 2014.

Elisabeth Honka, Ali Hortaçsu, and Maria Ana Vitorino. Advertising, consumer awareness, and choice: Evidence from the us banking industry. *The RAND Journal of Economics*, 48(3):611–646, 2017.

Ali Hortaçsu and Chad Syverson. Product differentiation, search costs, and competition in the mutual fund industry: A case study of s&p 500 index funds. *The Quarterly journal of economics*, 119(2):403–456, 2004.

John A. Howard and Jagdish N. Sheth. *The Theory of Buyer Behavior*. John Wiley & Sons, New York, 1969.

Aljoscha Janssen and Johannes Kasinger. Shrinkflation and consumer demand. *Marketing Science*, 2025. doi: 10.1287/mksc.2024.0948. Working paper version available at SSRN: https://ssrn.com/abstract=4783491.

Zhenling Jiang, Tat Chan, Hai Che, and Youwei Wang. Consumer search and purchase: An empirical investigation of retargeting based on consumer online behaviors. *Marketing Science*, 40(2):219–240, 2021.

Joonhwi Joo. Rational inattention as an empirical framework for discrete choice and consumer-welfare evaluatio. *Journal of Marketing Research*, 60:278–298, 2023.

M. G. Kendall. A new measure of rank correlation. *Biometrika*, 30(1/2):81–93, 1938.

In Kyung Kim. Consumers' preference for downsizing over package price increases. *Journal of Economics & Management Strategy*, 33(1):25–52, 2024. doi: 10.1111/jems.12548.

Jun B. Kim, Paulo Albuquerque, and Bart J. Bronnenberg. Online demand under limited consumer search. *Marketing Science*, 29(6):1001–1023, 2010.

Jun B. Kim, Paulo Albuquerque, and Bart J. Bronnenberg. The probit choice model under sequential search with an application to online retailing. *Management Science*, 63(11): 3911–3929, 2017. doi: 10.1287/mnsc.2016.2545.

Anja Lambrecht and Bernd Skiera. Paying too much and being happy about it: Existence, causes, and consequences of tariff-choice biases. *Journal of Marketing Research*, 43:212–223, 2006.

Youngeun Lee. Shrinkflation: Evidence on product downsizing and consumer response. Working paper, Kilts Center at Chicago Booth Marketing Data Center, September 2024. URL https://ssrn.com/abstract=5053745.

John A. List, Ian Muir, Devin Pope, and Gregory Sun. Left-digit bias at Lyft. *The Review of Economic Studies*, 90(6):3186–3237, November 2023. doi: 10.1093/restud/rdad014.

Yihan Ma. Annual consumer packaged goods (cpg) e-commerce sales in the united states in 2017 and 2018. *Statisa*, November 2025. URL https://www.statista.com/statistics/609822/cpg-e-commerce-sales-forecast-usa/.

Kenneth C. Manning, David E. Sprott, and Anthony D. Miyazaki. Consumer response to quantity surcharges: Implications for retail price setters. *The Journal of Retailing*, 74(3): 173–399, 1998.

Filip Matejka and Alisdair McKay. Rational inattention to discrete choices: A new foundation for the multinomial logitmodel. *American Economic Review*, 105(1):272–298, 2015.

Nitin Mehta. Investigating consumers' purchase incidence and brand choice decisions across multipleproduct categories: A theoretical and empirical analysis. *Marketing Science*, 26 (2):196–217, 2007.

James Melton. Us online cpg sales grew 35.4 *Digital Commerce*, February 2019. URL https://www.digitalcommerce360.com/2019/02/08/online-cpg-sales-in-the-us-grew-35-4-in-2018-and-represented-11-of-the-total-market/.

Ilya Morozov, Stephan Seiler, Xiaojing Dong, and Liwen Hou. Estimation of preference heterogeneity in markets with costly search. *Marketing Science*, 40(5):871–899, 2021. doi: 10.1287/mksc.2021.1287.

Olivia R. Natan. Choice frictions in large assortments. *Marketing Science*, 44(3):491–732, 2025.

National Institute of Standards and Technology. Unit pricing guide: A best practice approach to unit pricing. Technical Report NIST Special Publication 1181, U.S. Department of Commerce, 2015. URL https://www.nist.gov/system/files/documents/2017/04/28/SP1181-Unit-Pricing-Guide.pdf.

NIST. Unit pricing guide: A best practice approach to unit pricing. Technical Report NIST Special Publication 1181, U.S. Department of Commerce, 2015. URL https://nist.gov.

John H. Roberts and James M. Lattin. Development and testing of a model of consideration set composition. *Journal of Marketing Research*, 28(4):429–440, 1991.

Murray Rosenblatt. Remarks on a multivariate transformation. *The Annals of Mathematical Statistics*, 23(3):470–472, 1952.

Stephen P Ryan. The costs of environmental regulation in a concentrated industry. *Econometrica*, 80(3):1019–1061, 2012.

David Sefcik. The unit price is right. NIST Taking Measure Blog, March 2017. URL https://www.nist.gov/blogs/taking-measure/unit-price-right. Published March 1, 2017.

Stephan Seiler. The impact of search costs on consumer behavior: A dynamic approach. *Quantitative Marketing and Economics*, 2013.

Herbert A Simon. A behavioral model of rational choice. *The quarterly journal of economics*, pages 99–118, 1955.

V. P. Singh, Karsten T. Hansen, and Sunil Gupta. Modeling preferences for common attributes in multicategory brand choice. *Journal of Marketing Research*, 42(2):196–209, 2005. doi: 10.1509/jmkr.42.2.196.62288.

David E. Sprott, Kenneth C. Manning, and Anthony D. Miyazaki. Grocery price setting and quantity surcharges. *Journal of Marketing*, 67:34–46, 2003.

George J. Stigler. The economics of information. *Journal of Political Economy*, 69(3): 213–225, 1961.

Avner Strulov-Shlain. More than a penny's worth: Left-digit bias and firm pricing. *The Review of Economic Studies*, 90(5):2612–2645, October 2023. doi: 10.1093/restud/rdac082.

Peter Stüttgen, Peter Boatwright, and Robert T Monroe. A satisficing choice model. *Marketing Science*, 31(6):878–899, 2012.

Raluca M. Ursu. The power of rankings: Quantifying the effect of rankings on online consumer search and purchase decisions. *Marketing Science*, 37(4):530–552, 2018. doi: 10.1287/mksc.2017.1072.

Lucio Wasserman and Pedro Pires. Shrinkflation: What happens when package sizes change? *Kilts Center at Chicago Booth Marketing Data Center Paper*, 2025.

Martin L. Weitzman. Optimal search for the best alternative. *Econometrica*, 47(3):641–654, 1979.

William D. Wells and Leonard A. LoSciuto. Direct observation of purchasing behavior. *Journal of Marketing Research*, 3(3):227–233, 1966. doi: 10.1177/002224376600300301.

Stanley M. Widrick. Measurement of incidents of quantity surcharge among selected grocery products. *The Journal of Consumer Affairs*, 13(1):99–107, 1979.

A Data Appendix

A.1 Product Attributes

We analyze six product categories: beer, cereal, cracker, chewing gum, laundry detergent, and toilet paper. For each category, we classified each UPC into a unique product by partly matching with Syndigo database and also by manually querying external UPC databases.

A.2 Construction of Choice Panel

1. We pulled and matched Nielsen Homescan and RMS using store_code_uc and the week of transaction.

2. Purchase price is taken from RMS

For the structural estimation, we drop trips that have less than three UPCs. This allows us to estimate choice probabilities as a function of the attributes of three options (good 1, the other option we use to take our cross derivatives and a composite option that summarizes the other UPCs available). This filter removes a very small fraction of trips for each category. Specifically:

- Beer: 7 (0.015%)
- Cereal: 31 (0.014%)
- Crackers: 15 (0.011%)
- Detergent: 79 (0.174%)
- Gum: 0 (0%)
- Toilet paper: 545 (0.995%).

Further, for beer, we remove kegs and drafts, and for gum, we remove one very large UPC that has 912 sticks and is an outlier relative to the rest of the UPCs in the category.

B Further Descriptive Evidence

B.1 Descriptive Analysis With Other QS Modes

We now provide descriptive statistics for QS under three alternative definitions than the one used in the body of the text:

Table 14: Retention of Choice-Panel Trips during Sample Construction

Category	Selection stage	# trips (all years)	Trips retained (% of loaded)
beer	Loaded, pre-trim	942,481	100.0
	After single-UPC-trip filter	838,294	88.9
	After audited-store filter	169,115	17.9
cereal	Loaded, pre-trim	3,734,541	100.0
	After single-UPC-trip filter	2,455,223	65.7
	After audited-store filter	771,225	20.7
cracker	Loaded, pre-trim	2,468,241	100.0
	After single-UPC-trip filter	1,934,283	78.4
	After audited-store filter	438,486	17.8
gum	Loaded, pre-trim	723,886	100.0
	After single-UPC-trip filter	608,191	84.0
	After audited-store filter	115,826	16.0
laundry detergent	Loaded, pre-trim	1,040,462	100.0
	After single-UPC-trip filter	965,724	92.8
	After audited-store filter	154,322	14.8
toilet paper	Loaded, pre-trim	1,656,633	100.0
	After single-UPC-trip filter	1,618,833	97.7
	After audited-store filter	205,924	12.4

Notes. This table reports the number of trips and the fraction of trips retained at each stage of our sample-construction process. "Loaded, pre-trim" corresponds to all matched trips during which a positive quantity of at least one UPC is purchased in a given category. "After single-UPC-trip filter" retains only those trips during which a positive quantity of a single UPC is purchased in a given category. "After audited-store filter" retains only trips to stores that report their in-store promotion behavior. Percentages are calculated relative to the loaded pre-trim sample and rounded to one decimal place.

1. $\mathcal{QS}_{exact}$: The exact quantity of the QS'ed larger packsize could be obtained more cheaply by purchasing multiple units of a smaller packsize of the same commodity.

2. $\mathcal{QS}_{half}$: At least one smaller packsize less than half the size of the QS'ed larger packsize of the same commodity is available more cheaply.

3. $\mathcal{QS}_{any}$: At least one smaller packsize is available for a lower unit price than that of the larger packsize of the same commodity.

Below we replicate Tables 4 and 5 using each of our three alternative QS variable definitions.

Table 15: QS Incidence using $\mathcal{QS}_{exact}$

category	(1) at least one $\mathcal{QS}$	(2) $\mathcal{QS}$ no. items	(3) total no. items	(4) $\mathcal{QS}$ share (2÷3)	(5) $\mathcal{QS}$ no. top 5% items	(6) total no. top 5% items	(7) $\mathcal{QS}$ share of top 5% items (5÷3)
beer	0.717	3.84	64.38	0.060	3.80	53.49	0.059
cereal	0.060	0.07	83.11	0.001	0.07	47.23	0.001
cracker	0.029	0.03	47.06	0.001	0.03	29.40	0.001
gum	0.487	1.05	40.35	0.026	0.99	26.95	0.025
laundry detergent	0.259	0.58	43.93	0.013	0.53	23.59	0.012
toilet paper	0.345	0.62	17.32	0.036	0.52	7.06	0.030

Notes. This table summarizes $\mathcal{QS}_{exact}$ by product module in the 2018 RMS data using our granular item definitions. Col. (1) reports the observed proportion of store-weeks during which at least one $\mathcal{QS}_{exact}$ item is available. Cols. (2) to (4) report the average number of items on the shelf in a store-week, how many of them have at least one surcharged packsize, and the $\mathcal{QS}_{exact}$ item share. Cols. (5) to (7) report the analogous statistics for the subset of items comprising the top 5% on the basis of annual category-year dollar sales. All statistics are calculated using the entire sample of RMS stores in 2018.

Table 16: $\mathcal{QS}_{exact}$ Sales by Category, 2018

category	(1) \$ sales	(2) $\mathcal{QS}_{exact}$ \$ sales	(3) $\mathcal{QS}_{exact}$ opportunity cost	(4) Col (2)/(1)	(5) Col (3)/(1)
beer	\$ 9,200.6M	\$ 588.8M	\$ 54.5M	0.064	0.006
cereal	\$ 4,235.1M	\$ 4.7M	\$ 0.9M	0.001	0.000
cracker	\$ 2,450.8M	\$ 1.5M	\$ 0.3M	0.001	0.000
gum	\$ 976.0M	\$ 14.9M	\$ 3.2M	0.015	0.003
laundry detergent	\$ 2,856.6M	\$ 59.7M	\$ 8.2M	0.021	0.003
toilet paper	\$ 3,271.6M	\$ 96.3M	\$ 9.5M	0.029	0.003

Notes. This table reports the dollar-amount sales subject to $\mathcal{QS}_{exact}$ in the 2018 RMS data using our granular item definitions. Dollar figures are expressed in millions of dollars (\$ M). Shares are $\mathcal{QS}_{exact}$ dollar sales divided by total dollar sales.

Table 17: QS Incidence using $\mathcal{QS}_{half}$

category	(1) at least one $\mathcal{QS}$	(2) $\mathcal{QS}$ no. items	(3) total no. items	(4) $\mathcal{QS}$ share (2÷3)	(5) $\mathcal{QS}$ no. top 5% items	(6) total no. top 5% items	(7) $\mathcal{QS}$ share of top 5% items (5÷3)
beer	0.822	6.52	64.38	0.101	6.47	53.49	0.100
cereal	0.190	0.35	83.11	0.004	0.34	47.23	0.004
cracker	0.298	0.45	47.06	0.009	0.44	29.40	0.009
gum	0.652	1.84	40.35	0.045	1.77	26.95	0.044
laundry detergent	0.475	1.30	43.93	0.030	1.18	23.59	0.027
toilet paper	0.417	0.79	17.32	0.046	0.62	7.06	0.036

Notes. This table summarizes $\mathcal{QS}_{half}$ by product module in the 2018 RMS data using our granular item definitions. Col. (1) reports the observed proportion of store-weeks during which at least one $\mathcal{QS}_{half}$ item is available. Cols. (2) to (4) report the average number of items on the shelf in a store-week, how many of them have at least one surcharged packsize, and the $\mathcal{QS}_{half}$ item share. Cols. (5) to (7) report the analogous statistics for the subset of items comprising the top 5% on the basis of annual category-year dollar sales. All statistics are calculated using the entire sample of RMS stores in 2018.

Table 18: $\mathcal{QS}_{half}$ Sales by Category, 2018

category	(1) $ sales	(2) $\mathcal{QS}_{half}$ $ sales	(3) $\mathcal{QS}_{half}$ opportunity cost	(4) Col (2)/(1)	(5) Col (3)/(1)
beer	$ 9,200.6M	$ 1,129.3M	$ 123.2M	0.123	0.013
cereal	$ 4,235.1M	$ 25.2M	$ 4.1M	0.006	0.001
cracker	$ 2,450.8M	$ 60.0M	$ 12.3M	0.024	0.005
gum	$ 976.0M	$ 38.1M	$ 6.8M	0.039	0.007
laundry detergent	$ 2,856.6M	$ 104.1M	$ 16.4M	0.036	0.006
toilet paper	$ 3,271.6M	$ 119.6M	$ 15.4M	0.037	0.005

Notes. This table reports the dollar-amount sales subject to $\mathcal{QS}_{half}$ in the 2018 RMS data using our granular item definitions. Dollar figures are expressed in millions of dollars ($ M). Shares are $\mathcal{QS}_{half}$ dollar sales divided by total dollar sales.

Table 19: QS Incidence using $\mathcal{QS}_{any}$

category	(1) at least one $\mathcal{QS}$	(2) $\mathcal{QS}$ no. items	(3) total no. items	(4) $\mathcal{QS}$ share (2÷3)	(5) $\mathcal{QS}$ no. top 5% items	(6) total no. top 5% items	(7) $\mathcal{QS}$ share of top 5% items (5÷3)
beer	0.906	8.70	64.38	0.135	8.64	53.49	0.134
cereal	0.439	3.13	83.11	0.038	3.10	47.23	0.037
cracker	0.537	1.87	47.06	0.040	1.83	29.40	0.039
gum	0.711	2.30	40.35	0.057	2.24	26.95	0.055
laundry detergent	0.645	2.95	43.93	0.067	2.75	23.59	0.063
toilet paper	0.588	1.51	17.32	0.087	1.16	7.06	0.067

Notes. This table summarizes $\mathcal{QS}_{any}$ by product module in the 2018 RMS data using our granular item definitions. Col. (1) reports the observed proportion of store-weeks during which at least one $\mathcal{QS}_{any}$ item is available. Cols. (2) to (4) report the average number of items on the shelf in a store-week, how many of them have at least one surcharged packsize, and the $\mathcal{QS}_{any}$ item share. Cols. (5) to (7) report the analogous statistics for the subset of items comprising the top 5% on the basis of annual category-year dollar sales. All statistics are calculated using the entire sample of RMS stores in 2018.

Table 20: $\mathcal{QS}_{any}$ Sales by Category, 2018

category	(1) $ sales	(2) $\mathcal{QS}_{any}$ $ sales	(3) $\mathcal{QS}_{any}$ opportunity cost	(4) Col (2)/(1)	(5) Col (3)/(1)
beer	$ 9,200.6M	$ 1,956.2M	$ 217.7M	0.213	0.024
cereal	$ 4,235.1M	$ 282.1M	$ 54.8M	0.067	0.013
cracker	$ 2,450.8M	$ 157.9M	$ 46.3M	0.064	0.019
gum	$ 976.0M	$ 53.4M	$ 9.4M	0.055	0.010
laundry detergent	$ 2,856.6M	$ 237.7M	$ 26.1M	0.083	0.009
toilet paper	$ 3,271.6M	$ 248.6M	$ 29.1M	0.076	0.009

Notes. This table reports the dollar-amount sales subject to $\mathcal{QS}_{any}$ in the 2018 RMS data using our granular item definitions. Dollar figures are expressed in millions of dollars ($ M). Shares are $\mathcal{QS}_{any}$ dollar sales divided by total dollar sales.

B.2 RMS Descriptives, QS Identified Using Only NIQ Attributes

Table 21: $\mathcal{QS}$ Incidence Across All RMS Product Departments, 2018

department	(1) $ sales	(2) $\mathcal{QS}$ sales ($)	(3) $\mathcal{QS}$ opportunity cost ($)	(4) $\mathcal{QS}$ sales share (2÷1)	(5) $\mathcal{QS}$ opp. cost share (3÷1)
Health & Beauty Care	$ 38,656.7M	$ 3,220.0M	$ 1,174.0M	0.083	0.030
Dry Grocery	$ 124,190.4M	$ 18,318.7M	$ 5,180.6M	0.148	0.042
Frozen Foods	$ 22,916.2M	$ 2,735.9M	$ 633.8M	0.119	0.028
Dairy	$ 27,282.0M	$ 3,220.1M	$ 677.0M	0.118	0.025
Deli	$ 8,202.6M	$ 1,949.6M	$ 615.5M	0.238	0.075
Packaged Meat	$ 11,007.5M	$ 1,768.5M	$ 479.4M	0.161	0.044
Fresh Produce	$ 13,398.5M	$ 616.0M	$ 137.0M	0.046	0.010
Non-Food Grocery	$ 45,629.3M	$ 6,042.6M	$ 1,443.2M	0.132	0.032
Alcoholic Beverages	$ 22,043.1M	$ 1,673.7M	$ 200.7M	0.076	0.009
General Merchandise	$ 14,610.8M	$ 891.5M	$ 311.8M	0.061	0.021
Total	$ 327,937.0M	$ 40,436.5M	$ 10,853.0M	0.123	0.033

Notes. This table summarizes $\mathcal{QS}$ across product departments in the 2018 RMS data using the NIQ attributes for item definitions. Each row corresponds to a product department. Dollar figures are expressed in millions of dollars ($ M). Columns (4) and (5) normalize QS_{margin} dollar sales and QS_{margin} opportunity cost by total dollar sales.

Table 22: QS Calculated NIQ Attributes by Focal Category, 2018

category	(1) $ sales	(2) QS_{margin} $ sales	(3) QS_{margin} opportunity cost	(4) Col (2)/(1)	(5) Col (3)/(1)
beer	$ 9,200.6M	$ 1,340.7M	$ 152.7M	0.146	0.017
cereal	$ 4,237.4M	$ 325.7M	$ 91.1M	0.077	0.021
cracker	$ 2,471.4M	$ 253.7M	$ 54.4M	0.103	0.022
gum	$ 976.0M	$ 84.5M	$ 17.5M	0.087	0.018
laundry detergent	$ 2,856.7M	$ 289.3M	$ 47.7M	0.101	0.017
toilet paper	$ 3,273.8M	$ 760.5M	$ 266.7M	0.232	0.081

Notes. This table reports the extent of potential quantity surcharging under QS_{margin} in the RMS data across the focal categories. Dollar figures are expressed in millions of dollars ($ M). The final column reports QS_{margin} opportunity cost normalized by total dollar sales.

B.3 Price Versus Packsize Assortment Variation

In our structural analysis, we assume consumers must search a product's price to consider it for purchase, but already knows the packsize prior to search. We now show this assumption is consistent with the considerably higher degree of variation in an item's prices versus its packsize assortment in our data. We use the RMS data at the UPC-store-week level, grouping the UPCs by their item classification.

Table 23 shows that an item's available packsize assortments are far more stable than prices in every category. For a given item-store, we observe a change in the set of packsizes available for only 0.5-3.8% of weeks, about once in six months (toilet paper) to once in four years (crackers). At least one packsize's price changes in 33-51% of weeks, about once in two to three weeks. Prices change more often than sizes for 77-95% of item-stores, so the gap holds across most brands rather than being driven by a few.

These results are consistent with a consumer knowing the packsizes at the start of a trip, but still needing to search and browse the prices.

Table 23: Price vs. Packsize Assortment Variation

	(1) beer	(2) cereal	(3) cracker	(4) gum	(5) laundry detergent	(6) toilet paper
Panel A: Weekly frequency of change						
Prob. set of sizes changes	0.026	0.009	0.005	0.019	0.020	0.038
Prob. any size's price changes	0.326	0.406	0.445	0.365	0.459	0.509
Difference (price - sizes)	0.300	0.397	0.440	0.346	0.439	0.471
	(0.001)	(0.001)	(0.001)	(0.001)	(0.001)	(0.001)
Share of brand-stores with price > sizes	0.770	0.920	0.914	0.822	0.938	0.948
Brand-stores	3,667,507	6,004,534	3,605,480	3,406,544	3,055,694	1,220,853
Brand-store-weeks	82,401,878	144,541,273	88,420,605	70,147,295	57,314,161	30,517,557
Panel B: Share of variance in Δ log price per unit						
Within-size prices (1-week change)	0.942	0.994	0.998	0.987	0.982	0.978
Within-size prices (13-week change)	0.919	0.988	0.998	0.979	0.980	0.965
Pack-size mix (1-week change)	0.058	0.006	0.002	0.013	0.018	0.022
Pack-size mix (13-week change)	0.081	0.012	0.002	0.021	0.020	0.035

Notes. This table uses the NIQ RMS data at the UPC-store-week level for 2018. The item groups UPCs with identical attributes that differ only in their packsize and price. Since RMS only reports data for UPC-store-weeks with non-zero sales, we assume a packsize missing for four weeks or less was still available if the item had positive sales of any UPC. Panel A reports weekly probabilities that an item-store's set of packsizes changes and that any packsize's per-unit price changes by more than 1%. These are averaged across item-stores, weighted by number of weeks. Standard errors, clustered by store, are in parentheses. Panel B decomposes changes in the item's log price per unit, using fixed sample-volume weights across packsizes, into within-size price changes and packsize mix (sizes added or dropped). It reports each component's share of the variance; the two shares sum to one.

C Descriptive Tests for QS and QD

C.1 Descriptive tests for rarity of QS and pervasiveness of QD

For both tests, we use a rank-based probability integral transform (PIT) test. Both tests are defined relative to an exhangeability null hypothesis that packsize is uninformative about a given item's prices. We use a randomized PIT test that is exact even with a discrete, rank-based variable.

The randomized PIT. For a statistic T with discrete null c.d.f. F_0, the randomized probability integral transform

$$U \;=\; F_0(T^-) + V\big[F_0(T) - F_0(T^-)\big], \qquad V \sim \mathrm{Unif}(0,1), \qquad U \perp\!\!\!\perp V, \tag{19}$$

satisfies $U \sim \mathrm{Unif}(0,1)$ exactly under the null (Rosenblatt, 1952; Brockwell, 2007; Czado et al., 2009). Here $F_0(T^-)$ is the null probability of a value strictly below the observed T, and $F_0(T) - F_0(T^-)$ is the null mass at T. Randomization is what makes the transform exact despite ties. Inference can be conducted using a block-bootstrap that re-samples store-week-items with replacement.

Quantity Surcharging For each item, let the unit of analysis be a store-week-item-packsize combination c. Denote the focal large-packsize product's unit price by u_c^L and denote by $\{u_{c,1}^S, \ldots, u_{c,k_c}^S\}$ for its k_c feasible smaller substitutes. Define the continuous (log) surcharge margin as follows:

$$\Delta_c \;=\; \log u_c^L - \operatorname*{logmin}_j u_{c,j}^S.$$

We test the sharp null that packsize is uninformative about the unit price: the k_c+1 unit prices are *exchangeable* for any c. Under the null, the rank R_c of the large pack's unit price among its k_c+1 comparables is uniform on $\{1,\ldots,k_c+1\}$. We then apply the *randomized* probability integral transform (PIT):

$$U_c \;=\; \frac{R_c - 1 + V_c}{k_c + 1}, \qquad V_c \sim \mathrm{Unif}(0,1), \qquad U_c \perp\!\!\!\perp V_c,$$

which satisfies $U_c \sim \mathrm{Uniform}(0,1)$ under the null. Quantity discounting pushes the large pack toward the cheapest rank ($U_c \to 0$); a surcharge is the opposite, upper tail ($U_c \to 1$).

Pooling offerings c, the empirical CDF of the margin locates the *break-even percentile* $q^\star = \widehat{\Pr}(\Delta_c \le 0)$, the quantile corresponding to equal unit pricing. The surcharge prevalence

consists of the upper-tail mass $1-q^\star = \widehat{\Pr}(\Delta_c > 0)$. We benchmark against the rate a price-neutral (exchangeable) null would generate, $\mathbb{E}[k_c/(k_c+1)]$, and test uniformity of $\{U_c\}$ with the tail-weighted Anderson-Darling statistic (Anderson and Darling, 1954). For inference, we block bootstrap on the store-week-item clusters. Quantity surcharging is a tail event when $q^\star$ is large (the break-even point sits high in the margin distribution) and the observed tail mass falls well below the exchangeability benchmark.

Quantity Discounting In contrast with the surcharge test which isolates a single dominated large pack, quantity discounting pertains to the *entire* size ladder of an item. So we test discounting at the level of the store-week-item, using all packsizes present. Within a cell, let the distinct sizes be $x_1 < \cdots < x_K$ with unit prices $u_1, \ldots, u_K$. We measure the concordance between size and unit price with Kendall's rank correlation Kendall (1938),

$$\tau_c \;=\; \frac{\#\{\text{concordant pairs}\} - \#\{\text{discordant pairs}\}}{\binom{K}{2}},$$

using the tie-corrected form τ_b if unit prices coincide. Quantity discounting corresponds to $\tau_c < 0$ (larger packs cheaper per unit), and a perfectly monotone volume-discount schedule is $\tau_c = -1$.

Analogous to the surcharge test, we test the sharp null that packsize is uninformative about the unit price: within a cell, the K unit prices are *exchangeable* across sizes. Under the null, the permutation distribution $F_{0,K}$ of τ_c is symmetric around zero and known (exactly, by enumerating the $K!$ assignments; or by Monte Carlo for large K). Since packsizes are distinct, only unit prices tie and τ_b is a monotone transform of Kendall's score $S = \#\text{concordant} - \#\text{discordant}$. So we apply the randomized probability integral transform directly to S,

$$U_c \;=\; F_{0,K_c}(S_c^-) + V_c\big[F_{0,K_c}(S_c) - F_{0,K_c}(S_c^-)\big], \qquad V_c \sim \text{Unif}(0,1), \qquad U_c \perp\!\!\!\perp V_c$$

which satisfies $U_c \sim \text{Unif}[0,1]$ under the null. Systematic quantity discounting drives S_c toward its minimum and hence $U_c \to 0$: discounting arises in the *lower* tail of the transform, the mirror image of the surcharge test's upper tail.

Pooling cells, we report the discount-ordered share $\widehat{\Pr}(\tau_c < 0)$ and the strictly monotone share $\widehat{\Pr}(\tau_c = -1)$, each benchmarked against the exchangeability rate $\mathbb{E}[\,\Pr_0(\tau_c < 0)\,]$ that a size-uninformative null would produce, with confidence intervals from a block bootstrap clustered on the item. We again test uniformity of $\{U_c\}$ with the tail-weighted Anderson-Darling statistic. Quantity discounting is pervasive and systematic—rather than a tail event—when the discount-ordered share sits well above the exchangeability benchmark and the mass of

$\{U_c\}$ concentrates near zero.

C.2 Predictability of QS

We provide two sets of evidence to confirm that QS is not persistently concentrated within specific products so that consumers of those products could potentially learn when to anticipate QS at the shelf. First, we collapse alternatives to unique UPC-store-week shelf cells and retain UPC-store pairs observed in at least 52 distinct eligible weeks, which need not be consecutive. We then classify each UPC-store pair as "always QS," "never QS," or "both states," respectively.

Table 24 reports the results. Among 227,691 UPC-store pairs, 1.62% are always QS'd, 41.30% are never QS'd, and 57.08% exhibit both states. Only 1.62% of the UPCs are always QS'd.

Table 24: Persistence of QS Status

category	N	always QS	never QS	both states
beer	64,137	0.0422	0.2182	0.7396
cereal	109,473	0.0004	0.5521	0.4475
cracker	33,296	0.0279	0.4522	0.5199
gum	7,600	0.0007	0.2158	0.7836
laundry detergent	7,186	0.0000	0.1187	0.8813
toilet paper	5,999	0.0010	0.3409	0.6581
Overall	227,691	0.0162	0.4130	0.5708

Notes. Entries are proportions of qualifying identities classified as always quantity surcharged, never quantity surcharged, or observed in both states.

Next, we conduct a variance decomposition of QS by item, packsize, store, and household identity. Denote by j UPC and by t trip. Define $\mathbf{1}\left\{\mathrm{QS}_{jt}\right\}$ as the indicator that UPC j is QS'd at that store and week, and we do not trim data for this analysis. Within each category, we estimate the following linear probability specification:

$$\mathbf{1}\left\{\mathrm{QS}_{jt}\right\} = \mathbf{D}'_{hjt}\boldsymbol{\eta} + \varepsilon_{hjt}, \tag{20}$$

where $\mathbf{D}_{hjt}$ denotes the set of fixed-effects included in each specification.

Table 25: R^2 by Fixed-effects

category	N	Item	Item × Store	UPC	UPC × Store	UPC + Store + Household	UPC × Store + Household	UPC × Household + Store
beer	17,926,314	0.0778	0.2502	0.2677	0.5659	0.3009	0.5707	0.6406
cereal	41,703,158	0.0372	0.1027	0.1887	0.3235	0.1974	0.3259	0.4289
cracker	10,108,740	0.1030	0.1411	0.4965	0.6120	0.5081	0.6168	0.6889
gum	3,388,109	0.0344	0.1920	0.1130	0.3238	0.1551	0.3436	0.5192
laundry detergent	3,794,507	0.0424	0.2169	0.1026	0.3030	0.1841	0.3331	0.5053
toilet paper	3,505,078	0.0284	0.1309	0.2046	0.4171	0.2339	0.4274	0.5687
Overall	80,425,906	0.0641	0.1824	0.2611	0.4716	0.2875	0.4783	0.5719

Notes. Column labels refer to the fixed-effects included in each specification. Entries are total in-sample R^2 from category-specific linear probability models of QS status. Overall R^2 equals $1 - \sum_c RSS_c / \sum_c TSS_c$, with each category's TSS centered at its own sample mean.

We report the in-sample R^2s in Table 25. The Overall row pools category-specific RSS and TSS, using each category's own QS mean as the benchmark, and reports $R^2 = 1 - \sum_c RSS_c / \sum_c TSS_c$.

The overall R^2 ranges from 0.0641 to 0.5719. UPC, store, and household identity explain substantial variation. Nevertheless, UPC×household+store FE leaves 42.81% of within-category QS variation unexplained. Thus, the fixed-effects do not fully determine QS status of a UPC in these specifications.

D Proof of Lemma 1

We prove the lemma in three steps.

Step 1: the second-derivative ratio identifies the MRS even under partial information.
This follows from recognizing that the derivation in Appendix A.1 in ACZ does not hinge on linearity of utility or beliefs in $\mathbf{x}$. Thus, defining $\tilde{u}_j \equiv u(\mathbf{x}_j) + \beta p_j$ and $\tilde{eu}_j \equiv u(\mathbf{x}_j) + \beta f(\mathbf{x}_j)$, the same argument yields

$$\frac{\frac{\partial^2 \sigma_1}{\partial p_1 \partial p_j}}{\frac{\partial^2 \sigma_1}{\partial p_1 \partial size_j}} = \frac{\beta}{\frac{\partial u(\mathbf{x}_j)}{\partial size_j}}. \tag{21}$$

Step 2: compare first- and second-derivative ratios under full information.
If consumers are fully informed, then

$$\sigma_j = P(\tilde{u}_j + \epsilon_{ij} \geq \tilde{u}_k + \epsilon_{ik} \ \forall k),$$

so by the chain rule,

$$\frac{\partial \sigma_j}{\partial p_j} = \frac{\partial \sigma_j}{\partial \tilde{u}_j}\beta, \qquad \frac{\partial \sigma_j}{\partial size_j} = \frac{\partial \sigma_j}{\partial \tilde{u}_j}\frac{\partial u(\mathbf{x}_j)}{\partial size_j}.$$

Hence

$$\frac{\frac{\partial \sigma_j}{\partial p_j}}{\frac{\partial \sigma_j}{\partial size_j}} = \frac{\beta}{\frac{\partial u(\mathbf{x}_j)}{\partial size_j}} = \frac{\frac{\partial^2 \sigma_1}{\partial p_1 \partial p_j}}{\frac{\partial^2 \sigma_1}{\partial p_1 \partial size_j}},$$

which proves equation (8) of the lemma.

Step 3: show attenuation of first-derivatives ratio under partial information.
Following the same derivations as in the proof of Lemma 3 of ACZ, we can show that, when consumers have partial information,

$$\frac{\frac{\partial \sigma_j}{\partial p_j}}{\frac{\partial \sigma_j}{\partial size_j}} = \frac{\beta T_1}{\frac{\partial u(\mathbf{x}_j)}{\partial size_j} T_1 + T_2 + T_3 \left(\frac{\partial u(\mathbf{x}_j)}{\partial size_j} + \beta \frac{\partial f(\mathbf{x}_j)}{\partial size_j} \right)}, \tag{22}$$

where T_1, T_2, T_3 are positive terms under the maintained assumptions, which implies that $\left| \frac{\frac{\partial \sigma_j}{\partial p_j}}{\frac{\partial \sigma_j}{\partial size_j}} \right| \leq \frac{|\beta|}{\frac{\partial u(\mathbf{x}_j)}{\partial size_j}}$. Combining this with the result in step 1 yields equation (11) in the lemma.

E Extensions of Lemma 1

E.1 Proof of Proposition 1

Increasing τ_k in the common scale family generates a mean-preserving spread. By Weitzman (1979), reservation utility weakly increases, holding expected utility and search cost fixed. Since $\beta < 0$, this gives

$$m_i'(\tau_k) \leq 0, \qquad \beta m_i'(\tau_k) \geq 0. \tag{23}$$

Further, by the eventual-purchase result of Choi et al. (2018), the consumer chooses the product with the largest value of

$$V_{ik} = \min\{U_{ik}, z_{ik}^s\} = \min\{\tilde{u}_k, e\tilde{u}_k + \beta m_i(\tau_k)\} + \epsilon_{ik}, \tag{24}$$

where we have defined

$$\tilde{u}_k = u(\mathbf{x}_k) + \beta p_k, \qquad e\tilde{u}_k = u(\mathbf{x}_k) + \beta f(\mathbf{x}_k), \tag{25}$$

so that $U_{ik} = \tilde{u}_k + \epsilon_{ik}$ and $EU_{ik} = e\tilde{u}_k + \epsilon_{ik}$.

For brevity within this proof, write R_1 and R_2 for the absolute first- and second-derivative ratios, respectively, in the proposition.

Part (i). Let D_{i1} be the indicator that good 1 has the largest value in (24), so that $\sigma_1 = \mathbb{E}[D_{i1}(\mathbf{V}_i)]$. Fix $j \neq 1$. Suppressing the arguments for the remaining products, we have

$$D_{i1}(V_{i1}, V_{ij}) = D_{i1}(V_{i1}, U_{ij}) + D_{i1}(z_{i1}^s, V_{ij}) - D_{i1}(z_{i1}^s, U_{ij}). \tag{26}$$

To see this, consider two cases.

Case 1: $U_{i1} < z_{i1}^s$. This implies $V_{i1} = U_{i1}$. In turn, condition (15) implies $U_{ij} < z_{ij}^s$, and hence $V_{ij} = U_{ij}$. Thus, $D_{i1}(z_{i1}^s, V_{ij}) - D_{i1}(z_{i1}^s, U_{ij}) = 0$ and (26) holds.

Case 2: $U_{i1} \geq z_{i1}^s$. This implies $V_{i1} = z_{i1}^s$. Hence, $D_{i1}(V_{i1}, U_{ij}) - D_{i1}(z_{i1}^s, U_{ij}) = 0$ and (26) holds.

Taking expectations in (26) gives $\sigma_1 = P_A + P_B - P_C$. Neither P_B nor P_C depends on p_1, and P_A depends on the attributes of j only through $\tilde{u}_j$. Thus, by the chain rule,

$$\begin{aligned} \frac{\partial^2 \sigma_1}{\partial p_1 \partial p_j} &= \beta^2 \frac{\partial^2 P_A}{\partial \tilde{u}_1 \partial \tilde{u}_j}, \\ \frac{\partial^2 \sigma_1}{\partial p_1 \partial size_j} &= \beta \frac{\partial u(\mathbf{x}_j)}{\partial size_j} \frac{\partial^2 P_A}{\partial \tilde{u}_1 \partial \tilde{u}_j}. \end{aligned} \tag{27}$$

Taking ratios proves the result.

Finally, good 1 has the lowest surprise price, so $\tilde{p}_1 \leq \tilde{p}_j$. If $\tau_1 \leq \tau_j$, equation (23) gives $m_i(\tau_1) \geq m_i(\tau_j)$, which implies (15).

Part (ii). Define the size derivative of the reservation value by

$$q_{ij} \equiv \frac{\partial z_{ij}^s}{\partial size_j} = \frac{\partial EU_{ij}}{\partial size_j} + \beta m_i'(\tau_j)\tau'(size_j).$$

Equation (23) establishes $\beta m_i'(\tau_j) \geq 0$ for every consumer. Together with (16), this gives $q_{ij} \geq 0$. Economically, larger size weakly increases expected utility and the value of the option to inspect a product. When dispersion is locally constant, the reservation-premium derivative drops out entirely. These primitive conditions are sufficient, rather than equivalent, to reservation-value monotonicity: the latter can also hold through offsetting terms.

Conditional on search heterogeneity, write the choice probability as $\sigma_j^i(\tilde{\mathbf{u}}, \mathbf{r}_i)$, where $r_{ik} = \tilde{eu}_k + \beta m_i(\tau_k)$. Define

$$T_{1i} = \frac{\partial \sigma_j^i}{\partial \tilde{u}_j} \geq 0, \qquad T_{2i} = \frac{\partial \sigma_j^i}{\partial r_{ij}} \geq 0.$$

Both inequalities follow because increasing either index weakly increases V_{ij} and cannot change a choice of j to another product. By the chain rule,

$$\begin{aligned} \frac{\partial \sigma_j}{\partial p_j} &= \beta \mathbb{E}[T_{1i}], \\ \frac{\partial \sigma_j}{\partial size_j} &= \frac{\partial u(\mathbf{x}_j)}{\partial size_j}\mathbb{E}[T_{1i}] + \mathbb{E}[q_{ij}T_{2i}]. \end{aligned} \tag{28}$$

Under (16), the defined denominator is positive, and

$$R_1 = \frac{|\beta|\mathbb{E}[T_{1i}]}{\dfrac{\partial u(\mathbf{x}_j)}{\partial size_j}\mathbb{E}[T_{1i}] + \mathbb{E}[q_{ij}T_{2i}]} \leq \frac{|\beta|}{\partial u(\mathbf{x}_j)/\partial size_j} = R_2.$$

This proves part (ii).

E.2 Search and Discovery Model

Assumption SD1 (Search and discovery)**.** Consumers follow Greminger's stationary SD model with empty initial consideration and awareness sets. Discovery of product k reveals $(\mathbf{x}_k, \epsilon_{ik})$ and hence $EU_{ik} = u(\mathbf{x}_k) + \beta f(\mathbf{x}_k) + \epsilon_{ik}$; inspection reveals its price, so that $U_{ik} = EU_{ik} + \beta\tilde{p}_k$. We maintain the baseline SD assumptions on costs, beliefs about undiscovered products, and the discovery process (Greminger, 2022, Sections 3–4).

Assumption SD2 (Common beliefs about price surprises)**.** For a given consumer i, the pre-inspection distribution F_i of $\beta\tilde{p}_k$ is common across products and does not depend on discovered attributes, utility shocks, or information acquired about other products. Together with consumer i's inspection cost c_{si}, common across products, it yields a finite inspection premium ξ_i, also common across products, satisfying

$$z_{ik}^s = EU_{ik} + \xi_i, \qquad \int_{\xi_i}^{\infty} [1 - F_i(t)]\, dt = c_{si}.$$

Assumption SD3 (Regularity)**.** The relevant choice and auxiliary probability functions used below are twice continuously differentiable locally in their deterministic utility and expected-utility indices. The discovery, search and utility values exhibit ties with probability zero.

Assumption SD4 (Marginal utility of size)**.** $\partial u(\mathbf{x}_j)/\partial size_j > 0$.

Proof of Proposition 2

Proof. Let

$$\tilde{u}_k = u(\mathbf{x}_k) + \beta p_k, \qquad e\tilde{u}_k = u(\mathbf{x}_k) + \beta f(\mathbf{x}_k). \tag{29}$$

Then $U_{ik} = \tilde{u}_k + \epsilon_{ik}$, $EU_{ik} = e\tilde{u}_k + \epsilon_{ik}$, and

$$V_{ik} \equiv \min\{U_{ik}, z_{ik}^s\} = \min\{\tilde{u}_k, e\tilde{u}_k + \xi_i\} + \epsilon_{ik}. \tag{30}$$

By Theorem 2 of Greminger (2022), the event that good j is chosen can be written as a function of the vector of indices $\mathbf{V}_i \equiv (V_{ik})_{k=1}^J$. Letting D_{ij} be the associated indicator function, we write:

$$\sigma_j = \mathbb{E}[D_{ij}(\mathbf{V}_i)], \tag{31}$$

where the expectation is over utility shocks, search heterogeneity, and discovery positions. The dependence of D_{ij} on the fixed discovery environment and positions is suppressed in the notation.

Part (i). We decompose D_{i1} as follows:

$$\begin{aligned} D_{i1}(V_{i1}, \mathbf{V}_{i,-1}) &= D_{i1}(V_{i1}, \mathbf{U}_{i,-1}) + \Big[D_{i1}(V_{i1}, \mathbf{V}_{i,-1}) - D_{i1}(V_{i1}, \mathbf{U}_{i,-1})\Big] \\ &= D_{i1}(V_{i1}, \mathbf{U}_{i,-1}) + \Big[D_{i1}(z_{i1}^s, \mathbf{V}_{i,-1}) - D_{i1}(z_{i1}^s, \mathbf{U}_{i,-1})\Big]. \end{aligned}$$

The second equality follows from the fact that, since good 1 has the best price surprise,

$$U_{i1} < z^s_{i1} \quad \Longrightarrow \quad \beta\tilde{p}_k \leq \beta\tilde{p}_1 < \xi_i \quad \Longrightarrow \quad \mathbf{V}_{i,-1} = \mathbf{U}_{i,-1}, \tag{32}$$

so that $D_{i1}(V_{i1}, \mathbf{V}_{i,-1}) - D_{i1}(V_{i1}, \mathbf{U}_{i,-1})$ vanishes when $U_{i1} < z^s_{i1}$. Thus, taking expectations,

$$\sigma_1 = P_A(\tilde{\mathbf{u}}, e\tilde{u}_1) + P_B(\tilde{\mathbf{u}}_{-1}, \mathbf{e\tilde{u}}) - P_C(\tilde{\mathbf{u}}_{-1}, e\tilde{u}_1), \tag{33}$$

where

$$\begin{aligned} P_A &= \mathbb{E}[D_{i1}(V_{i1}, \mathbf{U}_{i,-1})], \\ P_B &= \mathbb{E}[D_{i1}(z^s_{i1}, \mathbf{V}_{i,-1})], \\ P_C &= \mathbb{E}[D_{i1}(z^s_{i1}, \mathbf{U}_{i,-1})]. \end{aligned}$$

Neither P_B nor P_C depends on p_1, and P_A depends on the attributes of $j \neq 1$ only through $\tilde{u}_j$. Thus, by the chain rule,

$$\frac{\partial \sigma_1}{\partial p_1} = \beta \frac{\partial P_A}{\partial \tilde{u}_1}, \tag{34}$$

and

$$\begin{aligned} \frac{\partial^2 \sigma_1}{\partial p_1 \partial p_j} &= \beta^2 \frac{\partial^2 P_A}{\partial \tilde{u}_1 \partial \tilde{u}_j}, \\ \frac{\partial^2 \sigma_1}{\partial p_1 \partial size_j} &= \beta \frac{\partial u(\mathbf{x}_j)}{\partial size_j} \frac{\partial^2 P_A}{\partial \tilde{u}_1 \partial \tilde{u}_j}. \end{aligned} \tag{35}$$

Taking ratios proves part (i) of the proposition.

Part (ii). By the chain rule,

$$\begin{aligned} \frac{\partial \sigma_j}{\partial p_j} &= \beta T_1, \\ \frac{\partial \sigma_j}{\partial size_j} &= \frac{\partial u(\mathbf{x}_j)}{\partial size_j} T_1 + \left(\frac{\partial u(\mathbf{x}_j)}{\partial size_j} + \beta \frac{\partial f(\mathbf{x}_j)}{\partial size_j} \right) T_2, \end{aligned} \tag{36}$$

where

$$T_1 \equiv \frac{\partial \sigma_j}{\partial \tilde{u}_j} \geq 0, \qquad T_2 \equiv \frac{\partial \sigma_j}{\partial e\tilde{u}_j} \geq 0. \tag{37}$$

The last display follows from the fact that the indicator D_{ij} is nondecreasing in V_{ij}, and thus in $\tilde{u}_j$ and $e\tilde{u}_j$ (Theorem 2 in Greminger (2022)). Under Assumption SD4 and Condition (9), we have

$$R_1 = \frac{|\beta| T_1}{\dfrac{\partial u(\mathbf{x}_j)}{\partial size_j} T_1 + \dfrac{\partial EU_{ij}}{\partial size_j} T_2} \leq \frac{|\beta|}{\partial u(\mathbf{x}_j)/\partial size_j} = |MRS|. \tag{38}$$

Together with part (i), this proves part (ii). □

F Estimation Details

We approximate the choice propensity σ_1 using Legendre polynomials of ($\widetilde{size}_1$, $\widetilde{size}_2$, $\widetilde{size}_{-1,2}$, $p_{1t}, p_{2t}, p_{-1,2t}$) and the parametric terms involving $\mathbf{w}$ discussed above. We restrict σ_1 to be weakly increasing in $\widetilde{size}_1$ and $p_{2t}, p_{-1,2t}$, and weakly decreasing in $\widetilde{size}_2$, $\widetilde{size}_{-1,2}$ and p_{1t}.

Further, we use the following Lemma to directly enforce the constraint that the ratio of second derivatives $\frac{\partial^2 \sigma_1}{\partial p_1 \partial p_j} / \frac{\partial^2 \sigma_1}{\partial p_1 \partial \widetilde{size}_j}$ be constant across values of ($\widetilde{size}_1$, $\widetilde{size}_2$, $\widetilde{size}_{-1,2}$, $p_{1t}, p_{2t}, p_{-1,2t}$). Under the assumption that utility is linear in price and size, this constraints follows directly from the fact that the second derivative ratio is equal to the (constant) MRS between price and size.

Lemma 2 (Representation under the derivative restriction)**.** *Let $\Omega \subset \mathbb{R}^4$ be a nonempty open rectangle and*

$$f = f(x_1, x_2, z_1, z_2)$$

be twice continuously differentiable on Ω. Suppose that, for some $k \in \mathbb{R}$,

$$\frac{\partial^2 f}{\partial z_1 \partial z_2} = k \frac{\partial^2 f}{\partial z_1 \partial x_2} \qquad \text{for all } (x_1, x_2, z_1, z_2) \in \Omega.$$

Then there exist functions

$$a = a(x_1, x_2, z_2) \qquad \text{and} \qquad b = b(x_1, z_1, w),$$

where

$$w = x_2 + k z_2,$$

such that

$$f(x_1, x_2, z_1, z_2) = a(x_1, x_2, z_2) + b(x_1, z_1, x_2 + k z_2).$$

Proof. Define

$$u(x_1, x_2, z_1, z_2) \equiv \frac{\partial f}{\partial z_1}(x_1, x_2, z_1, z_2).$$

Because $f \in C^2(\Omega)$, we have $u \in C^1(\Omega)$. The maintained restriction can be written as

$$u_{z_2}(x_1, x_2, z_1, z_2) = k u_{x_2}(x_1, x_2, z_1, z_2),$$

or equivalently

$$u_{z_2} - k u_{x_2} = 0$$

for some k.

Now fix (x_1, z_1). Conditional on this fixed pair, the preceding equation is a first-order linear transport equation in the variables (x_2, z_2):

$$u_{z_2} - k u_{x_2} = 0.$$

By the standard transport-equation characterization, solutions are constant along characteristic curves. See, for example, Section 2.1 of Evans (2022), where the solution to

$$v_t + c v_x = 0$$

is shown to be constant along the lines

$$x_2 - (-k) z_2 = x_2 + k z_2.$$

Hence, for each fixed (x_1, z_1), the function

$$u(x_1, x_2, z_1, z_2)$$

can depend on (x_2, z_2) only through the scalar index

$$w = x_2 + k z_2.$$

Therefore, there exists a function M such that

$$u(x_1, x_2, z_1, z_2) = M(x_1, z_1, x_2 + k z_2).$$

Equivalently,

$$\frac{\partial f}{\partial z_1}(x_1, x_2, z_1, z_2) = M(x_1, z_1, x_2 + k z_2).$$

Next, fix an arbitrary baseline value $\bar{z}_1 \in \mathcal{Z}_1$. By the fundamental theorem of calculus,

$$f(x_1, x_2, z_1, z_2) - f(x_1, x_2, \bar{z}_1, z_2) = \int_{\bar{z}_1}^{z_1} \frac{\partial f}{\partial z_1}(x_1, x_2, t, z_2)\, dt.$$

Using the representation of f_{z_1} derived above,

$$\frac{\partial f}{\partial z_1}(x_1, x_2, t, z_2) = M(x_1, t, x_2 + k z_2).$$

Therefore,

$$f(x_1, x_2, z_1, z_2) = f(x_1, x_2, \bar{z}_1, z_2) + \int_{\bar{z}_1}^{z_1} M(x_1, t, x_2 + kz_2)\, dt.$$

Now define

$$a(x_1, x_2, z_2) \equiv f(x_1, x_2, \bar{z}_1, z_2),$$

and define

$$b(x_1, z_1, w) \equiv \int_{\bar{z}_1}^{z_1} M(x_1, t, w)\, dt.$$

Then

$$f(x_1, x_2, z_1, z_2) = a(x_1, x_2, z_2) + b(x_1, z_1, x_2 + kz_2). \tag{39}$$

This proves the claimed representation. □

In our model, the function f is the choice probability of good 1, σ_1, x is log pack size, z is negative price, and k is the (constant) ratio of second derivatives R_2. Lemma 2 has two benefits: (i) we can approximate σ_1 using the sum of two functions each of which depends on strictly fewer arguments than σ_1 (helping alleviate the curse of dimensionality); and (ii) we can directly target the ratio R_2 in estimation as opposed to computing it from the polynomial approximation of the function.

More specifically, for any candidate value of R_2, we approximate σ_1 using (39)

$$\tilde{\sigma}_{1t}(\theta) = \theta_a' \mathbf{a}(\widetilde{size}_1, \widetilde{size}_2, p_{2t}, \widetilde{size}_{-1,2}, p_{-1,2t}) + \theta_b' \mathbf{b}(\widetilde{size}_1, p_{1t}, \widetilde{size}_2 + R_2 p_{2t}, \widetilde{size}_{-1,2}, p_{-1,2t}) + \theta_w' \mathbf{w}_{1t}, \tag{40}$$

where $\theta \equiv (\theta_a, \theta_b, \theta_w)$ are the parameters to be estimated and $(\mathbf{a}, \mathbf{b})$ is the Legendre polynomials approximation. We include all polynomials of total degree up to 4. For each value of R_2, we compute sum of squared residuals (RSS) between the approximation $\tilde{\sigma}_{1t}$ and the observed outcome y_{1t} (a dummy for whether product 1 is chosen in trip t), and then pick the value of R_2 on a grid that delivers the lowest RSS.

G Estimation of Utility Parameters

Under the linear utility assumption in (18), the R_2 ratio from the test in §6.4 identifies $MRS \equiv \frac{\beta}{\alpha_{\widetilde{size}}}$. Further, as argued in the main text, fitting a nested logit model on trips with choice sets exhibiting low variation in $\tilde{p}$ across goods yields an estimate of the impact of $\widetilde{size}$ on expected utility: $\delta_{\widetilde{size}} \equiv \alpha_{\widetilde{size}} + \beta\gamma_{\widetilde{size}}$. This formula reflects the fact that size affects expected utility directly ($\alpha_{\widetilde{size}}$) and indirectly via price expectations ($\beta\gamma_{\widetilde{size}}$). Given

our estimates of price beliefs $\hat{\boldsymbol{\gamma}}$ from S6.1, we can then combine the estimates of R_2 and $\delta_{\widetilde{size}}$ to obtain separate estimates of the utility coefficients as follows:

$$\begin{aligned} \hat{\beta} &= \frac{\hat{R}_2 \cdot \hat{\delta}_{\widetilde{size}}}{\hat{\gamma}_{\widetilde{size}} \hat{R}_2 - 1} \\ \hat{\alpha}_{\widetilde{size}} &= \frac{\hat{\beta}}{\hat{R}_2}. \end{aligned}$$

Similarly, we can compute the remaining coefficients in $\boldsymbol{\alpha}$ from (18) by combining our expected utility coefficient estimates $\hat{\boldsymbol{\delta}}$ with the beliefs estimates $\hat{\boldsymbol{\gamma}}$.

We cluster-bootstrap the entire procedure at the household level to obtain standard errors that reflect the uncertainty in each step and allow for correlation across trips within households. We draw households with replacement from the original sample, and for each of these bootstrap samples we estimate the beliefs parameters $\hat{\boldsymbol{\gamma}}$, the nested logit model on the limited $\tilde{p}$ variation trips and the R_2 ratio.[32] We then combine these bootstrap estimates as described above to obtain separate estimates of the utility coefficients. Finally, we compute confidence intervals by taking the 2.5^{th} and 97.5^{th} percentiles of the bootstrap distribution for each parameter.

H Distribution of Utility Parameters

Figures 7–12 show the household-weighted CDFs of the estimated utility coefficients from Section 7.1. Tables 26–31 show the corresponding estimates of the correlation matrices.

H.1 Beer

Table 26: Utility Coefficient Correlations: Beer

	$\alpha_{\widetilde{size}}$	β	α_{promo}	$\alpha_{\#\text{ of packs}}$	λ
$\alpha_{\widetilde{size}}$	1.000				
β	-0.998	1.000			
α_{promo}	0.272	-0.218	1.000		
$\alpha_{\#\text{ of packs}}$	-0.452	0.402	-0.981	1.000	
λ	0.403	-0.351	0.990	-0.999	1.000

[32] As always, we perform the estimation separately for each household cluster.

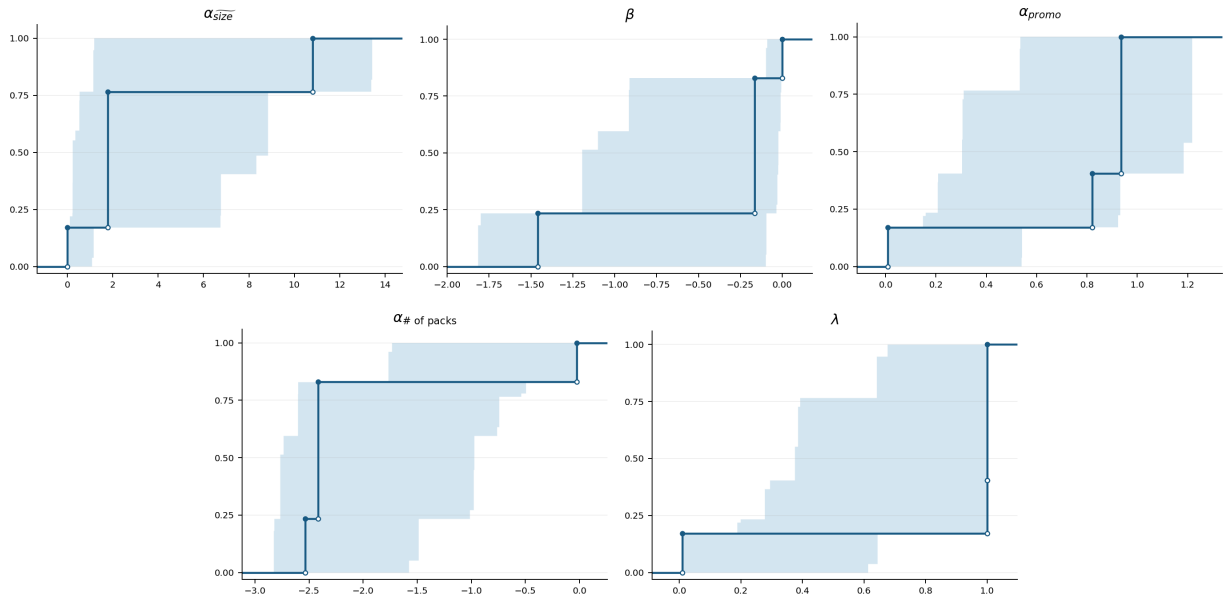


Figure 7: Utility Coefficient CDFs and 95% Confidence Bands

H.2 Cereal

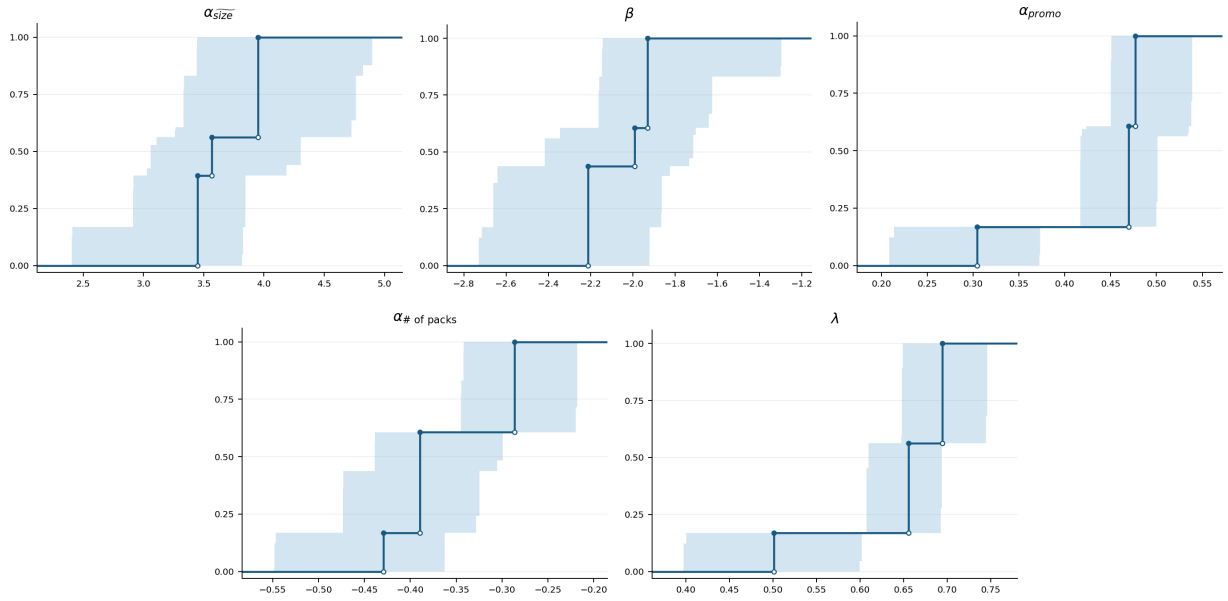


Figure 8: Utility Coefficient CDFs and 95% Confidence Bands

Table 27: Utility Coefficient Correlations: Cereal

	$\alpha_{\widetilde{size}}$	β	α_{promo}	$\alpha_{\#\text{ of packs}}$	λ
$\alpha_{\widetilde{size}}$	1.000				
β	-1.000	1.000			
α_{promo}	0.184	-0.198	1.000		
$\alpha_{\#\text{ of packs}}$	-0.658	0.648	0.619	1.000	
λ	0.480	-0.492	0.951	0.345	1.000

H.3 Crackers

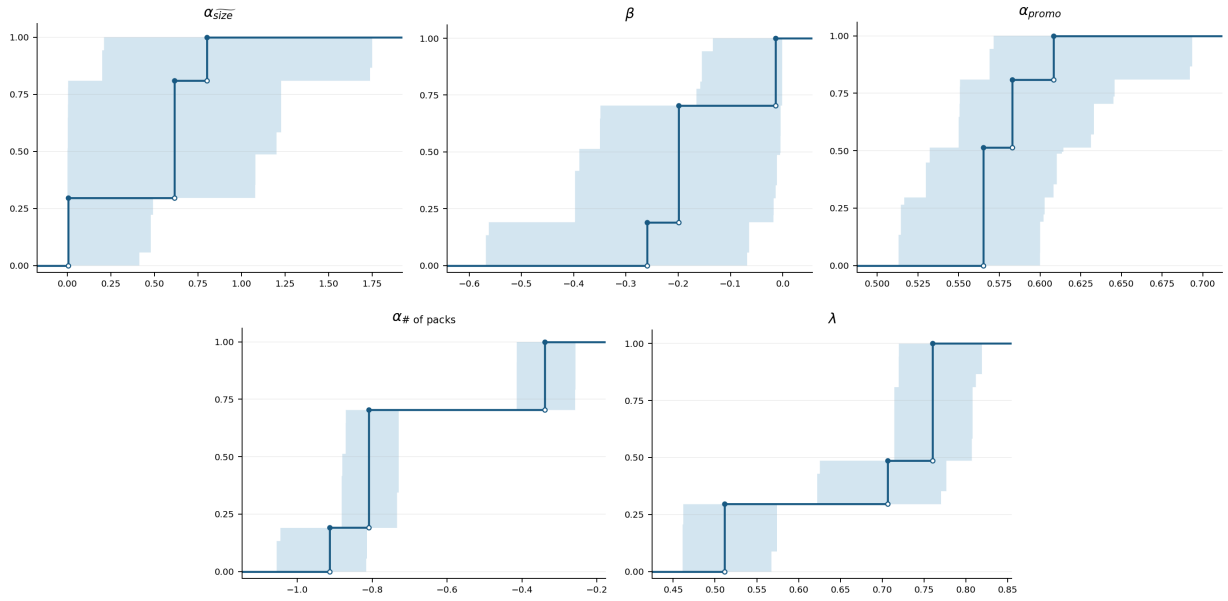


Figure 9: Utility Coefficient CDFs and 95% Confidence Bands

Table 28: Utility Coefficient Correlations: Crackers

	$\alpha_{\widetilde{size}}$	β	α_{promo}	$\alpha_{\#\text{ of packs}}$	λ
$\alpha_{\widetilde{size}}$	1.000				
β	-1.000	1.000			
α_{promo}	0.058	-0.071	1.000		
$\alpha_{\#\text{ of packs}}$	-0.998	0.998	-0.002	1.000	
λ	0.917	-0.912	-0.346	-0.938	1.000

H.4 Laundry Detergent

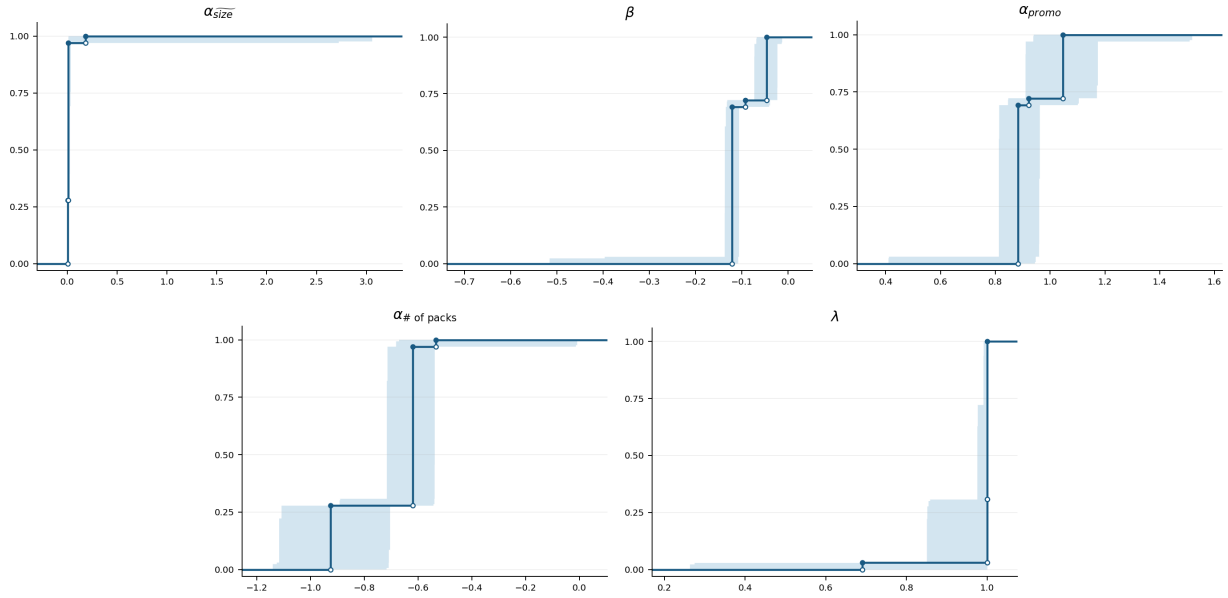


Figure 10: Utility Coefficient CDFs and 95% Confidence Bands

Table 29: Utility Coefficient Correlations: Laundry detergent

	$\alpha_{\widetilde{size}}$	β	α_{promo}	$\alpha_{\#\text{ of packs}}$	λ
$\alpha_{\widetilde{size}}$	1.000				
β	-0.070	1.000			
α_{promo}	-0.129	0.998	1.000		
$\alpha_{\#\text{ of packs}}$	0.319	-0.968	-0.981	1.000	
λ	-0.994	-0.040	0.019	-0.212	1.000

H.5 Gum

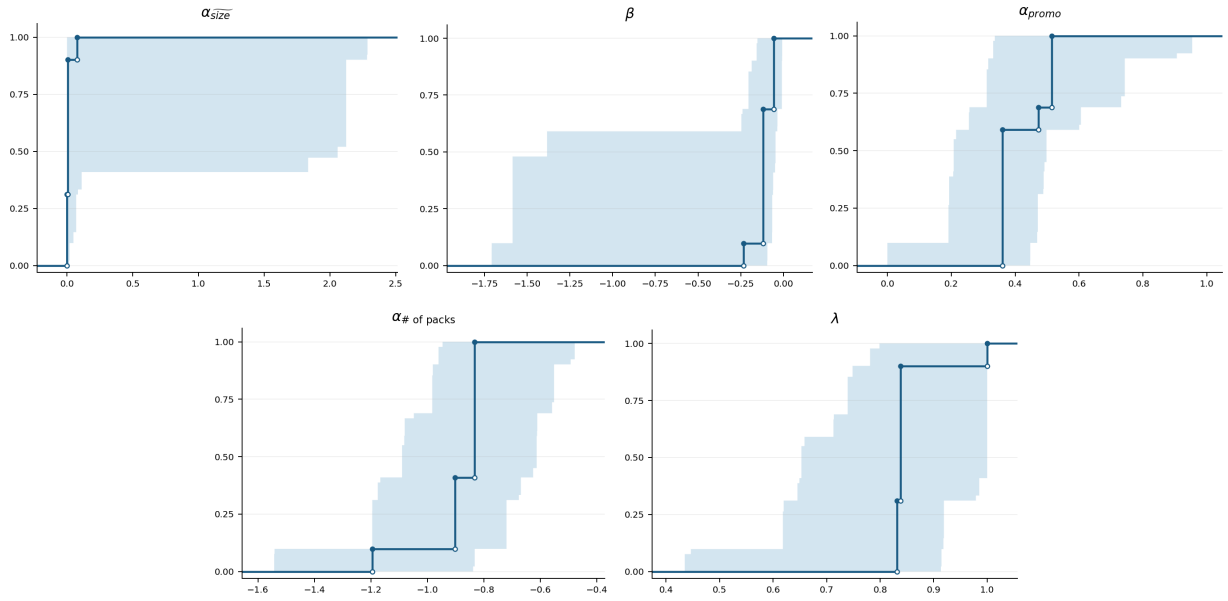


Figure 11: Utility Coefficient CDFs and 95% Confidence Bands

Table 30: Utility Coefficient Correlations: Gum

	$\alpha_{\widetilde{size}}$	β	α_{promo}	$\alpha_{\#\text{ of packs}}$	λ
$\alpha_{\widetilde{size}}$	1.000				
β	-0.883	1.000			
α_{promo}	0.121	0.360	1.000		
$\alpha_{\#\text{ of packs}}$	-0.910	0.608	-0.521	1.000	
λ	0.998	-0.850	0.186	-0.935	1.000

H.6 Toilet Paper

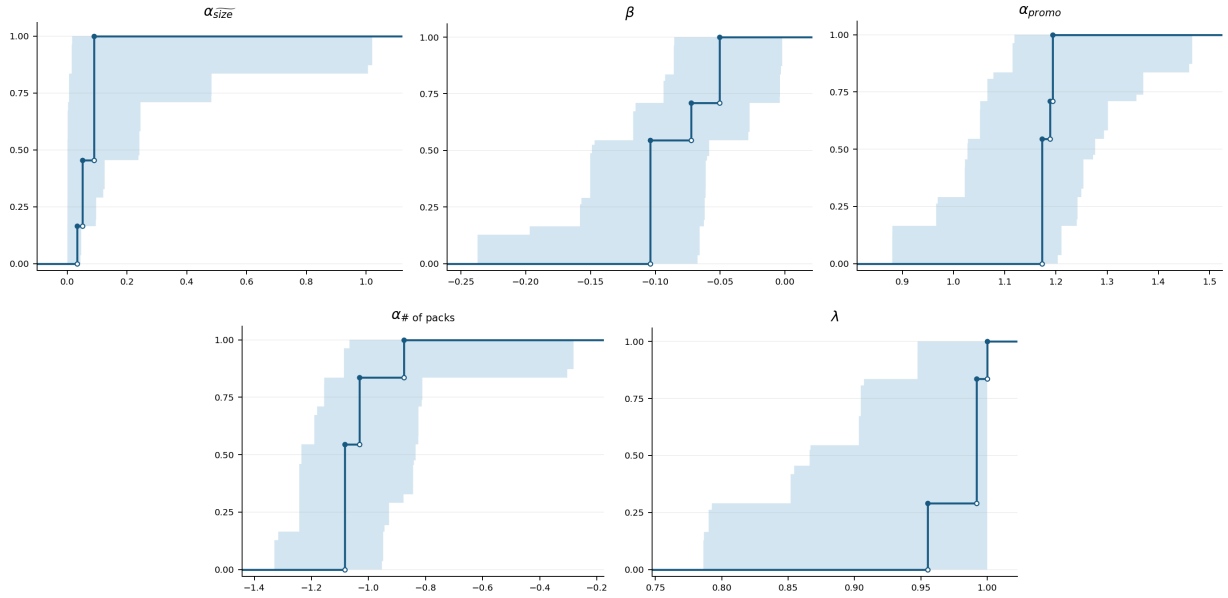


Figure 12: Utility Coefficient CDFs and 95% Confidence Bands

Table 31: Utility Coefficient Correlations: Toilet paper

	$\alpha_{\widetilde{size}}$	β	α_{promo}	$\alpha_{\#\text{ of packs}}$	λ
$\alpha_{\widetilde{size}}$	1.000				
β	-0.852	1.000			
α_{promo}	-0.915	0.991	1.000		
$\alpha_{\#\text{ of packs}}$	-0.874	0.490	0.602	1.000	
λ	0.356	-0.792	-0.703	0.144	1.000

I QS Decomposition Implementation

For every trip with at least one QS'd UPC, we proceed as follows. We draw a large number of ϵ draws from the generalized extreme-value distribution (parameterized by λ) and retain the draws such that $U_{ij}^{sq} \geq U_{ik}^{sq}$ for all k, where j is one of the QS'd options in the choice set and U_{ij}^{sq} denotes j's utility with the price coefficient replaced by the (smaller in magnitude) coefficient implied by the first-derivative ratio.[33] As in the full-information counterfactual of Section 7.3, these U^{sq} utilities are used to approximate the status quo where consumers are less price sensitive relative to the case with full price information. As such these are not normative utilities, but rather just devices to simulate the status quo. The same approach was also used by ACZ. Conditional on these ϵ draws, the probability of choosing a QS'd option in the status quo is mechanically 100%. We then, compute the probability of choosing a QS'd option under each of the counterfactual scenarios in Section 7.2 as follows:

1. π_1: Calculate the full-information utilities U_{ij} (using the price coefficient based on the second-derivative ratio from Table 11) and then compute $\sum_j P\left(U_{ij} \geq U_{ik} \ \forall k\right)$, where the sum is taken over the QS'd options in each trip.

2. π_2: The same as in scenario #1, but with the number of packs set to 1 for the multi-pack options dominating a QS'd option (i.e., no disutility from multiple packs).

3. π_3: The same as in scenario #2, but with the promo variable set to zero for all QS'd options.

4. π_4: The same as in scenario #3, but with the size of the cheaper dominating bundle increased to match the size of the QS'd bundle whenever the former is below the latter.

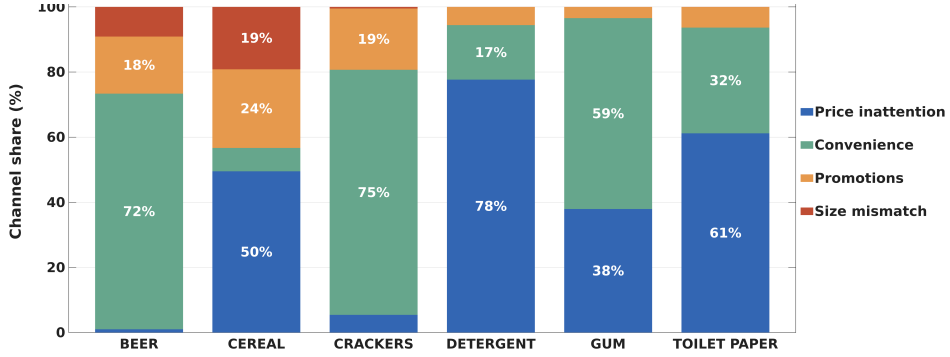
We then take $1 - \pi_4$ as the total probability that is explained by the model (as opposed to idiosyncratic shocks) and decompose it as follows:

1. Price inattention: $\frac{1-\pi_1}{1-\pi_4}$
2. Convenience: $\frac{\pi_1-\pi_2}{1-\pi_4}$
3. Promotions: $\frac{\pi_2-\pi_3}{1-\pi_4}$
4. Size mismatch: $\frac{\pi_3-\pi_4}{1-\pi_4}$.

[33] We ensure that for each trip we have at least 100 retained draws.

J Alternative Decomposition of QS Purchases

Figure 13: Decomposition of QS Purchases



Notes. For each category, this figure shows a decomposition of the probability of a QS purchases into various channels. We follow the same steps as in I except that we condition on ϵ draws such that consumers choose an option from the same item in scenario 4 as they do in the status quo.

K Derivations for Welfare Calculations

To quantify the welfare improvement from providing consumers with full price information, we use the following formula:

$$\Delta CS = CS_{full} - CS_{inatt}, \tag{41}$$

where

$$CS_{full} = \frac{1}{\beta} \log \sum_{g} \left(\sum_{j \in g} e^{\delta_j / \lambda} \right)^{\lambda} \tag{42}$$

and

$$CS_{inatt} = \frac{1}{\beta} \sum_{j} s_j(\delta_{inatt}) \log \left(\sum_{g} \left(e^{\delta_j/\lambda} \mathbb{I}\{j \in g\} + \sum_{k \in g, k \neq j} e^{(\delta_{k,inatt} + \delta_j - \delta_{j,inatt})/\lambda} \right)^{\lambda} \right) \tag{43}$$

In the expressions above, λ is the nesting parameter ($\lambda = 1$ corresponds to standard logit and the lower λ is the more correlated utilities are within each nest), $\delta_j = \alpha \mathbf{x}_j + \beta p_j$, $\delta_{j,inatt} = \alpha \mathbf{x}_j + \beta_{inatt} p_j$, where β_{inatt} is the (attenuated) price coefficient capturing consumers' behavior in the status quo where they may be partially inattentive to price, and $s_j(\delta_{inatt})$ is product j's choice probability evaluated at the vector of mean utilities δ_{inatt}.

The full information expression (42) is the standard nested logit consumer surplus formula, whereas (43) is derived by adapting the derivation in Appendix G of ACZ to the nested logit model. Note that CS_{inatt} collapses to CS_{full} when $\beta_{inatt} = \beta$, i.e. when consumers are attentive in the status quo. In the case where $\beta_{inatt} \neq \beta$, CS_{inatt} captures the consumer surplus associated with the possibly suboptimal choices they make under price inattention.