

Sales, Quantity Surcharge, and Consumer Inattention*

Sofronis Clerides[†]

Pascal Courty[‡]

February 2015

Abstract

Quantity surcharges occur when retailers carry a product in two sizes and offer a promotion on the small size: the large size then costs more per unit than the small one. When quantity surcharges occur, sales of the large size decline only slightly even though the same quantity can be purchased for less. We document this behavior in two datasets and four different product categories. It is consistent with the notion of passive shoppers found in the industrial organization literature and the notion of rational inattention in macroeconomics. We discuss implications for consumer decision making, demand estimation, and firm pricing.

Keywords: quantity surcharge, sales, promotions, consumer inattention, quantity discounts, nonlinear pricing.

JEL Classification: L12, L13, D4.

*We thank Evangelos Constantinou, Oege Dijk and Machi Michalopoulou for research assistance, *The Nielsen Company* for providing the data and Antonis Michis for invaluable help in understanding them. We are solely responsible for any errors. Participants at several conferences and at departmental seminars at McGill, Oklahoma, Victoria, Drexel, Northwestern, Zurich, CREST, HEC, Piraeus, ALBA, and Hebrew provided many helpful comments. We especially thank Catherine Thomas for very useful feedback. Much of this research was carried out while Clerides was visiting Northwestern University on a European Commission Marie Curie fellowship.

[†]University of Cyprus and CEPR; s.clerides@ucy.ac.cy.

[‡]University of Victoria and CEPR; pcourty@uvic.ca.

1 Introduction

A quantity surcharge is the opposite of a quantity discount; it occurs when the same physical product is sold in two packs of different sizes and the large size has a higher price per unit than the small one. Quantity discounts find strong theoretical justification in models of nonlinear pricing and are indeed ubiquitous in everyday life and in most datasets, including the two we analyze in this paper. Crucially however, in our data we also observe many instances of quantity surcharges. This typically occurs when a supermarket promotes a small size while leaving the price of the large size unchanged. In such an instance, the unit price of the small size drops well below the unit price of the large size.

Quantity surcharges (QS) give rise to a substitution opportunity. Since the small size has the lowest unit price, all buyers of the large size should switch to the small size, provided they can approximate the desired quantity by purchasing multiple small packs. If consumers do not buy dominated options, we would expect the sales of the large size during a QS to be zero. Stated differently, sales should decrease by 100% relative to what they would have been in the absence of a QS. Instead, in our primary dataset of Dutch laundry detergents we find a decrease of only 27% in the median case relative to the sales the week preceding the surcharge. The figure is similar if we use the median sales during non-promotion periods as a comparison benchmark. We obtain comparable sales declines in a second dataset covering detergent, razor blades, shampoo and diapers in US stores.

We argue that limited substitution away from the large size in the face of a QS is evidence that many consumers purchase a dominated option. This claim rests on two assumptions: (i) at the time of purchase, both ways of purchasing a given amount of detergent (in one large container or in two – or more – smaller containers) are available to the consumer; and (ii) the two options are (almost) perfect substitutes. The first assumption relates to data quality. We go to great lengths to ensure that it holds by focusing on high quality data and by exploiting multiple datasets. The second assumption is more subjective. Consumers may have a strong preference for the large size, perhaps because it is easier to transport, because it fits nicely in their cupboard or because it is environmentally friendlier (requires less packaging material). These are all plausible, but there are also arguments to the contrary. Two small packages may be easier to store than one large one, or they could be used in different locations in the house. Our thesis is that the two options are as close to being perfect substitutes as one could expect two different products to be, because they deliver the same quantity of the same physical product. The fact that a large majority of consumers do not purchase the cheaper option suggests a lack of price awareness. Borrowing terminology from the industrial organization and macroeconomics

literatures, we call these purchases passive or inattentive.

We also investigate whether the incidence of passive purchases is greater for premium brands than for value brands. The answer is positive. A QS for a value brand is associated with a 41% decline in the sales of the large size relative to the median regular sales. By contrast, a QS for a premium brand decreases the sales of the large size by only 16%. Passive purchases are less common among buyers of value brands, who are likely to be relatively price sensitive. Below we argue that this is a key assumption of the theory literature on retail sales.

How widespread are passive purchases? Can we say that 73% of the purchases of the large pack during non-promotion periods are passive purchases? Answering these questions with non-experimental data requires care. A contribution of this paper is to discuss the inference problems associated with grocery data that often aggregate time periods, consumers within a store, and even stores within a chain. We also develop a simple econometric framework based on the distinction between regular and promotion buyers to estimate the fraction of regular purchases that are passive. When pooling all products together we obtain an estimate of the fraction of inattentive of 0.601 in our Dutch data. Allowing for different degrees of inattention for premium versus value brands gives estimates of 0.659 and 0.560 respectively but the difference is not statistically significant. With our US data, the mean inattention parameter estimates are 0.658, 1.09, 0.910 and 0.789 for detergent, blades, shampoo and diapers respectively.

Models of store sales in industrial organization have used the concept of passive consumers to explain why retailers can sell the same product at a regular price and also (in different periods) at promotional prices that are substantially lower. It is difficult to come up with robust evidence of passive purchases because one has to rule out the alternative explanation of strong attachment to a brand/product/size. Using the unusual temporary price patterns that happen during QS events, we offer unique evidence of passive consumer behavior. Our evidence also contributes to the macroeconomic literature on information frictions that has argued that passive purchases may in fact be rational and has developed the notion of rational inattention.¹ Simple computations using the time cost of checking a product's price show that inattentive behavior leading to the choice of dominated options may indeed be rational.

The only empirical works on limited attention that we are aware of are behavioral studies of information opacity.² Two papers have considered consumer goods as we do. They show

¹Sims (1998, 2003), Reis (2006), Mackowiak and Wiederholt (2009), Mondria (2010).

²See DellaVigna (2009) for a review. Violation of rationality is also found when decision-making involves intertemporal payoffs. DellaVigna and Malmendier (2006) show that some gym users could save by switching from a monthly subscription to pay-per-visit. They rationalize this by questioning rational expectations and intertemporal commitment. They do not consider consumer attention. Lacetera, Pope, and Syndor (2012) find

that consumers respond less to opaque information that is more difficult to process, such as nontransparent taxes (Chetty, Looney, and Kroft, 2009) or shipping costs (Hossain and Morgan, 2006). In our application there is no opaque component to the price that may confuse consumers. On the contrary, manufacturers and retailers go to great lengths to draw consumers’ attention to promotional offers and the fact that consumers still miss them seems striking.³

This work also contributes to the literature that tries to understand how consumers process price information. There is some indirect evidence in support of passive consumers. Behavioral and marketing research has produced many examples – mainly through survey evidence – where consumers are not fully aware of available products and their prices. Cognitive research on decision-making in grocery purchases has demonstrated that the majority of consumers are not aware of point-of-purchase prices (Monroe and Lee, 1999). These considerations are important in the modeling of consumers decision making in demand models as we argue in the discussion section.

Our work is related to the large empirical literature that has tried to explain the source of the large demand response during promotions (Pesendorfer, 2002; Erdem, Imai, and Keane, 2003; Hosken and Reiffen, 2004; Hendel and Nevo, 2006a,b). The focus of our work is radically different from that literature. We look at whether consumers take advantage of the opportunities offered by QS. We are not aware of any work that has looked at within-product-line (across pack size) substitution responses during QS as we do here. We explain how our evidence on consumer inattention complements the promotion literature in section 6.3.

The rest of the paper is organized as follows. The next section describes the market and the dataset. Section 3 presents the empirical framework and section 4 presents the findings of our econometric analysis with our main dataset and section 5 repeats the analysis using data from a different country, products and retailers. Section 6 outlines a model of consumer behavior that is consistent with our findings and discusses its implications. Section 7 concludes.

2 Data and stylized facts

Our analysis will utilize scanner data from grocery stores. Such datasets have been widely used in economic and marketing research in recent years. Two different data sources will be used. Our

evidence of bias from heuristic information processing even in markets with large stakes and easily observed information.

³This literature argues that, among other things, opacity depends negatively on salience. Salience increases during promotions due to advertising and more visible price tags. This should increase substitution, not limit it. Change in salience cannot explain our results.

primary dataset covers the market for laundry detergents in the Netherlands. The dataset has been provided by *The Nielsen Company* and covers a period of 120 weeks between September 2002 to December 2004. The Dutch detergent market is dominated by three multinational manufacturers (Henkel, Procter & Gamble, Unilever) but private labels also have a significant presence. Each manufacturer markets several brands and each brand name is carried by several products. We observe every size (pack) of every brand sold in each of the country's four major chains. For each product and size, we observe the total quantity sold, the sales-weighted average price, and the fraction of stores that carry that size.

Our Dutch dataset has two shortcomings. One is that it only features a single product in a single country. The other is that the data are at the chain level rather than the store level. It would appear that stronger results could be obtained by exploiting one of the richer datasets that are available to researchers. We do use one of those datasets, the well-known IRI dataset covering a large number of product categories and thousands of stores across the United States over several years (Bronnenberg, Kruger, and Mela, 2009). Despite its richness, this dataset has some of its own shortcomings that make the analysis we wish to conduct less straightforward than we would like. The Dutch dataset has several features (that we document in detail below) that make it uniquely appropriate for our purpose. We therefore treat the Dutch data as the primary focus of our analysis and use the US data to provide supporting evidence (see section 5).

2.1 Data issues

Table 1 summarizes the empirical challenges of identifying inattentive purchases using non-experimental data. Some issues are common to the entire literature while others are specific to particular datasets. A shortcoming of the Dutch data is that the information is at the level of the chain, not the individual store. Sales are aggregated over all stores and the price is a weighted average. Chain level data could be problematic to study consumer substitution across packs of the same product offered in the same store (Table 1, point 1-a). To start, stores within a chain may not use the same prices. We eliminate this problem by limiting our analysis to a single chain that we know follows a national pricing policy. Stores have no discretion in setting regular or promotional prices. It is one of the leading chains and it carries the major multinational detergent brands as well as its own private label.

Each product is typically offered in either one or two pack sizes (containing the same physical product), respectively 45% and 54% of the times. Sales are roughly balanced between them; sales of the small size are 61.5% of total sales in the median case when two sizes are offered

Table 1: Challenges in measuring the fraction of inattentive regular buyers of large packs

Nature of problem	Discussion/solution
1-Chain level data (store aggregation)	
(a) QS may not apply in some stores if prices are decided locally	We focus on a single chain with a national price and promotion policy
(b) Large pack may not be available in some stores thus preventing substitution	We focus on stores that carry both packs at least 99 percent of the time
2-Weekly data (time aggregation)	
(a) Store may stock out of the small product during a QS	No solution. For our supermarket chain, stock out may happen only as forecasting error or inventory accidents
(b) Promotions may end in the middle of a week thus eliminating substitution during the remaining days	For QS events lasting more than a week, we only consider the first week and drop the next observation
3-Product packs	
(a) Small and large packs are not exact multiples	We focus on products for which the large pack quantity is double the one of the small pack
(b) Consumers have a preference for the large pack	No solution. The substitution prediction does not apply for these consumers
4-Other	
(a) The small pack is not available next to the large pack but only elsewhere in the store	In our supermarket chain, promoted products are available next to non-promoted products
(b) Other marketing actions take place during QS	See discussion in Section 4

during regular weeks. However, not all stores carry the same set of pack sizes. This could be a problem when we study responses to QS: switching to the small pack is not possible in stores that carry only the large pack (Table 1, point 1-b). To be safe, we will consider only promotion events when substitution is possible in at least 99% of all stores.

Another issue, not related with the use of chain level data, is that the small and large packs may not be perfect substitutes (Table 1, point 3-a). In the introduction we argued that the two options are as close to being perfect substitutes as can be practically feasible. This claim requires that the pack sizes are exact multiples of each other, so that the consumer can get the exact same quantity of product with either choice. To ensure this, we select products for which the large pack is exactly twice the size of the small pack. We initially describe the promotions in our sample period using all products sold by the retailer. We then construct our subsample of QS events that satisfy these additional two conditions (universal availability and exact multiple).

This sample of QS events is designed to rule out obvious reasons against the hypothesis of full substitution. We think our approach is conservative. The fact that in our Dutch dataset most products come in two packs and the large pack is usually exactly twice the size of the small pack is a significant advantage over similar datasets. If we want to test for inattention, we would like to do it in an environment where the price comparison is as simple as possible. Our data and our sample selection process allow us to do that.

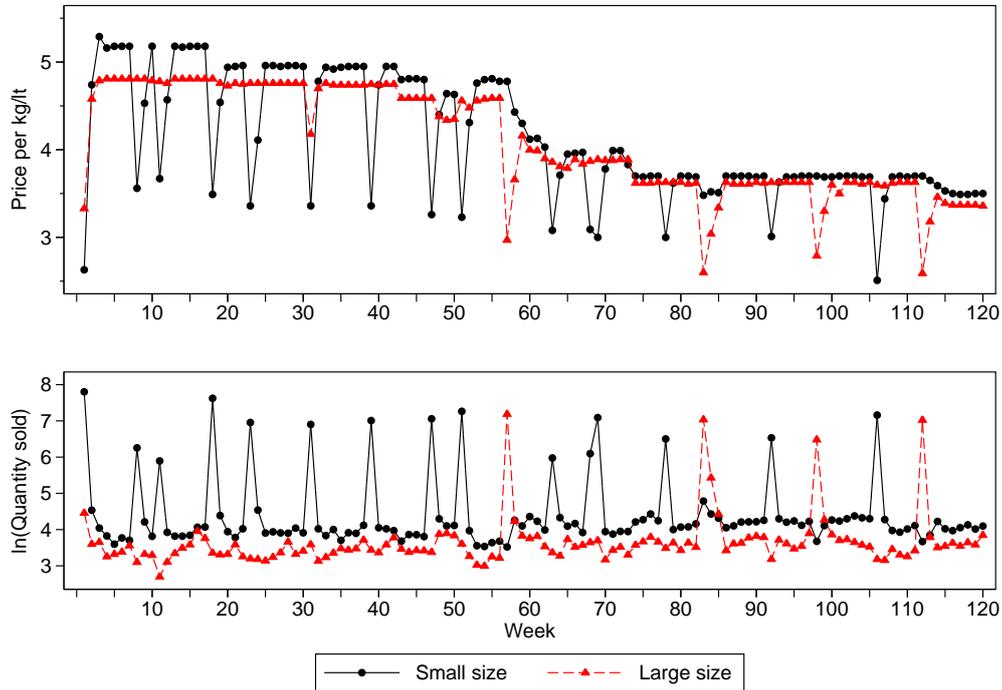
A final issue is that a store may stock out of product in the middle of a promotion. Switching to the small pack is not possible if this option is not available (Table 1, point 2-a). Ruling out this possibility would require store level data on availability at a higher frequency level than is typically available. This problem is also relevant for the entire literature that studies consumer responses to promotion and we do not know of any study that has identified this as a serious issue. We are not aware that stockouts are a recurrent problem for this retailer. A small number of store managers we interviewed assured us that they are rare in their stores. This seems reasonable as frequent stockouts would give the store a bad reputation. A related issue is that promotions may end in the middle of a week (Table 1, point 2-b). This does seem to happen frequently in our data. To be safe, we eliminate observations for the second week of a promotion because they may contain promotion and non-promotion days.

2.2 Promotion

Figure 1 illustrates the nature of the data variation we seek to leverage. The top graph plots the temporal evolution of price per unit for each of the two pack sizes of a selected product.⁴ Promotions are easy to identify as large and temporary downward deviations from the regular unit price. In each promotion the unit price drops sharply for a week, partially recovers in the following week and returns to its original level the week after that. Thus promotions last for between one and two weeks. In the first week all units are sold at the discounted price while in the second week some units are sold at the discounted price and others at the regular price, leading to a sales-weighted price lying somewhere between the two. The bottom graph plots the volume of sales for each size (in logarithms). The impact of promotions is quite striking, with a large spike in sales for the promoted pack. Sales increase by a factor of 5-10 relative to periods with regular price.

The general pattern of a relatively constant “regular” price punctuated by significant, temporary, and fairly regular price reductions resulting in large demand responses is consistent with

⁴All prices in this study are measured per unit of product. Because there is no risk of confusion, we use interchangeably the terms price per unit, unit price, or more simply price.



Product line sold in two sizes throughout sample period. Prices are weekly averages across all stores. Quantity sold is total weekly sales in all stores.

Figure 1: Prices and quantities sold for packs of a selected product line

previous works in the literature on store promotion.⁵ The novel aspect of Figure 1 is that it separately plots the price and quantity for each size of the same product. We are not aware of any work that has disaggregated products by size to study the response to QS. Previous studies have focused on a single size (Pesendorfer, 2002; Hosken and Reiffen, 2004), which amounts to plotting a single time series in each panel of Figure 1, or have aggregated products at the brand level (van Heerde, Leeflang, and Wittink, 2004), which amounts to taking averages. The authoritative studies of Erdem, Imai, and Keane (2003) and Hendel and Nevo (2006a) also do not examine within-product substitution during QS as we do here. Instead, these two studies as well as the rest of the literature attempt to explain the large demand responses to promotions as inter-temporal substitution (through inventory or demand acceleration) and intra-temporal substitution (from other brands and retailers).

⁵See Hendel and Nevo (2006b) and Berck, Brown, Perloff, and Villas-Boas (2008) for detailed analyses of promotions.

Table 2: Descriptive statistics on promotions

	Total proms	Promotions per pack			% time on prom.		% prom. discount	
		Mean	Median	Max.	Mean	Median	Mean	Median
Manufacturer A	152	5.5	4	14	6.9	5.4	26.5	28.9
Manufacturer B	105	7.4	8	12	7.3	7.0	28.2	30.3
Manufacturer C	110	5.5	5	11	6.6	7.4	25.7	28.0
Private label	32	3.7	4	7	5.1	3.9	25.9	27.7
Overall	399	5.8	5	14	6.7	6.7	26.7	29.1

The 399 promotions occur in a total of 5,927 pack-week observations. % time on promotion is the number of weeks on promotion divided by weeks the pack is available. % prom. discount is the discount relative to the median price in the six-week window.

In order to proceed with our analysis we need to provide an operational definition of what constitutes a promotion. We identify a promotion as a temporary decrease in price of at least 10%.⁶ In practice we implement this by looking at a six-week window around any given price. If the price in the current period is at least 10% lower than the modal price during the six-week window, then the current period is labeled as a promotion period. Promotions lasting more than one week are counted as a single event. The use of the six-week window to define promotions means that we cannot identify promotions in the first and last three weeks of the sample, so we are left with 114 weeks of data.

This procedure identifies 399 promotions, the properties of which are summarized in Table 2. In the rest of this paper, we refer to weeks without a promotion as regular weeks/periods for that pack. Products are on promotion on average 6.7% of the weeks. Firm 4 (the private label) is the least frequent promoter. The depth of promotional discounts is in the range 25-30% and is very similar across firms. These price patterns are consistent with those reported elsewhere using detergent in other countries (United States, UK) or other grocery products (Hosken and Reiffen, 2004, among others). Out of the 399 promotions we identify, 197 involve single-item product lines. Of the 202 instances when two sizes are available, the large size is promoted 76 times and the small size 126 times. On 26 occasions both sizes are promoted, which leaves us with 50 solo promotions of the large size and 100 solo promotions of the small size.

⁶Many authors use a 5% threshold. We prefer to be more conservative on what constitutes a promotion. In our analysis of the US data later on we use a 5% threshold because discounts are smaller.

Table 3: Quantity surcharge and promotions

	Regular periods	Promotion periods	All periods
Quantity discount	1,230	66	1,296
Quantity surcharge	90	110	200
Total	1,320	176	1,496
% of time with quantity discount	93.2%	37.5%	86.6%

Table uses only products coming in exactly two sizes. It classifies the 1,496 product-weeks by whether there was a quantity discount or surcharge (rows) and by whether there was a promotion of either product (columns).

2.3 Quantity surcharge

During regular periods, the price of the small pack is almost always higher than the price of the large pack. This corresponds to quantity discounting, a practice that is consistent with standard nonlinear pricing theory. When the small pack is promoted, however, the price order is reversed: the price is lower for the small pack. This is what we call a quantity surcharge.

The frequency of QS events is summarized in Table 3. The pricing schedule displays quantity discounts 93.2% of the time during regular periods, which is broadly consistent with nonlinear pricing theory. When at least one pack size is promoted, this percentage drops to 37.5%. This is quite striking, and it demonstrates that QSs are a frequent side effect of promotions. Most importantly, out of the 100 solo promotions of the small size, 97 result in a QS.

The large discounts on offer during QS periods suggest that there is substantial incentive to substitute. When the small pack is promoted, for example, the median surcharge for the large pack is 28.2%. In absolute terms, this corresponds to a median saving of €0.83 per kilogram. The savings exceed one euro per kilogram in 45% of promotion-induced QS events. QSs occurring during non-promotion periods are substantially smaller, with a median of 5.0% or €0.21 per kilogram. Because QSs outside promotions are small and infrequent, we do not consider them further and focus on demand response during promotion-induced QSs only.

Quantity surcharges have been documented in the marketing literature. Studies have found that “quantity surcharges occur in 16% to 34% of supermarket brands that are available in two or more package sizes” (Sprott, Manning, and Miyazaki, 2003). The marketing literature distinguishes QSs that are a deliberate and systematic policy used in place of quantity discount (not the case for detergent in our sample) and those that are a by-product of promotional

activities (explains events of quantity surcharge during large pack promotions in our sample). The marketing literature has tried to explain deliberate and systematic QS because it challenges nonlinear pricing theory. Little attention has been dedicated to the second type of QS (the type that interests us in this study) because it is less frequent and it occurs as “as an unintentional by-product of common price setting processes” (Sprott, Manning, and Miyazaki, 2003).

As stated earlier, we will restrict our attention to QS events meeting certain criteria in order for our evidence to be as clean as possible. Starting with the 110 events reported in Table 3, we first remove 13 cases where a QS occurs when both products are promoted because consumer behavior might be different in those cases. Second, we restrict our sample to products where the large pack is exactly twice the size of the small pack; this leads to the removal of 21 events. Finally, we remove 8 events where availability of both pack sizes during a QS event might be an issue. This leaves 68 events that will be the focus of the rest of the analysis.

2.4 Substitution during quantity surcharges

Do consumers switch from the large to the small pack during QSs? Figure 1 offers some preliminary eyeball evidence. The small pack is promoted 14 times (drop in the upper time series in the upper panel). If most consumers would substitute from the large to the small pack, we would expect a large drop in the sales of the large pack during QS. This does not appear to be the case. Sales of the large pack (the lower time series in the lower panel) seem to drop in some cases but the decrease is small. Note that unit price comparison can be done by reading product labels because information on unit pricing is posted together with the product’s price.⁷

Table 4 reports the median ratio of sales during a QS to sales in the previous week. Under complete substitution, we would expect the sales of the large pack to be zero in these cases. When we look at all 68 instances of QS (first column), this figure is 73.1 percent, clearly much higher than zero. Sales of the large pack decrease by only 26.9 percent relative to the previous period. This figure is consistent with the evidence from Figure 1. Even though purchasing a large pack is a dominated option during QSs, the majority of consumers do not substitute from the large to the small pack.

It may be argued that the previous week’s sales is not a good benchmark because sales may vary in systematic ways in between promotion periods. In order to account for this, we also calculated the ratio using the median quantity sold during regular weeks in the denominator.

⁷In 1998 the European Parliament and the European Council established directive 98/6/EC on consumer protection, compelling stores to display unit prices in an unambiguous, easily identifiable, and clearly legible way.

Table 4: Ratios of promotion week sales to “regular” sales

	All	Premium	Value
Relative to previous period sales	.731	.843	.589
Relative to median regular sales	.735	.859	.523
Observations	68	46	22

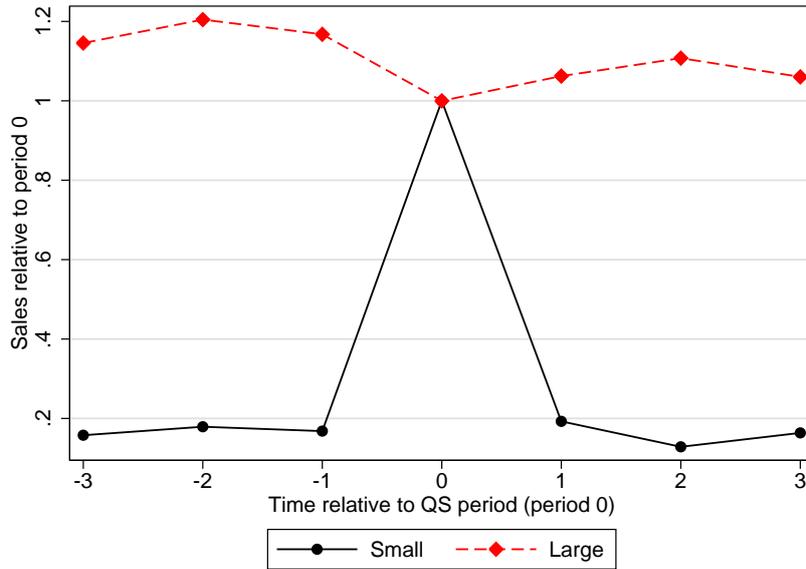
For each QS event we divide sales of the large size by the relevant base (previous period sales or median regular sales). We report the median of these ratios for each case. We use only the 68 QS events that satisfy our strict inclusion criteria.

The number is reported in the second row and is very similar to the ratio reported above.

Does the extent of substitution during QSs vary across products? We split the brands in our sample into value products and premium brands on the basis of their prices. Value products are the store brand (private label) plus Henkel’s Witte Reus and Unilever’s Sunil. The latter two brands sell at a substantial discount relative to other brands sold by the three multinationals. The last two columns of Table 4 report the ratios of interest. The ratios are much higher for premium brands than for value products. A premium large size keeps 84.3 percent of its sales during a QS; a value large size only keeps 58.9 percent. The latter figure drops to 52.5 percent when median regular sales are used. These figures suggest that premium brand buyers respond less to promotional discounts than buyers of value products. This is in line with our intuition, as we expect value brand buyers to more price sensitive than premium brand buyers and therefore more likely to search for bargains.⁸

Figure 2 provides a visual summary of sales patterns before and after a QS event. For each of the 68 QS events we normalize the week of the QS to be time 0 and calculate the fraction q_t/q_0 for $t \in [-3, 3]$. We take the simple average of these fractions and plot it in Figure 2. Sales of the small size before and after the QS are about 20% of sales in the promotion week (meaning that a promotion of the small size increases its sales by a factor of five). Sales of the large size drop on average about 15-20% during a QS relative to the previous periods and recover afterwards,

⁸We have also tried to compare substitution responses for small and large promotion discounts (size of the drops in the upper panel in Figure 1). One might have expected the substitution response to increase with the promotion discount. We did not find evidence in support of this hypothesis.



For each of the 68 QS events we normalize the week of the QS to 0 and consider the three preceding weeks and the three following weeks. We divide the sales of the small (large) size in each period by the sales of the small (large) size in period 0 (the QS week). We report the median of these fractions for each period.

Figure 2: Normalized sales paths of small and large size around a QS

though not fully.

3 Empirical framework

How widespread is the lack of substitution? Can we say that 73.1 percent of purchases of the large pack fail to take advantage of QSs? Doing so assumes that in the absence of a QS, the sales of the large pack would be the same as the sales we observe in the week preceding the QS. This may not be a bad assumption. At least, it is supported by Figures 1 and 2: regular period sales of both packs are fairly stable over time. But the issue requires further scrutiny because promotion is a firm decision that may be correlated with demand and supply.

We must develop the appropriate empirical framework to match the nature of our data. We have time series data on several products coming in two sizes. For each product, in every period t we observe three key variables: the sales of the small and large size Q_t^S and Q_t^L and whether

there is a QS or not, denoted by the indicator variable S_t (equals 1 for QS periods). The sales variables Q_t^S and Q_t^L are count data as they have a discrete, nonnegative support.

3.1 A simple model

We call the purchases of the large pack during QSs inattentive purchases. We want to estimate the fraction of regular week purchases of the large pack that are inattentive. To do so, we present a simple framework that explains the sales of the small and large pack during regular and QS periods. Each product-week observation in our data can be classified into one of five categories: (i) weeks with no promotion are regular weeks; (ii) weeks with a promotion of the small pack resulting in a QS; (iii) weeks with a promotion of the small pack that do not result in a QS; (iv) weeks with a promotion of the large pack; and (v) weeks with both packs on promotion. For the sake of exposition, we ignore cases (iii) and (v) because they occur infrequently (2 and 26 occurrences respectively) and case (iv) because these observations are not necessary to estimate the inattention parameter. We are left with just regular weeks and QS weeks.

We start by outlining a simple theoretical model that will form the basis of our empirical framework. Consider a single product that is sold in two sizes. During regular weeks, both sizes are sold at their regular price. During QS weeks, the small size is discounted. In any given week, there is a regular stream of consumers who walk into a store with the intention of buying a specific item. During regular weeks these visits translate into sales with means θ^L and θ^S respectively. For example, one can think of regular sales as a Poisson process with rates of arrival θ^L and θ^S respectively.

During QS weeks, a regular large pack buyer may switch to the small pack option or she may continue to purchase the large pack. When the latter occurs, we say that the buyer is inattentive. We want to estimate the fraction of the regular buyers of the large pack who are inattentive. Denote this fraction, our structural parameter of interest, α .⁹ Expected sales of the large size during a QS is $\alpha\theta^L$; the remaining $(1 - \alpha)\theta^L$ will switch to the small size.

In addition to the switchers, the small size will attract increased demand emanating from a variety of responses that the literature has documented (stockpiling or cross-period effects, substitution across brands, category expansion explained by bargain hunters). We denote the mean of the new demand by $\theta^{qs,S}$. There is no reason to expect a similar increase in demand for the large size, unless some of the new consumers are inattentive and inadvertently purchase the

⁹Note that we could have called our structural parameter the fraction of regular purchases of the large pack that are inattentive. The reader should keep in mind that we associate purchases to buyers to simplify the exposition.

large size. Since these are consumers who are responding to the promotion, they are unlikely to be inattentive; it is therefore reasonable to assume that $\theta^{qs,L} = 0$.¹⁰ Expected sales for each size in regular and QS weeks are summarized in equations (1)–(2):

$$E [Q_t^L] = \begin{cases} \theta^L & \text{if } S_t = 0 \\ \alpha\theta^L & \text{if } S_t = 1 \end{cases} \quad (1)$$

$$E [Q_t^S] = \begin{cases} \theta^S & \text{if } S_t = 0 \\ \theta^S + (1 - \alpha)\theta^L + \theta^{qs,S} & \text{if } S_t = 1 \end{cases} \quad (2)$$

The means $(\theta^L, \theta^S, \theta^{qs,S})$ are parameters to be estimated. This is a highly stylized representation of demand that aims to capture the notion of “regular” sales. In particular, we have assumed that the means θ^L and θ^S are fixed. We could have assumed instead that they depend on control variables, $\theta^L(Z)$ and $\theta^S(Z)$, where Z accounts for the fact that regular sales could vary over time and as a function of price, as is sometimes done in the literature. We chose not to do so for two reasons. From an expositional viewpoint, carrying control variables does not shed additional insight into the identification problem at hand. Second, during our sample period the sales of each pack remain relatively constant in the absence of a promotion (see Figure 1). Thus, we do not expect control variables to play an important role.

Note that from equation (1) it follows that

$$\alpha = \frac{E[Q^L|S = 1]}{E[Q^L|S = 0]}. \quad (3)$$

This means that a valid estimate of the fraction of regular buyers who are inattentive (under the assumption $\theta^{qs,L} = 0$) can be obtained by taking the ratio of mean sales in surcharge periods to mean sales in regular periods. This is close to the ratios presented in Table 4 and those displayed in Figure 2 in the previous section. Those simple ratios can therefore be interpreted as quick-and-ready estimates of α . The econometric framework allows us to formalize the estimation process and facilitates hypothesis testing. Equations (1) and (2) can be written more concisely

¹⁰In an earlier version of the paper we discussed how to identify α when $\theta^{qs,L} > 0$. We did not find any evidence that $\theta^{qs,L} > 0$.

as:

$$E [Q_t^L | S_t] = \theta^L - \theta^L(1 - \alpha)S_t \quad (4)$$

$$E [Q_t^S | S_t] = \theta^S + ((1 - \alpha)\theta^L + \theta^{qs,S})S_t \quad (5)$$

These are two econometric models that can be taken to the data. In fact we only need to estimate equation (4) because, as is clear from equation (3), α is identified from that equation alone. Since we are not particularly interested in θ^S and $\theta^{qs,S}$, we will ignore equation (5) and use only observations on the large size to estimate equation (4). This also helps us avoid possible endogeneity problems (see next subsection). Note that (4) is a nonlinear equation as the two structural parameters θ^L and α enter multiplicatively.

3.2 Exogenous variation in S_t

One issue with interpreting model (4) as a regression equation is that the explanatory variable S_t could be correlated with unobserved demand shifters in the disturbance term. It could be the case, for example, that promotions are initiated in order to boost sales of products losing market share. If so, the error term in (4) will be correlated with S_t .

To address this point we refer back to the empirical literature on promotion that faces the same endogeneity issue. The promotion literature views promotions as supply side decisions rather than a response to changes in preferences. After a thorough review of the issue, Erdem, Imai, and Keane (2003) conclude that “We find it extremely implausible that these deals are the result of manufacturer, wholesaler or retailer responses to aggregate taste shocks, for several reasons. Why would demand for a good like ketchup or yogurt suddenly jump every several weeks and then return to normal? And how could sellers detect such a jump quickly enough to incorporate it into daily or weekly price setting? [...] A more plausible explanation for the observed price variation is some sort of inter-temporal price discrimination.”¹¹ Our data are consistent with this view. The timing and frequency of promotions do not appear to be correlated with the demand for the product. Figure 1, for example, shows that the sales of the large pack remains fairly constant during non-promotion weeks, and in particular in weeks adjacent to a promotion. This pattern generalizes to other promotion events (and therefore QS events) in our sample.

¹¹See also the discussion in Hendel and Nevo (2006a), page 1645. Hendel and Nevo (2013) present and estimate a model of promotions as intertemporal price discrimination.

That being said, it is also true that the event of a QS may be correlated with the level of consumer inventory in the market. This endogeneity issue has received much attention in the empirical promotion literature. It matters for the interpretation of the parameter estimates if we were to jointly estimate (4) and (5). Keane et al. state the problem clearly: “On the other hand, an important reason for endogeneity of prices in demand models is the failure to account for consumer inventories [...]. If prices are persistent over time and consumers engage in stockpiling behavior, then inventories will be correlated with current prices. This causes price to be econometrically endogenous due to the omitted variables problem, even though price fluctuations are exogenous from the point of view of consumers.”¹² But this is not an issue for us because we are not interested in measuring the own response to promotion, that is equation (5), which depends on the level of consumer inventory in the market. We are interested instead in measuring substitution responses (sales of the large pack during QS) which are not influenced by the level of consumer inventory in the market. Stated differently, while S_t may well be endogenous in equation (5), there is no reason why it should be endogenous in equation (4). This is because the level of inventory in the market influences only the response to promotion (consumers who buy the small pack); not the regular consumers who buy the large pack.

4 Estimation and results

Given that the dependent variable is a count variable, model (4) should in principle be estimated using count data techniques. However, our dependent variable has a very large mean because our data at the national chain level. This makes least squares a valid approximation. We will therefore estimate our model with conventional least squares methods.¹³ The regression equation implied by model (4) is:

$$Q_t^L = \theta^L - \theta^L(1 - \alpha)S_t + \eta_t. \quad (6)$$

Either ordinary or nonlinear least squares (NLS) can be used to estimate equation (6). The error term η_t is unlikely to be iid across observations; in particular the variance will likely be smaller when $S_t = 1$. It is therefore important to use robust standard errors that correct this violation of the iid assumption. Despite the simplicity of OLS, we prefer to estimate (6) using NLS. The appeal of NLS is that it delivers direct estimates of our structural parameters and

¹²For example, the literature reports that demand of the promoted product during promotion periods increases with the duration since last sale; e.g. Pesendorfer (2002); Hendel and Nevo (2006b).

¹³Note that this will not be the case in the next section when we use store-level data.

Table 5: Estimates of equation (6)

	Premium products				Value products			
θ^L	26.35** (2.922)	42.97** (7.604)	32.22** (4.212)	2.553** (0.310)	3.610** (0.523)	4.087** 0.645	58.30** (5.442)	45.85** (3.399)
α	0.696** (0.108)	0.630** (0.119)	0.745** (0.122)	1.550** (0.330)	0.610* (0.207)	0.461** (0.093)	0.559** (0.121)	0.561** (0.114)
N	16	37	33	14	23	21	18	31
R^2	0.916	0.600	0.697	0.897	0.727	0.704	0.896	0.881
Test of $H_0 : \alpha = 1$ (p-values)								
	0.014	0.004	0.045	0.122	0.073	0.000	0.002	0.001

Estimates obtained using nonlinear least squares. Robust standard errors are reported. Significance levels: † : 10%, * : 5%, ** : 1%.

makes hypothesis testing straightforward.¹⁴ It can also be easily extended to the case of multiple products, where α might be common across products or could vary with product characteristics. This cannot be easily done with OLS, especially because products have different market shares and frequency of QS occurrences.

From our full dataset we keep observations that meet the following criteria: (i) the product is available in two sizes, the large size is exactly double the small size, and both are available in at least 99% of stores; (ii) there are at least 14 observations of the product; and (iii) Qs are the result of a promotion of the small size with no simultaneous promotions of the large size. This leaves us with 193 observations; 141 of those are regular weeks and 52 are QS weeks (compared to 68 QS weeks in the analysis in the previous section). There are eight different products for which we have enough observations to support inference; four are premium and four are value products.

In Table 5 we present results from estimating equation (6) separately for each product. In seven out of eight cases the estimates are very reasonable. The α parameter is estimated in the range (0.461, 0.745), suggesting that the fraction of inattentive varies across products but within a modest range. Premium products have a higher α , consistent with the notion that buyers of premium products are less price elastic and therefore more likely to be inattentive. In one case

¹⁴We have also run the OLS regression $Q_t^L = \beta_0 + \beta_1 S_t + \eta_t'$. The structural parameters can be recovered as $\hat{\theta}^L = \hat{\beta}_0$ and $\hat{\alpha} = 1 + \hat{\beta}_1 / \hat{\beta}_0$ and they are identical as with NLS. It should also be noted that regardless of estimation method, nonlinearity implies that estimates are biased but weakly consistent. From Jensen's inequality, $E\hat{\alpha} = E((\hat{\beta}_0 + \hat{\beta}_1) / \hat{\beta}_0) \geq 1 + E\hat{\beta}_1 / E\hat{\beta}_0$ with the inequality approaching an equality as $\hat{\beta}_0$ converges in probability to $E(Q^L | S = 0)$. Thus $\hat{\alpha}$ over-estimates α when the sample used to estimate $\hat{\beta}_0$ is small. This might be an issue with the Dutch data but not with the US data, where we typically have 200+ observations per product.

we get an estimate of α that exceeds 1 but not statistically significantly so.¹⁵ The simple mean of the eight α parameters is .727, very close to the ratio presented in Table 4.

It may be desirable to pool data on different products in order to either get a single estimate of α or to allow α to vary with product characteristics. For example, by pooling the data we could formally test whether α is greater for premium products than for value products. Letting k index products, the fully general pooled version of (6) has both parameters (θ_k^L, α_k) be product specific, giving the results in Table 5. The case for product-specific θ_k^L parameters is clear: they represent regular sales, which vary substantially across products. On the other hand, there is no reason to think that the inattention parameter should vary substantially across products, except perhaps along the premium/value distinction; the estimates in Table 5 are consistent with this view. We therefore estimated two models with the pooled data. In one we fix $\alpha_k \equiv \alpha$ for all k . In the other we let α vary along the premium/value dimension; in other words, we estimate α_p and α_v for premium and value products respectively.

Estimated parameters of interest are presented in Table 6 (estimates of the θ_k^L 's are omitted). The pooled data give an estimate of $\alpha = 0.601$, somewhat smaller than .727, the simple mean of the individual estimates in Table 5. Part of the difference can be explained by the fact that the simple mean does not account for the number of observations; the weighted mean of the estimates in Table 5 is .683.¹⁶ The estimated α_p (for premium products) is about ten percentage points lower than α_v (for value products), but the difference is not statistically significant. Perhaps with more data we would be able to draw stronger conclusions about differences in the inattention parameter across different products.

What have we learned about inattentive purchases? The goal was to estimate the fraction of regular buyers of the large pack (the consumers who buy the large pack in the absence of promotion) who are inattentive. We started from the 73.1% figure of inattentive consumers computed as the fraction of sales during QS relative to regular periods. We have shown that this is not a bad estimate under the assumption that Qs do not happen simultaneously with other marketing actions that influence the number of regular consumers (point 4-b in Table 1).¹⁷

¹⁵This product has the smallest number of observations among the eight we analyze. It would have been justifiable to drop this product but we chose to report it to illustrate where our method might go wrong. In the 14 periods that we observe this product, there are three QS events. In two of the three events, sales of the large size rise. This is why we get $\alpha > 1$. It could be that sales of the large size rise in those two periods because of additional marketing activity or for other exogenous reasons. With such a small number of events, it is certainly conceivable that this is just a random outcome.

¹⁶Some of the difference may also be explained by the fact that the frequency of promotion differs across products.

¹⁷We also obtained qualitatively similar estimates in an earlier version of the paper using the full dataset and a different methodology (Clerides and Courty, 2010).

Table 6: Pooled estimates of equation (6)

α	0.601** (0.059)	
α_v		0.560** (0.082)
α_p		0.659** (0.089)
$H_0 : \alpha_v = \alpha_p$ (p-val)		0.411
N	193	193
R ²	0.747	0.748

Estimates obtained using nonlinear least squares.
 Robust standard errors are reported. Significance
 levels: † : 10%, * : 5%, ** : 1%.

The econometric analysis produced a lower figure of 60.1%. There is some evidence that buyers of premium products are less attentive but richer data are needed for more robust conclusions.

5 Generality: different country, stores and products

We repeat the analysis with data from a different source, covering a different country, different retail chains and stores, and different product categories. The main objective of this exercise is to demonstrate that inattention is not specific to Dutch consumers, to detergent, or to the way the price and quantity variables are constructed and recorded in the Nielsen dataset. In the process, we will highlight the lessons we learned about the challenges involved with documenting inattention with store-level rather than chain-level data.

We use the well-known IRI database for the period 2001-2005.¹⁸ The type of information provided by IRI is very similar to our Nielsen data: weekly data on unit price and sales at the pack level. But there are some important differences with respect to the size and scope of the data. The Nielsen data are for a single product category (detergent) and are recorded at the chain level. The IRI database covers 31 different product categories and the information is recorded at the store level for a wide range of stores belonging to different chains and located in different local markets. We selected four product categories: detergent for comparability, and three additional categories of products that are storable and have constant consumption:

¹⁸For details on the database see Bronnenberg, Kruger, and Mela (2009).

Table 7: Observation counts and pack sales

	Detergents	Blades	Shampoos	Diapers
Observations	1,819,284	164,894	114,168	96,366
Weeks	260	260	260	260
Stores	1001	414	128	138
Products per store	11.1	1.6	6.1	6.5
Median pack sales	6	7	6	7
25-75 percentiles pack sales	2-12	4-13	3-10	4-12

Years 2001-2005. Includes product-store pairs with at least two packs and weekly average sales across packs of at least 10 units.

razor blade cartridges (henceforth blades), shampoos and diapers.¹⁹ Another difference with the Dutch data is that the stores in the sample may not follow a chain pricing policy. Because the timing, length and depth of promotion may vary across stores, we cannot aggregate sales across stores.

Having rich, store-level data is very useful but has a downside in that most items sell just a few units per store every week. Median sales across all packs of detergent in our sample is just 6 units. In 25 percent of the week-store-pack observations at most 2 units are sold and in 75 percent of the observations at most 12 units are sold. Similar patterns hold for the other three product categories. This feature of the data – “small unit sales” – requires some adjustments to our estimation strategy. The support of our dependent variable is restricted to be nonnegative and discrete. This was also true for the Dutch chain-level data, but in that case the mean level of regular sales (θ^L) was high enough to make the normal approximation for the error term reasonable. This is not the case here because θ^L at the store level is small. To accommodate this we estimate our equation using Poisson regression, which is appropriate for count data.²⁰

For maximum comparability we selected only product-store pairs that sell in exactly two sizes. We dropped products that averaged fewer than 10 units sold per week (summed over the two sizes). Table 7 presents the basic features of the extracted dataset. The unit of observation is at the level of product, size and store. Detergent is the product with the most observations. We observe detergents in 1001 distinct stores and there are on average 11 detergent products per store. In contrast with the Dutch data, the large size is usually not an exact multiple of the small size; for this reason we did not impose this requirement with the US data. Price are per ounce for detergent and shampoo and per physical unit for blades and diapers.

¹⁹Selecting products with constant demand over time avoids problems associated with demand seasonality and peak-demand periods that could be correlated with promotions (Chevalier, Kashyap, and Rossi, 2003).

²⁰In practice it turns out that OLS with robust standard errors yields very similar estimates to Poisson.

Table 8: Promotion and Quantity Surcharge

	Detergent	Blades	Shampoo	Diapers
Promotion frequency	8.43%	2.87%	9.96%	9.41%
Small pack promotion: frequency	9.3%	3.9%	12.3%	10.9%
Small pack promotion: frequency of QS	71.6%	77.5%	45.0%	69.4%
Small pack promotion: price reduction	25.6%	12.3%	21.7%	11.1%
Median price ratio (p_s/p_l) during QS	.79	.93	.93	.94

The last two lines use the sample of QS observations (counting only the first week if a QS lasts more than one week and eliminating QS events if both sizes are promoted). The price reduction is computed using the price in the previous week.

Table 8 presents similar information on promotion and QS as Tables 2 and 3. We define a promotion as a 5 percent reduction in price relative to the median price within a window of ± 3 weeks.²¹ Razor blades are much less likely to be on promotion than the other products. Consistently with earlier results, the small pack is more likely to be promoted and price reduction during a small pack promotion is substantial, between 11% and 26% depending on the product category. As before, we say that a QS surcharge occurs if the small pack is promoted and if the per unit price is lower for the small pack than for the large pack. QSs are very likely when the small pack is promoted (between 45% and 78% of the time). The impact of a quantity surcharge on the price ratio is largest for detergent. For the other three products, the savings from switching pack size during a QS is a modest 6-7 percent.

Table 9 replicates Table 4. As in the previous analysis, we count only the first week if a QS lasts more than one week and we consider only events where only the small size is promoted (thus we eliminate events where both sizes are promoted). The first line reports the median sales ratio across all QS events (across all products and stores). As before, the ratio is computed as the sales of the large pack during a QS week to the sale in the previous week. The results do not change if we use the median sales in periods without promotions. The median ratio for detergent, .727, is very close to the figure of .731 reported in Table 4. The figures for the other three product categories are higher. This could be because purchasing behavior varies across products or because the price reduction is smaller for these products.

As a robustness check, the lower panel in Table 9 restricts the sample to product-store observations that have median sales of the large pack of at least 10 units. Doing so dramatically reduces the number of QS events, but the estimated sales ratios do not alter the picture in a significant way.

²¹We chose a lower threshold (5% here instead of 10% in the Dutch dataset) because two of the added categories (blades and diapers) have lower promotional price reductions (see Table 8).

Table 9: Median ratio of promotion week sales to sales of previous week

	Detergents	Blades	Shampoos	Diapers
All QS events	.727	1.00	.818	.80
Count of QS events	26,081	1,530	1,979	2,152
QS events for products with median sales of large pack ≥ 10	.900	1.00	1.09	.818
Count of QS events	412	254	45	523

In the first two lines, the median sales across both packs must be greater than 10. In the last two lines the median sales of the large pack must be greater than 10.

Table 10: Product-store estimates of α with at least 3 QSs

	Detergents	Blades	Shampoos	Diapers
Number of estimated α 's	786	216	199	247
Mean	0.658	1.09	0.91	0.789
Median	0.664	1.04	0.77	0.788
Standard Deviation	0.295	0.352	0.776	0.234
5th percentile	0.193	0.569	0.494	0.365
10th percentile	0.277	0.676	0.545	0.475
25th percentile	0.463	0.878	0.666	0.63
75th percentile	0.832	1.259	0.913	0.94
90th percentile	1.005	1.549	1.136	1.08

We restrict to product-store with at least 25 observations (weeks). For detergent we require 200 weeks.

We obtain econometric estimates of product-store specific α 's by estimating model (4) using Poisson regression. To be on the safe side, we select only product-stores with at least three QS events. Table 10 reports statistics on the distribution of the thousands of α 's estimated this way. The mean and median of the estimated α 's are similar in magnitude to the ratios reported in Table 9. There is substantial dispersion in the estimated inattention parameters. The 5th percentile value of α for detergent is .193. When we restrict to products with at least 10 QS events, that figure increases to .277. About 10% of the estimated α 's are greater than one.

The estimates of α for a given product category have a fairly wide support. This could be because inattention varies across products and stores or because the estimates of α are not precise. Imprecise estimates is an issue if the hypothesis $\alpha = 0$ cannot be rejected. This is not the case: more than 97% of product-stores estimates are statistically different from zero and this holds for all four product categories. Inattention prevails for most product-stores: 75% of the product-stores in the sample have at least 46% inattentive buyers for detergent, 88% for razor blades, 67% for shampoos and 63% for diapers. Interestingly, inattention is highest for

razor blades and lowest for detergent. We cannot reject the hypothesis that all consumers are inattentive ($\alpha = 1$) for 40% of detergent estimates, 80% of razor blades, 59% of shampoos and 45% of diapers. Inattention parameters greater than one are not consistent with assumption $\theta^{qs,L} = 0$ (see footnote 10). We test the hypothesis that $\alpha \leq 1$ for each product category. The hypothesis is rejected in less than 3% of the cases for detergent, shampoo and diapers; for blades it is rejected almost 13% of the time.

6 Discussion

6.1 Theories of grocery sales and nonlinear pricing

Our evidence of passive purchases relates directly to the industrial organization literature on grocery sales. A central assumption in many models of retail pricing is that consumers differ in their willingness to search for the best deal. Some consumers compare prices while others do not. They may compare prices across stores (Varian, 1980), over time (Sobel, 1984), or across products (Lal and Matutes, 1994; Hosken and Reiffen, 2007). This is an essential ingredient in models that explain why retailers can sell the same product at a regular price and also at promotional prices that are substantially lower. The literature has labeled the consumers who do not respond to temporary price differences as passive, naive, loyal, captive, uninformed, non-strategic or regular, depending on the model.

The distinction between strategic and passive has proven useful in empirical studies of promotion. It is consistent, for example, with the finding that sales are fairly constant during non-promotion periods (passive consumers) and with the large response during promotion (strategic consumers). But, as pointed out by (Pesendorfer, 2002), this is not definite evidence of passive purchases. It is also consistent, for example, with heterogeneous consumer inventory costs: high inventory costs consumers buy at fixed intervals while low inventory cost consumers stock out during promotions (Hendel and Nevo, 2013). The finding that many consumers buy dominated options during Qs offers unique evidence in support of the assumption that some consumers do not compare prices.

The theory literature also assumes that the consumers who actively compare prices have a lower willingness to pay than the passive consumers. Passive purchases are associated with loyal consumers of branded products while value products purchases are made by price sensitive consumers who are more likely to make price comparisons. This is consistent with the finding in Table 4: passive purchases are less common for value products than for the premium brands

sold by the big multinationals.

Our results are also related to nonlinear pricing theory. This theory explains quantity discounts as a mechanism for exercising price discrimination (Maskin and Riley, 1984; Stole, 2007).²² The existence of QS during promotion periods presents a challenge to nonlinear pricing theory: why do firms violate the concavity of their pricing schedule? In the standard model firms would not want to do that because it would cannibalize the sales of the large size. This challenge disappears, however, once one recognizes that inattentive consumers do not respond to temporary promotion opportunities. The assumption at the heart of nonlinear pricing theory, that the buyers of the large pack would substitute to the low pack in the event of QS, does not hold for temporary price reductions. Thus, one should be careful in the conduct of empirical analysis of nonlinear pricing and should make a point of distinguishing promotion and non-promotion periods. The derivation of the optimal nonlinear price schedule has to separately account for consumer sorting across pack sizes during promotion and non-promotion periods because short and long term demand substitution elasticities across packs could be very different.

Finally, the grocery sales literature will have to develop models that integrate promotion and nonlinear pricing decisions. Everything else equal, firms should promote less often products for which there is more substitution. Interestingly, value brands in the Dutch data, for which there is more substitution, are promoted less often than premium brands. The median value product is on promotion 4.4% of the time while the median premium product is on promotion 8.3% of the time, almost twice as often. This indicates that firms internalize the fact that promotions can come at the cost of losing sales from competing packs – the cannibalization effect.

6.2 Rational inattention

The fact that grocery shoppers are not aware of all prices at all times has been recognized for a long time in the marketing literature. Cognitive research has shown that the majority of consumers cannot correctly quote the price of items they have just placed in their shopping cart (Monroe and Lee, 1999). Dickson and Sawyer (1990) find that “more than half of the shoppers who purchased an item that was on special were unaware that the price was reduced.” In addition, consumers consider few alternatives for each item they purchase and spend little

²²Both the economic and marketing literatures have applied the theory of nonlinear pricing to consumer packaged goods and most empirical studies use grocery data. McManus (2007) applies the theory of nonlinear pricing to the market for coffee drinks, Cohen (2008) to the paper towel market, and Allenby, Shively, Yang, and Garratt (2004) to the beer market. In a review of nonlinear pricing for the *Handbook of Pricing Research in Marketing*, Iyengar and Gupta (2009) quote consumer packaged goods as one of the leading applications of nonlinear pricing.

time doing so (Hoyer, 1984).²³

Could it be rational that many consumers do not make simple unit price comparisons? It is difficult to answer this question without some knowledge of the costs of processing information. The cost of visually scanning the supermarket aisle for possible promotions seems quite small and the savings of close to one euro in the case of Dutch laundry detergent are not negligible. To rule out the possibility that there could be a QS, a consumer has to find out whether a smaller pack exists, locate it, make sure it is the same physical content, and then compare the unit prices. This may take several seconds. Hoyer's (1984) study of consumer in-store decision making for laundry detergents reports that the median consumer considered only 1.19 packages when making a purchase and spent 4.77 seconds doing so. Taking this literally, it takes about 5 seconds to check a price. Engaging in a price comparison requires checking two prices, which takes 10 seconds.

The rational inattention hypothesis says that consumers take into account the time cost of keeping track of QS. For the sake of illustration, consider a consumer purchasing a large pack. Say that the probability of finding a QS is p and the savings from substituting from large to small pack is $\text{€}x$. The consumer's opportunity cost of time is $\text{€}w$ per hour. A w higher than the hourly wage is not implausible if consumers get disutility from engaging in price comparisons. Finally, say it takes s seconds on average to seize a QS opportunity. This could include the time of keeping track of QS and the time of comparing unit prices. The consumer rationally ignores QS altogether if $s \frac{w}{3600} > px$. For example, if $p = 10\%$, $x = \text{€}1$, and $w = \text{€}20$ a consumer will check for QS only if doing so takes less than 18 seconds.

Costly information processing could be a reasonable explanation for the finding that some consumers do not substitute. This is consistent with the rational inattention literature in macroeconomics that argues that consumers re-optimize decisions only periodically when the processing of price information is costly. This is related to the old notion of search costs (Stigler, 1961) and is therefore not entirely new. An important difference is that price information is freely available in models of rational inattention. Terms like consumer inattention or limited attention have been used to describe the fact that consumers may not be fully attentive to (freely available) prices.

²³Lennard, Mitchell, and McGoldrick (2003) report that only half of the consumers use unit pricing as an information source to find the best option. The other consumers find that the information is too complicated to use or report that they do not have the time to compare prices. The marketing notion of "consideration set", that goes back to Howard and Sheth (1969), argues that consumers only look at a selected subset of products. More recently, Piccione and Spiegler (2009) have assumed that consumers follow procedural rules to decide what product to purchase (e.g. first choose a pack size and then select a product of that size).

6.3 Empirical demand models, consumer behavior, and promotion literature

Most demand models used in empirical studies of nonlinear pricing (Cohen, 2008; Thomas, 2011) or promotions (Hendel and Nevo, 2006a) typically assume that consumers know all prices relevant to their purchase, and most importantly, that they purchase the item with the lowest price.²⁴ Our evidence suggests the possibility that some consumers do not compare prices even when search costs are tiny. In that light, our findings present a direct challenge to demand models that assume that consumers are fully informed about all products' prices. All demand models discussed above predict that consumers will exploit substitution opportunities. We find that a substantial fraction will not.

Take the empirical literature on promotion. Many studies have documented large consumer responses to promotions (the positive spikes in the lower panel of Figure 1). Clearly, the consumers who buy mostly during promotion periods are attentive and represent a significant fraction of total sales. Our evidence shows that other consumers are inattentive to temporary price reductions. This demonstrates the existence of important differences in how consumers respond to promotion opportunities. This novel dimension of consumer heterogeneity – degree of attentiveness to short term price changes – has been largely ignored in the empirical literature on promotions. Discrete choice models assume that all consumers process price information uniformly. The literature has struggled to explain the complex ways in which a product's sales depend on the time elapsed since the last own promotion, on whether other products are being promoted and on whether other stores offer promotions. To do so, studies have allowed for heterogeneity in price sensitivity, storage cost, willingness to substitute, and other dimensions. Heterogeneity in how consumer process price information, which is largely absent from the literature, offers new opportunities.

Another important implication is that one has to be careful in dealing with items sold in multiple pack sizes in estimating demand responses to promotions. We are not aware of any promotion study that separately models the demand for different pack sizes. Studies focus on a single pack size or aggregate all pack sizes for a given product. Our evidence shows that the sales of a given pack size responds only partially to the price of the other pack. Not accounting for incomplete substitution across pack sizes is restrictive.

Ignoring passive consumers attributes the limited degree of substitution to other factors that are included in the model, with the possible implication, for example, of overestimating the

²⁴Seiler (2013) makes the point that people do not know all prices in the store but only the prices of those products that they have on their shopping list. But he assumes that once they have, say, detergent on their shopping list, then they check all detergent prices in the store.

impact of intangible variables such as brand loyalty or storage cost. Consumer inattention may wrongly lead one to conclude that there are strong preferences for specific pack sizes. One may also wrongly attribute the limited extent of cross-brand substitution during promotions to brand loyalty. The common finding in the promotion literature that short-run elasticities are smaller than long-run elasticities has given rise to a number of explanations such as storage capacity, commitment, adjustment frictions, or switching costs. Our evidence suggests that demand substitution may be small in the short run for entirely different reasons than adjustment friction or switching cost. Consumers do not switch because they are passive shoppers.

7 Summary and conclusions

Entering to a store with a shopping list and buying what is on the list without checking prices is not an unusual thing to do. For the large items on the list, such behavior that is not attentive to prices will miss a saving opportunity when the small pack of the same item is offered on promotion. This paper leverages QS events associated with promotions to measure the extent of inattentive purchases. We use two sources of non-experimental data from different countries to estimate the fraction of regular buyers (the consumers who buy during non-promotion periods) who make inattentive purchases during QS periods. We highlight the inference challenges that arise from using scanner data and argue that the unresolved challenges could be addressed only with experimental data.

Dutch consumers do not fully exploit substitution opportunities arising when a small pack size of a detergent product is heavily discounted. Looking specifically at value brands, we found that the large size of a value brand loses 44% of its sales when its smaller counterpart is promoted. This means that the rest of regular buyers of large size value products choose a dominated option during promotions of the small size. With premium brands, we pick up a smaller impact (34%) on the sales of the large size in a similar situation. Similar patterns are found for American consumers and for three other products.

These findings come from estimating consumer responses to temporary price reductions. The fact that we find only partial substitution is somewhat surprising but not shocking. Keeping track of the prices of every product and its close substitutes on a weekly basis is a time-consuming and, for many people, unpleasant task. Back of the envelope calculations suggest that it is quite plausible that the cost of time associated with the process of collecting and processing information might exceed the expected benefit from it, at least for some consumers. These consumers will choose to not be attentive to temporary price changes. Information complexity

does not have to be great to generate significant deviations from the predictions of the standard rational, full-information model.

References

- Allenby, G., T. Shively, S. Yang, and M. Garratt, 2004, “A Choice Model for Packaged Goods: Dealing with Discrete Quantities and Quantity Discounts,” *Marketing Science*, 23, 95–108.
- Berck, P., J. Brown, J. M. Perloff, and S. B. Villas-Boas, 2008, “Sales: Tests of Theories on Causality and Timing,” *International Journal of Industrial Organization*, 26, 1257–1273.
- Bronnenberg, B. J., M. W. Kruger, and C. F. Mela, 2009, “The IRI Marketing Data Set,” *Marketing Science*, 27, 745–748.
- Chetty, R., A. Looney, and K. Kroft, 2009, “Salience and Taxation: Theory and Evidence,” *American Economic Review*, 99, 1145–1177.
- Chevalier, J. A., A. K. Kashyap, and P. E. Rossi, 2003, “Why Don’t Prices Rise During Periods of Peak Demand? Evidence from Scanner Data,” *American Economic Review*, 93, 15–37.
- Clerides, S., and P. Courty, 2010, “Sales, Quantity Surcharge and Consumer Inattention,” *CEPR Discussion Paper No. 8115*.
- Cohen, A., 2008, “Package Size and Price Discrimination in the Paper Towel Market,” *International Journal of Industrial Organization*, 26, 502–516.
- DellaVigna, S., 2009, “Psychology and Economics: Evidence from the Field,” *Journal of Economic Perspectives*, 47, 315–372.
- DellaVigna, S., and U. Malmendier, 2006, “Paying Not to Go to the Gym,” *American Economic Review*, 96, 694–719.
- Dickson, P. R., and A. G. Sawyer, 1990, “The Price Knowledge and Search of Supermarket Shoppers,” *Journal of Marketing*, 54, 42–53.
- Erdem, T., S. Imai, and M. Keane, 2003, “Brand and Quantity Choice Dynamics under Price Uncertainty,” *Quantitative Marketing and Economics*, 1, 1–64.
- Hendel, I., and A. Nevo, 2006a, “Measuring the Implications of Sales and Consumer Inventory Behavior,” *Econometrica*, 74, 1637–1673.

- Hendel, I., and A. Nevo, 2006b, "Sales and Consumer Inventory," *RAND Journal of Economics*, 37, 543–561.
- Hendel, I., and A. Nevo, 2013, "Intertemporal Price Discrimination in Storable Goods Markets," *American Economic Review*, 103, 2722–2751.
- Hosken, D., and D. Reiffen, 2004, "Patterns of Retail Price Variation," *RAND Journal of Economics*, 35, 128–146.
- Hosken, D., and D. Reiffen, 2007, "Pricing Behavior of Multiproduct Retailers," *The B.E. Journal of Theoretical Economics (Topics)*, 7, Article 39.
- Hossain, T., and J. Morgan, 2006, "...Plus Shipping and Handling: Revenue (Non) Equivalence in Field Experiments on eBay," *Advances in Economic Analysis & Policy*, 6, Article 3, available at: <http://www.bepress.com/bejeap/advances/vol6/iss2/art3>.
- Howard, J. A., and J. N. Sheth, 1969, *Theory of Buyer Behaviour*, John Wiley & Sons.
- Hoyer, W. D., 1984, "An Examination of Consumer Decision Making for a Common Repeat Purchase Product," *Journal of Consumer Research*, 11, 822–829.
- Iyengar, R., and S. Gupta, 2009, "Nonlinear Pricing," in Vithala R. Rao (ed.), *Handbook of Pricing Research in Marketing*, Edward Elgar Publishing.
- Lacetera, N., D. G. Pope, and J. R. Syndor, 2012, "Heuristic Thinking and Limited Attention in the Car Market," *American Economic Review*, 102, 2206–2236.
- Lal, R., and C. Matutes, 1994, "Retail Pricing and Advertising Strategies," *Journal of Business*, 67, 345–370.
- Lennard, V., D. Mitchell, and P. McGoldrick, 2003, "Consumer Awareness, Understanding and Usage of Unit Pricing," *British Journal of Management*, 14, 173–87.
- Mackowiak, B., and M. Wiederholt, 2009, "Optimal Sticky Prices under Rational Inattention," *American Economic Review*, 99, 769–803.
- Maskin, E., and J. Riley, 1984, "Monopoly with Incomplete Information," *The RAND Journal of Economics*, 15, 171–196.
- McManus, B., 2007, "Nonlinear Pricing in an Oligopoly Market: The Case of Specialty Coffee," *RAND Journal of Economics*, 38, 512–532.

- Mondria, J., 2010, "Portfolio Choice, Attention Allocation, and Price Comovement," *Journal of Economic Theory*, 145, 1837–1864.
- Monroe, K. B., and A. Y. Lee, 1999, "Remembering Versus Knowing: Issues in Buyers' Processing of Price Information," *Journal of the Academy of Marketing Science*, 27, 207–225.
- Pesendorfer, M., 2002, "Retail Sales: A Study of Pricing Behavior in Supermarkets," *Journal of Business*, 75, 33–66.
- Piccione, M., and R. Spiegler, 2009, "Framing Competition," unpublished manuscript, London School of Economics.
- Reis, R., 2006, "Inattentive Consumers," *Journal of Monetary Economics*, 53, 1761–1800.
- Seiler, S., 2013, "The Impact of Search Costs on Consumer Behavior: A Dynamic Approach," *Quantitative Marketing and Economics*, 11, 155–203.
- Sims, C. A., 1998, "Stickiness," *Carnegie-Rochester Conference Series On Public Policy*, 49, 317–356.
- Sims, C. A., 2003, "Implications of Rational Inattention," *Journal of Monetary Economics*, 50, 665–690.
- Sobel, J., 1984, "The Timing of Sales," *The Review of Economic Studies*, 51, 353–368.
- Sprott, D., K. Manning, and A. Miyazaki, 2003, "Grocery Price Setting and Quantity Surcharges," *Journal of Marketing*, 67, 34–46.
- Stigler, G., 1961, "The Economics of Information," *Journal of Political Economy*, 69, 213–225.
- Stole, L. A., 2007, "Price Discrimination and Competition," in Mark Armstrong, and Robert Porter (ed.), *The Handbook of Industrial Organization*, Vol. 3 . pp. 2221–2299, North-Holland.
- Thomas, C., 2011, "Too Many Products: Decentralized Decision Making in Multinational Firms," *American Economic Journal: Microeconomics*, 3, 280–306.
- van Heerde, H. J., P. S. J. Leeflang, and D. R. Wittink, 2004, "Decomposing the Sales Promotion Bump with Store Data," *Marketing Science*, 23, 317–334.
- Varian, H. R., 1980, "A Model of Sales," *American Economic Review*, 70, 651–659.