

Hendel, Igal; Nevo, Aviv

Working Paper

Sales and consumer inventory

CSIO Working Paper, No. 0061

Provided in Cooperation with:

Department of Economics - Center for the Study of Industrial Organization (CSIO), Northwestern University

Suggested Citation: Hendel, Igal; Nevo, Aviv (2004) : Sales and consumer inventory, CSIO Working Paper, No. 0061, Northwestern University, Center for the Study of Industrial Organization (CSIO), Evanston, IL

This Version is available at:

<https://hdl.handle.net/10419/38644>

Standard-Nutzungsbedingungen:

Die Dokumente auf EconStor dürfen zu eigenen wissenschaftlichen Zwecken und zum Privatgebrauch gespeichert und kopiert werden.

Sie dürfen die Dokumente nicht für öffentliche oder kommerzielle Zwecke vervielfältigen, öffentlich ausstellen, öffentlich zugänglich machen, vertreiben oder anderweitig nutzen.

Sofern die Verfasser die Dokumente unter Open-Content-Lizenzen (insbesondere CC-Lizenzen) zur Verfügung gestellt haben sollten, gelten abweichend von diesen Nutzungsbedingungen die in der dort genannten Lizenz gewährten Nutzungsrechte.

Terms of use:

Documents in EconStor may be saved and copied for your personal and scholarly purposes.

You are not to copy documents for public or commercial purposes, to exhibit the documents publicly, to make them publicly available on the internet, or to distribute or otherwise use the documents in public.

If the documents have been made available under an Open Content Licence (especially Creative Commons Licences), you may exercise further usage rights as specified in the indicated licence.

THE CENTER FOR THE STUDY
OF INDUSTRIAL ORGANIZATION
AT NORTHWESTERN UNIVERSITY

Working Paper #0061

Sales and Consumer Inventory^{*}

By

Igal Hendel

University of Wisconsin, Madison and NBER

And

Aviv Nevo

University of California, Berkeley and NBER

June 2004

* We wish to thank David Bell for the data, Iain Cockburn, Ken Hendricks, Chad Jones, John Kennan, Rody Manuelli, Ariel Pakes, Rob Porter, Joel Waldfogel and several seminars audiences for comments and suggestions. We gratefully acknowledge the support from the NSF (SES-0093967 and SES-0213976) and the Food System Research Group. The second author wishes to thank the Center for the Study of Industrial Organization at Northwestern University, for hospitality and support.

ABSTRACT

Temporary price reductions (sales) are common for many goods and naturally result in a large increase in the quantity sold. We explore whether the data support the hypothesis that these increases are, at least partly, due to dynamic consumer behavior: at low prices consumers stockpile for future consumption. This effect, if present, has broad economic implications. We construct a dynamic model of consumer choice, use it to derive testable predictions and test these predictions using two years of scanner data on the purchasing behavior of a panel of households. The results support the existence of household stockpiling behavior and suggest that static demand estimates may overestimate price sensitivity.

1. Introduction

For many non-durable consumer products prices tend to be at a modal level with occasional short-lived price reductions, namely, sales. Naturally, during sales the quantity sold is higher. Quantity purchased may increase due to a consumption effect if consumption is price sensitive, and a stockpiling effect if dynamic considerations lead consumers to accumulate inventory for future consumption. For example, in our sample the quantity of laundry detergents sold is 4.7 times higher during sales than during non-sale periods (provided there was no sale the previous week). Instead if there was a sale in the previous week, then the quantity sold is only 2.0 times higher. This pattern suggests not only that demand increases during sales, but that demand accumulates between sales.² Demand accumulation has been documented by Pesendorfer (2002) using store level data of ketchup purchases (see also Blatteberg and Neslin, 1990). Our goal is to study whether stockpiling is behind the demand accumulation documented by Pesendorfer. Is the observed behavior consistent with dynamic forward looking consumer behavior? In order to address these questions we derive and test the implications of a consumer inventory model. In contrast to Pesendorfer, we explicitly model the dynamic problem of a consumer who can store the product (see a comparison of the models in Section 2), which we test using household level data.

There are several reasons to study and quantify consumers' stockpiling behavior. First, most of the work in industrial organization, from theoretical models to demand estimation, assume away demand dynamics. In contrast, the purchase of most products involves some sort of intertemporal substitutability. The substitutability may arise because the product is durable (cars or furniture), storable (groceries), or because of consumption is intertemporally substitutable (like a vacation or a golf game). Scanner data present an opportunity to study whether these potential dynamics are observed in household behavior. The pronounced price changes, observed in some of the products we study (i.e., laundry detergents, yogurt and soft-drinks), create incentives for consumers to stockpile at low prices, for future consumption. From our

²For differentiated products there is another potential effect: brand switching. We discuss later how to distinguish stockpiling from other responses to prices.

data we can proxy the potential gains from dynamic behavior. One such measure is given by comparing the actual amount paid by the household to what they would have paid (for the same bundle) if prices were drawn at random from the price distributions observed at the same locations over time. In our data the average household pays 12.7 percent less than if they were to buy the exact same bundle at the average price for each product.³ Some households save little, they are essentially drawing prices at random, while others save more (e.g., the 90th percentile saves 23 percent). Assuming savings in these 24 categories are representative, the total amount saved by the average household over two years in the stores we observe is 500 dollars (with 10th and 90th percentiles of 150 and 860 dollars, respectively). The observed price movements provide incentives for storage.

The second reason to look at stockpiling is to quantify the implications of the frequent price reductions (present in the typical scanner data) on demand estimation. In principle, the presence of frequent sales is a blessing for demand estimation, as they provide price variability needed to identify price sensitivities. However, when the good in question is storable, there is a distinction between the short run and long run reactions to a price change. Standard static demand estimation would capture short run reactions to prices, which reflect both the consumption and stockpiling effects. In contrast, for most demand applications (e.g., merger analysis or computation of welfare gains from introduction of new goods) we want to measure long run responses.

Third, consumer stockpiling has implications for how sales should be treated in the consumer price index. If consumers stockpile, then ignoring the fact that consumers can substitute over time will yield a bias similar to the bias generated by ignoring substitution between goods as relative prices change (Feenstra and Shapiro, 2001).

A final motivation to study stockpiling behavior, is to gain some understanding of the forces that determine sellers' incentives when products are storable. Although this paper does not address the question

³ This is for the 24 products we have in our data set. These 24 products account for 22 percent of the total grocery expenditure of the sampled households.

of the optimal seller behavior, our findings provide some basic numbers that may guide theoretical modeling.⁴

Assessing whether consumers stockpile in response to price movements would be straightforward if we observed consumers' inventories. For instance, we could test if after sales end-of-period inventories are higher. However, consumption and therefore inventory, is unobservable. We could assume a consumption rate which jointly with the observed purchases would enable us to infer inventories. While this approach might be reasonable for some products (those with no consumption effects), it would not help disentangle long run from short run effects for those products for which the distinction really matters.⁵

We take an alternative route. We propose a dynamic model of consumer choice and use it to derive implications about the variables we observe. For example, using household purchase data we test the link between current prices and duration to next purchase, instead of testing the (negative) relation between end-of-period inventories and current price.

We concentrate on those predictions of the model that stem from the stockpiling effect, but would not be expected under static behavior. In the model the consumer maximizes the discounted expected stream of utility by choosing in each period how much to buy and how much to consume. She faces uncertain future prices and in any period decides how much to purchase for inventory and current consumption. Optimal behavior is characterized by a trigger and target level of inventory which depend on current prices.

In order to test the model we use store and household-level data. The data were collected using scanning devices in nine supermarkets, belonging to five different chains, in two sub-markets of a large mid-west city. The store level data includes weekly prices, quantities, and promotional activities. The

⁴In ongoing work we study the behavior of a storable good monopolist. Most of the literature on sales is based on the Sobel (1984, 1991) model, which is a model of durable goods. Preliminary results show the main forces at play are quite different when the good is instead storable. The numbers in the current paper can shed some light on the degree and sources of consumer heterogeneity, the typical length of the inter-purchase period, and a rough estimate on the storage costs to help in the theoretical modeling.

⁵We report in Section 4 evidence that suggests that the consumption effect is important for some products.

household-level data set follows the purchases of about 1,000 households over two years. We know when each household visited a supermarket, how much was spent in each visit, which product was bought, where it was bought, how much was paid and whether a coupon was used.

Since the model deals with a single homogenous product purchased in a single store, while the data includes multiple varieties purchased in several stores, we need a practical way to link model and data. Under the maintained assumption that visits to the different stores are exogenous to the needs of the goods in question, the multiplicity of stores presents no problem. Each visit, regardless of the store, is just a draw from the price distribution prevailing at the frequented stores. The multiplicity of products is more delicate. It requires a definition of what a product is. We take a broad product definition (unless otherwise stated) treating whole categories as a single product. How close substitutes different brands (or UPCs) are, is an empirical matter beyond the scope of this paper. As we discuss later, a broad product definition seems natural for our descriptive purposes for several reasons. First, the observed purchasing behavior of each household defines, and narrows down, the product. Second, purchases of a product are likely to be affected by the duration from the purchase of an imperfect substitute as well. Finally, this treatment of the products is consistent with the model which abstracts from product differentiation, treating all goods as perfect substitutes. The cost of treating different product varieties as a single product is that it imposes duration dependence within categories, while their might not be such a link. Thus, we introduce measurement error in the definition of duration.

We test the implications of the model regarding both household and aggregate behavior. Our results support the model's predictions in the following ways. First, using the aggregate data, we find that duration since previous sale has a positive effect on the aggregate quantity purchased, both during sale and non-sale periods. Both these effects are predicted by the model since the longer the duration from the previous sale, on average, the lower the inventory each household currently has, making purchase more likely. Second, we find that indirect measures of storage costs are negatively correlated with households' tendency to buy on sale. Third, both for a given household over time and across households we find a

significant difference between sale and non-sale purchases, in both duration from previous purchase and duration to next purchase. The duration effects are a consequence of the dependence of the trigger and target inventory level on current prices. In order to take advantage of the low price, during a sale a household will buy at higher levels of current inventory. Furthermore, during a sale a household will buy more and therefore, on average, it will take more time till the next time the inventory crosses the threshold for purchase. Fourth, even though we do not observe the household inventory, by assuming constant consumption over time we construct a measure of implied inventory. We find that this measure of inventory is negatively correlated with the quantity purchased and with the probability of buying. Finally, we find that the pattern of sales and purchases during sales across different product categories is consistent with the variation in storage costs across these categories.

There are several potential stories on consumption patterns that explain why demand increases during sales. It would be virtually impossible to rule them all out (especially, given that consumption is unobserved). The main alternative hypothesis we consider is that consumers behave in a static fashion, buying more during sales, purely for consumption. Another alternative hypothesis to stockpiling, is that price sensitive consumers accumulate in the market until they find a sale (as in Sobel, 1984). Although some of the patterns in the data are consistent with Sobel-type models, others are not. In particular, household level behavior is inconsistent with that model (see next section).

The closest paper to ours is Boizot et. al. (2001). They present a dynamic inventory model which they test using consumer dairy data. The main difference between the papers is in the data. The key advantage of our data is detailed information about the product purchased (for example, brand and exact size). Such data is necessary to distinguish sales from substitution towards cheaper brands. Absent exact brand information it is impossible to distinguish if a consumer pays a lower than average price because she purchased a different, cheaper, brand or because she bought the usual brand on sale. This problem not only introduces measurement error but it also renders impossible to figure out the extent and depth of the sales faced by the households in their sample. While Boizot et. al. clearly point this out, due to data limitations,

they cannot overcome with this problem. Another advantage of our study is that our data enables a richer descriptive analysis of household heterogeneity in their shopping behavior. The models, although different in several ways (price process, demand uncertainty, their model assumes away consumption effects) deliver similar testable implications. We discuss later the key overlap with their findings.

There are several studies in the marketing literature that examine the effects of sales, or more generally the effects of promotions (see Blatteberg and Neslin, 1990, and references therein).⁶ In Section 4 we discuss how our results relate to this literature. Erdem et al (2003) and Hendel and Nevo (2002a) also look at sales from the inventory perspective. In contrast to this paper that is primarily descriptive, their starting point is a dynamic forward looking model, which they structurally estimate.

The rest of this paper is organized as follows. Next, we present a formal model of consumer inventory and use it to derive testable predictions. In Section 3 we present our data and display some preliminary analysis describing the three categories we focus on. Section 4 presents the results of the tests. We conclude by discussing how the findings relate to our motivation.

2. The Model

We present a simple inventory model, which we use to generate testable predictions about both observable household purchasing patterns and aggregate (store level) demand patterns. We depart from most of the literature on sales which is based on Sobel's model. For example, Pesendorfer's model is a Sobel-type one where consumers accumulate in the market until they buy.

Sobel's model is a good starting point for studying sales, but it does not capture the main feature of the goods in question: storability. In order to generate predictions testable with household level data, the model explicitly captures consumers' endogenous decision to return to the market. Predictions will differ at the aggregate level as well. At the aggregate level our model predicts that quantity sold is a function of the

⁶For example, see Ward and Davis (1978), Shoemaker (1979), Wilson, Newman and Hastak (1979), Blattberg, Eppen and Lieberman (1981), Neslin, Henderson and Quelch (1985), Gupta (1988), Chiang (1991), Grover and Srinivasan (1992) and Bell, Chiang and Padmanabhan (1999). Gonul and Srinivassan (1996) study the use of coupon in the context of a dynamic inventory model.

duration since the previous sale during both sale and non-sale periods (while Pesendorfer's predictions are only for sale periods). The model also delivers comparative statics across products, for which we find anecdotal support.

The model abstracts from important dimensions of the problem, like non-linear pricing and brand choice. Our goal here is to test the main implications of stockpiling in a simple setup. In Hendel and Nevo (2002a) we impose more structure in order to deal with the additional dimensions ignored here.

2.1 The Basic Setup

Household h obtains the following flow utility in period t

$$u(c_{ht}, v_{ht}) + \alpha m_{ht}$$

where c_{ht} is the quantity consumed, v_{ht} is a shock to utility and m_{ht} is the outside good consumption. The utility function is assumed increasing and concave. v_{ht} captures demand shocks unobserved by the researcher. For simplicity we assume the shock is additive in consumption, $u(c_{ht}, v_{ht}) = u(c_{ht} + v_{ht})$, affecting the marginal utility from consumption. Low realizations of v_{ht} increase the household's need, increasing demand and making it more inelastic.⁷ We also assume $\frac{\partial u(c_{ht} + v_{ht})}{\partial c} \geq \alpha p \forall p$ and $\forall v$, which is sufficient for positive consumption every period. This assumption has no impact on the predictions of the model, while it avoids having to deal with corner solutions.

Facing random prices, p_t , the consumer at each period has to decide how much to buy, denoted by x_{ht} , and how much to consume. Since the good is storable, quantity not consumed is kept as inventory for future consumption. We could assume consumption is exogenously determined, either at a fixed rate or randomly distributed (independently of prices). Both these alternative assumptions, which have been made by previous work, are nested within our framework. All the results below hold, indeed the proofs are simpler. However, it is important to allow consumption to vary in response to prices since this is the main alternative explanation for why consumers buy more during sales. We want to make sure that our

⁷Notice that v affects consumption but not the slope of the demand curve since it appears additively ($c^*(v,p)=u^{-1}(p)-v$) at each p . Hence, for any price, the elasticity increases in v .

results are not driven by assuming it away.

We assume the consumer visits stores at an exogenously given frequency, i.e., the timing of shopping is assumed to be determined by overall household needs (a bundle). Each of the products is assumed to be a minor component of the bundle, hence, need for these products does not generate a visit to the store.

After dropping the subscript h , to simplify notation, the consumer's problem can be represented as

$$\begin{aligned} V(I(0)) = \max_{\{p_t, x_t\}} \sum_{t=0}^{\infty} \delta^t E \left[u(c_t + v_t) - C(i_t) - \alpha p_t x_t \mid I(t) \right] \\ \text{s.t.} \quad 0 \leq i_t, \quad 0 \leq x_t \\ i_t = i_{t-1} + x_t - c_t \end{aligned} \quad (1)$$

where u is the marginal utility from income, δ is the discount factor, and $C(i_t)$ is the cost of storing inventory, with $C(0) = 0$, $C' > 0$ and $C'' > 0$.⁸

The information set at time t , $I(t)$, consists of the current inventory, i_{t-1} , current prices, and the current shock to utility from consumption, v_t .⁹ Consumers face two sources of uncertainty: utility shocks and future prices. We assume that shocks to utility, v_t , are independently distributed over time.

Prices are set according to a first-order Markov process. We assume there are two states, sale and non-sale. While it is easy to spot a modal (non-sale) price in the data, sales occur at a whole range of different prices. Thus, we assume there is a single non-sale price, \bar{p} ; but many sale prices. Conditional on a sale prices are drawn from the distribution $F(p)$, with support $[p_L, p_H]$ and $p_L \leq p_H < \bar{p}$. We assume that each sale gets a new draw from $F(p)$ independent of the current price. Let \bar{q} be the transition probability from non-sale to sale, and q be the transition probability from sale to sale. In the data the

⁸Notice we do not need to impose $c > 0$ since we assumed $\partial u / \partial c$ is such that there is always positive consumption.

⁹It is reasonable to assume that at the time of purchase the current utility shock has still not been fully realized. This will generate an additional incentive to accumulate inventory – to avoid the cost of a stock out. Since this is not our focus we ignore this effect, but it can easily be incorporated.

transition between sale prices is less frequent than the transition from the regular price to a sale price, hence we assume $1 > \bar{q} > q > 0$.

2.2 Consumer Behavior

In each period the consumer weights the costs of holding inventory against the (potential) benefits from buying at the current price instead of future expected prices. She will buy for storage only if the current price and her inventory are sufficiently low. At high prices the consumer might purchase for immediate consumption, depending on her inventory and the realization of the random shock to utility. The consumer's behavior is described by two thresholds S and s that respectively determine the target inventory, in case of a purchase, and a trigger inventory below which the consumer buys. We now formalize this result. Notice that the model is not an (s, S) one, in spite of some resemblances.

The solution of the consumer's inventory problem is characterized by the following Lagrangian

$$\max_{\{c_t, x_t, i_t\}} E \left[\sum_{t=0}^{\infty} \delta^t \{u(c_t + v_t) - C(i_t) - \alpha p_t x_t + \lambda_t (i_{t-1} + x_t - c_t - i_t) + \psi_t x_t + \mu_t i_t\} \mid I(t) \right] \quad (2)$$

where μ_t , ψ_t and λ_t are the Lagrange multipliers of the constraints in equation (1). From equation (2) we derive the first order conditions with respect to consumption,

$$u'(c_t + v_t) = \lambda_t, \quad (3)$$

purchase,

$$\alpha p_t = \lambda_t + \psi_t, \quad (4)$$

and inventory,

$$C'(i_t) + \lambda_t = \delta E(\lambda_{t+1} \mid i_t, p_t) + \mu_t. \quad (5)$$

Let $c^*(p_t, v_t)$ be the consumption level such $u'(c^*(p_t, v_t) + v_t) = \alpha p_t$ and let $S(p)$ be the inventory level such $C'(S(p)) + \alpha p_t = \delta E(\lambda_{t+1} \mid S(p), p_t)$. Manipulating the first order conditions we get the main result. Proofs are provided in the Appendix.

Proposition 1 In periods with purchases, $x_t > 0$, the target level of inventory, i_t , equals $S(p_t)$, a decreasing

function of p_t , independent of the other state variables i_{t-1} and v_t . Moreover, the inventory level that triggers a purchase is $s(p_t, v_t) = S(p_t) + c^*(p_t, v_t)$, which is decreasing in both arguments.

Optimal consumer behavior is characterized by a trigger s , and a target inventory S . The target, S , is a decreasing function of current price. On the other hand the trigger, s , which is the sum of the target and current consumption, depends on prices and the utility shock.¹⁰ It depends on the utility shock, since current consumption depends on the shock to utility.

Proposition 2 The quantity purchased, $x(i_{t-1}, p_t, v_t)$, declines in the three arguments.

2.3 Testable Implications

In this section we present the testable implications of the model. We focus on those predictions that help us distinguish the model from a static one, where all the reactions to sales stem purely from a consumption effect. For example, Proposition 2 predicts that purchases decline in prices. Such a relation is testable but is implied by static consumer behavior as well. Quantity purchased can increase during sales simply because consumption increases or because of stockpiling. Since we do not know the magnitude of the consumption effect, showing that quantity purchased increases during sales does not necessarily imply stockpiling. However, an immediate implication of Propositions 1 and 2, not predicted by the static model, is:

Implication I1: *Quantity purchased and the probability of purchase decline in the current inventory.*

¹⁰Optimal behavior is characterized by two functions that determine trigger and target. The two levels differ by current consumption. Hence, there is a single cut-off which determines both the target and (post-consumption) trigger inventory. The inaction region is not dictated by $s < S$, as it is in (s, S) models (see Arrow et. al. (1951)), rather by the movement in prices; which determine when a purchase is triggered.

The reason why quantity purchased and probability of purchase decline in inventories is that purchases are triggered by an inventory threshold and the end of period target inventory is independent of the initial inventory, hence the higher the initial inventory the lower the purchase needed to reach the target. Proof sketches can be found in the appendix.

Since we do not observe inventories, we cannot directly test this implication. We use two alternative strategies. In Section 4.3 we assume that consumption is fixed, which allows us to compute a proxy for the unobserved inventory. This is not an attractive assumption (and seems to be inconsistent with some of our findings that point to a consumption effect) since it assumes away the main alternative. Therefore, for most of the paper we resort to predictions on other aspects of consumer behavior, which indirectly testify on stockpiling. The following predictions follow this approach. They exploit the fact that s and S are decreasing functions of price. A decreasing $S(p)$ means a higher end of period inventory during sales. All else equal, this implies a longer duration until the next time the consumer's inventory crosses the threshold for purchase, s .

Implication I2: *Duration until the following purchase is longer during a sale.*

For deriving the next two implications we will make the additional assumption: $\bar{q} = q$. This condition is sufficient for the validity of implications I3 and I4 but not necessary. For highly persistent price process the implications may fail to hold. It is difficult to derive an analytic cut off on the transition probabilities that guarantee the validity of implications I3 and I4. To be on the safe side we assume no persistence, and discuss in the appendix why lack of persistence is sufficient but not necessary. Namely, for not too persistent price process we expect implications I3 and I4 to hold.

Implication I3: *Duration from previous purchase is shorter at sale periods.*

This prediction is a consequence of Proposition 1 which shows that the inventory that triggers a purchase, s , is lower at non-sale prices. Other things equal, crossing the lower trigger threshold implies that the previous purchase was further back in time. To see what would go wrong with this prediction if the price process is highly persistent, notice that it would be meaningless to say ‘other things equal’ while comparing sale and non-sale events, since the two events come from different histories.

The next implication is based on the same reasoning. If the previous purchase was on sale then, all else equal, end of period inventory would have been higher (since S declines in p). Then consumer’s inventory would be higher today, relative to her inventory if the previous purchase was not during a sale. Therefore, conditional on purchasing on non-sale today, it is more likely that the previous purchase was not during a sale. Intuitively, since s declines in p , a lower initial inventory generates non-sale purchases, while a lower initial inventory is more likely if the previous purchase was not on sale. As with the previous implication we assume $\bar{q} = q$.

Implication I4: *Non-sale purchases have a higher probability that the previous purchase was not on a sale, namely: $Pr(NS_{t-1}|S_t) < Pr(NS_{t-1}|NS_t)$, where S = sale purchase and NS = non-sale purchase.*

We now move to Implication I5 which holds in the aggregate. Store level demand increases with duration since the last sale, both on sale and no-sale periods. Which is a consequence of both quantity purchased and probability of purchase increasing in duration since the last purchase.

Implication I5: *Aggregate demand increases in the duration from the previous sale.*

A couple of caveats on implications I2 and I3. First, these implications are similar to those derived in Proposition 7 of Boizot et. al. (2001) coming from a slightly different model. Second, duration effects may be present also in a model without storage, but with duration dependence of sales. For example,

suppose there is no storage, consumers have a random reservation utility for a single unit of the good and there is positive duration dependence of sales. In that case, a non sale purchase is more likely to be followed by a sale period, which is likely to generate a purchase leading to an equivalent to implication I2, namely, that duration is shorter after a sale. However, notice that positive duration dependence of sales actually imply the opposite of Implications I3 and I4. If positive duration dependence of sales instead of stockpiling was the driving force behind implications I2 we should expect that the probability of a purchase in a non-sale period would be higher if the previous purchase was on sale, the opposite of implication I4. Moreover, we would expect that duration backwards would be shorter under a non-sale, contradicting implication I3.

3. The Data, Product Categories and Preliminary Analysis

3.1 Data

We use data collected by IRI using scanning devices in nine supermarkets, belonging to five different chains, in two separate sub-markets in a large mid-west city during the period of June 1991 to May 1993. The data has two components, one with store and the other with household-level information. The first contains prices, quantities sold and promotional activities, for each product (brand-size) in each store, in each week. The second component of the data set is a household-level sample of the purchases of 1,039 households over a period of 104 weeks. We know when a household visited a supermarket and how much they spent each visit. The data includes detailed information on 24 different product categories about which we know exactly which product each household bought, where it was bought, how much was paid, and whether or not a coupon was used. Table 1 presents some basic household demographics and their purchasing habits, with regards to the different stores they frequent.

3.2 The Product Categories

We focus here on three product categories available in the data: laundry detergents, soft-drinks and yogurt. We chose these three categories because of the simplicity of the choice set. A handful of brands

have a significant market share in each of these categories. In addition, the differences between the characteristics of these products allow us to examine cross-category implications.

Laundry detergents come in two main forms: liquid and powder. Liquid detergents account for 70 percent of the quantity sold. Unlike many other consumer goods there are a limited number of brands offered. The top eight (six) brands account for 75 percent of the liquid (powder) volume sold. The leading firms are Procter and Gamble (which produces Tide and Cheer) and Unilever (All, Wisk and Surf). Detergents can be stored for a long time before and after they are initially used. However, they probably require a designated area for storage.

The yogurt category is very concentrated at the brand level with the top two brands, Dannon and Yoplait, accounting for roughly 78 percent of the quantity sold. These brands are offered in many different varieties, which are differentiated along two main dimensions: fat contents and flavor (plain, vanilla and various fruit flavors, which can be blended or on the bottom). Unlike detergent, yogurt can be stored for a limited time only (several weeks). Nevertheless, for the relevant time horizon, which is a function of the frequency of visits to the store (at least once a week for most of the households in the sample), yogurt is still a storable product. Once the container is initially opened yogurt has a much shorter life.

The soft-drinks category combines several sub-categories: cola, flavored soda and club soda/mixer, all of which can be divided into regular and low calorie. The club soda/mixer sub-category is the smallest and for much of the analysis below will be excluded. The cola and low-calorie cola sub categories are dominated by Coke, Pepsi and Rite, which have a combined market share of roughly 95 percent. The flavored soda sub categories are much less concentrated with both more national brands and also a larger share of generic and private labels.

In all three categories, the prices for brand-size combinations have a clear pattern: they are steady at a “regular” price, which might vary by store, with occasional temporary reductions. While this pattern is easy to spot it is less easy to put in practice because modal prices also change. The first possibility we explore is to define the regular price as the modal price for each brand-size-store over the entire period, and a sale as any price below this level. This definition can miss changes in the regular price and therefore

mis-classify sale and non-sale periods. We check the robustness of the analysis to the definition of sales in two ways. First, we explore defining a sale as any price at least 5, 10, 25 or 50 percent below the regular price (defined as above). Second, we define the regular price as the max price in the previous three weeks, and a sale as any price at least 0, 5, 10, 25 or 50 percent below this price.

For the purpose of this section, which is purely descriptive, the exact definition is less important. Although for the most part all quantitative results reported below are robust to the different definitions, we must keep in mind that none of the definitions is perfect.

Using different definitions of a sale we display in Table 2 for each category the percent of weeks the product was on sale and the percent of the quantity sold during those weeks. The figures are averaged across all products at all stores. It is not surprising that for any definition of a sale the percent of quantity sold on sale is larger than the percent of weeks the sale price is available. Notice, in support of the theory, that for any definition of a sale, despite the fact that sales are less frequent for laundry detergents the quantity sold on sale is higher than that sold for yogurt. Since laundry detergent is more storable than yogurt this is consistent with stockpiling behavior. The main alternative explanation is that consumers simply increase their consumption in response to a price reduction. If anything it is more likely that the response of consumption to price is higher in yogurt, which makes this result even stronger. We return to this point in Section 4.

The products we examine come in different sizes. Consumers can stockpile by buying more units or by buying larger containers.¹¹ In Table 3 we display statistics for the major sizes in each product category. The sizes displayed account for 97 and 99 percent of the quantity sold of liquid detergent and yogurt, respectively. Soft-drinks are sold in either cans or various bottle sizes (16 oz. 1, 2 and 3 liter). For the purpose of this table we focus on cans, which can be sold as singles or bundled into 6, 12 or 24-unit packs.

¹¹ Indeed, size discounts are consistent with price-discrimination based on consumer storage costs. Size discounts are also consistent with varying costs by size, and therefore one could claim are not at all due to price discrimination (Lott and Roberts, 1991).

The first column in Table 3 displays the quantity discounts. Since not all sizes of all brands are sold in all stores reporting the average price per unit for each size could potentially be mis-leading. Instead we report the ratio of the size dummy variables to the constant, from a regression of the price per 16 ounce regressed on size, brand and store dummy variables. The results show quantity discounts in all three categories, but more so in detergents and soft-drinks.

The next three columns document the frequency of a sale, quantity sold on sale and average discount during a sale, for each size. We define a sale as any price at least 5 percent below the modal price. In all three categories there is an interaction between size and both the frequency of a sale and the quantity sold. The figures suggest that for both detergents and soft-drinks the larger sizes have more sales, and more quantity is sold on sale in the larger sizes. For yogurt, however, the pattern is reversed. There are more sales, and a larger fraction sold on sale, for the smaller sizes. In Section 4 we discuss this finding, which is consistent with storability.

Our data records two types of promotional activities: *feature* and *display*. The *feature* variable measures if the product was advertised by the retailer (e.g., in a retailer bulletin sent to consumers that week.) The *display* variable captures if the product was displayed differently than usual within the store that week.¹² Defining as a sale any price at least 5 percent below the modal price we find that conditional on being on sale, the probability of being featured (displayed) is 19 (18), 31 (7) and 30 (14) percent for detergents, yogurt and soft-drinks, respectively. While conditional on being featured (displayed) the probability of a sale is 88 (47), 87 (83) and 78 (53) percent, respectively. The probabilities of being featured/displayed conditional on a sale increase as we increase the percent cutoff that defines the sale.

3.3 Preliminary Analysis: the effect of duration from previous sales at the Store Level

In this section we study how quantity sold, conditional on price, increases with duration since the last sale. We look for patterns beyond those reported by Pesendorfer (2002), that may hint what is behind

¹²These variables both have several categories (for example, type of display: end, middle or front of aisle). We treat these variables as dummy variables.

demand accumulation. Pesendorfer (2002, Tables 9 and 10) finds that demand during sales increases with duration from previous sale, for the two items he studies: 32oz. Heinz and Hunt ketchup bottles. In his model, as in Sobel (1984), duration has no impact on non-sale periods. In accordance with his model, he only tests the effect of duration on quantity sold during sale periods. Using the store level data we find that the quantity purchased increases with duration from the previous sale, not only during sales periods but also in non-sale periods. Models in the spirit of Sobel (1984), explain the accumulation of demand as being a consequence of shoppers' waiting in the market for price reductions. Such models, predict that accumulation should occur during sale periods, but during non-sale periods demand should be independent of duration since the previous sale. The patterns of accumulation we find are consistent with an inventory model.

Table 4 presents the results of regressing the log of quantity sold, measured in 16 ounce units, on price, measured in dollars per 16 ounce, current promotional activity and duration since previous promotional activity.

We are still defining a sale as a price at least 5% below the modal price for each product/store and each UPC is a product. Hence, duration is product (UPC) and store specific. The implication of the latter is that two varieties of the same product, are treated as non-substitutes. We are not allowing the duration since last sale of Sprite to affect the demand of Coke. This narrow definition clearly introduces errors, but (i) there is no better definition of a product that could be implemented without fully estimating demand to reveal which products are indeed substitutes, and (ii) our current purpose is purely descriptive, even if imperfectly so, the regressions in Table 4 show the role of duration on quantity sold.

Different columns present the results for the different product categories, each of the categories divided into sale and non-sale periods.

The results in the first column (of each category), show the coefficient on duration since previous sale is positive and significant, for all three products.¹³ As already recorded in the literature, for other

¹³Duration is measured in weeks/100. In all the columns, even in cases where the coefficient on duration squared is significant, the implied marginal effect will be of the same sign as the linear term for the range of duration

products, demand accumulates between sales. The second column of each category shows the effect of duration on demand during non-sale periods. Quantity sold absent sales also increases in the duration from previous sale. In sum, duration has a positive effect on quantity sold in both sale and non-sale periods. Moreover, as we would expect the effect of duration is stronger during sales than non-sale periods, for all three categories. The larger coefficient on sale periods implies a larger impact of duration during sales in percentage terms. According to the model during non-sale periods consumers purchase exclusively for consumption. Thus, we expect duration to have a larger impact during sale periods. Note, that the duration from previous feature and duration from previous display have a negative effect. There are potential explanations to this result. For example, feature or display generate awareness of the product, and could impact demand (positively) even after they are over. In Hendel and Nevo (2002b) we discuss this issue further and show how this could be used to address a “puzzle” in the marketing literature.

The numbers in this section highlight the potential relevance of the inventory model. We next move to the testable predictions of the model at the consumer level.

4. Results: Household Level Analysis

In this section we use household data to (1) study which household characteristics determine proneness to buy on sale; (2) characterize the difference between sale and non-sale purchases, both across households and for a given household over time; and (3) examine the purchase decision conditional on being in a store and the decision of how much to buy conditional on a purchase. We conclude this section by comparing the results across product categories.

For these tests we will need to define a sale. As we discussed in the previous section there is no unequivocal definition of a sale. Nevertheless, for the sake of consistency all the tests below were conducted defining a sale as any price at least 5 percent below the modal price, for that UPC in that store over the two years. We checked the robustness of the results to this definition by looking at different

values mostly observed in the data. Therefore, we limit the discussion to the linear coefficient on duration.

definitions of the “regular” price (e.g., the max over 3 or 4 previous weeks) and by varying the cutoff for a sale (from 0 to 25 percent below the regular price). Qualitatively the results are robust to the different definitions we examined.

Another measurement issue to keep in mind is the definition of a product. In Tables 5 and 6, we treat each category as a single product. A broad product definition captures the fact that different brands are substitutes. The duration since last sale of a specific yogurt brand is likely to affect another brand’s sales. How close substitutes different brands (or UPCs) are, is an empirical matter beyond the scope of this paper.

Although imperfect, a broad product definition seems natural in this section. First, in what follows we use household level data. For each household, the relevant category might not include all products but only those UPCs the household actually consumed. The observed purchasing behavior of each household defines, and narrows down, the product. Second, purchases of any product are likely to be affected by the duration from the purchase of a substitute, even an imperfect one. More specifically, if consumers’ behavior can be characterized as a sequence of discrete choices, then all the brands in the choice set should be included in the definition of the products (see details in Hendel and Nevo, 2002a). Finally, this treatment of the products is consistent with the model presented in the previous section. The model abstracts from product differentiation, treating all goods as perfect substitutes. The empirics mimics the model by lumping all varieties consumed by each household.

What type of errors is our product definition likely to create? For example, suppose a household’s consumption of diet sodas is independent from the non-diet soda consumption. By treating diet/non-diet as a single product we will impose duration dependence across these categories, while their might not be such a link. Thus, we will introduce (classical) measurement error in the definition of duration and therefore the effects we find probably under estimate the true effects.

4.1 Household sales proneness

We now turn to the household data to study households’ propensity to purchases on sale. For the

1,039 households we regress the fraction of times the household bought on sale, in any of the three categories we study, during the sample period on various household characteristics.¹⁴ The results suggest that demographics have little explanatory power. We found that households without a male tend to buy more on sale, as do households with a female working less than 35 hours a week. Households with higher per person income are less likely to buy on sale, and so are households with a female with post high school education. These effects are just barely statistically significant, and some not significant, at standard significance levels. Overall observed demographics explain less than 3 percent of the cross household variation. Both the direction and lack of significance of these results is consistent with previous findings (Blattberg and Neslin, 1990).

While the frequency a household buys on sale is not strongly correlated with standard household demographics it is correlated with two other household characteristics, relevant from the theory perspective. First, households that live in market 1 tend to buy less on sale. This is true even after controlling for demographic variables including income, family size, work hours, age and race, as seen in column (i) of Table 5. Market 1 has smaller homes with less rooms and bedrooms, relative to the other market. Under the assumption that home size proxies for storage costs, this finding is consistent with stockpiling. One would expect lower storage costs to be positively related to the frequency of purchasing on sale. Second, though we know nothing about the households' house, we know the number of dogs they own. Column (ii) shows that the having a dog is positively, and significantly, correlated with purchasing on sale, even after we control for other household characteristics. At the same time owning a cat is not. Assuming that dog owners have larger homes, while cat owners do not, this further supports our theory. Dog ownership is not just a proxy for the market since the effects persist once we also include a market dummy variable, as seen in column (iii).¹⁵

¹⁴We also looked at the fraction of quantity purchased on sale. The results are essentially identical.

¹⁵The dog dummy variable might, alternatively, be a proxy for spare time, which may reflect a higher propensity to search. However, if the dog dummy variable was capturing propensity to search it would lose importance once we control for measures that proxy for the propensity to search (e.g., frequency of visits and number of stores). In fact dog ownership is uncorrelated with those proxies, moreover, the significance of the dog dummy variable is not affected by

In the last three columns we explore the correlation between frequency of purchasing on sale and other shopping characteristics. The results in column (iv) show that households who bought in more than one store tend to buy more on sale. This finding relates to Pesendorfer (2002) who reports that consumers who buy at low prices tend to shop in more stores. These effects also hold once we control for the characteristics used in columns (i) - (iii).

4.2 Sale vs. non-sale purchases

We now turn to the implications of the model presented in Section 2. These are predictions at the household level. We focus on those predictions that stem from the inventory model but that would not be expected under the alternatives. In Table 6, we compare, for each product category, the averages of several variables comparing sale and non-sale purchases. The first column, in each category, displays the average during non-sale purchases. The next three columns display the averages during sale purchases minus the average during non-sale purchases. The columns labeled *Total* display the difference between the mean of all sale purchases and the number in the first column. The *Total* difference averages purchases over time and across households. Hence, it reflects two different components: (i) a given household's sale purchases are likely to differ from non-sale ones (a within effect), and (ii) the profiles of households purchasing more frequently on sale is likely to differ from those not purchasing on sale (a between effect).¹⁶ Actually, the model has predictions regarding both the within and between effects and therefore in some cases also regarding the total effect. However, since each effect has a different interpretation to test the theory one has to separate these effects.¹⁷ In order to do so, the next column, labeled *Within*, displays the difference between each household's sale and non-sale purchases, averaged across households. Finally, the last column, labeled *Between*, displays the coefficient, from a cross household regression, of the mean

controlling for search proxies (see column (vi)).

¹⁶For a precise definition of within and between estimates see Greene (1997).

¹⁷Previous work was not always careful in separating these effects (see, for example, Neslin and Schneider Stone, 1996; van Heerde, Leeflang and Wittink, 2000; and references therein.)

of the variable in question for each household, on the proportion of purchases on sale (namely, the mean of the sale dummy across purchases of that household).

The results in the first row of Table 6 suggest that when purchasing on sale households buy more quantity (size times number of units). This is true both when comparing between households (households that make a larger fraction of their purchases during sales tend to buy more quantity) and within a household over time (when buying during a sale a household will tend to buy more), as predicted by Proposition 2. There is a difference across the three categories in how the additional quantity is bought. When buying laundry detergents households buy both more units and larger sizes. When buying yogurt households buy smaller units, but more of them. When buying soft-drinks households buy less units but of larger size (e.g., a single 24 pack instead of 4 six packs). This relates to Table 3, which highlights the interaction of sales and non-linear prices.

While the effect that households buy more on sale is consistent with our theory it is also consistent with the main alternative theory: when prices go down households buy and consume more of the product. If one is willing to assume that increased consumption is less relevant for some of our products, then increased quantity would testify of stockpiling. Instead, we turn to predictions that allow us to separate the two theories. Rows 4 and 5 of Table 6 show that duration to next purchase is larger for purchases on sale, while duration from previous purchase is shorter for sale purchases. These findings match the within household duration predictions of Implications I2 and I3. The alternative, of a pure increase in consumption, cannot explain these results. Furthermore, a simple comparison of the quantity and duration effects suggests that consumption goes up after sales. The consumption effect is particularly clear for sodas where the within increase in quantity purchased is 60% while the duration forward increases roughly 15%.

Notice that both implications I2 and I3 are within household implications. However, they have between households counterparts, namely, those households that consume more buy more on sale. Indeed all the between affects are positive and quite large in economic terms. Households more prone to buy on sale buy larger quantities and less frequently. Although these figures do not rule out alternative theories, they are consistent with stockpiling, and possibly generated by heterogeneity in storage costs. A possible

explanation for the between differences is that high demand households have a larger incentive to search, as they spend a higher budget on the item, and also a higher incentive to store for future consumption once they find a low price; making them more prone to buy on sale, buy larger quantities, store and hence buy less frequently. This hypothesis is further supported by the findings of section 4.1 where we found a positive correlation between the propensity to purchase on sale and the shopping frequency. This is relevant because it shows that sale-prone consumers buy less frequently not because they shop less frequently (an alternative explanation) but in spite of shopping more often, which is consistent with our theory.

The large between effects suggest substantial heterogeneity across households in how responsive they are to sales, and perhaps in how much they store. Such heterogeneity provides sellers incentives to hold sales as a way of discriminating across types with different abilities to store or responsiveness to sales.

The magnitude of the within effects seem small, especially compared to the magnitude of the between effects. This could be driven by several factors. First, the between effects imply heterogeneity in the sensitivity to sales. Therefore the within effects, which average responses across all households, are likely to understate the responses of those households aware of sales. Second, the definition of sale probably introduces measurement error and biases the effect towards zero.¹⁸ Third, there might be consumption and (potentially) stockpiling of several products generating further measurement problems. For example, a household might buy diet colas for the parents and a flavored soda for the kids. These could be two separate processes or there could be substitution between them. The results in Table 6 implicitly assume that these are perfect substitutes, since duration is measured to any purchase in the category.

Finally, we find that the probability the previous purchase was not on sale, given that current purchase was not on sale is higher (Implication I4). The reasoning behind the prediction is that since non-sale purchases have a lower inventory threshold (namely, inventories have to be low for the buyer not to be willing to wait for a sale) a non-sale purchase informs us that inventories are low which in turn means,

¹⁸For example, if we use a higher cutoff for the definition of a sale (i.e., a sale is any price at least 10, instead of 5, percent below the regular price) then the effects increase both in statistical significance and in economic magnitude.

other things equal, that the last purchase was not on sale. As before, the large between effects suggest a large cross-household heterogeneity in sales proneness, as those households buying today on sale are a lot more likely to have purchased last time on sale as well.

Implication I4 helps us distinguish whether the duration effects were caused by stockpiling or positive duration of sales. If positive duration dependence of sales instead of stockpiling was the driving force behind implication I2 we should expect that the probability of a purchase in a non-sale period would be higher if the previous purchase was on sale, the opposite of implication I4.

The findings regarding the within quantity and duration effects, relate to Boizot et. al. (2001). They test the dependence of duration and quantity sold on current and past prices using a marked failure time model. Like us they find significant effects. Our findings, in this part, differ from theirs in two ways. First, as they point out they cannot separate, due to data limitations, if the consumer is paying a low price due to a sale or because they bought a cheaper brand (or a larger size, which is cheaper per quantity). In contrast, we have the detailed data that contains information on the brand (and size) purchased necessary to define a sale.

A second difference with their results is that we are able to decompose the total difference, into between and within effects. The former measures cross household differences in behavior, while the latter tells us of consumers' responses to prices. They focus only on the latter. However, in order to characterize the incentives for sellers to hold sales, one has to understand and take account of both (within) consumers' responses to sales, as well as (between) consumer heterogeneity. Specially, since the above results seem to suggest that the between effects dominate.

It is worth mentioning that if the marginal utility from consuming a product depends negatively on previous period consumption, then sales could generate dynamics similar to those discussed in this section even absent stockpiling. The stock of past consumption would affect behavior (i.e., purchases) in a way similar to the physical stock help in storage under stockpiling. We have no way of separating the stories. However, we find the same patterns across products while it is reasonable to assume that non-time-separable preferences is not an issue for detergents. Thus, under the alternative story we should not expect

the dynamic effects for detergents. Moreover, inter-purchase durations are around a month, which is probably too long for last periods consumption to impact current marginal utility. For these two reasons stockpiling sounds like a more reasonable explanation.

4.3 Inventories, purchases and promotional activities

Up to now the results focused on testing the implications of our model assuming we cannot observe inventories. In this section we take an alternative approach. We assume constant consumption, compute a proxy for inventory and use it to study in Table 7: (i) the decision to purchase conditional on being in a store and (ii) the quantity purchased by a household conditional on a purchase, as a function of the price paid and promotional activities. The dependent variable in the first set of regressions is equal to one if the household purchased the product and zero if they visited the store but did not purchase. In the second set of regressions, the dependent variable is the quantity purchased, measured in 16 ounce units. The independent variables include the price and promotional variables for the brand-size purchased, household-specific dummy variables (as well as dummy variables for each store and for each, broadly-defined, product).¹⁹

We approximate the unobserved inventory in the following way. For each household we sum the total quantity purchased over the two year period. We divide this quantity by 104 weeks to get the average weekly consumption for each household. Assuming the initial inventory for each household is zero, we use the consumption variable and observed purchases to construct the inventory for each household at the beginning of each week. Since we include a household-specific dummy variable in the regressions assuming a zero initial inventory does not matter (as long as the inventory variable enters the regression linearly).

The results, presented in Table 7, are consistent with implication I1: the higher the inventory a household holds the lower the probability they buy and the less they buy, conditional on a purchase. To

¹⁹Notice there is no price associated with observation of the discrete choice regression (purchase/no purchase) when there was no purchase, that is why prices and promotions are not included.

get an idea of the magnitude of the coefficients consider the following. The average purchase of soft-drinks is roughly 7.25 units (116 oz.). Increasing the inventory by this amount, holding everything else constant, the probability of purchase conditional on being in a store decreases by 2.2 percentage points (relative to roughly 3 percent if inventory is zero). The effects for detergents and yogurt are 2.4 percentage points and 1.1, respectively. In the quantity regression the estimated coefficients suggest that each unit of (16 ounce) inventory reduces the quantity purchased by 0.72, 0.19 and 0.47 ounces, for the three categories respectively.

The effect of inventories on quantity purchased is statistically different than zero. The economic significance of these effects might seem small. It is hard to judge whether their magnitude is in line with the model's predictions (for example, Proposition 1) since these were derived assuming continuous quantities and linear prices. Simulations based on the preliminary results in Hendel and Nevo (2002a), where we model the discreteness of purchases and non-linear prices, suggest that the magnitude of the coefficients presented in Table 7 is consistent with stockpiling behavior that is economically significant.²⁰

4.4 A cross-category comparison

Unfortunately, none of the categories in our data is completely perishable. We were able to obtain data comparable to ours, but from a different city, on milk.²¹ The retail price exhibits a very different pattern than the one we find in the categories in our data set. Prices tend to change every 6-7 weeks and stay constant until the next change. There are essentially no temporary price reductions. Assuming that milk is not storable (and that the only reason for sales is to exploit consumer heterogeneity in storage costs), then there should be no sales for milk.

²⁰There are also two data related reasons why the estimated coefficient might be biased towards zero. First, the inventory variable was constructed under the assumption of constant consumption, which might be right on average but will yield classical measurement error and will bias the coefficient towards zero. As we noted in the previous section there is support in the data that consumption is not constant but reacts to prices. Second, we ignore differentiation in the definition of inventory. Once a quantity is bought we just add it to inventory. In reality, however, consumers might be using different brands for different tasks, which is also likely to bias the coefficient towards zero.

²¹We wish to thank Sachin Gupta, Tim Conley and Jean-Pierre Dube for providing us with these data.

Another cross-category comparison involves the difference between laundry detergents and yogurt. Since the average duration between supermarket visits is less than a week both these products are storable. However, there is a key difference between how one would store them. Unlike detergents, the storability of yogurt decreases once the container is opened. This suggests that for detergents we should see more sales for larger sizes and when consumers purchase on sale they buy larger units. For yogurt we should see the opposite: more sales for smaller sizes and purchase of smaller units on sale. Both these predictions hold and can be seen in Table 3 columns two and three and Table 6 second row.

Further evidence linking the relation between the easier-to-store size and sales is presented in the last column of Table 3, where we show the potential gains from stockpiling (defined in the Introduction) for the different sizes. Bigger savings are associated with the containers easier to store, namely larger sizes of detergents and soda, while small yogurt containers.

5. Conclusions and Extensions

In this paper we propose a model of consumer inventory holding. We use the model to derive several implications, which we take to the data. Our data consists of an aggregate detailed scanner data and a household-level data set. Using these data sets we find several pieces of evidence consistent with our model. (1) Aggregate demand increases as a function of duration from previous sale, and this effect differs between sale and non-sale periods. (2) Fraction of purchases on sale is higher in one market (the market that on average has larger houses) and if there is a dog in the house. Both of these could potentially be correlated with lower storage costs. (3) When buying on sale households tend to buy more quantity (either by buying more units or by buying larger sizes), buy earlier and postpone their next purchase. (4) Inventory constructed under the assumption of fixed consumption over time, is negatively correlated with quantity purchased and the probability of purchase. (5) The patterns of sales across different product categories is consistent with the variation in storage costs across these products.

Calculations based on our findings suggest that in the presence of stockpiling standard, static, demand estimation may be misleading. Static demand estimates, which neglect dynamics, may overestimate

own-price elasticities.

References

- Aguirregabiria, V. (1999) "The Dynamics of Markups and Inventories in Retailing Firms," *The Review of Economic Studies*, 66, 275-308.
- Arrow, K, T. Harris, J. Marschak (1951). "Optimal Inventory Policy," *Econometrica*, Vol. 19, No 3, pp. 250-272
- Bell D., J. Chiang and V. Padmanabhan (1999), "The Decomposition of Promotional Response: An Empirical Generalization," *Marketing Science*, 18, 504-26.
- Blattberg R., G. Eppen and J. Lieberman (1981), "A Theoretical and Empirical Evaluation of Price Deals in Consumer Nondurables," *Journal of Marketing*, 45, 116-129.
- Blattberg, R. and S. Neslin (1990), *Sales Promotions*, Prentice Hall.
- Blattberg R., R. Briesch and E. Fox (1995), "How Promotions Work," *Marketing Science*, 14, G122-132.
- Boizot, C., J.-M. Robin, M. Visser (2001), "The Demand for Food Products. An Analysis of Interpurchase Times and Purchased Quantities," *Economic Journal*, 111(470), April, 391-419.
- Chiang J. (1991), "A Simultaneous Approach to the Whether, What and How Much to Buy Questions," *Marketing Science*, 10, 297-315.
- Erdem, T., S. Imai and M. Keane (2003), "Consumer Price and Promotion Expectations: Capturing Consumer Brand and Quantity Choice Dynamics under Price Uncertainty," *Quantitative Marketing and Economics*, 1, 5-64.
- Feenstra, R. and M. Shapiro (2001), "High Frequency Substitution and the Measurement of Price Indexes," NBER Working Paper no. 8176.
- Gonul, F. and K. Srinivasan (1996) "Estimating the Impact of Consumer Expectations of Coupons on Purchase Behavior: A Dynamic Structural Model," *Marketing Science*, Vol. 15, No. 3, pp. 262-279.

- Grover R. and V. Srinivasan (1992), "Evaluating the Multiple Effects of Retail Promotions on Brand Loyal and Brand Switching Segments," *Journal of Marketing Research*, 29, 76-89.
- Gupta S. (1988), "Impact of Sales Promotions on When, What, and How much to Buy," *Journal of Marketing Research*, 24, 342-55.
- van Heerde H., Leeflang P. and D. Wittink (2000), "The Estimation of Pre- and Postpromotion Dips with Store-Level Scanner Data," *Journal of Marketing Science*, 37, 383-395.
- Hendel, I. and A. Nevo (2002a), "Measuring the Implications of Sales and Consumer Stockpiling Behavior". University of California at Berkeley, mimeo.
- Hendel, I. and A. Nevo (2002b), "The Post-Promotion Dip Puzzle: What Do the Data Have to Say?" University of California at Berkeley, mimeo.
- Hosken, D., D. Matsa, and D. Reiffen (2000) "How do Retailers Adjust Prices: Evidence from Store-Level Data," working paper.
- Lott Jr., J. and R. Roberts (1991), "A Guide to the Pitfalls of Identifying Price Discrimination," *Economic Inquiry* 29: 14-23.
- MacDonald, J. (2000), "Demand, Information, and Competition: Why Do Food Prices Fall At Seasonal Demand Peaks?," *Journal of Industrial Economics*, 48 (1), 27-45.
- Narasimhan, C. (1988), "Competitive Promotional Strategies," *Journal of Business*, 61 (4), 427-49.
- Neslin S., C. Henderson and J. Quelch (1985), "Consumer Promotions and the Acceleration of Product Purchases," *Marketing Science*, 4(2), 147-165.
- Neslin, S. and L. Schneider Stone (1996), "Consumer Inventory Sensitivity and the Postpromotion Dip," *Marketing Letters*, 7, 77-94.
- Pesendorfer, M. (2002), "Retail Sales. A Study of Pricing Behavior in Supermarkets," *Journal of Business*, 75, 33-66.
- Rao, R. (1991), "Pricing and Promotions in Asymmetric Duopolies," *Marketing Science*, 10(2), 131-44.
- Shoemaker R. (1979), "An Analysis of Consumer Reactions to Product Promotions," in *Educators' Conference Proceedings*, Chicago: American Marketing Association, 244-248.

Sobel, J. (1984), "The Timing of Sales," *Review of Economic Studies*, 51, 353-368.

Sobel, J. (1991), "Durable Goods Monopoly with Entry of New Consumers," *Econometrica*, 59(5), 1455-85.

Ward, R. and J. Davis (1978), "A Pooled Cross-Section Time Series Model of Coupon Promotions," *American Journal of Agricultural Economics*, 60, 393-401.

Wilson D., L. Newman and M. Hastak (1979), "On the Validity of Research Methods in Consumer Dealing Activity: An Analysis of Timing Issues," in *Educators' Conference Proceedings*, Chicago: American Marketing Association, 41-46.

Appendix

Proof of Proposition 1: If $x_t > 0$ then $\psi_t = 0$. If $i_t = 0$, there is nothing to show, simply $S(p_t) = 0$. In the complementary case, $i_t > 0$, we know $\mu_t = 0$. Using equation (4) and $\mu_t = \psi_t = 0$, equation (5) becomes: $C'(i_t) + \alpha p_t = \delta E(\lambda_{t+1} | i_t, p_t)$, which shows the end-of period inventory, i_t , is independent of the states variables i_{t-1} and v_t . Furthermore, notice that conditional on a sale, the specific realization of the current price conveys no information about future prices. Thus, conditional on a sale the right hand side is independent of prices, while non-increasing in i_t (that is, if we start the next period with a higher inventory the marginal utility from consumption must be weakly lower). Hence, since $C'' < 0$ the end of period inventory, i_t , declines in price. This shows that $S(p)$ is a declining function for all sale prices. To finish the argument we have to make sure $S(\bar{p}) \leq S(p_H)$. The inequality trivially holds since the consumer would never buy for storage at a non-sale price. By not buying she saves the storage cost, moreover, the price in the future cannot be higher. That is, $S(\bar{p}) = 0$, which completes the argument that $S(p)$ is a decreasing function.

To show that the inventory level that triggers a purchase is $S(p_t) + c^*(p_t, v_t)$, assume first that the consumer is willing to buy when she has an initial inventory $i_{t-1} > S(p_t) + c^*(p_t, v_t)$. In such a case, $i_t > S(p_t)$, which violates equation (5) since it would hold with equality for $i_t = S(p_t)$, but the left-hand side is bigger and the right-hand side smaller for $i_t > S(p_t)$. Now suppose the consumer does not want to purchase when $i_{t-1} < S(p_t) + c^*(p_t, v_t)$. Since $\psi_t \geq 0$ by equation (3) $c_t \geq c^*(p_t, v_t)$. Hence, $i_t < S(p_t)$, which implies equation (5) cannot hold. By definition, it holds for $S(p_t)$, but for $i_t < S(p_t)$ the left-hand side is lower than the right-hand side. We conclude that the inventory $i_{t-1} = S(p_t) + c^*(p_t, v_t)$ triggers purchases. #

Proof of Proposition 2: There are two cases to consider. Case 1: $x_t > 0$ and $i_t = 0$. In this case purchases equal consumption minus initial inventories: $x(i_{t-1}, p_t, v_t) = c(i_{t-1}, p_t, v_t) - i_{t-1}$. Since $x_t > 0$ we can combine equations (3) and (4) to get $u'(c_t + v_t) = \alpha p_t$, which implies that $c(i_{t-1}, p_t, v_t)$ declines in v_t and p_t , and is independent of i_{t-1} . Thus, $x(i_{t-1}, p_t, v_t)$ declines in v_t , p_t , and i_{t-1} . Case 2: $x_t > 0$ and $i_t > 0$. From Proposition 1 we know $x(i_{t-1}, p_t, v_t) = S(p_t) + c(i_{t-1}, p_t, v_t) - i_{t-1}$. The result follows from Case 1 and

Proposition 1, which showed $S(p)$ declines in p . #

Sketch of Proofs of the Implications:

I1: A purchase is triggered by $I_{t-1} < s(p, v)$. The lower the initial inventory the larger the range of v 's that generate a purchase. Hence, for any distribution of v the probability of purchase declines in I_{t-1} . The quantity purchased is given by $S(p) + c(p, v) - I_{t-1}$, thus quantity purchased also declines in the initial inventory.

I2: The longer duration after a sale is an immediate consequence of $S(\bar{p}) = 0$, namely that in our model non-sale purchases are only for consumption. Thus, duration after a non-sale purchase is one. In contrast after a sale purchase households consume from storage.²²

The result is also valid in a richer model in which duration after a non-sale purchase is not necessarily one. For instance, with indivisibilities, fixed costs of purchase or several non-sale prices, the duration after a non-sale purchase need not be one. The same forces will be at play, leading to a larger post sale duration due to $S(p)$ being a declining function of price. Namely, a larger end of the period inventory.

I3: If prices are *i.i.d.* ($\bar{q} = q$) it is immediate to show that duration backwards is shorter during sales by virtue of the inventory cut off that generates (triggers) a purchase being a decreasing function of price. For non *i.i.d.* prices a lower cut off is not enough because a sale and non-sale purchases are in principle preceded by different histories, which may impact the duration since the previous purchase. In other words, there are two pieces of information on today's purchase having being on sale. First, that inventories are high, relative to those if the purchase was not on sale. This first piece of information pushes towards a shorter duration backwards. However, there is a second piece of information contained on today being a

²²There are parameters of the model, e.g. high storage costs, for which the consumer purchases with probability one after a sale also. However, they are of little interest as they neutralize stockpiling.

sale period, in case there is persistence in the price process. This effect may go in the opposite direction of the previous one, depending on the type of persistence. If the persistence is very high, it may overturn the impact of the former, while it vanishes as we approach $\bar{q} = q$, namely, as the persistence disappears. In our application in which both \bar{q} and q are small (i.e., sales are not high probability events) we do not expect the second effect to overcome the first one. However, it is an empirical matter whether the price process exhibits sufficient persistence to overturn I3.

I4 Notice that a non-sale purchase signals that inventories are low (at least lower than consumption) which make more likely that the previous sale did not generate storage. Which in turn makes more likely that the previous purchase was not on sale. As in the previous implication the result is immediate for the $\bar{q} = q$ case, but not necessarily valid for any transition matrix. The proximity of \bar{q} and q (or low persistence) is needed for the validity of this prediction.

I5 First consider a non-sale period, t , and a consumer whose last sale purchase was in period $t-j$. Since purchases in non-sale periods are only for consumption, the consumer will buy in period t only if $I_{t-1} < c^*(v, p) = u'^{-1}(p) - v$, namely if the initial inventory is lower than intended consumption at current price and shock. Notice that as j grows, inventory declines (both because at non-sale price the consumer does not buy for storage, and because consumption is positive every period). Notice that the magnitude of a non-sale purchase is determined by the shock, and the leftover inventories. Hence, expected purchase conditional on having purchased at $t-j$ increases in j .²³

Now consider a sale period, t . The consumer will buy only if $I_{t-1} < s(v, p)$. The consumer has been consuming out of inventories since the last purchase on sale, at $t-j$. Namely, I_{t-1} declines in j . Moreover, since the target inventory, $S(p)$ is independent of I_{t-1} we can conclude that expected purchase conditional

²³Once the inventory is exhausted the consumer buys every period, so after some point duration has no further impact on consumers' demand.

on having purchased at $t-j$ increases in j . (As above, once the inventory has been exhausted duration has no further impact on demand).

Aggregating over consumers we get, I5, that demand accumulates during both sales and non-sale periods.

Table 1
Summary Statistics of Household-level Data

	mean	median	std	min	max
Demographics					
income (000's)	35.4	30.0	21.2	<10	>75
size of household	2.6	2.0	1.4	1	6
live in suburb	0.53	–	–	0	1
Store Visits					
number of stores visited over the 2 years	2.38	2	1.02	1	5
store HHI	0.77	0.82	0.21	0.27	1.00

Each household is an observation. store HHI is the sum of the square of the expenditure share spent in each store by each household.

Table 2
Percent of Weeks on Sale and Quantity Sold on Sale,
by Category for Different Definitions of Sale

	Laundry Detergents		Yogurt		Soft-drinks	
	weeks on sale	quantity sold	weeks on sale	quantity sold	weeks on sale	quantity sold
regular price equals modal price and a sale is any price less than:						
< regular price	18.6	39.0	22.8	33.2	30.9	63.2
< .95*regular price	12.6	32.3	16.9	25.2	28.0	60.9
< .9*regular price	7.5	26.9	13.0	20.5	23.1	54.8
< .75*regular price	1.8	14.9	4.4	10.4	12.0	36.6
< .5*regular price	0.04	1.4	0.4	1.8	2.0	4.3
regular price equals max in previous 3 periods and a sale is any price less than:						
< regular price	12.9	33.8	16.2	26.2	29.9	61.1
< .95*regular price	8.9	28.6	13.4	21.5	26.3	58.5
< .9*regular price	5.9	24.8	10.0	17.4	22.4	54.5
< .75*regular price	1.7	13.9	4.0	9.7	11.5	36.1
< .5*regular price	0.05	1.4	0.5	1.9	1.8	3.8

Table 3
Quantity Discounts and Sales

	price/ discount (\$ / %)	quantity sold on sale (%)	weeks on sale (%)	average sale discount (%)	quantity share (%)	saving (%)
Detergents						
32 oz.	1.08	2.6	2.0	11.0	1.6	4.3
64 oz.	18.1	27.6	11.5	15.7	30.9	1.3
96 oz.	22.5	16.3	7.6	14.4	7.8	10.0
128 oz.	22.8	45.6	16.6	18.1	54.7	18.6
256 oz.	29.0	20.0	9.3	11.8	1.6	–
Yogurt						
6 oz.	1.39	37.8	23.6	19.7	27.4	13.7
6*4.4 oz.	7.8	19.4	15.2	18.5	12.4	8.9
8 oz.	9.3	25.3	14.4	21.9	40.4	7.2
16 oz.	9.9	1.1	1.8	16.6	5.7	1.3
32 oz	28.3	15.9	10.8	13.0	12.9	3.0
Soft-drinks						
1 can	1.07	24.3	19.4	21.9	6.8	6.3
6 cans	2.3	59.5	34.3	35.4	16.8	21.8
12 cans	14.7	72.8	43.9	22.0	21.8	17.2
24 cans	34.4	78.3	41.7	20.8	54.5	17.6

All cells are based on data from all brands in all stores. The column labeled *price/discount* presents the price per 16 oz. for the smallest size and the percent quantity discount (per unit) for the larger sizes, after correcting for differences across stores and brands (see text for details). The columns labeled *quantity sold on sale*, *weeks on sale* and *average sale discount* present, respectively, the percent quantity sold on sale, percent of weeks a sale was offered and average percent discount during a sale, for each size. A sale is defined as any price at least 5 percent below the modal. The column labeled *quantity share* is the share of the total quantity (measured in ounces) sold in each size. The column labeled *savings* is the average percent increase in the amount consumers would pay if instead of the actual price they paid the average price for each product they bought.

Table 4
Demand as a Function of Duration from Previous Promotional Activity

Variable	Detergents		Yogurt		Soft-Drinks	
	sale	non-sale	sale	non-sale	sale	non-sale
log(price per 16 oz)	-2.42 (0.06)	-2.40 (0.02)	-1.46 (0.05)	-1.83 (0.03)	-2.59 (0.05)	-1.50 (0.05)
duration from previous sale	1.30 (0.58)	0.67 (0.13)	1.62 (0.80)	0.33 (0.24)	2.99 (0.70)	1.72 (0.28)
(duration from previous sale) ²	-1.90 (1.89)	-1.44 (0.26)	6.50 (3.59)	-2.03 (1.05)	-4.13 (1.71)	-3.18 (0.40)
feature	0.44 (0.03)	0.56 (0.07)	0.31 (0.02)	0.66 (0.04)	-0.01 (0.02)	0.56 (0.04)
display	1.12 (0.03)	1.19 (0.02)	0.67 (0.03)	0.96 (0.06)	1.48 (0.02)	1.32 (0.04)
duration from previous feature	-0.30 (0.23)	-0.84 (0.11)	-3.01 (0.43)	-1.46 (0.21)	0.55 (0.19)	-0.50 (0.17)
(duration from previous feature) ²	-0.09 (0.27)	1.41 (0.14)	8.48 (1.50)	5.51 (0.72)	-0.35 (0.24)	-0.69 (0.23)
duration from previous display	-1.21 (0.19)	-0.37 (0.08)	-0.68 (0.19)	-0.29 (0.09)	-2.01 (0.21)	-1.39 (0.15)
(duration from previous display) ²	1.05 (0.23)	0.04 (0.12)	1.74 (0.35)	0.36 (0.19)	1.89 (0.29)	0.90 (0.22)
N =	6,681	35,314	9,297	41,226	14,889	22,135

The dependent variable in all regressions is the natural logarithm of quantity purchased (measured in 16 ounce units). Each observation is a brand-size combination in a particular store. Duration from previous sale/feature/display is measured as number of weeks, divided by 100, from previous sale/feature/display for that brand in that store for any size. All regressions include brand and store dummy variables. The regressions in the soft-drinks category are for the sub-sample of cans and include a dummy variables for high demand holiday weeks (July 4, labor day, Thanksgiving and Christmas).

Table 5
Correlation Between Households Fraction of Purchases on Sale
and Household Characteristics

Variable	(i)	(ii)	(iii)	(iv)	(v)	(vi)
constant	0.50 (0.02)	0.50 (0.02)	0.49 (0.02)	0.39 (0.02)	0.51 (0.02)	0.44 (0.03)
market 1	-0.05 (0.01)	-	-0.05 (0.01)	-	-	-0.04 (0.01)
dog dummy variable	-	0.04 (0.01)	0.04 (0.01)	-	-	0.04 (0.01)
cat dummy variable	-	-0.001 (0.016)	0.005 (0.016)	-	-	0.005 (0.016)
# of stores	-	-	-	0.033 (0.006)	-	0.027 (0.006)
avg days b/ shopping	-	-	-	-	-0.008 (0.002)	-0.005 (0.002)
R-squared	0.045	0.037	0.051	0.059	0.042	0.080

The dependent variable is the fraction of purchases made during a sale averaged across the three categories: laundry detergents, yogurt, and soft-drinks. A sale is defined as a price at least 5 percent below the modal price. There are 1039 observations, where each household is an observation. All regressions also include per person HH income and dummy variables for a male head of HH, female works less than 35 hours and if she works more than 35 hours (excluded category is retired/unemployed), female post high school education and if head of HH is Latino. See text for discussion of the effect of these variables.

Table 6
Differences in Purchasing Patterns Between Sale and Non-Sale Purchases

Variable:	Laundry Detergents				Yogurt				Soft-drinks			
	Avg during non-sale	Difference during sale			Avg during non-sale	Difference during sale			Avg during non-sale	Difference during sale		
		Total	Within	Between		Total	Within	Between		Total	Within	Between
		households				households				households		
Quantity (16 oz.)	4.79 (0.04)	1.55 (0.07)	1.14 (0.07)	2.22 (0.27)	1.60 (0.01)	0.16 (0.02)	0.20 (0.02)	0.22 (0.08)	5.00 (0.26)	5.04 (0.31)	3.01 (0.34)	6.44 (0.61)
Units	1.07 (0.01)	0.09 (0.01)	0.08 (0.01)	0.12 (0.03)	2.63 (0.03)	0.99 (0.05)	0.80 (0.04)	1.24 (0.16)	4.18 (0.17)	-2.34 (0.20)	-1.75 (0.26)	-1.70 (0.29)
Size (16 oz.)	4.50 (0.03)	0.91 (0.05)	0.63 (0.05)	1.28 (0.20)	0.80 (0.01)	-0.19 (0.01)	-0.11 (0.01)	-0.23 (0.04)	2.82 (0.13)	4.31 (0.15)	2.73 (0.16)	5.05 (0.27)
Days from previous	44.38 (0.68)	6.70 (1.12)	-2.01 (1.03)	29.85 (8.11)	27.35 (0.59)	6.25 (1.10)	-1.27 (1.03)	6.87 (8.85)	24.71 (2.30)	8.85 (2.75)	-2.47 (2.07)	23.64 (7.66)
Days to next	43.75 (0.67)	8.56 (1.14)	1.95 (1.04)	28.91 (8.46)	26.08 (0.59)	9.87 (1.09)	2.78 (1.03)	21.64 (8.53)	21.49 (2.31)	12.89 (2.77)	2.50 (1.99)	29.74 (8.00)
Previous purchase not on sale	0.75 (0.01)	-0.29 (0.01)	-0.05 (0.01)	-0.77 (0.02)	0.78 (0.01)	-0.31 (0.01)	-0.13 (0.01)	-0.66 (0.03)	0.53 (0.02)	-0.26 (0.02)	-0.07 (0.03)	-0.36 (0.04)

A sale is defined as any price at least 5 percent below the modal price, of a UPC in a store over the observed period. The column labeled *Within households* controls for a household fixed effect, while the column labeled *Between households* is the regression of household means. Standard errors are provided in parentheses.

Table 7
Purchase Conditional on Store Visit and Quantity Purchased Conditional on Purchase by Household as a Function of Price and Promotional Activities

Dep variable: variable	Laundry Detergents		Yogurt		Soft-drinks		
	= 1 if purchase	quantity	= 1 if purchase	quantity	= 1 if purchase	quantity	
constant	0.08 (0.001)	8.98 (0.64)	0.06 (0.0005)	0.94 (0.44)	0.03 (0.0004)	4.96 (1.14)	
inventory/100	-0.43 (0.01)	-4.45 (0.27)	-0.63 (0.0002)	-1.16 (0.16)	-0.30 (0.007)	-2.91 (0.43)	
price	-	-3.79 (0.15)	-	-0.27 (0.08)	-	-5.09 (0.55)	
price*sale	-	-1.53 (0.15)	-	0.70 (0.09)	-	-5.83 (0.78)	
sale	-	1.39 (0.16)	-	-0.08 (0.12)	-	3.22 (0.62)	
feature	-	0.14 (0.09)	-	0.11 (0.03)	-	-0.76 (0.16)	
display	-	0.18 (0.08)	-	0.14 (0.04)	-	0.38 (0.15)	
	N =	149,802	12,731	149,802	10,457	149,802	4,768

All results are from linear regressions. The dependent variable in the regressions in columns 1, 3, and 5 is equal to one if the HH bought and zero if visited the store and did not buy. In all other regressions the dependent variable is the quantity purchased (measured in 16 oz units), conditional on purchasing a strictly positive quantity. All regressions also include household, product and store dummy variables. Prices (\$/16 oz) and promotional variables are for the product purchased. A sale is defined as any price at least 5 percent below the modal price. The sample for soft-drinks includes only purchases of cans of low calorie colas. Standard errors are provided in parentheses.