

NBER WORKING PAPER SERIES

SALES AND CONSUMER INVENTORY

Igal Hendel
Aviv Nevo

Working Paper 9048
<http://www.nber.org/papers/w9048>

NATIONAL BUREAU OF ECONOMIC RESEARCH
1050 Massachusetts Avenue
Cambridge, MA 02138
July 2002

We wish to thank David Bell for the data, Iain Cockburn, Ken Hendricks, Chad Jones, John Kennan, Ariel Pakes, Rob Porter, Joel Waldfogel and seminar participants in several seminars for comments and suggestions. The second author wishes to thank the Center for the Study of Industrial Organization at Northwestern University, for hospitality and support, and gratefully acknowledges support from the NSF (SES-0093967). The views expressed herein are those of the authors and not necessarily those of the National Bureau of Economic Research.

© 2002 by Igal Hendel and Aviv Nevo. All rights reserved. Short sections of text, not to exceed two paragraphs, may be quoted without explicit permission provided that full credit, including © notice, is given to the source.

Sales and Consumer Inventory
Igal Hendel and Aviv Nevo
NBER Working Paper No. 9048
July 2002
JEL No. L0, L4, D1, D4

ABSTRACT

Temporary price reductions (sales) are common for many goods and naturally result in 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, renders standard static demand estimates misleading, which 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, which neglect dynamics, may overestimate price sensitiveness by up to a factor of 2 to 6.

Igal Hendel
Department of Economics
Social Science Building
University of Wisconsin
1180 Observatory Drive
Madison, WI 53706-1393
and NBER
igal@ssc.wisc.edu

Aviv Nevo
Department of Economics
University of California
549 Evans Hall #3880
Berkeley, CA 94720-3880
and NBER
nevo@econ.berkeley.edu

1. Introduction

For many non-durable consumer products prices tend to be at a modal level with occasional short-lived price reductions, namely, sales. Unsurprisingly, during sales the quantity sold is higher. Price reductions may have two effects on quantity bought: first, a consumption effect if consumption is price sensitive, second, 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. These effects, which have been documented in the economics and marketing literature (see Section 1.1), are consistent with stockpiling. We want to analyze what is behind these patterns of demand. Do consumers stockpile? Does stockpiling explain the observed demand accumulation? Is the observed behavior consistent with dynamic forward looking behavior? In order to address these questions we derive and test the implications of a consumer inventory model.

There are several reasons to study and quantify consumers' stockpiling behavior. First, suppose the data available for demand estimation presents frequent price reductions (as is the case with scanner data). 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.

Second, detailed data, such as the household-level sample we describe below, present an opportunity to study whether a dynamic model of forward looking agents fits household behavior. The pronounced price changes, observed in some of the products we study, create incentives for

consumers to stockpile. Our analysis will focus on grocery products, in particular, laundry detergents, yogurt and soft-drinks. From our data we can compute 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 if the price was drawn at random from the distribution of prices for the same product at the same location over time. This is a lower bound on the potential gross gains from optimizing behavior. 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.² Some households save little, i.e., they are essentially drawing prices at random, while others save more (the 90th percentile save 23 percent). Assuming savings in these 24 categories represent saving in groceries in general, 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). Hence, the price movements provide incentives for storage and dynamic behavior.

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 understand sellers' incentives when products are storable. Although this paper does not answer this question, our findings provide some guidance on how to model the problem.³

If we observed the consumers' inventories then determining whether consumers stockpile in response to price movements would be straightforward. For instance, we could test if after a sale end-of-period inventories are higher. However, consumption and therefore inventory, is unobservable. We could make an assumption about consumption, for example, that it is constant.

² This is for the 24 products we have in our data set. For the households we observe these products account for 22 percent of their total grocery expenditure.

³In 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.

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.⁴

We, therefore, 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, as we show below the model predicts that when purchasing on sale the duration to next purchase should increase (relative to non-sale purchases by the same household). Furthermore, if there is a consumption effect this increase should be less than the increase in quantity purchased. Therefore, 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. This is just one example of how the model helps us circumvent the problem of not observing inventories, by providing implications about behavior we do observe.

We concentrate on those predictions of the model that stem exclusively from the stockpiling effect, but would not be expected under static behavior (i.e., when only the consumption effect is present). In our 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 follows a (conditional) *s-S* type behavior: if inventory is low enough the consumer buys and fills her inventory to a target level.

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 different chains, in two sub-markets of a large mid-west city. The store level data includes weekly prices, quantities, and promotional activities. The 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

⁴Indeed we will see in Section 4 that the consumption effect is important for some products.

used.

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 levels (s and S respectively) 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, i.e., s is higher. Furthermore, during a sale a household will buy more and therefore, on average, it will take more time till the next time inventory crosses the threshold for purchase. Fourth, even though we do not observe the household inventory, by assuming constant consumption over time we can 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. All these findings are consistent with the predictions of the model.

The findings suggest that in the presence of stockpiling standard, static, demand estimation may be misleading. The full estimation of short vs long run demand elasticities is beyond the scope of this paper (see Hendel and Nevo, 2002). However, the results can give an approximate idea of the relevant order of magnitude of the impact of stockpiling on demand estimates. We find that static demand estimates that neglect dynamics may overestimate price sensitiveness by up to a factor of 2 to 6.

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). Our goal is to study the implications of stockpiling. The main alternative hypothesis we consider is that consumers behave in a static fashion, buying more during sales, purely for consumption reasons. Another hypothesis, alternative to stockpiling, that we consider 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's model, others are not. In particular, household level behavior is inconsistent with that model (see next section).

The rest of this paper is organized as follows. We next survey the relevant literature and explain how it relates to our paper. In Section 2 we describe our data and display some preliminary analysis describing the three categories we focus on. Next, we present a formal model of consumer inventory and use it to derive testable predictions. Section 4 presents the results of the tests. We conclude by discussing how the findings relate to our motivation.

1.1 *Literature Review*

Pesendorfer (2002) studies sales of ketchup. He proposes a model in which a fixed number of consumers appears every period. These consumers differ in the willingness to pay (as in the Sobel (1984)) and in their store loyalty. He shows that in equilibrium the decision to hold a sale is a function of the duration since the last sale. His empirical analysis shows that both the probability of holding a sale and the aggregate quantity sold (during a sale) are a function of the duration since the last sale. The latter is evidence of demand accumulation, which he interprets as support for a Sobel-type model where consumers accumulate in the market until they buy.

This paper is motivated by Pesendorfer's findings. The analysis below differs from Pesendorfer's work in both theory and empirics. 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. Our goal is to study whether stockpiling is behind the demand accumulation documented by Pesendorfer (see also Blatteberg and Neslin, 1990). We model the behavior of a consumer who can store the product. In order to generate predictions testable with household level data, the model explicitly captures

consumers' endogenous decision to return to the market (rather than having an exogenously given number of consumers arrive each period.) Sobel's model is mute about individuals' behavior over time; they buy once and leave the market.⁵ Furthermore, at the aggregate level our model predicts that quantity sold is a function of the duration since the previous sale during both sale and non-sale periods (while Pesendorfer's predictions are only for sale periods). The data supports these predictions. Demand accumulation during non-sale periods is consistent with consumers' following a *s-S* type policy, which is not predicted by a Sobel type model. A longer duration from previous sale, makes more consumers cross their inventory threshold *s*, even during non-sale periods. Our model also delivers comparative statics across products, for which we find anecdotal support in the data. From the empirical perspective, the key distinction with Pesendorfer is that we test for stockpiling using household level data. The testable predictions delivered by the model pertain to within household differences across sales and non-sale periods.

Boizot et. al. (2001) study dynamic consumer choice with inventory. They test whether (i) duration from previous purchase increases in current price and declines in past price, and (ii) quantity purchased increases in past prices. The papers, although similar in spirit, differ both in the theoretical implications and data. On the theory side, they assume that consumption does not respond to prices, which are assumed i.i.d. over time. While this is a good starting point, it assumes away an alternative explanation of some of the results. Furthermore, consumption effects are present in our data (see section 4.3). On the data side, the key advantage of our data is detailed information about the product purchased (for example, brand and exact size). Such data is important for distinguishing sales from substitution towards cheaper brands (for example, if a consumer pays a lower than average price it may be due to purchase of a different, cheaper, brand and not due to a sale). While Boizot et. al. clearly point this out, due to data limitations, they cannot deal with this problem. The key overlap with their findings is in Section 4.3 (more precisely, in our Table 5), see discussion there for a comparison of results.

⁵Even if we enriched Sobel's model by allowing exogenous depreciation, so that consumers do return to the market, the model still would not provide testable implications at the individual level. Low types would buy exclusively on sale and high types would buy immediately upon return to the market.

Erdem et al (2000), look at sales from the inventory perspective, constructing a structural model, closer to Hendel and Nevo (2002). Besides several modeling assumptions we differ from them in focus. Our main goal is primarily descriptive, testing the most general predictions of the theory, while their starting point is a dynamic forward looking model, which they structurally estimate. Hosken et al. (2000) study the probability of a product being put on sale as a function of its attributes. They report that sales are more likely for more popular products and in periods of high demand. Warner and Barsky (1995), Chevalier, et al. (2000) and MacDonald (2000), study the relation between seasonality and sales. The effect we study complements the seasonality they focus on. The same is also true for Aguirregabiria (1999), who studies retail inventory behavior. His paper is about firm's inventory policy and its effect on prices, while our focus is on consumers' inventory policies given the prices they face. Hong, McAfee and Nayyar (2000) present a storable good model that generates demand dynamics, namely, a (negative) link between current prices and future demand. They characterize suppliers' behavior in a competitive context. They show that there exist equilibria where firms use random pricing. Moreover, prices are negatively correlated over time.

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 use household-level data to show that when purchasing during a promotion households tend to buy more units, larger sizes and in shorter duration to their previous purchase. We discuss, in somewhat more detail, how our findings relate to this literature in Section 4 (as we present our results).

2. The Data, Product Categories and Preliminary Analysis

2.1 Data

We use data collected by IRI using scanning devices in nine supermarkets, belonging to different chains, in two separate sub-markets in a large mid-west city. The data has two components, one with store and the other with household-level information. The first contains prices, quantities

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).

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.

2.2 The Product Categories

We focus here on three product categories available in the data: laundry detergents, soft-drinks and yogurt. We focus on only three categories due to space limitations. We chose these three categories because of the simplicity of the choice set. As we describe below, a handful of brands have a significant market share in each of these categories. Furthermore, we believe each of these categories represents a class of similar products. 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 consumption and 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 can still be stored although for a shorter period.

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 these 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 define exactly what is a sale price. 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, non-sale, 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.

None of these definitions is perfect. For the purpose of this section, which is purely descriptive, the exact definition is less important. As we show in the next section the theory provides an exact definition of what is a sale from the consumer’s perspective, but this provides little insight as to what definition we should use. 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, hence any of them will introduce measurement error into the analysis in Section 4.

Using these different definitions of a sale we display in Table 1 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. Already in this table we can foresee some of the support for our theory. The comparison between the figures for detergent and yogurt provide support for stockpiling behavior. 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. Furthermore, the main alternative explanation is that consumers simply increase their consumption in response to a price reduction (i.e., they buy more because they consume more, not in order to stockpile). 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 sizes. Indeed, size discounts are consistent with price-discrimination based on consumer storage costs.⁷ In Table 2 we display statistics for the major sizes in each 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 sizes bottles (the main size bottles are 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 packs.

The first column in Table 2 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 misleading. 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 that there are 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

⁷Size 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 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 opposite. There are more sales, and a larger fraction sold on sale, for the smaller sizes. In Section 4 we discuss this finding, which we think 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.⁸ 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.

2.3 Preliminary Analysis: Demand Accumulation

In this section we study the patterns of demand accumulation. We first confirm the presence of demand accumulation during sales for the products we study, and then, we look for additional patterns in the data to understand what is behind demand accumulation. More precisely, using the store level data we show that the quantity purchased during sales increases with duration from the previous sale. However, in addition, we find that accumulation is present during non-sale periods as well. 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.

⁸These variables both have several categories (for example, type of display: end, middle or front of aisle). We treat these variables as dummy variables.

Table 3 presents the results of regressing the log of quantity sold, measured in 16 ounce units, as a function of price, measured in dollars per 16 ounce, current promotional activity and duration since previous promotional activity.⁹ 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 pervious sale is positive and significant, for all three products.¹⁰ As already recorded in the literature, for other products, demand during sales accumulates. 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. However, the effect of duration is stronger during sales than non-sale periods, for all three categories.

Pesendorfer (2002, Tables 9 and 10) also 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. Our findings support a richer model, consistent with accumulation in non-sale periods. The model we present in the next section, predicts that duration from previous sale matters both during sale and non-sale periods. Individuals accumulate inventory according to an (S, s) rule, hence, duration affects demand, no matter what the prevailing price is. Moreover, the model also predicts the duration effects to be more pronounced for sale periods (see Implication 5 in Section 3.3).

The numbers in this section highlight the relevance of the inventory model. We next derive other testable predictions. We are particularly interested in predictions at the consumer level. In

⁹We will discuss later how to define a sale and what a product is. For the moment we define a sale as a price at least 5% below the modal price for each product/store and a each UPC as a product. 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 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 3 show the role of duration on quantity sold.

¹⁰Duration 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 values mostly observed in the data. Therefore, we limit the discussion to the linear coefficient on duration.

Sobel-type models the product is treated as durable, buyers leave the market forever after purchasing one unit. Such a model, as in Pesendorfer (2002), does not generate testable predictions at the individual level. Our goal is to look into the detailed data to test what is behind the demand behavior, already (partially) documented at the aggregate level.

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

3.1 The Basic Setup

Consumer i obtains the following flow utility in period t

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

m_{it} where c_{it} is the quantity consumed, v_{it} is a shock to utility and is the utility from consumption of the outside good. The stochastic shock, v_{it} , captures demand shocks unobserved to the researcher. For simplicity we assume the shock is additive in consumption, $u(c_{it}, v_{it}) = u(c_{it} + v_{it})$, affecting the marginal utility from consumption. Low realizations of v_{it} increase the household's need, increasing demand and making it more inelastic. We also assume $\frac{\partial u(c_{it} + v_{it})}{\partial c} \geq \alpha p \forall p$ and $\forall v$, which is sufficient for positive consumption every period. This assumption has no major 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_{it} , 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), instead of endogenously

determined. 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. It is important to allow consumption to vary in response to prices since this is the main alternative explanation to why consumers buy more during sales. We want to make sure that our 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 these 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 i , to simplify notation, the consumer's problem can be represented as

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

where α 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 $C'' > 0$.¹¹

The information set at time t , $\mathbf{I}(t)$, consists of the current inventory, i_{t-1} , current prices, and the current shock to utility from consumption, v_t .¹² 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.¹³ We assume $F(p_{t+1} | p_t) < F(p_{t+1} | p'_t)$ for all $p'_t > p_t$. Namely, the probability of a sale at time $t+1$ is higher if there is no sale

¹¹Notice we do not need to impose $c \geq 0$ since we assumed $\partial u / \partial c$ is such that there is always positive consumption.

¹²It is quite 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 – the cost of a stock out. Since this is not our focus we ignore this effect, but it can easily be incorporated.

¹³A Markov process fits the observed prices reasonably well. Qualitatively, we believe our results continue to hold with a more general price process, but the analysis is more complicated.

at time t .

3.2 Consumer Behavior

In each period a 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 a (conditional) S - s type rule. We now formalize this result.

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,

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

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 $C'(S(p)) + \alpha p = \delta E(\lambda_{t+1} \mid I(t))$. Manipulating the first order conditions we get the main result.

Proposition 1 In periods with purchases, $x_t > 0$, the target level of inventory, i_t , equals $S(p)$, a decreasing function of p , independent of the other state variables i_{t-1} and v_t . Moreover, the inventory level that triggers a purchase is $s(p, v_t) = S(p) + c^*(p, v_t)$, which is decreasing in

both arguments.

Proof: All proofs are provided in the Appendix.

Consumers behave according to an S - s rule, where the upper band, S , is a decreasing function of current price and the lower band, s , declines both on prices and the utility shock.¹⁴

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

3.3 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 when the price decreases (see equation (3)) or because of stockpiling.¹⁵ Since we do not know the size of the consumption effect, showing that quantity purchased increases during sales does not necessarily imply stockpiling.

An immediate implication of Propositions 1 and 2, not predicted by the static model, is that

Implication II: *Quantity purchased and the probability of purchase decline in inventories.*

Since we do not observe inventories, we cannot directly test this implication. We use two different strategies to test the model in light of not observing inventories. In Section 4.4 we assume that consumption is fixed, which allows us to compute a proxy for the unobserved inventory. As we

¹⁴If only discrete quantities are available or prices are non-linear in quantities then the target inventory $S(\cdot)$ becomes a function of i_{t-1} and v_t .

¹⁵Corollary 1, in the appendix, describes the stockpiling effect for every price less than p^r , which we have implicitly equated with a sale. In the empirical analysis we do not observe p^r and therefore we will experiment with several definitions of a sale.

mentioned above this is not an attractive assumption (and seems to be inconsistent with some of our findings) 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 both bands, s and S , are decreasing functions of price.

According to Proposition 1 $S(p)$ is a decreasing function, which 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, $s(\cdot)$. Absent stockpiling, duration would be unaffected by sales. This gives the following implication, which indirectly testifies to the presence of stockpiling.

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

From Proposition 1 we know that the inventory that triggers a purchase, $s(\cdot)$, is lower at non-sale prices. Hence, according to our model, since during a non-sale purchase initial inventory is lower, on average the duration from previous purchase will be longer.

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

Furthermore, 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 purchase, while a lower initial inventory is more likely if the previous purchase was not on sale.

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 turn to implications on aggregate demand. The aggregation of Implication I2 over a population of buyers who visit the supermarket at different periods leads to Implication I5, namely, that store level demand increases with duration since the last sale. Moreover, since at non-sale periods consumers only demand for current consumption, while, on sale they hoard inventories we expect duration to have stronger effects during sales.

Implication I5: *Aggregate demand increases in the duration from the previous sale. Furthermore, duration effects are stronger during sales.*

4. Results

In this section we test the implications derived in the previous 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. Two issues arise. First, sales have to be distinguished from changes in the “regular” price, for example, due to seasonality. In addition we have to define what is the highest price the consumers treat as a sale. Corollary 1, in the appendix, defines this cutoff but gives us little guidance as to how we should define a sale in the non-structural tests we perform below. Both these problems suggest that a correct definition of a sale will vary across households and across products. Any definition that is held constant across households and products will introduce noise into the analysis (and most likely bias the estimated effects towards zero). 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 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 4 through 6, we treat each category as a single product. A broad product definition captures the fact that different brands are substitutes. Naturally, 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 for our descriptive purposes for two reasons. 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, the broad product definition (including several UPCs) does not mean we treat all the UPCs as perfect substitutes. 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. In that case the duration from last purchase, regardless of the brand, determine the current inventory (see details in Hendel and Nevo, 2002).

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.

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.

4.1 Aggregate data: the effect of duration from previous sales

According to implication I5, aggregate demand should increase with the duration from the previous sale (i.e., as consumers run out of the inventory stockpiled during the last sale). Moreover, the effect of duration while stronger during sales, should also be present during non-sale periods.

We already discussed demand accumulation in section 2.3. The numbers presented in Table 3 show duration effects are present at the aggregate level. In accordance with I5, duration matters during both non-sale and sale periods and the effect is more pronounced during the sale periods.

4.2 Household sales proneness

We now turn our attention to the household data described in Section 2. In this section we study the factors that impact a household's fraction of 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 observed period on various household characteristics.¹⁶ 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 variation, across households, in the fraction of purchases on sale. 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

¹⁶We also looked at the fraction of quantity purchased on sale. The results are essentially identical.

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 4. 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 our model that predicts lower storage costs are correlated with purchasing more frequently 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).¹⁷

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. We provide a broader profile, by looking at additional characteristics of the shoppers and by looking at a much wider set of products. Column (v) shows that households who shop more frequently tend to buy more on sale. These effects also hold once we control for the characteristics used in columns (i) - (iii).

4.3 Sale vs. non-sale purchases

In this section we discuss the main predictions of the model at the level of the individual. We focus on those predictions that stem from the inventory model but that would not be expected under the alternatives. In Table 5, we compare, for each product category, the averages of several

¹⁷The 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 controlling for search proxies (see column (vi)).

variables between 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).¹⁸ Actually, our theory 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 we believe that in order to rigorously test the theory one has to separate these effects.¹⁹ 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 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 5 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, while 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 2, which highlights the interaction of sales and non-linear prices.

¹⁸For a precise definition of within and between estimates see Greene (1997).

¹⁹Previous 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.)

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 will 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 we could use the increased quantity as proof for stockpiling. Instead of making such an assumption we turn to predictions that allow us to separate the two theories (and indirectly say something about this assumption). Rows 4 and 5 of Table 5 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 back of the envelope 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. We believe that the alternative explanations, of the between differences, are likely to involve a stockpiling component and therefore the between differences also point to the relevance of stockpiling. For example, a possible alternative explanation is that households with large demand are likely to 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.

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 within effects might 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 be more pronounced for households who are sensitive to sales, while close to absent (zero) for non-deal prone households. Hence, the within effect probably understate the responses of those households aware of sales. Second, the definition of sale probably introduces measurement error and biases the effect towards zero.²⁰ Third, there might be consumption and (potentially) stockpiling of several products. 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 5 implicitly assume that these are perfect substitutes, since duration is measured to any purchase in the category. If we take the other extreme and compute the results for the sub categories separately the effects generally increase.

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, 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.

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). We, on the other hand, have detailed data that contains information on the

²⁰For 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.

brand (and size) purchased, which we use to define a sale. We can therefore deal with the problem they raise. As we discussed above, there are some potential problems also with our definition of a sale. We therefore find it encouraging that a study using different data and methods finds similar effects.

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. Especially, since the above results seem to suggest that the between effects dominate.

4.4 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 6, (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).

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 6, 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 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. Intuitively, 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 (2002), where we model the discreteness of purchases and non-linear prices, suggest that the magnitude of the coefficients presented in Table 6 is consistent with stockpiling behavior that is economically significant.²¹

4.5 A cross-category comparison

The last set of tests of our theory involve a comparison across products. Unfortunately, none of the categories in our data is completely perishable. We were able to obtain data comparable to

²¹There 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.

ours, but from a different city, on milk.²² 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 till 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 according to our model there should be no sales for milk. Indeed, that seems to be the case.

Another cross-category comparison involves the difference between laundry detergents and yogurt. Since the average duration between supermarkets visits is less than a week both these products are storable. However, there is a key difference between how one would store them. It is reasonable to assume that the storage costs for yogurt, holding total quantity fixed, are lower for smaller sizes of yogurt. It is easier to store in the refrigerator four small 8 ounce containers rather than one large 32 ounce container. Furthermore, unlike detergents, the storability of yogurt decreases once a container is opened. This suggests that for detergents we should see more sales for larger sizes and when consumers purchase on sales 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 2 columns two and three and Table 5 second row .

Further evidence linking the relation between the easier-to-store size and sales is presented in the last column of Table 2, 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. Implications for Demand Estimation: Short vs long run elasticities

In this section we attempt to quantify the bias in demand elasticities that would arise from neglecting dynamics. Short run elasticity estimates are likely to overstate consumers' long run price responses, which involve consumption responses but no stockpiling. While static demand estimates are likely to capture both the consumption and stockpiling effects. A full demand estimation is

²²We wish to thank Sachin Gupta, Tim Conley and Jean-Pierre Dube for providing us with these data.

beyond the scope of this paper.²³ We present a simple exercise, which approximates short and long run estimates using the numbers from Table 5. The results are presented in Table 7.

We proxy the short run elasticities by the increase in demand associated with a sale. We divide the percent change in *Total* quantity sold by the average percent price decline during sales reported in Table 2. For example, quantity demanded of detergents is 32.4% higher during sales (see *Total* quantity effects in Table 5) while the average price discount during sales is 16.2% (see Table 2). Leading to a short run price elasticity of 2 (third row in Table 7).

On the other hand, we proxy the long run elasticity using inter-purchase duration data to purge the elasticities from the stockpiling effect. The idea is to compute a duration-corrected change in demand, we do so by spreading the purchases after a sale, over the longer durations to next purchase that take place after a sale. That way we get a measure of the extra consumption induced by the lower prices, net of stockpiling. Absent consumption effects, the duration adjusted consumption after a sale and a non-sale period should be equal. We interpret higher quantity adjusted purchases during sales as evidence of consumption effects. For detergents, we compute the consumption rate after a sale by dividing the quantity sold during sales, $4.79+1.14$, by duration after sales, $43.75+1.95$.²⁴ We then divide this number by the consumption rate after a non-sale purchase: $4.79/43.75$. The ratio of consumption rates is: 1.185. Namely, after a sale consumption increases by 18.5%. Which translates into a long run demand elasticity of 1.15. Thus, neglecting dynamics would have lead us to conclude that demand is 74% more elastic that it really is.

Similar numbers for the other two products are presented in Table 7. They show even larger biases than detergents. Yogurt shows the lowest short run elasticity, 0.6, but at the same time the largest stockpiling bias. We would attribute demand to be almost 6 times more elastic than it really is if we neglect dynamics. Finally, soft drinks show the largest response to prices, but at the same time the largest stockpiling effect, which would lead us to a to infer elasticities are more than twice

²³See Hendel and Nevo (2002) for such estimation.

²⁴We use the *Within* figures from Table 5, in order to concentrate on pure consumption effects, namely, to capture quantity responses due to price changes.

larger than they really are.

6. 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. We find that static demand estimates, which neglect dynamics, may overestimate own-price elasticities by up to a factor of 2 to 6. While these calculations can give an idea of the order of magnitude of the impact of stockpiling on demand estimates, they are only approximations. The full estimation of short vs long run demand elasticities is beyond the scope of this paper.

We are currently exploring extensions along several dimensions. First, in Hendel and Nevo (2002) we estimate a structural model that addresses some of the issues we ignore here. Obviously, the cost is that we have to impose more structure on the data. The structural model provides interpretable estimates and enables us to perform counterfactual experiments. Preliminary results, using data from the laundry detergents category, suggest that ignoring the dynamic effects can substantially bias the estimates of own- and cross-price elasticities and have profound effects on their

implications. We are extending our theoretical analysis to include the supply side. This, jointly with the structural estimates, will allow us to examine questions like what are the optimal patterns of sales and why are they profitable to sellers.

References

- Aguirregabiria, V. (1999) "The Dynamics of Markups and Inventories in Retailing Firms," *The Review of Economic Studies*, 66, 275-308.
- 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.
- Chevalier, J., A. Kashyap and P. Rossi (2000), "Why Don't Prices Rise During Peak Demand Periods? Evidence from Scanner Data," NBER Working Paper no. 7981.
- Chiang J. (1991), "A Simultaneous Approach to the Whether, What and How Much to Buy Questions," *Marketing Science*, 10, 297-315.
- Conlisk, J., E. Gerstner and J. Sobel (1984), "Cyclic Pricing by a Durable Goods Monopolist," *Quarterly Journal of Economics*; 99(3), 489-505.
- Erdem, T., M. Keane and S. Imai (2000), "Consumer Price and Promotion Expectations: Capturing Consumer Brand and Quantity Choice Dynamics under Price Uncertainty," University of California at Berkeley, mimeo.
- Feenstra, R. and M. Shapiro (2001), "High Frequency Substitution and the Measurement of Price Indexes," NBER Working Paper no. 8176.
- van Heerde H., Leeftang 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 (2002), "Measuring the Implications of Sales and Consumer Stockpiling

Behavior”.

Hong, P., P. McAfee and A. Nayyar (2000), “Equilibrium Price Dispersion with Consumer Inventories,” working paper.

Hosken, D., D. Matsa, and D. Reiffen (2000) “How do Retailers Adjust Prices: Evidence from Store-Level Data,” working paper.

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.

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.

Salop, S. and J. Stiglitz (1982) “The Theory of Sales: A Simple Model of Equilibrium Price Dispersion with Identical,” *American Economic Review*; 72(5), 1121-30.

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.
- Varian, H. (1980), "A Model of Sales," *American Economic Review*; 70(4),651-59.
- Ward, R. and J. Davis (1978), "A Pooled Cross-Section Time Series Model of Coupon Promotions," *American Journal of Agricultural Economics*, 60, 393-401.
- Warner, E. and R. Barsky (1995), "The Timing and Magnitude of Retail Store Markdowns: Evidence from Weekends and Holidays," *Quarterly Journal of Economics*;110(2), 321-52.
- 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: $C'(i_t) + \alpha p_t = \delta E(\lambda_{t+1} | I(t))$, which shows the end-of period inventory, i_t , is independent of the states variables i_{t-1} and v_t . Moreover, since $F(p_{t+1} | p_t)$ increases in p_t , by equation (3) we get that the right hand side of the last equality declines in p_t . Hence, since $C'' > 0$ the end of period inventory, i_t , declines in price.

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 $x_t = 0$ we know $\psi_t > 0$, which in turn, by equation (3), implies $c_t > c^*(p_t, v_t)$. 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 $u'(c_t + v_t; \theta) = \alpha p_t$, 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 $S(p_t)$ declines in p_t . ■

Corollary 1: There is a price $p^r < p^m$, where p^m is the highest (non-sale) price, such that at any price $p \geq p^r$ if consumers buy they do so for current consumption exclusively.

Proof of Corollary 1: At p^m if $x_t = 0$ there is nothing to show. If $x_t > 0$ we can combine equations (4) and (5) to get $C'(i_t) + \alpha p_t = \delta E(\lambda_{t+1} | I(t)) + \mu_t$. Moreover, from equation (4) we know $E(\lambda_{t+1} | I(t)) \leq \alpha E(p_{t+1} | I(t))$. The right hand side of the last inequality is strictly lower than αp^m (as long as prices lower than p^m , arise with positive probability). Hence, since $\delta \leq 1$ we know $C'(i_t) + \alpha p^m > \delta E(\lambda_{t+1} | I(t))$ for any $i_t \geq 0$. Therefore, equation (5) can hold with equality only if $\mu_t > 0$, i.e., when $i_t = 0$. Since the inequality is strict it holds also for some $p^r < p^m$. Concluding the proof that if any quantity is purchased, it is for consumption only, since no inventories will be left at the end of the period. ■

Table 1
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 2
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 3
Demand as a Function of Duration from Previous Promotional Activity

Variable						
	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 4
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 5
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 households	Between households		Total	Within households	Between households		Total	Within households	Between 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 6
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.

Table 7
Consumption vs Stockpiling Effects

	Laundry Detergents	Yogurt	Soft-drinks
Total change in quantity bought	32.4%	10.0%	100.0%
Elasticity (short run)	2.0	0.6	4.2
Within change in quantity bought	23.8%	12.5%	60.0%
Duration corrected consumption change	18.5%	1.8%	43.5%
Elasticity (long run or duration adjusted)	1.2	0.1	1.8
Bias (ratio Elast/Adjusted Elast)	1.7	5.6	2.3

Based on the figures reported in Tables 2 and 5, and authors calculations. See discussion in text for exact definitions.