

1680

Discussion
Papers

Consumer Stockpiling and Sales Promotions

Opinions expressed in this paper are those of the author(s) and do not necessarily reflect views of the institute.

IMPRESSUM

© DIW Berlin, 2017

DIW Berlin
German Institute for Economic Research
Mohrenstr. 58
10117 Berlin

Tel. +49 (30) 897 89-0
Fax +49 (30) 897 89-200
<http://www.diw.de>

ISSN electronic edition 1619-4535

Papers can be downloaded free of charge from the DIW Berlin website:
<http://www.diw.de/discussionpapers>

Discussion Papers of DIW Berlin are indexed in RePEc and SSRN:
<http://ideas.repec.org/s/diw/diwwpp.html>
<http://www.ssrn.com/link/DIW-Berlin-German-Inst-Econ-Res.html>

Consumer Stockpiling and Sales Promotions

Anna Lu*

Abstract

In retailing markets of storable goods, consumer behavior is typically characterized by stockpiling. While existing research has developed rich models for such strategic consumer behavior, little is known about how sellers should ideally respond to it. In this paper, we provide insights into how frequency and depth of promotions affect consumer purchases and seller revenues in the long run. We show an application to the U.S. market for laundry detergent. We use estimates from a structural dynamic demand model to simulate different pricing policies and find that in the detergent market, an increase in promotion depth is more effective than a change in promotion length. Our results suggest that this finding can be translated to markets with a large heterogeneity in storage costs and steady consumption rates.

JEL: M31, D22, L11

Keywords: Promotion, Frequency, Depth, Stockpiling, Storable Goods

*Deutsches Institut für Wirtschaftsforschung (DIW Berlin) and Düsseldorf Institute for Competition Economics (DICE). I received valuable advice from Jaap Abbring, Ulrich Doraszelski, Jean-Pierre Dubé, Anthony Dukes, Tomaso Duso, Jana Friedrichsen, Germain Gaudin, Matthew Osborne, Stephan Seiler, and Hannes Ullrich. I also thank seminar participants at DIW Berlin and DICE for their helpful comments and Mike Kruger from Information Resources Inc. for his help with understanding the data.

1 Introduction

The design of promotions is a key concern for both marketing practitioners and marketing research. Since the 1970s, price promotions have become the main marketing instrument in many industries (Currim and Schneider, 1991; Blattberg, Briesch and Fox, 1995). Promotions are particularly important in grocery retailing. In 2016, they accounted for almost 66% of the marketing budget of consumer packaged goods (Bhardwaj et al., 2016), by far exceeding advertising expenditures.

It is well known that price promotions in storable-goods markets induce consumer stockpiling: When retailers offer promotions, consumers strategically buy large amounts at low prices and store them for future consumption. Previous research developed rich models to describe and estimate such strategic consumer behavior (e.g. Erdem, Imai and Keane, 2003; Hendel and Nevo, 2006*a*, 2013; Su, 2010). What is much less well understood is how firms *should* design promotions in response to consumer stockpiling.

In this paper, we are the first to study the effect of promotion length vs. promotion depth on firms' long-term revenue in a structural framework. To do so, we use rich panel data from the U.S. market for laundry detergent. Detergent can be stored long before and after its first use, it comes in large packs associated with substantial storage cost and, thus, is a typical product in which to study the dynamics of stockpiling (Bell, Iyer and Padmanabhan, 2002; Hendel and Nevo, 2006*b*; Seiler, 2013; Pires, 2016).¹ We use a dynamic discrete-choice model of strategic consumer stockpiling. In our model, consumers are

¹ There are many other product categories that also have these characteristics, e.g. soda (Hendel and Nevo, 2013), ketchup (Pesendorfer, 2002; Erdem, Imai and Keane, 2003; Sun, Neslin and Srinivasan, 2003), coffee (Neslin, Henderson and Quelch, 1985), as well as razors and razor blades (Hartmann and Nair, 2010).

forward-looking with rational expectations over future prices. They can choose to stockpile but incur storage costs. We estimate consumer preferences, price sensitivities and storage costs. Using these estimates, we can simulate how changes in promotion depth and length affect consumer purchase decisions and seller revenues.

Our results suggest that in the detergent market, shorter but deeper promotions are preferred over longer, shallower promotions. We find that the revenue elasticity with respect to promotion depth is, *ceteris paribus*, about four times higher than with respect to promotion length. Our findings provide general insights into promotion design that can prove helpful for industry practitioners: Compared to longer promotions, deeper promotions tend to perform better in markets with substantial heterogeneity in storage costs, large heterogeneity in price sensitivity, and steady consumption rates.

The rest of the paper is organized as follows: In Section 2, we give a brief overview of the literature. In Section 3, we describe the data. We lay out the model in Section 4 and describe our identification and estimation strategy in Section 5. We discuss the estimation results in Section 6 and the counterfactual simulations in Section 7. Finally, we conclude in Section 8.

2 Literature

The core contribution of this paper is to quantify and compare the revenue effects of promotion length and promotion depth. In doing so, we add to a large literature on the impact of promotional price cuts. Since the mid-1990s, the economics and marketing literature has used structural models to investigate how sellers should set their prices (Kadiyali, 1996; Besanko, Gupta and Jain, 1998; Sudhir, 2001; Chintagunta, 2002; Verboven, 2002;

Besanko, Dubé and Gupta, 2003; Draganska and Jain, 2006; Pancras and Sudhir, 2007; Richards, 2007). The majority of these pricing models are based on static demand models in which consumers are not forward-looking and remain unaware of the fact that their present-day decisions will affect their future payoffs.

However, it is well-known that consumers *do* behave in a forward-looking fashion in a vast array of situations and markets.² In particular, researchers have found strong evidence that consumers stockpile in storable-goods markets when they face temporary price cuts, anticipating that prices will increase shortly after (Blattberg, Eppen and Lieberman, 1981; Neslin, Henderson and Quelch, 1985; Mela, Jedidi and Bowman, 1998; Pesendorfer, 2002). One well-documented example is the market for laundry detergent (Bell, Iyer and Padmanabhan, 2002; Hendel and Nevo, 2006*a,b*; Seiler, 2013; Pires, 2016).

Previous work in economics and marketing research has developed sophisticated structural models to estimate consumer stockpiling (Erdem, Imai and Keane, 2003; Sun, Neslin and Srinivasan, 2003; Hendel and Nevo, 2006*a*; Osborne, 2010; Hartmann and Nair, 2010; Seiler, 2013; Pires, 2016). These studies model forward-looking strategic stockpiling by incorporating a storage cost parameter and consumer price expectations. The main strength of these structural models is that they allow us to perform counterfactual simulations. Consequently, the literature looks at various demand-side counterfactuals, for example how household purchases change when storage costs are lowered.

There is limited empirical work on pricing for markets with dynamic consumer demand. One reason for this is that, even in toy models, it is

² For example, consumers display forward-looking behavior when booking flight tickets (Li, Granados and Netessine, 2014), buying college textbooks (Ching and Osborne, 2015), and making career decisions (Keane and Wolpin, 1997).

computationally challenging to solve for optimal prices in a framework with both a dynamic demand side and a dynamic supply side. We are aware of only one paper doing this: Hendel and Nevo (2013) develop a simple dynamic demand model in order to empirically quantify the impact of intertemporal price discrimination on profits in the soda market. They find that sales capture 25-30% of the gap between non-discriminatory profits and third-degree price discrimination profits. Their model relies on simplifying assumptions, e.g. storage cost is assumed to be zero and products are assumed to be perishable. These assumptions allow them to derive an optimal pricing strategy, but make the model unsuitable for many markets.

In related work, Nair (2007) solves for optimal prices in a durable goods market with forward-looking consumers. More specifically he investigates different pricing policies in the market for video games: After the introduction of a new game, consumers may choose to wait for the price of the game to drop. Nair solves for optimal prices in a market with one monopolist and two types of consumers who differ in their product valuation. Such an approach could theoretically be taken to storable goods data. However, product storability comes with additional computational challenges: It requires keeping track of households' inventories because, unlike in markets of durable goods, consumers do not drop out of the market after making a purchase. In practice, it is computationally challenging if not infeasible to incorporate this in the estimation.

Instead of solving for optimal prices, Osborne (2010) simulates different pricing regimes. This is the paper closest to our work. Osborne studies how changes in frequency and depth of promotions affect revenues in the canned tuna category. He finds that increasing promotion depth significantly increases

sold quantity whereas an increase in promotion frequency has a much smaller effect. Our work differs from his in that we investigate promotion length instead of promotion frequency. This difference is important because, unlike a change in frequency, a change in promotion length will not affect the fixed costs of running promotions.

Our work is also related to a large literature on sales. This literature proposes a multitude of rationales for sales. One important explanation is that sales take advantage of some dimension of consumer heterogeneity. This may be heterogeneity in information about prices (Varian, 1980), preferences and brand loyalty (Sobel, 1984; Narasimhan, 1988; Raju, Srinivasan and Lal, 1990; Hendel and Nevo, 2013), storage levels (Hong, McAfee and Nayyar, 2002; Pesendorfer, 2002), storage costs (Blattberg, Eppen and Lieberman, 1981; Jeuland and Narasimhan, 1985), or heterogeneity in the stores that shoppers visit (Salop and Stiglitz, 1982).³ In our paper, we investigate heterogeneity in storage costs and storage levels as the main rationale for sales.

3 Data

3.1 IRI Panels

The market research company Information Resources Inc. provided us with U.S. retail data on households' laundry detergent purchases in the years 2001 to 2004. Laundry detergent comes in two main forms: liquid and powder. In order to be able to compare pack sizes, we have to choose one form of detergent. In the following, we study liquid laundry detergent because it

³ There is also a vast business and operations research literature investigating seller-side rationales for sales, e.g. inventory management (e.g. Whitin, 1955; Petruzzi and Dada, 1999) or loss-leading (e.g. Mason and Mayer, 1984; Lal and Matutes, 1994). In this paper, we abstract from such motives for sales.

has a market share of 94.9% in our sample period. The household panel contains 6,000 to 10,000 households (the number of panelists varies by year) in Eau Claire, Wisconsin, and Pittsfield, Massachusetts. These households use handheld scanners to scan their purchases after every shopping trip. We observe all detergent purchases of the panelists, and for each purchase, we observe the paid price, the date of the purchase, the store identifier, the chain identifier, and the characteristics of the household. The average household purchases detergent every six weeks, buys only one pack per trip, and switches between two different brands.

Table 1 displays brand summary statistics. The market is relatively concentrated, with the largest three firms capturing almost half of the market. We focus our analysis on the 12 best-selling brands, which have an accumulated market share of 82.5%. All residual brands are collected into one composite brand called “OTHER”. In the supermarket retailing context, detergent is a relatively expensive product. More than half of the brands in our sample have at least one variety that retails at more than 10 USD.

We also use an auxiliary store panel to supplement our estimation with information on prices that we do not observe in the household panel (see Section 5.1). The store panel contains check-out scanner data from 1,588 U.S. supermarkets. We observe the weekly sales quantity for every laundry detergent brand and size that was sold at least once, together with prices, promotions, and information on the stores’ location.⁴

⁴ For more details on the IRI data set see Bronnenberg, Kruger and Mela (2008).

Table 1: Summary Statistics: Brands

Brand	Price				Market Share
	Mean	Std.Dev.	Min	Max	
ALL ¹	4.974	1.534	.65	15.4	9.794
CHEER	6.115	1.558	3.5	27.96	1.713
DYNAMO	4.119	1.909	1.49	9.99	7.529
ERA	5.002	1.922	1.48	15.13	8.481
FAB	4.891	1.344	1.98	6.79	1.067
GAIN	5.776	1.245	2.44	10.59	.86
OTHER	3.844	1.654	.34	14.89	17.499
PUREX	4.096	1.549	1.8	11.69	12.39
SURF	5.656	.891	3.09	9.98	.466
TIDE	7.774	3.19	1.28	24.59	19.468
TREND	2.797	.156	2.69	3.75	.933
WISK	6.087	1.946	1.5	18.51	6.504
XTRA	2.842	.74	.5	5.99	12.323
YES	4.334	1.693	1.99	6.99	.972

¹ ALL is the name of a brand.

3.2 Preliminary Analysis of Stockpiling

We observe a distinctive price pattern in the store level data. Prices are not kept constant, but vary across time. Figure A.1 and A.2 show how the average prices for 100-ounce and 200-ounce packs change over time across all stores. It shows that promotions do not occur in predictable intervals and that consumers consequently face uncertainty with respect to future prices and promotions.

In the following, we provide preliminary evidence that consumers stockpile strategically, i.e. they purchase for *future* consumption. In contrast, myopic consumers make purchases only with current consumption in mind. The household panel shows that during promotions, more ounces of detergent are sold. Imagine that this sales spike is the result of consumers only caring about

the present. Then present-day purchase decisions should be independent of previous or following purchases. Interestingly, we find that a) the duration since the previous purchase is shorter during a promotion; and b) the duration before the next purchase is longer during a promotion. Both effects are significant at the 10% level (see Table A.1). This finding is inconsistent with the hypothesis that consumers do not have strategic stockpiling motives. Instead, it suggests that households build stock during promotions by purchase acceleration, i.e. by making earlier purchases than initially planned.⁵

Furthermore, we find that households rarely buy multiple packs per shopping trip; instead, the vast majority of households (84.7%) buys only one pack of detergent per trip. We also find that consumers stockpile by buying *larger* packs: The share of large packs increases considerably during promotional periods (see Table A.2). Consumers switch frequently between pack sizes; 54.0% of all households buy more than one pack size, and on 12.6% of all shopping trips, the household chooses a different pack size than the pack size it purchased last time. All in all, our preliminary findings suggest that consumers stockpile primarily through purchase acceleration and by buying larger pack sizes.

4 Model

In each period t , household $i = 1, \dots, I$ can buy one pack of laundry detergent.⁶ We define a product j as a unique combination of brand $b = 1, \dots, B$ and size $z = 1, \dots, Z$, for example “Tide 50-ounce pack” or “Xtra 100-ounce pack”.⁷

⁵ For a more detailed investigation into the evidence of stockpiling in the market for laundry detergent see Hendel and Nevo (2006b) and Pires (2016).

⁶ We make this assumption since the vast majority of households in our sample never buys more than one pack per trip. For a further discussion see Section 5.2.

⁷ In this paper, we do not model retailer choice because it adds significantly more complexity to the consumer decision (Farley, 1968; Fotheringham, 1988; Rhee and Bell, 2002; Smith,

The household can also choose to buy no detergent ($j = 0$). In the following, we denote a household's decision to buy product j in period t as $d_{it} = j$. The flow utility of $d_{it} = j$ at the time of the purchase decision is:

$$U_{ijt} = \begin{cases} \underbrace{v(c_i) - C(k_t) + \alpha p_{jt} + x_{jt}\beta + \varepsilon_{ijt}}_{=u_{ijt}} & \text{if } j = 1, \dots, J \\ v(c_i) - C(k_t) + \varepsilon_{i0t} & \text{if } j = 0, \end{cases} \quad (1)$$

where p_{jt} is the retail price, x_{jt} is a vector of the observable characteristics, and ε_{ijt} is an individual-product-time-specific demand shock. The parameter α is the disutility from price, β is the taste parameter. The term u_{ijt} denotes the flow utility minus the idiosyncratic error. Households consume the product at a household-specific constant rate c_i and receive utility $v(c_i)$ from consumption. Any units that are not consumed enter the household's inventory. An inventory of k_t units creates storage costs $C(k_t)$. Both the consumption rate and the storage costs do not depend on the composition of brands in stock, i.e. product differentiation occurs at the moment of purchase, not at the moment of consumption.⁸

In the following, we denote by s_{it} a vector that captures the current level of inventory and the current prices of all brand-size combinations at time t . The vector $\varepsilon_{it} = (\varepsilon_{ijt})_{j=1, \dots, J}$ stacks the household-time-specific shocks for all brands and pack sizes. Together, s_{it} and ε_{it} describe the so-called *state*, i.e. all the information that is relevant for a household's decision. In each

2004; Lu, 2016). We assume that during any given shopping trip, a household's choice set includes only the products at the visited store.

⁸ We borrow this model specification from Hendel and Nevo (2006a). It simplifies our estimation because it implies that not the brand composition of storage but only the total quantity in stock matters for consumption.

period, the household fully observes its state and forms expectations over future states.

We assume that households form rational expectations, i.e. their expectations are correct in equilibrium. This is a standard assumption in the literature (for a survey of the dynamic discrete-choice literature, see Aguirregabiria and Mira (2010)). In general, the rational-expectations assumption is made for identification purposes because observed choices may be explained by multiple specifications of expectations and beliefs. This assumption can be relaxed if the researcher has data on elicited beliefs (e.g Van der Klaauw and Wolpin, 2008; Arcidiacono, Hotz and Kang, 2012).

Households understand that future inventory depends on how much detergent they buy, how much they consume, and whether they run out of stock in the meantime. Consequently, inventory evolves according to

$$k_t = \max(k_{t-1} - c_i + q_{it}, 0), \quad (2)$$

where q_{it} is the amount of new detergent that enters storage if household i buys a pack of detergent in period t . While the evolution of inventory is deterministic, consumers face uncertainty with regards to future prices and future utility shocks. This implies that, in order to compute expected future values, we have to integrate over prices and utility shocks. Since this is computationally burdensome, we make two assumptions to reduce the dimensionality of the problem.

Firstly, we make assumptions on the error structure. Without any assumptions, we have to numerically integrate over the errors for every future value term. To avoid this, we make a simplifying assumption that was ini-

tially proposed by Rust (1987) and since then has become a standard in the literature:

Assumption 1 *The demand shocks ε_{ijt} are independently and identically extreme-value I distributed.*

This assumption is popular in the literature on dynamic discrete-choice models (see Aguirregabiria and Mira (2010)) because it delivers a closed-form solution of expectations of future utility conditional on the states and, thus, significantly reduces the computational burden.

Secondly, we aim to simplify consumer expectations of future prices. A model in which price expectations depend on the (infinitely long) history of previous prices is neither realistic nor tractable. Instead, the literature makes simplifying assumptions of various degrees, for example assuming that price expectations are conditional on a small number of price lags (Hendel and Nevo, 2006a) or the identity of the store (Seiler, 2013). In particular, we follow Pires (2016) and make the following assumption:

Assumption 2 *Prices are identically and independently distributed.*

This assumption implies that consumers do not condition their price expectations on past prices. Figures A.1 and A.2 show that prices indeed do not display clear promotional patterns. Therefore, it seems difficult for consumers to form sophisticated price expectations based on lagged prices. This is supported by the fact that, in general, consumers have been found to have poor knowledge of the prices that they pay; even seconds after selecting a product, only about 50% of shoppers are able to correctly recall its price (Dickson and Sawyer, 1990; Wakefield and Inman, 1993).

Furthermore, a number of important theoretical studies suggests that prices can be intertemporally independent since the optimal decision to conduct a sale may involve randomization (e.g. Varian, 1980; Salop and Stiglitz, 1982; Narasimhan, 1988; Raju, Srinivasan and Lal, 1990; Pesendorfer, 2002; Su, 2010).

In each period, the household makes a decision according to a decision rule r_{it} which assigns a decision – whether to purchase anything and, if so, which product – to each possible state $(s_{it}, \varepsilon_{it}) \in \mathcal{S}$. Let a policy be a sequence of decision rules for each point in time. For a given policy $\pi_i = (r_{i1}, r_{i2}, \dots)$, the discounted value of utilities is

$$V(\pi_i, s_{i0}, \varepsilon_{i0}) = E \left[\sum_{t=1}^{\infty} \tau^{t-1} U_{ijt}(s_{it}, \varepsilon_{it}, r_{it}(\pi_i, s_{it}, \varepsilon_{it})) | s_{i0}, \varepsilon_{i0} \right], \quad (3)$$

where τ is the discount factor. Since the time horizon remains infinite in every period and the transition probabilities are stationary, we do not have to find the optimal policy but only the optimal decision rule. For the same reason, we can drop the time subscript in the following. For notational simplicity we also drop the household subscript, i.e. the following solution of the dynamic discrete choice problem relates to a specific household i . Given that the set of decision rules is finite, there is an optimal decision rule $r^*(s, \varepsilon)$. The associated value function satisfies

$$\begin{aligned} V(s, \varepsilon) &= \max_{m \in \mathcal{R}} \{u(s, m) + \varepsilon(m) + \tau E(V(s', \varepsilon') | s, m)\} \\ &= \max_{m \in \mathcal{R}} \{u(s, m) + \varepsilon(m) + \tau \int V(s', \varepsilon') f(s' | m) g(\varepsilon') ds' d\varepsilon'\}, \end{aligned} \quad (4)$$

where \mathcal{R} denotes the set of all possible decision rules. s' and ε' denote s and ε in the next period, respectively. $f(\cdot)$ and $g(\cdot)$ are probability distribution

functions. We reformulate the problem in order to characterize this optimality condition as a function of s only. Before the demand shock ε is realized, the consumer expects

$$W(s) = \int V(s, \varepsilon) g(\varepsilon) d\varepsilon. \quad (5)$$

If we plug this expression into Equation (4), we obtain

$$V(s, \varepsilon) = \max_{m \in \mathcal{R}} \left\{ u(m, s) + \varepsilon(m) + \tau \int W(s') f(s'|s, m) ds' \right\}. \quad (6)$$

Finally, if we take the expectation on both sides of Equation 6, we obtain the integrated Bellman equation which depends only on s :

$$W(s) = \int \max_{m \in \mathcal{R}} \left\{ u(m, s) + \varepsilon(m) + \tau \int W(s') f(s'|s, m) ds' \right\} g(\varepsilon) d\varepsilon. \quad (7)$$

The right-hand side of this equation defines a contraction mapping $\Gamma : \mathcal{B} \rightarrow \mathcal{B}$. According to the Banach fixed-point theorem, the equation has a unique solution that must equal the expected value function $W = \Gamma(W)$. Once we have solved for $W(s)$, we can define the choice-specific value function

$$V_k(s) = u_k(s) + \tau \int W(s') f(s'|s, k) ds' + \varepsilon_k, \quad (8)$$

where the subscript k denotes the choice of k in the present period. Due to Assumption 1, the choice probabilities have the simple logit form. Choosing option k has the probability

$$\Pr(k|s) = \frac{\exp(V_k(s))}{\sum_{j=0}^J \exp(V_j(s))}. \quad (9)$$

In general, dynamic discrete-choice models with infinite-time horizons do not have closed-form solutions but must be solved numerically. We use the nested fixed-point algorithm proposed and made popular by Rust (1987). This is an iterative gradient search method to obtain the maximum-likelihood estimates of the structural parameters. It nests two loops: The inner loop solves the dynamic optimization problem for a given set of structural parameters, and the outer loop solves for the parameters that maximize the likelihood.

5 Estimation

5.1 Estimating Non-Observed Data

Prices In order to obtain purchase probabilities, we need to know what price a household *would have* paid for any other alternative at the store it was shopping at. To fill in these prices, we use the IRI store-level panel data. It records weekly prices for all products that were sold at least once during each week. However, not all prices can be matched like this because not every product is sold at every supermarket at least once a week. For 16,905 out of 35,132 brand-pack-size-store-week combinations we are able to retrieve the corresponding store prices. For the remaining cases – mostly smaller brands and exotic varieties – we replace the missing price with the corresponding weekly or monthly average price across all stores of the chain. We are able to fill in all missing prices this way.

Inventory Consumer inventory and consumption are unobserved in our data. However, if we knew both the consumption rate and the initial inventory of a household, we could use the observed purchases to construct the series of inventories. Because the consumption of laundry detergent is of a relatively

stable nature, we can compute the weekly consumption rate of a household as the total sum of purchases during the sample period divided by the total number of weeks (e.g. Erdem, Imai and Keane, 2003).

Then, we follow Seiler (2013) and assume that households start with zero inventory before the first observed purchase in our sample.⁹ The impact of the initial inventory will fade over time because consumption is not constant; instead, it drops to zero when stocks are depleted. In the estimation, we follow Hendel and Nevo (2006*a*) and Pires (2016) and drop the first ten observations of each household in order to mitigate the effect of the initial inventory. Finally, in order to reduce the state space in the inventory level dimension, we keep only those observations with inventories that are less or equal to 500 liquid ounces, i.e. about 14.785 liters. This assumption is not too restrictive, as we lose only 0.67% of observations.¹⁰ Figure 1 shows the final distribution of estimated inventories. There is a spike at zero inventory - 1.9% of the households face stock-outs. Higher levels of inventory are increasingly less likely to occur. In the estimation, we discretize inventory in units of 2 liquid ounces. We do so in order to make the problem tractable; it does not imply that serving size equals 2 ounces.


5.2 Reducing Dimensionality

Applications to real-world data often have the problem of dimensionality of the state space. Consumers (theoretically) face infinitely many possible

⁹ We run robustness checks of this and use an initial inventory equal to three, five, or ten times the consumption rate. Our results are robust to these changes.

¹⁰ If a household at one point in time has an inventory above 500 ounces, we drop only the observations during that time.

Figure 1: Estimated Inventory Distribution 2001-2004



inventory levels¹¹ and many possible brand-size combinations that are all offered at a variety of prices. This extends the state space so much that the estimation may become computationally infeasible. To reduce the state space, we adopt an approach proposed by Hendel and Nevo (2006a) that simplifies the estimation by decomposing the problem into a static discrete-choice part and a dynamic discrete-choice part. We further reduce the state space by dropping a selection of households that show outlier behavior.

Decomposition Approach In the following, we detail the steps of the decomposition approach of Hendel and Nevo (2006a). We break the problem

¹¹ We only have to track the total amount of inventory and not the brand-size composition of it. This is because in our model product differentiation takes place at the time of purchase and not at the time of consumption (e.g Hendel and Nevo, 2006a).

down into 1) optimal brand choice and 2) optimal size choice. This means that the probability of choosing brand b and size z can be written as the probability of choosing a pack size times the probability of choosing a brand given this pack size, i.e.

$$\Pr(d_t = (b, z)|p_t, k_t) = \Pr(d_t^{brand} = b|p_t, z_t, k_t) \cdot \Pr(d_t^{size} = z|p_t, k_t), \quad (10)$$

where k_t is the inventory carried in period t and p_t is the price, z_t denotes the pack size chosen in period t , and d_t^{brand} and d_t^{size} denote the choice of brand and size, respectively.

Firstly, we estimate a static discrete-choice model in which we restrict the choice of options to products of the same size as the one that was actually purchased. This estimation yields the static demand parameters, namely the price sensitivity and taste preferences. We then use the estimates from this first step to compute for each pack size the inclusive value, i.e. the expected utility from choosing that pack size:¹²

$$\omega_{zt} = \log \left(\sum_{b=1}^B \exp(\alpha p_{zbt} + x_{zbt}\beta) \right). \quad (11)$$

Since we collapse brand valuation into the inclusive value, the inclusive value serves as a size-specific adjusted price-index. In the following, we only need to track one inclusive value per pack size instead of tracking one price for each brand-size combination. This considerably reduces the state space. We use the empirical distribution of inclusive values as the probability distribution with which households face a certain inclusive value. This implies that consumer

¹²The inclusive value originates from the nested logit (e.g. McFadden, 1980) but has gained popularity in many applications with a nested structure.

expectations over future inclusive values are time-invariant; this assumption is analogous to Assumption 2 of the original problem (see Section 4). The utility function of the simplified dynamic problem is

$$\tilde{U}_{izt} = \begin{cases} v(c_i) - C(k_{it}) + \omega_{zt} + \varepsilon_{izt} & \text{if } z = 1, \dots, Z, \\ v(c_i) - C(k_{it}) + \varepsilon_{i0t} & \text{if } z = 0. \end{cases} \quad (12)$$

where $z = 0$ corresponds to the no-purchase option. This modified utility function corresponds to the utility function in Equation 1 and can be solved analogously.

Household Selection We drop households that make excessively many or few purchases because our model may not be able to properly describe their shopping behavior.¹³ More specifically, we drop households that, on average, make less than one purchase every six months or more than one purchase every two weeks. We conduct robustness checks with different thresholds and find that our results do not change substantially.

In order to reduce the state space and, thus, the computational burden, we limit households' choice sets. We assume that a household's choice set includes only the pack sizes that it buys over the total sample period, i.e. from 2001 to 2004. We further limit our analysis to households that consider only the two dominant pack sizes: 100-ounce packs (≈ 2.957 liters) and 200-ounce packs (≈ 5.915 liters) which have a combined market share of 68.95%. In other words, we drop a household from our analysis if it purchases another pack size at least once during our sample period.

¹³For example, households with extremely few detergent purchases may regularly visit a laundromat. Households with extremely frequent purchases may be buying not only for private consumption but also for resale or on behalf of someone else, such as a relative, friend, or organization.

The large majority of households (84.68%) buys only one pack on each trip, conditional on buying detergent (see Table A.3). This implies that, in the detergent market, consumers do not stockpile by increasing the number of purchased packs per trip. Instead, they keep buying only one pack per trip but increase their purchase frequency. In the following, we restrict our analysis to households that buy at most one pack per shopping trip. We drop all households that at some point in time buy more than one pack.

Our final sample contains 243 households over a span of 208 weeks. This is a very standard sample size in the literature (see for example Hendel and Nevo (2006a) with 218 households or Osborne (2010) with 299 households). Table A.4 details how much each selection step reduces the sample size. Table A.5 shows the summary statistics of the households in the final sample.

5.3 Identification

We informally discuss the empirical identification of both static and dynamic parameters.¹⁴ The parameters to be estimated are the price coefficient α , the taste parameter β , the consumption rate c_i , and the parametrized function $C(k)$. The utility of consumption $v(c_i)$ is not well-identified since households in our model always consume a constant amount of detergent, unless they face a stock-out. Therefore, we follow Seiler (2013) and define $v(c_i) = c_i$. Previous research establishes that the discount factor is difficult to identify (e.g. Rust, 1994; Magnac and Thesmar, 2002). Therefore, we follow the literature and set τ equal to 0.975, which corresponds to an annual interest rate of about 2.56%.

¹⁴For a formal discussion of identification in dynamic discrete-choice models, see Rust (1996) and Magnac and Thesmar (2002).

The identification of the static parameters is standard. Price sensitivity is identified by variation in prices. Brand and size preferences are identified by variation in shares of products. Heterogeneity in the price sensitivity is driven by variation in relevant household characteristics (here: family size) and heterogeneity in consumer response to promotions.

Identification of the dynamic parameters comes from variation in inter-purchase duration and the extent to which consumers exploit price cuts: Higher storage costs decrease the consumer’s ability to benefit from sales and decreases interpurchase duration. Imagine two households that face the same prices, have identical consumption rates, and buy the same total amount of detergent over a given period of time. One household buys only small packs, the other household buys only large packs. The household that buys small packs is then characterized by higher storage cost. For the storage costs, we assume the functional form $C(k_t) = \theta_1 k_t + \theta_2 k_t^2$. Doing so, we follow the literature that typically assumes a linear or quadratic cost function, for example Erdem, Imai and Keane (2003) in the market for ketchup, Osborne (2010) for canned tuna, as well as Hendel and Nevo (2006*a*), Seiler (2013), and Pires (2016) for liquid detergent.¹⁵

Jointly, the two sets of estimates from stages one and two determine consumer response. Consequently, both sets of estimates are crucial to the simulation of pricing counterfactuals. The static parameters from the first stage determine the short-term effects of how consumers substitute between brands. The storage cost parameters from the second stage determine long-term effects of how consumers substitute in quantities.

¹⁵ We are aware that storage costs may depend not only on the total amount of liquid but also on the total number of detergent bottles. However, modelling this would require either further data or strong assumptions on how households consume their stock, i.e. whether they consume bottle by bottle or spread consumption evenly over the bottles in stock.

6 Results

Table 2 shows a selection of parameters from the static estimation of brand choice conditional on chosen pack size. In all four specifications, price coefficients are negative and significant at the 1% level. We find that price sensitivity is heterogeneous; in particular, it tends to be larger for families with a larger income. We include brand fixed-effects and brand-size fixed-effects and find that they are almost all significant (see Table A.6 for the complete table of results). In the following, we continue with the estimates from Model 4, the richest specification.

Storage cost can be interpreted as the opportunity cost of the storage space. This cost is likely to decrease as total available storage space increases. Housing size should therefore be an important determinant of households' storage costs. However, we do not observe direct measures of housing size in our data. Instead, we use the existence of children as a proxy because families with children tend to live in larger homes. In our sample, 34.16% of the households have at least one child in their home (see Table A.5). We sort households into two types, those without children (type 1) and those with children (type 2).

Table 2: Estimation Results: Static Parameters

Variables	Model 1	Model 2	Model 3	Model 4
Price	-0.0560***	-0.482***	-0.525***	-0.436***
× Household Income				yes
Brand Dummies		yes		
Brand-Size Dummies			yes	yes
Observations	104,730	104,730	104,730	104,730

The symbols *, ** and *** denote significance at the 1%, 5%, and 10% level, respectively.

Table 3: Estimation Results: Dynamic Parameters

Variable	Mean	Standard Error	P-Value
TYPE 1: NO CHILDREN			
Storage Cost (linear)	-0.0182	0.0091	0.03421
Storage Cost (quadratic)	-1.9128e-05	0.0018	0.4958
TYPE 2: CHILDREN			
Storage Cost (linear)	-0.0179	0.0077	0.0119
Storage Cost (quadratic)	-1.9114e-05	0.0018	0.4958

Table 3 shows the results from the dynamic choice problem, split by household type. We find that storage costs increase in a linear way because the quadratic cost term is not significantly different from zero. Furthermore, households with children, i.e. households that are more likely to have larger homes, incur a lower storage cost. To provide an idea of the economic relevance of these estimates, consider the storage cost for a 100-ounce pack at zero inventory. It is given by 50 (units of 2 ounces) $\times 0.017925 = 0.895$ USD for a household with children and $50 \times 0.018208 = 0.91$ USD for a household without children. Buying a 100-ounce pack doubles this storage cost.

7 Counterfactual Promotion Policies

In this section, we simulate how counterfactual promotional pricing would affect consumer decisions and seller revenue.¹⁶ We look at two types of changes in pricing: Firstly, we study an increase in promotion length, i.e. existing promotions are extended in time. Secondly, we study an increase in promotion depth, i.e. promotional prices are further lowered. In both counterfactual simulations, consumers adjust their price expectations and face new dynamic programming problems. We compute counterfactual seller revenues as follows:

1. Change prices from observed p to counterfactual \tilde{p} .
2. Compute new inclusive values using the new prices and the previously estimated preference parameters $\hat{\alpha}$ and $\hat{\beta}$:

$$\tilde{\omega}_{zt} = \log \left(\sum_{b=1}^B \exp(\hat{\alpha}\tilde{p}_{zbt} + x_{zbt}\hat{\beta}) \right). \quad (13)$$

3. Compute empirical probability of each inclusive value.
4. Compute value function $\tilde{W}(s)$ given new inclusive values $\tilde{\omega}$ and new empirical probability of inclusive values.
5. Use $\tilde{W}(s)$ to simulate household choices of quantity. Note that for every household and in each period this choice is affected by a random shock.

¹⁶We do not study profits because we do not observe marginal costs. Theoretically, we could back out marginal costs from a structural model – in this case, a model with both a dynamic demand and a dynamic supply side. However, this is so technically challenging that we are not aware of any paper in the literature that does this. One alternative is to make assumptions on marginal costs. For example, Nair (2007) assumes a constant marginal cost of 12 USD per video game. But since the packaging and production of detergent are less standardized than those of CD-ROM disks, it is much more difficult to make assumptions on marginal costs of detergent. Therefore we decide to focus, like Osborne (2010), on revenues instead of profits.

Therefore, we simulate each household decision 1000 times, each time drawing a random shock from an extreme-value I distribution of shocks.

6. Simulate household brand choice conditional on quantity choice. Again, we use $n = 1000$ draws of the extreme-value I distributed error term. For each draw n , household i , and chosen pack size z , we compute :

$$\tilde{U}_{ibt n} = c_i - (\hat{\theta}_1 k_t + \hat{\theta}_2 k_t^2) + \tilde{\omega}_{zt} + \varepsilon_{ibt n} \quad \forall b \in B^z, \quad (14)$$

where B^z are all brands that are available in pack size z . A household i chooses brand b if $\tilde{U}_{ibt n} > \tilde{U}_{ilt n} \quad \forall l \in B^z$. We average household choices over $n = 1000$ draws.

7. Compute retailer revenue across products and time.

7.1 Promotion Length

We first simulate how a change in promotion length affects quantities and revenues. Since we observe prices and purchases on a weekly level, we can only vary promotion length in steps of one week. In the following, we simulate an extension of all promotions in our sample period (2001-2004) by one week. Table 4 displays how this affects purchase probabilities. Note that our counterfactuals are simulated for a time span of four years, i.e. they show long-term effects. We see that with longer promotions fewer households choose not to buy anything. This is because an extension of promotion length affects price expectations. When consumers expect lower product prices, the outside option of not buying anything becomes relatively less attractive. We find that, with longer promotions, consumers buy more packs of both sizes. In the baseline as well as in the extended-promotions scenario, type 2 households with children buy more of both pack sizes than childless type 1 households.

This is because type 2 households tend to be larger and therefore consume more detergent.

7.2 Promotion Depth

We now simulate a change in promotion depth. In order to ensure that we can compare counterfactuals in promotion length and depth, we compare counterfactuals of equal promotion value. The promotion value is defined as the sum of the price discounts across products. We find that the promotion value of prolonging all sales by one week corresponds to the promotion value of an additional 2.307% price cut on all sales prices. Table 4 displays the results for this increase in promotion depth. Similar to the previous counterfactual – and following the same logic – we find that increased promotion depth leads to fewer households choosing the outside option. Again, type 2 households generally buy more detergent than type 1 households.

Table 4: Promotion Depth vs. Length: Quantities

Model	Small Pack (%)	Large Pack (%)	No Pack (%)
TYPE 1: NO CHILDREN			
Baseline	0.1276	0.017702	0.8546
1 week longer	0.1281	0.017703	0.8541
2.307% off	0.1308	0.017209	0.8519
TYPE 2: CHILDREN			
Baseline	0.1344	0.019367	0.8461
1 week longer	0.1354	0.019733	0.8448
2.307% off	0.1392	0.018987	0.8419

7.3 Comparison: Length Vs. Depth

When we compare the two counterfactuals, we see that consumers react more strongly to an increase in promotion depth. We compute the revenues (see Table 5) and find that the elasticity of revenue¹⁷ is 0.0314 for a change in

¹⁷ The elasticity of revenue with respect to sales value is computed by $\frac{\Delta \text{revenue}}{\Delta \text{sales value}} \cdot \frac{\text{sales value}}{\text{revenue}}$.

promotion depth and 0.0082 for a change in promotion length, i.e. making a promotion deeper such that the promotion value increases by 1% will lead to a 3.14% increase in revenue. This effect is about four times larger than for an increase in promotion length. Our estimated revenue elasticities are of the same magnitude as those by Osborne (2010).

In the following, we discuss why a change in promotion depth is more effective than a change in promotion length. We have different dimensions of heterogeneity among households. Firstly, households carry different levels of inventory. Secondly, households differ in price sensitivity. Thirdly, we have two types of households with inherently different storage costs per unit. Lastly, households experience an idiosyncratic demand shock ε_{ijt} such that, even under identical conditions, a household may make different decisions on two days. These four dimensions of heterogeneity affect consumer behavior differently in the two pricing counterfactuals.

Table 5: Promotion Depth vs. Length: Revenues

Model	Revenue Type 1	Revenue Type 2	Total Revenue	Elasticity (of Revenue)
Baseline	20535.26	11305.69	31840.95	-
1 week longer	20687.36	11415.45	32102.81	0.0082
2.307% discount	21144.61	11729.06	32873.67	0.0314

When promotions are made deeper by 2.307%, price $p_{\text{promo},t}$ is replaced by $\tilde{p}_{\text{promo},t} = p_{\text{promo},t} \cdot (1 - 0.02307)$. Now, there are households that would not buy at price $p_{\text{promo},t}$ but would buy at price $\tilde{p}_{\text{promo},t}$. Note that both the level of household inventory and the idiosyncratic demand shock ε_{ijbt} are unaffected by the price change. Instead, the influx of new buyers is driven by households who were previously too price-sensitive and/or had too

high storage costs to make a purchase but can afford the purchase after the additional cut to sale prices.

Now look at the case in which promotions are extended by one week. Consider households that do not buy at price $p_{\text{promo},t}$ but do buy in the following week at price $p_{t+1} = p_{\text{promo},t}$. This influx of new buyers can be explained by two dimensions of heterogeneity and their interplay: A purchase may suddenly become attractive in period $t+1$ if the demand shock for a purchase is sufficiently large compared to the previous period and/or if household inventory i_{t+1} has sufficiently dropped due to mean-time consumption and, thus, lowered total storage costs.

Importantly, an increase in promotion depth is not generally more effective than an increase in promotion length: The relative effectiveness of the two promotion policies varies from market to market. In general, a change in promotion depth will be effective in a market in which storage costs and price sensitivities are relatively heterogeneous. A change in promotion length will be more effective in markets in which idiosyncratic shocks vary heavily and in which inventory can drastically change from one period to the next. The latter typically applies to markets that are strongly affected by demand shocks. Examples include the market for ice-cream (with weather-specific demand spikes), baking powder (rarely consumed on a daily basis, instead used irregularly for baking), and champagne (with demand spiking due to festive events).

8 Conclusion

In this paper, we target a core question of both marketing researchers and industry practitioners: How should one design promotions? In particular,

we are the first to investigate how the length of a promotion affects its effectiveness and how this compares to changes in promotion depth. We study this in the context of a storable-goods market that is characterized by forward-looking consumers who strategically stockpile. We develop a dynamic, structural model of consumer stockpiling and apply it to the U.S. market for laundry detergent.

We find that in this market, shorter but deeper promotions generate more revenue than longer, shallower promotions. However, this is not a general result; instead, marketers need to tailor promotion policies to product markets. Our results suggest that shorter, deeper promotions are generally preferable in markets with relatively heterogeneous storage costs and price sensitivities. Longer, shallower promotions are better suited for markets with strong demand shocks and unsteady consumption rates.

9 Bibliography

References

- Aguirregabiria, Victor, and Pedro Mira.** 2010. “Dynamic Discrete Choice Structural Models: A Survey.” *Journal of Econometrics*, 156(1): 38–67.
- Arcidiacono, Peter, V. Joseph Hotz, and Songman Kang.** 2012. “Modeling College Major Choices Using Elicited Measures of Expectations and Counterfactuals.” *Journal of Econometrics*, 166(1): 3–16.
- Bell, David R., Ganesh Iyer, and V. Padmanabhan.** 2002. “Price Competition Under Stockpiling and Flexible Consumption.” *Journal of Marketing Research*, 39(3): 292–303.
- Besanko, David, Jean-Pierre Dubé, and Sachin Gupta.** 2003. “Competitive Price Discrimination Strategies in a Vertical Channel Using Aggregate Retail Data.” *Management Science*, 49(9): 1121–1138.
- Besanko, David, Sachin Gupta, and Dipak Jain.** 1998. “Logit Demand Estimation under Competitive Pricing Behavior: An Equilibrium Framework.” *Management Science*, 44(11): 1533–1547.
- Bhardwaj, Bhanu, Olga Casabona, Joy Joseph, and Howard Shimmel.** 2016. “Drive Growth with Media Parity.” IRI Technical Paper.
- Blattberg, Robert C., Gary D. Eppen, and Joshua Lieberman.** 1981. “A Theoretical and Empirical Evaluation of Price Deals for Consumer Nondurables.” *Journal of Marketing*, 45(1): 116–129.

- Blattberg, Robert C., Richard Briesch, and Edward J. Fox.** 1995. "How Promotions Work." *Marketing Science*, 14(3): 122–132.
- Bronnenberg, Bart J., Michael W. Kruger, and Carl F. Mela.** 2008. "Database Paper-The IRI Marketing Data Set." *Marketing Science*, 27(4): 745–748.
- Ching, Andrew, and Matthew Osborne.** 2015. "Identification and Estimation of Forward-Looking Behavior: The Case of Consumer Stockpiling." Working Paper.
- Chintagunta, Pradeep K.** 2002. "Investigating Category Pricing Behavior at a Retail Chain." *Journal of Marketing Research*, 39(2): 141–154.
- Currin, Imran S., and Linda G. Schneider.** 1991. "A Taxonomy of Consumer Purchase Strategies in a Promotion Intensive Environment." *Marketing Science*, 10(2): 91–110.
- Dickson, Peter R., and Alan G. Sawyer.** 1990. "The Price Knowledge and Search of Supermarket Shoppers." *The Journal of Marketing*, 54(3): 42–53.
- Draganska, Michaela, and Dipak C. Jain.** 2006. "Consumer Preferences and Product-Line Pricing Strategies: An Empirical Analysis." *Marketing Science*, 25(2): 164–174.
- Erdem, Tülin, Susumu Imai, and Michael P. Keane.** 2003. "Brand and Quantity Choice Dynamics Under Price Uncertainty." *Quantitative Marketing and Economics*, 1(1): 5–64.

- Farley, John U.** 1968. “Dimensions of Supermarket Choice Patterns.” *Journal of Marketing Research*, 5(2): 206–208.
- Fotheringham, A. Stewart.** 1988. “Consumer Store Choice and Choice Set Definition.” *Marketing Science*, 7(3): 299–310.
- Hartmann, Wesley R., and Harikesh S. Nair.** 2010. “Retail Competition and the Dynamics of Demand for Tied Goods.” *Marketing Science*, 29(2): 366–386.
- Hendel, Igal, and Aviv Nevo.** 2006a. “Measuring the Implications of Sales and Consumer Inventory Behavior.” *Econometrica*, 74(6): 1637–1673.
- Hendel, Igal, and Aviv Nevo.** 2006b. “Sales and Consumer Inventory.” *RAND Journal of Economics*, 37(3): 543–562.
- Hendel, Igal, and Aviv Nevo.** 2013. “Intertemporal Price Discrimination in Storable Goods Markets.” *American Economic Review*, 103(7): 2722–2751.
- Hong, Pilky, R. Preston McAfee, and Ashish Nayyar.** 2002. “Equilibrium Price Dispersion with Consumer Inventories.” *Journal of Economic Theory*, 105(2): 503–517.
- Jeuland, Abel P., and Chakravarthi Narasimhan.** 1985. “Dealing-Temporary Price Cuts by Seller as a Buyer Discrimination Mechanism.” *Journal of Business*, 58(3): 295–308.
- Kadiyali, Vrinda.** 1996. “Entry, its Deterrence, and its Accommodation: A Study of the US Photographic Film Industry.” *RAND Journal of Economics*, 27(3): 452–478.

- Keane, Michael P., and Kenneth I. Wolpin.** 1997. "The Career Decisions of Young Men." *Journal of Political Economy*, 105(3): 473–522.
- Lal, Rajiv, and Carmen Matutes.** 1994. "Retail Pricing and Advertising Strategies." *Journal of Business*, 67(3): 345–370.
- Li, Jun, Nelson Granados, and Serguei Netessine.** 2014. "Are Consumers Strategic? Structural Estimation from the Air-Travel Industry." *Management Science*, 60(9): 2114–2137.
- Lu, Anna W.** 2016. "Inference of Consumer Consideration Sets." Working Paper.
- Magnac, Thierry, and David Thesmar.** 2002. "Identifying Dynamic Discrete Decision Processes." *Econometrica*, 70(2): 801–816.
- Mason, J. Barry, and Moris L. Mayer.** 1984. *Modern Retailing : Theory and Practice*. Business Publications, Plano, TX.
- McFadden, Daniel.** 1980. "Econometric Models for Probabilistic Choice Among Products." *Journal of Business*, 53(3): 13–29.
- Mela, Carl F., Kamel Jedidi, and Douglas Bowman.** 1998. "The Long-Term Impact of Promotions on Consumer Stockpiling Behavior." *Journal of Marketing Research*, 35(2): 250–262.
- Nair, Harikesh.** 2007. "Intertemporal Price Discrimination with Forward-Looking Consumers: Application to the US Market for Console Video-Games." *Quantitative Marketing and Economics*, 5(3): 239–292.
- Narasimhan, Chakravarthi.** 1988. "Competitive Promotional Strategies." *Journal of Business*, 61(4): 427–449.

- Neslin, Scott A., Caroline Henderson, and John Quelch.** 1985. "Consumer Promotions and the Acceleration of Product Purchases." *Marketing Science*, 4(2): 147–165.
- Osborne, Matthew.** 2010. "Frequency Versus Depth: How Changing the Temporal Process of Promotions Impacts Demand for a Storable Good." Working Paper.
- Pancras, Joseph, and K. Sudhir.** 2007. "Optimal Marketing Strategies for a Customer Data Intermediary." *Journal of Marketing Research*, 44(4): 560–578.
- Pesendorfer, Martin.** 2002. "Retail Sales: A Study of Pricing Behavior in Supermarkets." *Journal of Business*, 75(1): 33–66.
- Petruzzi, Nicholas C., and Maqbool Dada.** 1999. "Pricing and the Newsvendor Problem: A Review with Extensions." *Operations Research*, 47(2): 183–194.
- Pires, Tiago.** 2016. "Costly Search and Consideration Sets in Storable Goods Markets." *Quantitative Marketing and Economics*, 14(3): 157–193.
- Raju, Jagmohan S., Venkatesh Srinivasan, and Rajiv Lal.** 1990. "The Effects of Brand Loyalty on Competitive Price Promotional Strategies." *Management Science*, 36(3): 276–304.
- Rhee, Hongjai, and David R. Bell.** 2002. "The Inter-Store Mobility of Supermarket Shoppers." *Journal of Retailing*, 78(4): 225–237.
- Richards, Timothy J.** 2007. "A Nested Logit Model of Strategic Promotion." *Quantitative Marketing and Economics*, 5(1): 63–91.

- Rust, John.** 1987. “Optimal Replacement of GMC Bus Engines: An Empirical Model of Harold Zurcher.” *Econometrica*, 55(5): 999–1033.
- Rust, John.** 1994. “Structural Estimation of Markov Decision Processes.” In *Handbook of Econometrics.* , ed. Robert F. Engle and Daniel L. McFadden, 3081–3143. North-Holland, Amsterdam.
- Rust, John.** 1996. “Numerical Dynamic Programming in Economics.” In *Handbook of Computational Economics.* , ed. H. M. Amman, D. A. Kendrick and J. Rust, 619–729. North-Holland, Amsterdam.
- Salop, Steven, and Joseph E. Stiglitz.** 1982. “The Theory of Sales: A Simple Model of Equilibrium Price Dispersion with Identical Agents.” *American Economic Review*, 72(5): 1121–1130.
- Seiler, Stephan.** 2013. “The Impact of Search Costs on Consumer Behavior: A Dynamic Approach.” *Quantitative Marketing and Economics*, 11(2): 155–203.
- Smith, Howard.** 2004. “Supermarket Choice and Supermarket Competition in Market Equilibrium.” *Review of Economic Studies*, 71(1): 235–263.
- Sobel, Joel.** 1984. “The Timing of Sales.” *Review of Economic Studies*, 51(3): 353–368.
- Sudhir, K.** 2001. “Competitive Pricing Behavior in the Auto Market: A Structural Analysis.” *Marketing Science*, 20(1): 42–60.
- Sun, Baohong, Scott A. Neslin, and Kannan Srinivasan.** 2003. “Measuring the Impact of Promotions on Brand Switching When Consumers are Forward-Looking.” *Journal of Marketing Research*, 40(4): 389–405.

- Su, Xuanming.** 2010. “Intertemporal Pricing and Consumer Stockpiling.” *Operations Research*, 58(4): 1133–1147.
- Van der Klaauw, Wilbert, and Kenneth I. Wolpin.** 2008. “Social Security and the Retirement and Savings Behavior of Low-Income Households.” *Journal of Econometrics*, 145(1): 21–42.
- Varian, Hal R.** 1980. “A Model of Sales.” *American Economic Review*, 70(4): 651–659.
- Verboven, Frank.** 2002. “Quality-Based Price Discrimination and Tax Incidence: Evidence from Gasoline and Diesel Cars.” *RAND Journal of Economics*, 33(2): 275–297.
- Wakefield, Kirk L., and J. Jeffrey Inman.** 1993. “Who Are the Price Vigilantes? an Investigation of Differentiating Characteristics Influencing Price Information Processing.” *Journal of Retailing*, 69(2): 216–233.
- Whitin, Thomson M.** 1955. “Inventory Control and Price Theory.” *Management Science*, 2(1): 61–68.

10 Appendix

Table A.1: Duration Since Last Purchase and Till Next Purchase

Promotion	Obs	Mean	Std. Err.	Std. Dev.	95% Conf. Interval	
DURATION TILL NEXT PURCHASE						
0	51,564	7.6009	.0455	10.3280	7.5118 7.6900	
1	39,032	7.0751	.0516	10.1905	6.9740 7.1762	
Diff=mean(0)-mean(1)		.5258	.0689		.3908 .6608	
DURATION SINCE LAST PURCHASE						
0	51,328	7.2769	.04469	10.1257	7.1892 7.3645	
1	39,268	7.5019	.0528	10.4595	7.3984 7.6053	
Diff=mean(0)-mean(1)		-.2250	.0689		-.3560 -.0900	

Table A.2: Shares of Pack Sizes

Pack Size (ounces)	Regular Price	Promotional Price
50	20.13%	2.09 %
100	71.30 %	86.06%
200	8.57 %	11.85%
Total	100.00 %	100.00%

This table shows the distribution of different pack sizes during regular-price periods (column 2) and promotional periods (column 3). The entries in the second and third column are column percentages.

Table A.3: Number of Purchased Packs Per Shopping Trip

Number of Packs per Trip	Number of Trips	%
1	83204	84.678
2	11209	11.408
3	3,010	3.063
4	392	0.399
5	113	0.115
6	221	0.225
7	7	0.007
8	18	0.018
9	34	0.035
10	14	0.014
12	11	0.011
13	1	0.001
14	1	0.001
15	21	0.021
16	1	0.001
21	1	0.001
30	1	0.001
Total	98259	100

Table A.4: Household Selection and Sample Size

Treatment: Keep households that...	Number Households
	8289
... make >7 but <105 purchases in 2001-2004	4397
... buy only 100- and 200-ounce packs	534
... never buy more than one pack	243

Table A.5: Summary Statistics: Household Characteristics

Group	Number	%
HOUSEHOLD SIZE		
1 person	50	20.58
2 people	110	45.27
3 people	31	12.76
4 people	34	13.99
5 people	14	5.76
6 people	4	1.65
CHILDREN		
No children	160	65.84
At least one child	83	34.16
ANNUAL HOUSEHOLD INCOME (IN USD)		
< 9,999	6	2.47
10,000 to 11,999	9	3.70
12,000 to 14,999	10	4.12
15,000 to 19,999	13	5.35
20,000 to 24,999	22	9.05
25,000 to 34,999	20	8.23
35,000 to 44,999	35	14.40
45,000 to 54,999	27	11.11
55,000 to 64,999	19	7.82
65,000 to 74,999	27	11.11
75,000 to 99,999	33	13.58
≥100,000	22	9.05
Total	243	100




Figure A.1: Price Series 100-Ounce Packs

Price development over sample period across stores.




Figure A.2: Price Series 200-Ounce Packs
Price development over sample period across stores.

Table A.6: Estimation Results: Static Parameters (Full Table)

Variables	(1)	(2)	(3)	(4)
Price	-0.0560*** (0.00806)	-0.482*** (0.0139)	-0.525*** (0.0284)	-0.436*** (0.0547)
Brand 2		-2.125*** (0.283)		
Brand 3		-4.063*** (0.321)		
Brand 4		-0.207*** (0.0741)		
Brand 5		-3.836*** (0.381)		
Brand 6		-1.601*** (0.163)		
Brand 7		-0.971*** (0.0881)		
Brand 8		-2.703*** (0.158)		
Brand 9		-1.878*** (0.161)		
Brand 10		2.435*** (0.0626)		
Brand 11		-0.965*** (0.124)		
Brand 12		-5.576*** (0.709)		
Brand 13		-4.205*** (0.412)		
Brand-Size 2			-1.048*** (0.195)	-1.069*** (0.196)
Brand-Size 3			-2.183*** (0.289)	-2.184*** (0.289)
Brand-Size 4			-4.270*** (0.323)	-4.271*** (0.323)
Brand-Size 5			-0.243*** (0.0772)	-0.240*** (0.0775)
Brand-Size 6			-1.805*** (0.318)	-1.824*** (0.318)
Brand-Size 7			-3.990*** (0.382)	-3.990*** (0.382)
Brand-Size 8			-1.716*** (0.165)	-1.709*** (0.165)
Brand-Size 9			-0.989*** (0.0952)	-1.010*** (0.0957)
Brand-Size 10			-1.919*** (0.228)	-1.930*** (0.228)
Brand-Size 11			-2.767*** (0.166)	-2.792*** (0.166)
Brand-Size 12			-4.099*** (0.712)	-4.102*** (0.712)
Brand-Size 13			-1.923*** (0.166)	-1.921*** (0.166)
Brand-Size 14			-2.815*** (0.721)	-2.866*** (0.721)
Brand-Size 15			2.308*** (0.0790)	2.310*** (0.0792)
Brand-Size 16			3.137*** (0.219)	2.978*** (0.221)

continued on next page

continued from previous page

Brand-Size 17			-1.054***	-1.045***
			(0.130)	(0.130)
Brand-Size 18			-1.187**	-1.295**
			(0.533)	(0.534)
Brand-Size 19			-5.773***	-5.786***
			(0.710)	(0.710)
Brand-Size 20			-4.378***	-4.377***
			(0.412)	(0.412)
Income Group 2 × Price				-0.272***
				(0.0832)
Income Group 3 × Price				-0.181**
				(0.0823)
Income Group 4 × Price				-0.246***
				(0.0693)
Income Group 5 × Price				0.0293
				(0.0582)
Income Group 6 × Price				-0.203***
				(0.0632)
Income Group 7 × Price				-0.162***
				(0.0557)
Income Group 8 × Price				-0.0114
				(0.0557)
Income Group 9 × Price				-0.0132
				(0.0562)
Income Group 10 × Price				-0.0892
				(0.0561)
Income Group 11 × Price				-0.00125
				(0.0536)
Income Group 12 × Price				-0.0312
				(0.0584)
Observations	104,730	104,730	104,730	104,730

Standard errors are in parentheses. The symbols *, ** and *** denote significance at the 1%, 5%, and 10% level, respectively.