

A Simple Model of Demand Anticipation*

Igal Hendel Aviv Nevo

September 11, 2009

Abstract

In the presence of intertemporal substitution, static demand estimation yields biased estimates and fails to recover long run price responses. Our goal is to present a computationally simple way to estimate dynamic demand using aggregate data. Previous work on demand dynamics is computationally intensive and relies on (hard to obtain) household level data. We estimate the model using store level data on soft drinks and find: (i) a disparity between static and long run estimates of price responses, and (ii) heterogeneity consistent with sales being driven by discrimination motives. The model's simplicity allows us to compute mark-ups implied by dynamic pricing.

*We are grateful to seminar participants for helpful comments. This research was funded by a cooperative agreement between the USDA/ERS and Northwestern University, but the views expressed herein are those of the authors and do not necessarily reflect the views of the U.S. Department of Agriculture. Contact info: Department of Economics at Northwestern University. igal@northwestern.edu and nevo@northwestern.edu.

1 Introduction

Demand estimation plays a key role in many applied fields. A typical exercise is to estimate a demand system and use it to infer conduct, simulate the effects of a merger, evaluate a trade policy or compute cost pass-through.¹ While for the most part the demand models used are static, there is evidence that product durability or storability may generate dynamics, which could contaminate estimates. Focusing on storable products, a number of papers (Erdem, Imai and Keane, 2003, and Hendel and Nevo, 2006b) use household level data to structurally estimate consumer inventory models and simulate long run price responses. The computational burden and (household level) data requirement have limited the use of these dynamic demand models.

We propose an alternative model to incorporate demand dynamics. Our goal is to present a computationally simple way to estimate dynamic demand for storable products, or test for its presence, using aggregate, rather than household level, data. In many studies dynamics are not the essence. A test for the presence of dynamics may help rule them out. If dynamics are present their impact can be quantified by comparing static estimates to estimates from our model.

The model allows us to separate purchases for current consumption from purchases for future consumption. That way we can relate consumption and prices, to recover preferences (clean of storage decisions); and translate short run responses to prices, observed in the data, into long run reactions. The latter are the object of interest in most applications. The way we impute purchases for storage is quite simple but intuitive. Its advantage is that it does not require solving the value function of the consumer and the estimation is straightforward.

A key to the simplicity of the model is in the storage technology: consumers are assumed to be able to store for a pre-specified number of periods. This assumption simplifies the solution to the consumer's problem. The intuition of the model can best be demonstrated by a simple example. Suppose there is a single variety of a product with (1) prices that take on two values: a sale and a non-sale price; and (2) some consumers can store the product for one period (while others cannot store). Given these assumptions the model defines four states depending on the current and previous period price. The states determine whether there are purchases for storage or not, and whether consumption comes out of storage. Thus,

¹See, for example, Berry, Levinsohn, and Pakes (1995, 1999), Goldberg (1995), Hausman, Leonard and Zona (1994).

for each period there is a well defined demand curve, which is a function of the state, the (long run) demand parameters, and the fraction of consumers who can store. For example, consider a non-sale preceded by a non-sale. All consumers purchase for consumption. Since it is not a sale there is no incentive to store, and since the previous period was also not a sale none of the consumers have any inventory. Similarly, during a non-sale that follows a sale only consumers that cannot store will purchase, since those who can store bought in the previous period.

Using the illustrative example we see that the parameters of the model can be estimated in one of two ways. We could restrict attention to periods where dynamics do not play a role. In the above example these are two consecutive non-sale periods, as well as two consecutive periods with sales. Alternatively, we could look at all periods and use the model to predict purchases, accounting for the fraction of consumers who may be stockpiling.

The model we propose builds on the intuition of the example. We formalize the required assumptions and show that they simplify the state space. The problem remains dynamic, but easy to characterize. Solving the value function is not necessary. The approach can be seen as an alternative model of storage (where storage is based on periods of consumption as opposed to physical units) or as an approximation to a complex dynamic inventory decision. We present simulations to evaluate how well the approximation works.

Using the model we describe the biases generated by neglecting dynamics. Estimated own price responses are upward biased. The reason is that estimates reflect a weighted average of long run price responsiveness (dictated by the underlying preferences) and short run inventory (intertemporal) considerations. In addition, the model suggests that standard static estimation controls for the wrong price of the competing goods. The "effective" price (the actual opportunity cost of consumption) in this dynamic setup might differ from current price. We show that the consequences of using the wrong price is to bias the estimated cross price effect downward. For most antitrust applications the interest lies in long run elasticities. For example, in assessing unilateral effects in merger analysis, both biases, the upward bias in own price effect and downward bias in cross price effect, attenuate the computed unilateral effect.

We apply the model to weekly store-level data on purchases of 2-liter bottles of colas. The estimates using our model deliver, as expected, lower own and higher cross price responses

than the static estimates. The order of magnitude of the bias is comparable to what Hendel and Nevo (2006b) find when they estimate a dynamic inventory model for laundry detergents.

We discuss alternative approaches in Section 7. Alternatives to dealing with dynamics include aggregating the data from weekly to monthly and quarterly frequency, or approximating the missing inventory by including lagged prices/quantities (and computing long run effects using impulse response). We show these alternatives perform poorly, yielding negative cross price effects. We argue that the alternative methods also require a model to translate the estimated coefficients into preferences.

Another advantage of the simplicity of the model is to make the supply side tractable. In principle, the presence of demand dynamics makes the pricing problem quite difficult to solve. Especially so when there are multiple products sold by different sellers. In contrast, the demand framework we propose leads to a simple solution to the sellers' pricing problem.

Studying the supply side is interesting in its own right, but it is particularly important in many applications. Demand elasticities are typically used in conjunction with static first order conditions to infer market power. Demand dynamics render static first order conditions irrelevant. A supply framework consistent with demand dynamics is needed. We show that sellers' optimal behavior can still be characterized by first order conditions. Interestingly, the demand estimates show that consumers who store are significantly more price sensitive than non-storers, which is consistent with price discrimination being the motive behind sales. We use the estimated demand elasticities and the dynamic first order conditions to infer markups.

Section 2 presents motivating facts and reviews the literature. The model is presented in Section 3 and the estimation in Section 4. Section 5 presents an application to soft drinks. Extensions of the model are presented in Section 6.

2 Evidence of Demand Accumulation

2.1 Motivating Facts

Several papers (discussed in the next sub-section) have documented demand dynamics. We first look at typical scanner data for direct evidence on the relevance of intertemporal demand effects.

Figure 1 shows the price of a 2-liter bottle of Coke in a store over a year. The pattern is typical of pricing observed in scanner data: regular prices and occasional sales, with return to the regular price. Since soft-drinks are storable, pricing like this creates an incentive for consumers to anticipate purchases: buy during a sale for future consumption.




Figure 1: A typical pricing pattern

Quantity purchased shows evidence of demand accumulation. Table 1 displays the quantity of 2-liter bottles of Coke sold during sale and non-sale periods (we present the data in more detail below). During sales the quantity sold is significantly higher (623 versus 227, or 2.75 times more). More importantly, the quantity sold is lower if a sale was held in the previous week (399 versus 465, or 15 percent lower).

The impact of previous sales is even larger if we condition on whether or not there is a sale in the current period (532 versus 763, or 30 percent lower, if there is a sale and 199 versus 248, or 20 percent lower in non sale periods).

We interpret the simple patterns present in Table 1 as evidence that demand dynamics are important and that consumers' ability to store detaches consumption from purchases. Table 1 shows that purchases are linked to previous purchases, or at least, to previous prices.

Table 1: Quantity of 2-Liter Bottles of Coke Sold

	$S_{t-1} = 0$	$S_{t-1} = 1$	
$S_t = 0$	247.8	199.4	227.0
$S_t = 1$	763.4	531.9	622.6
	465.0	398.9	

Note: The table presents the average across 52 weeks and 729 stores of the number of 2-liter bottles of Coke sold during each period. As motivated below, a sale is defined as any price below 1 dollar.

2.2 Related Literature

Numerous papers in Economics and Marketing document demand dynamics, specifically, demand accumulation (see Blattberg and Neslin (1990) for a survey of the Marketing literature). Boizot et al. (2001) and Pesendorfer (2002) show that demand increases in the duration from previous sales. Hendel and Nevo (2006a) document demand accumulation and demand anticipation effects, namely, duration from previous purchase is shorter during sales, while duration to following purchase is longer for sale periods. Erdem, Imai and Keane (2003), and Hendel and Nevo (2006b) estimate structural models of consumer inventory behavior.

Several explanations have been proposed in the literature to why sellers offer temporary discounts. Varian (1980) and Salop and Stiglitz (1982) propose search based explanations which deliver mixed strategy equilibria, interpreted as sales. Sobel (1984), Conlisk Gerstner and Sobel (1984), Pesendorfer (2002), Narasimhan and Jeuland (1985) and Hong, McAfee and Nayyar (2002) present different models of intertemporal price discrimination. Our estimates show that sellers have incentives to intertemporally price discriminate, suggesting that sales are probably driven by discrimination motives.

3 The Model

In order to convey the main ideas we start with the simplest model of product differentiation with storage. We later show the model can be generalized in several dimensions. For example, the proposed estimation can be applied to more flexible demand systems, e.g., Berry, Levinsohn and Pakes (1995).

3.1 The Main Assumptions

Assume quadratic preferences:

$$U(q, m) = Aq - q' Bq + m \tag{1}$$

where $q = [q_1, q_2, \dots, q_N]$ is the vector of quantities consumed of the different varieties of the product (colas in our application) and m is the outside good. Absent storage, quadratic preferences lead to a linear demand system:

$$q_i^t(p) = \alpha - \beta_i p_i^t + \sum_j \gamma_{ij} p_j^t \tag{2}$$

In a multi-period set up with storage, consumers can anticipate purchases for future consumption. We make the following assumptions:

A1: two prices: sale and non-sale ($p^S < p^N$)

The transition between sale and non-sale periods could be random or deterministic. In the next section we discuss exactly how consumers form expectations regarding future prices.

Assumption A1 is stronger than needed. A key for our method to work is to be able to define a sale price, i.e., a price at which (some) consumers store for future consumption. When prices take on two values the definition of a sale is immediate. There can be several sale prices (as well as non-sale ones), all we need is to correctly define periods at which consumers store. Namely, periods of demand anticipation.

Next we make assumptions on the storage technology.

A2: storage is free

A3: inventory lasts for T periods (depreciates afterwards)

Initially we will focus on the $T = 1$ case, which makes the analysis more transparent. Allowing inventory to last for a single period reduces the state space: we only have to consider whether there was a sale in the previous period. In Section 3.4 we show how to modify the estimation for $T > 1$. Notice the data can guide what the relevant T is, for example, by examining the effect of lagged sales on the quantity purchased.

A4: a proportion ω of consumers do not store.

We assume that a proportion of customers do not have access to the storing technology. This assumption helps us explain why we see purchases in a non-sale period after a sale.

If everyone stored in the previous period our model would predict no purchases. With two prices this assumption is not very restrictive, but as we add more prices it will have bite since it assumes that the fraction of non-storers does not change with price.

In Section 3.4 we discuss the assumptions, their limitations, and possible generalizations.

3.2 Purchasing Patterns

We now characterize consumer behavior. To ease exposition we ignore discounting. The application involves weekly data, and therefore discounting does not play a big role.

Consumers who store, purchase for storage at p^S , and never store at p^N . When they store, they do so for one period. Thus, to predict consumer behavior we only need to define 4 events (or types of periods): a sale preceded by a sale (SS), a sale preceded by a non-sale (NS), a non-sale preceded by a sale (SN), and two non-sale periods (NN). We assume for now perfect price foresight, and later discuss (in section 6) behavior under rational price expectations.

Assume, for a moment, that only product i is stored. Given Assumptions A1-A4 and perfect foresight, product i aggregate purchases, $x_i(p)$, are:

$$x_i(p^t) = \begin{cases} q_i(p_i^t, p_{-i}^t) & p_i^{t-1} = p_i^t = p_i^N \\ \omega q_i(p_i^t, p_{-i}^t) & p_i^S = p_i^{t-1} < p_i^t = p_i^N \\ \omega q_i(p_i^t, p_{-i}^t) + (1 - \omega)(q_i(p_i^t, p_{-i}^t) + q_i(p_i^t, p_{-i}^{t+1})) & p_i^N = p_i^{t-1} > p_i^t = p_i^S \\ \omega q_i(p_i^t, p_{-i}^t) + (1 - \omega)q_i(p_i^t, p_{-i}^{t+1}) & p_i^{t-1} = p_i^t = p_i^S \end{cases} \quad \text{if} \quad (3)$$

where $q_i(p_i^t, p_{-i}^t)$ is the long run demand (we are after). Here we assume identical preferences for all consumers, but it is easy to allow storers and non-storers to have different demands, which we do in the application.

The four rows represent events NN , SN , NS , and SS , respectively. For the fraction ω of non-storers demand is equal to consumption and therefore in all states contributes $\omega q_i(p_i^t, p_{-i}^t)$ to aggregate demand. For storers, demand is dictated by the model. At high prices there are no incentives to store, in which case purchases equal either: consumption, defined by the long run demand, or zero, if there was a sale in the previous period (i.e., in SN consumption is out of storage). During sales preceded by a non-sale period purchases include current consumption as well as inventory. During periods of sale preceded by a sale,

current consumption comes from stored units, so purchases are for future consumption only, and the contribution to aggregate demand is $(1 - \omega)q_i(p_i^t, p_{-i}^{t+1})$.²

Notice the difference in the second argument of the anticipated purchases relative to purchases for current consumption (i.e., during NN). Purchases for future consumption take into account the expected consumption of products $-i$. Here, for simplicity, we assume perfect foresight of future prices and therefore future demand is a function of p_{-i}^{t+1} . Alternatively, under rational price expectations the consumer would purchase based on the expected future price (see Section 6).

The key observation, regardless of price expectations, is the following: if a product is currently on sale we know its effective next period price is p^S (since the product will be stored today for consumption tomorrow). In other words, the way to incorporate the dynamics dictated by storage is to consider the effective cost (or price) of consumption, which does not necessarily coincide with current price. In an inventory model, the effective or shadow price is a complicated creature that requires solving the value function. In our framework effective prices is just the minimum of current and previous prices.

When all products are storable, the case we consider from here onward, accounting for the storability is no more complicated. We just need to control for the effective cross price. For example, consider the event NN (product i is not on sale at t or at $t - 1$) and assume that product $-i$ was on sale at $t - 1$ (but is not on sale at t). The demand from consumers who store is $q_i(p_i^t, pe_{-i}^t)$ instead of $q_i(p_i^t, p_{-i}^t)$, where pe_{-i}^t is the effective price, in this case p_{-i}^{t-1} . A similar adjustment is needed in all other states.

An important implication is that current prices of other products are the wrong prices to control for in the estimation. Controlling for current price generates a bias in the estimated cross price effect.

3.3 Predicted Biases

We now explore the model's implications for biases that may arise from neglecting dynamics. Suppose we observe several price regimes, with constant prices within each regime. Since

²When a sale follows a sale under our assumptions the consumer is indifferent between storing for consumption next period or buying now. We break the tie by assuming she buys now. This is justified with even a small amount of uncertainty about coming to the store in the next period or about future prices.

prices are constant within each regime, there is no reason to store and therefore the difference in purchases (and consumption) across regimes helps recover preference parameters β and γ .

Instead of observing long lasting price differences we may observe high frequency price changes, like in the case of sales. Consider for simplicity just three periods, and suppose product 1's price decreases during the second period: $p_1^1 = p_1^3 = p_1^N$ and $p_1^S = p_1^2 < p_1^1$, while product 2's price remains constant at p_2 . Denote by $\Delta p_1 = p_1^2 - p_1^1 = p_1^S - p_1^N < 0$.

Since storing is free, consumers (who store) will purchase all of period 3 consumption, $q_1(p_1^S, p_2)$, in period 2. Notice the effective price of product 1 in period 3 is actually the lowest of periods 2 and 3 prices, $\min\{p_1^2, p_1^3\} = p_1^S$. The consumer can time her purchases to minimize expenses. In this case, period 3 consumption is determined by p_1^2 .

Quantities purchased by a storing consumer over the three periods (according to equation 3) are:

	$t = 1$	$t = 2$	$t = 3$
$p_1^t =$	p^N	p^S	p^N
$x_1 =$	q_1^1	$2(q_1^1 - \beta \Delta p_1)$	0
$x_2 =$	q_2^1	$q_2^1 + \gamma \Delta p_1$	$q_2^1 + \gamma \Delta p_1$

where $q_1^1 = q_1(p_1^N, p_2)$ and $q_2^1 = q_2(p_1^S, p_2)$. Should we estimate demand statically we would estimate the following price effects:

Own price

$$\| -2\beta + \frac{3q_1^1}{2\Delta p_1} \| > \| -\beta \|$$

Cross price reaction

$$\frac{\gamma}{2} < \gamma$$

There is an over estimation of long run own price reaction and underestimation of long run cross price effects.

The over estimation of own price effects is caused by attributing the response to a temporary price reduction as an increase in consumption, while the consumer is purchasing for storage. In addition, the price increase in period 3 coincides with a decline in purchases, which is also misconstrued as a decline in consumption.

While it is natural to expect an overestimation of own price responses, the impact of dynamics on cross prices responses is more delicate. Previous work has documented the

effect on cross price responses, but did not show the expected bias theoretically. The model predicts cross price effects are understated. In period 3 the observed and effective prices differ. The effective price, which dictates consumption of good 1, is the period 2 purchase price. In the estimation we would instead interpret the price increase (observed in period 3), which is not accompanied by an increase in purchases of product 2, as lack of cross price reactions.

3.4 Discussion of the Main Assumptions

The model greatly simplifies the consumer problem. We now discuss the key assumptions that deliver the simplicity. A good point of comparison is the dynamic inventory model where the consumer maximizes the discounted expected flow of utility from consumption minus the price of the product and the cost of holding inventory. If we assume that prices follow a first order Markov process then the state variables are current prices and a vector of inventories.³ We will refer to this model as the (standard) inventory model.

Assumption A1 simplifies the determinants of demand anticipation. If current price is high then future prices will be (weakly) lower and there is no incentive to store. If instead current price is low then future prices will be (weakly) higher and consumers buy for inventory. The two-price support assumption is restrictive, and stronger than needed. What is really needed is that the periods of accumulation are well defined. The same logic applies even with a more general price process as long as we can properly define a sale (i.e., periods of demand anticipation). The simplicity comes with a cost. We need to define a sale, which is easy with 2 prices, but can be more difficult in general. In the inventory model a definition of a sale is not necessary. Instead consumption and purchases are determined by current price, and the other state variables.

Assumptions A2 and A3 define the storage technology. We assume that consumers can store for a pre-specified period of time. Taken literally this assumption fits perishable products like milk or yogurt. For more storable products the relevant constraint on advance purchases is probably storage space at home or transportation from the store. The assumption is quite convenient. First, it simplifies the dynamics. Under a capacity constraint, as in the inventory model, the researcher and the consumer need to keep track of how much is left

³The state space can be simplified to include a smaller number of inventories. See Erdem, Imai and Keane (2003) or Hendel and Nevo (2006b).

in storage in different states. Second, it helps detach the storage decision of different products. If the binding constraint is storage capacity, storing one product restrains the ability to store the other. Under A3 the optimal storage of one product depends on the effective price of the other product (as can be seen in equation 3), but not on the quantity stored. Such a link between products is quite easy to incorporate, we just control for effective prices of alternative products as oppose to their observed prices.

So far we assumed $T = 1$. We can increase the number of periods that consumers can store. If $T > 1$ predicted purchases depend on a larger number of lags. For example, if $T = 2$ we have to condition on 2 lags rather than one. As we see in equation 3, for $T = 1$ purchases in state NN are dictated by long-run needs, $q()$. When $T = 2$ we need to distinguish between NNN , where everyone buys according to the long-run demand $q()$, and SNN , where only non-storers buy. For $T = 2$ there are 8 states. Since some of the states involve similar predicted purchases (as shown in the Appendix) the model is easy to generalize to $T = 2$.

The number of states (in the state space) does not burden the estimation. But the more states the more tedious accounting is needed. In order to incorporate a longer horizon but keep the accounting to a minimum, we experimented with the following approximation:

$$x(p^t) = \begin{cases} q(p^t) & NN \\ \omega q(p^t) & SN \\ \omega + T(1 - \omega)q(p^t) & NS \\ q(p^t) & SS \end{cases} \quad \text{if} \quad (4)$$

The shortcut is based on equation 3 and predicted purchases for $T = 2$ (described in the Appendix). Some states (like SN and SS) are unaffected by longer histories, while others are. We are back to 4 states, knowing that predictions in states NN and NS are not exact. The shortcut seems to perform reasonably well.

Assumption A4 holds the fraction of storers constant. If Assumption A1 holds in the sample this is not restrictive for the purpose of estimation, but if the price takes on additional values this assumption has some bite. Even if in the data there are only two prices, this assumption might be problematic for counterfactuals. The fraction of consumers who store might be a function of price. The current model does not allow it, but in principle we could extend the model to allow ω to vary with prices.

Section 6 discusses extensions like discrete choice models instead of linear demand and rational expectations about future prices.

3.5 Simulation Results

The model can be taken literally or as an approximation to a standard inventory model. In this section we apply the proposed estimation to data generated from an inventory model. The goal is to assess how well the proposed estimation does in recovering preferences when applied to data generated from an inventory model.

The data was generated by computing the optimal (dynamic) behavior of a consumer with quadratic preference, linear inventory costs, and facing stochastic prices. We assume there are two prices, sales occur with probability 0.2, half (0.5) of the consumers store according to the standard inventory model while the other half have a static demand, and the discount factor is 0.995. The preference parameters used to generate the data imply $\beta = 4$. We compute the consumer's value function and optimal policy, and then simulate purchases for 100, 200, and 500 periods for different storage cost levels. We perform 1,000 repetitions of each estimation.

Figure 2 shows consumption and storage predicted by the inventory model as a function of the storage parameter. The different storage costs trace situations where storage is quite high (over twice the flow consumption) to no storage.




Figure 2: Optimal Dynamic Behavior as a Function of Storage Costs

Figure 3 displays the percent bias in the price coefficient for OLS estimates and our fix assuming $T = 1$ and $T = 2$. For moderate levels of anticipated purchases the proposed fix does well. On the other hand, OLS shows substantial bias, about 60%, even for modest levels of storage. For very low storage costs all estimates overstate price responses. However, while the $T = 1$ fix is off the mark by 40% the OLS estimate is over 160% off. As expected the $T = 2$ fix does better than the $T = 1$ fix for very low storage costs.




Figure 3: Percent Bias in Estimated Slope Parameter

Table 2 presents mean estimates and mean squared error of the different estimates by storage cost. It shows that the $T = 1$ fix does best when the average storage (conditional on holding storage) is in the ballpark of one period of consumption (i.e., $c = 0.29$ and 0.38), while the $T = 2$ fix is closest to target for $c = 0.17$ when the average storage is about twice the flow consumption. Both uniformly dominate OLS, unless storage is absent.

Table 2: Monte Carlo Simulations

Simulated Data			Mean β			MSE β		
			OLS	$T = 1$	$T = 2$	OLS	$T = 1$	$T = 2$
c	Consumption	Storage	N=100					
0.08	4.80	16.20	10.57	5.82	4.73	43.98	3.59	0.66
0.17	4.19	9.31	8.36	4.89	4.07	19.26	0.92	0.13
0.29	3.64	5.80	6.50	4.08	4.30	6.39	0.14	0.35
0.38	3.54	5.05	6.16	3.91	4.37	4.75	0.15	0.38
0.50	2.97	0	3.99	4.00	3.99	0.05	0.20	0.19
N=200								
0.08	4.80	16.20	10.50	5.71	4.68	42.59	3.02	0.52
0.17	4.19	9.31	8.35	4.87	4.07	19.06	0.83	0.07
0.29	3.64	5.80	6.50	4.07	4.29	6.32	0.07	0.21
0.38	3.54	5.05	6.14	3.90	4.35	4.63	0.08	0.23
0.50	2.97	0	4.00	4.01	3.99	0.03	0.09	0.09
N=500								
0.08	4.80	16.20	10.48	5.66	4.66	42.08	2.79	0.46
0.17	4.19	9.31	8.32	4.84	4.05	18.69	0.73	0.03
0.29	3.64	5.80	6.47	4.05	4.28	6.13	0.03	0.12
0.38	3.54	5.05	6.13	3.89	4.33	4.52	0.04	0.15
0.50	2.97	0	4.00	4.00	4.00	0.01	0.04	0.03

Note: Means and mean squared error of estimates of the slope coefficient, beta, computed based on 1,000 repetitions of each estimation. The data was generated using with a slope parameter of 4. The storage level is the average storage conditional on being positive, as oppose to Figure 2 that shows the unconditional average storage.

4 Identification and Estimation

4.1 How Do We Recover Preferences?

Before presenting the estimation we discuss intuitively how the model helps recover preferences. We offer two approaches, both are part of the full estimation, but discussing them

separately helps clarify what variation in the data identifies the parameters. The first approach is based on events without storage, while the second approach imputes storage and purges it from purchases.

For simplicity, assume a single product in which equation 3 suggests that during NN and SS demand is given by $q(p^t)$, while during SN demand is scaled down by ω and during NS it is scaled up by $2 - \omega$. This suggests two different ways to recover the model's parameters from the data. We will refer to the first as "timing" restrictions. According to the model during sale periods that follow a sale (event SS) purchases equal consumption: $x(p) = q(p)$. Basically, after purchasing for storage, the pantry is filled, consumers (whether they are a storer or not) purchase for a single consumption event. Since both NN and SS events involve purchases dictated by $q(p)$ we can rely on them to estimate preferences. Price variation across these states, across different stores during these states or within these states (if there are more than 2 prices), can identify long-run responses.

A different way to map purchases into preferences is to take advantage of the data from all periods but use the model to adjust the predicted purchases to account for storage. We will refer to the additional restrictions as "accounting" restrictions. For example, during SN we need to scale down demand because only non-storers purchase, while during NS we have to scale up purchases due to storing. This approach is more efficient, since it uses all the data, but it also adds additional parameters and imposes Assumption A4.

We note that the model is over-identified. For example, we can recover ω by looking at the ratio of purchases during SN to purchases during NN , or by looking at 2 minus the ratio of purchases during NS to purchases during SS . In principle we can use the additional degrees of over-identification to enrich the model somewhat.

To demonstrate how the different restrictions work we can use the numbers in Table 1 to recover the demand parameters. As a benchmark, we note that the static estimate of the slope coefficient is $\hat{\beta}^{Static} = \frac{x^S - x^N}{\Delta p} = \frac{623 - 227}{0.4} = 988$, where $\Delta p = 0.4$ is the price difference between sale and non-sale periods. Estimating the same slope using only the timing restrictions yields $\hat{\beta}^{Time} = \frac{x^{SS} - x^{NN}}{p^{SS} - p^{NN}} = \frac{532 - 248}{0.4} = 710$.

Since the model is over-identified, there are several ways to impose the accounting restrictions. One way is to use only the information from NN , SS , and SN and recover $\omega = x^{NN}/x^{SN} = 0.8$ and $\hat{\beta} = \frac{\omega x^{SS} - x^{SN}}{\omega \Delta p} = 708$. Another way is to use the information from NN , SS , and NS , which implies $\omega = 2 - x^{NS}/x^{SS} = 0.57$ and $\hat{\beta} = \frac{x^{NS} - (2 - \omega)x^{NN}}{(2 - \omega)\Delta p} = 713$. More

efficient estimators, coming in the next section, will combine all this information and further control for differences across stores, prices of other products, and promotional activities.

Notice that both restrictions render a lower price sensitivity than the one implied by the static estimates.

4.2 Estimation

We follow the two strategies described above to estimate preferences. The first strategy uses data only from the *NN* and *SS* periods, which involve no storage. The second approach uses data from all periods, and is therefore more efficient, but it requires non-linear estimation. Linear estimation allows us to recover all the parameters of the model, except the fraction of consumers who store. To obtain the exact estimating equations we combine equations 2 and 3, and allow for a panel structure (that exists in the data we use below). To account for the store level fixed effects we de-mean the data. For prices this is straightforward. For quantities we have to account for the re-scaling in different regimes. We show in the Appendix how to modify the estimating equation to account for this re-scaling.

We estimate all the parameters by least squares, linear or non-linear depending on the equation. In principle, we could use instrumental variables to allow for correlation between prices and the econometric error term. However, we do not think correlation between prices and the error term is a major concern in the example below.

5 An Empirical Application: Demand for Colas

The average numbers (from Table 1) used in the previous section do not exploit price variation across stores, or within a regime (for a given store). They also neglect to properly control for the prices of substitute products.⁴ We now estimate the model using all the events adding these additional controls.

⁴This is a serious concern since promotions of Coke and Pepsi are probably correlated, thus, a low Coke price may be also reflecting a high price of the closest substitute, thus contaminating the price reactions we infer.

5.1 Data

The data we use was collected by Nielsen and it includes store-level weekly observations of prices and quantity sold. The data set includes information at 729 stores that belong to 8 different chains throughout the Northeast, for the 52 weeks of 2004. We focus on 2-liter bottles of Coke, Pepsi and store brands, which have a combined market share of over 95 percent of the market.

There is substantial variation in prices over time and across chains. A full set of week dummy variables explains approximately 20 percent of the variation in the price in either Coke or Pepsi, while a full set of chain dummy variables explains less than 12 percent of the variation.⁵ On the other hand, a set of chain-week dummy variables explains roughly 80 percent of the variation in price. Suggesting similarity in pricing across stores of the same chain (in a given week), but prices across chains look quite different. As a first approximation it seems that all chains charge a single price each week. However, three of the chains appear to define the week differently than Nielsen. This results in a change in price mid week, which implies that in many weeks we do not observe the actual price charged just a quantity weighted average. In principle we could try to impute the missing prices. Since this is orthogonal to our main point we drop these chains.

We need a definition of a sale, or more precisely, we need to identify periods of advance purchases. Figure 4 displays the distribution of the price of Coke in the five chains we examine below. The distribution seems to have a break at a price of one dollar, which we use as the threshold to define a sale. Any price below a dollar is considered a sale, namely, a price at which storers purchase for future consumption. This is an arbitrary definition. A more flexible definition may allow for chain specific thresholds, or perhaps moving thresholds over time. For the moment we prefer to err on the side of simplicity. Using this definition we find that approximately 30 (36) percent of the observations are defined as a sale for Coke (Pepsi). Interestingly, sales are somewhat asynchronized with only 7 percent of the observations exhibiting both Pepsi and Coke on sale (compared to a 10.5 percent predicted if the sales were independent).

⁵These statistics are based on the whole sample, while the numbers in Table 2 below are based on only five chains as we explain next.




Figure 4: The Distribution of the Price of Coke

For the analysis below we use 24,674 observations from five chains. The descriptive statistics for the key variables are presented in Table 3.

Table 3: Descriptive Statistics

Variable	Mean	Std	% of variance explained by:		
			chain	week	chain-week
Q_{Coke}	446.2	553.2	5.6	20.4	52.5
Q_{Pepsi}	446.0	597.8	2.8	24.4	46.7
P_{Coke}	1.25	0.25	7.1	29.7	79.9
P_{Pepsi}	1.19	0.23	7.5	30.7	79.8
Coke Sale	0.30	0.46	6.4	30.0	86.6
Pepsi Sale	0.36	0.48	9.3	29.2	89.0

Note: Based on 24,674 observations for five chains, as explained in the text. As sale is defined as any price below one dollar.

5.2 Results

The estimation results are presented in Tables 4 and 5. All columns present least squares estimates of linear demand. The dependent variable is the number of 2-liter bottles of Coke or Pepsi sold in a week in a particular store. All the columns include the price of the store brand and store fixed-effects. The first column displays estimates from a static model, with store fixed effects. Column 2 presents our model estimated using the timing restriction only (namely, using the sub-sample with the events in which the model predicts no storage). Columns 3 and 4 present estimates of the full model.

The difference between columns 3 and 4 is that in column 3 we control for the current price of the competing products. Column 4 instead controls for the effective price. According to the model the effective price faced by a storer is the minimum of current and last period price.⁶

In column 5 we present the results from a model that allows different price sensitivity between storers and non-storers. Finally, in column 6 we present results where we replace the perfect foresight assumption with a rational expectation assumptions. We discuss this model and the results in the next section.

All the estimates from our model suggest lower (in absolute value) own price effects and higher cross price effects, for both Coke and Pepsi. The estimated proportion of consumers who do not stockpile is around half the population, and is slightly higher for Coke. Consistent with this estimate, the differences between the static and dynamic estimates are larger for Pepsi than for Coke.

The estimates in column 3 are of no interest on their own. According to the model, the cross prices controls are incorrect. The model prescribes the use of past prices during periods preceded by a sale (i.e., the effective price that dictates the consumption of storers is the lagged price). However, if the model is irrelevant, or demand dynamics absent, as we move from column 3 to column 4 we would be introducing noise in the price of the competing product. As such we would expect the coefficient of Pepsi in the Coke equation (and Coke's in the Pepsi equation) to be lower, due to measurement error (assuming the introduced noise will generate classical measurement error). Interestingly, both cross price

⁶For the non-storer the current price is the effective one. Thus, because of linearity of the demand curve, the aggregate effective price is the weighted average of the prices faced by storers and non-storer; weighted by the proportion of each type of buyer in the population. The price is recomputed as the estimation algorithm searches for the optimal ω .

effects increase substantially as we replace current price by effective price. Suggesting the latter is the correct control, and that indeed dynamics are present.

Table 4: Demand for Coke

	FE	Timing Only	All Restrictions		Different slopes	Rational Exp
	(1)	(2)	(3)	(4)	(5)	(6)
P_{Coke}	-1428.2 (11.1)	-743.4 (11.8)	-967.4 (11.2)	-938.9 (11.1)	-522.4 (35.1)	-767.8 (8.8)
P_{Pepsi}	66.5 (11.9)	191.6 (10.9)	82.8 (11.0)	150.1 (11.4)	-73.1 (36.4)	195.4 (12.3)
P_{Coke} storers					-1273.1 (16.7)	
P_{Pepsi} storers					145.01 (21.3)	
ω (fraction non- storers)			0.53 (0.01)	0.53 (0.01)	0.35 (0.01)	0.57 (0.01)
Cross price Corrections			No	Yes	Yes	Yes

Note: All estimates are from least squares regressions. The dependent variable is the quantity of Coke sold at a store in a week. The regression in column (1) includes store fixed effects. The regression in column (2) is the same as in column (1) but uses only the NN and SS periods. The regressions of columns (3)-(4) impose all the restrictions of the model using the actual and effective price. Column (5) allows for different slopes for consumers who store and those that do not. Column (6) assumes rational expectations rather than perfect foresight. Standard errors are reported in pharenthesis.

Table 5: Demand for Pepsi

	FE	Timing Only	All Restrictions	Different Slopes	Rational Exp	
	(1)	(2)	(3)	(4)	(5)	(6)
P_{Coke}	-20.9 (12.2)	71.8 (11.0)	62.8 (10.5)	106.6 (10.5)	-246.2 (26.9)	140.0 (12.1)
P_{Pepsi}	-1671.3 (13.1)	-994.0 (15.8)	-1016.5 (11.6)	-996.9 (11.6)	-341.3 (29.3)	-762.3 (8.9)
P_{Coke} storers					216.7 (13.3)	
P_{Pepsi} storers					-1255.0 (14.8)	
ω (fraction non-storers)			0.44 (0.01)	0.44 (0.01)	0.28 (0.01)	0.47 (0.01)
Cross Price Corrections			No	Yes	Yes	Yes

All estimates are from least squares regressions. The dependent variable is the quantity of Pepsi sold at a store in a week. The regression in column (1) includes store fixed effects. The regression in column (2) is the same as in column (1) but uses only the NN and SS periods. Column (5) allows for different slopes for consumers who store and those that do not. Column (6) assumes rational expectations rather than perfect foresight. Standard errors are reported in parenthesis.

The findings are consistent with the biases predicted in Section 3. Both proposed fixes lower own price effects by purging purchases for storage. At the same time the fixes raise cross price effects, by accounting for effective price variation (i.e., eliminating measurement problems in prices of substitute products).

In column 5 we allow for different price sensitivity between consumers who store and those who do not. We find that storers are substantially more price sensitive. This is consistent with price discrimination as being a motivation for the existence of sales. We return to this below.

5.3 Implications of the Estimates

The bias in the estimated elasticities has implications for almost all applications of the estimates. One common use of demand elasticities in the IO and trade literature is to plug

them into a first order condition and recover margins and implied marginal costs. For single product firms the price-cost margins are equal to the inverse of the price elasticity.

Using the static estimates from column 1 of Tables 4 and 5 implies a margin of 25 percent for Coke and 22 percent for Pepsi, which translates into a marginal cost of 0.94 for Coke and 0.92 for Pepsi.⁷ Repeating this computation but using instead the estimates from column 4, i.e., those estimated using all the restrictions of the model, the implied margins are 38 percent and the implied marginal costs are 0.77 for Coke and 0.73 for Pepsi.

A natural concern with this sort of calculation is that it ignores the dynamic aspect of demand. It would be unsatisfactory to resort to a static first order condition after we argued demand is dynamic.

One rationale for sales is to price discriminate between consumers who store and those who do not. The estimates in column 5 allow us to check this theory. Indeed, the estimates suggest that consumers who store are significantly more price sensitive than consumers who do not. So the estimated preferences render sales profitable, as a way to price discriminate.

Fully solving the dynamic pricing problem is beyond the scope of this paper. However, the case of a single product monopolist who commits to prices is simple to characterize given the demand structure we propose. Optimal behavior is characterized by two first order conditions, not that different from the static ones. The seller profits increase by holding sales rather than a constant price.

Let p_S^* and p_{NS}^* be the monopoly prices that maximize –static– profits from selling to the populations of storers and non-storers respectively, and p_{ND}^* the monopoly price of a non-discriminating monopolist. The monopolist will pick a pair of prices, \bar{p} and \underline{p} , to maximize

$$(\omega Q_{NS}(\underline{p}) + 2(1 - \omega)Q_s(\underline{p}))(\underline{p} - c) + \omega Q_{NS}(\bar{p})(\bar{p} - c)$$

where $Q_{NS}()$ and $Q_S()$ are the demands of non-storers and storers, and c is the (constant) marginal cost. The first term represents variable profits during sales, targeting storers who purchase for two periods, and non-storers for one period. The last term represents profits from non-storers, during non-sale periods. By repeating a two period pricing cycle the monopolist maximizes the present value of profits.

⁷These numbers are computed using the average prices and quantities for the sample presented in Table 3.

As long as $p_{NS}^* > p_S^*$ then $\bar{p} = p_{NS}^*$ while \underline{p} is the price charged by a non-discriminating monopolist who faces demand $\omega Q_{NS}(\underline{p}) + 2(1 - \omega)Q_S(\underline{p})$, namely, demand with additional weight on the storing population. It is easy to see that $p_S^* < \underline{p} < p_{ND}^*$.

Optimal pricing involves high prices targeting non-storers who are less price sensitive and sales, targeting storers. Under constant prices the seller would set a price that targets the average demand. The price cycle enables the seller to increase profits from non-storers without compromising profits from the storer. Naturally, the price cycle is of length $T + 1$.

We can use the estimates in column 5 to compute the margins implied by this pricing rule. The results suggest that Coke's margins during non-sale periods are 61 percent and the implied marginal cost is 0.54. During sale the margin is lower, at 43 percent, but the implied marginal cost is 0.53. Interestingly, the implied marginal cost is identical, even though we did not impose this in any way during the estimation or computation.

6 Extensions

6.1 Rational Expectations

In the base model we assumed consumers have perfect foresight. The simplicity of assuming perfect price foresight is that at any point in time the consumer maximizes $U(q, m)$ (where q are the quantities of the product and m the outside good) subject to the budget constraint. Anticipated demand is found by plugging the effective prices in equation 2. For example, on a period of Coke sale but no Pepsi sale the consumer buys $q^C(p_t^C, p_{t+1}^P)$ of Coke for consumption in the coming period.

Absent perfect information about future prices all the consumer can do is to maximize $E_t(U_{t+1}(q, m))$, that is, the $t + 1$ expected utility, given period t prices and the distribution of prices. Going back to the example, where Coke is on sale and Pepsi is not, the consumer has to decide how much Coke to purchase knowing Pepsi's price may end up at two different levels (assume A1 holds for Pepsi). The demand for Coke involves the solution of three first order conditions for q^C , \bar{q}^P and \underline{q}^P , where the last two quantities represent Pepsi consumption (at $t + 1$) if on sale, and consumption absent a sale. The demand for Coke, q^C , is still given by equation 2, but replacing the price of Pepsi by its expected price.

The quantities \bar{q}^P and \underline{q}^P are still linear in prices and have the same functional form, but they differ from the demand functions of the static problem (in equation 2). The reason is simple, in the static problem the consumer reacts to a Coke sale by adjusting both Coke and Pepsi quantities. Instead, \bar{q}^P and \underline{q}^P take q^C as given (since it was decided in the previous period before Pepsi prices were revealed). Thus, demand is slightly different from equation 2. Since demand for each good depends on whether the other was already purchased (on sale), we need a finer definition of the state.

The definition of the state involves the prices at t and $t-1$ of both (all) products. Demand depends on this finer state. For example, the demand for Coke from consumers who store if Coke did not have a sale at t and $t-1$, while Pepsi had a sale at $t-1$, is given by \bar{q}^C . Alternatively, if Coke has a sale at t but not at $t-1$, while Pepsi did not have a sale at either period, demand for Coke is $q(p_S^C, p_{NS}^P) + q^C$. The first quantity is for current consumption, while the second is for consumption at $t+1$. Similar expressions can be written for all states. In principle there are 16 states, but demand in some of them is identical so effectively there are 8 different states.

Column 6 of Tables 4 and 5 show demand estimates assuming rational expectations. As before the estimation is done by minimizing (in a least squares sense) the distance between the observed quantities and those predicted by the rational expectations model just described.

The results are not that different from the perfect foresight estimates: own prices elasticities are lower in absolute value while cross price effects are higher. If anything the results suggest that the bias from neglecting dynamics is larger under rational expectations. This shows that our previous results do not rely on, or are driven by, the perfect foresight assumption.

6.2 Discrete Choice Demand

In the above analysis we assume linear demand. However, the analysis goes through with more flexible functional forms. A popular model used in many recent applications, especially when dealing with many products, is the discrete choice model. We now show how this model fits into our framework.

Assume the utility consumer i gets from product j is given by

$$u_{ijt} = \delta_{ij} - \alpha_i p_{jt} + \varepsilon_{ijt}$$

where δ_{ij} is the utility from the attributes of the product both observed and unobserved⁸, α_i is the marginal utility of income and ε_{ijt} is a transitory shock. For now, we assume perfect foresight of both prices and individual shocks. We can think of ε_{ijt} as capturing transitory needs known in advance, like having guests the following week. As in the standard discrete choice model, we assume that in each period the consumer consumes at most one unit (but might consume none). However, the consumer can purchase additional units that can be stored up to T periods (as in Assumption A3).

Demand has a structure similar to equation 3. In each period, depending on their ε_{ijt+1} , consumers who store decide which brand they will consume next period. The choice is based on the standard discrete choice thresholds using $\min\{p_{jt}, p_{j,t+1}\}$ as the (effective) price. If the optimal choice is product j , and that product is on sale at period t it will be purchased then, otherwise it will be purchased at time $t + 1$. For example, assuming no heterogeneity in tastes or in the marginal utility of income and that ε_{ijt} is distributed *i.i.d.* extreme value (i.e., the simple Logit model) aggregate demand for product j is given by

$$x_i(p^t) = M \begin{cases} \omega \frac{e^{\delta_j - \alpha p_{jt}}}{\sum_k e^{\delta_k - \alpha p_{kt}}} + (1 - \omega) \frac{e^{\delta_j - \alpha p_{jt}}}{\sum_k e^{\delta_k - \alpha \min(p_{kt-1}, p_{kt})}} & NN \\ \omega \frac{e^{\delta_j - \alpha p_{jt}}}{\sum_k e^{\delta_k - \alpha p_{kt}}} & SN \\ \omega \frac{e^{\delta_j - \alpha p_{jt}}}{\sum_k e^{\delta_k - \alpha p_{kt}}} + (1 - \omega) \left(\frac{e^{\delta_j - \alpha p_{jt}}}{\sum_k e^{\delta_k - \alpha \min(p_{kt-1}, p_{kt})}} + \frac{e^{\delta_j - \alpha p_{jt}}}{\sum_k e^{\delta_k - \alpha \min(p_{kt}, p_{kt+1})}} \right) & NS \\ \omega \frac{e^{\delta_j - \alpha p_{jt}}}{\sum_k e^{\delta_k - \alpha p_{kt}}} + (1 - \omega) \frac{e^{\delta_j - \alpha p_{jt}}}{\sum_k e^{\delta_k - \alpha \min(p_{kt}, p_{kt+1})}} & SS \end{cases} \text{ if}$$

where M is the market size. Similar expressions can be written for more flexible demand models (e.g., the model proposed by Berry, Levinsohn and Pakes, 1995).

Allowing for rational price expectation, instead of perfect price foresight, is slightly more tedious but conceptually straightforward. Consumers, uncertain about future prices, compare the flow utility from purchasing today for future consumption to the option value of waiting to buy the best alternative tomorrow. Tomorrow's best alternative depends on uncertain prices. Computing predicted market shares may require simulation, but the dynamics involved are still immediate.

The option to wait is generated by uncertain future prices. On the other hand, we assumed the vector $\{\varepsilon_{ijt+1}\}$ is know while shopping at time t . The assumption that the ε_{ijt+1} are know in advance is quite handy, making the analysis simpler, but also appealing. It is reasonable to assume consumers anticipate future need, as well as the future ranking

⁸In many applications $\delta_{ij} = X_j \beta_i + \xi_j$ where X_j are the observable attributes of product j , β_i are the consumer specific taste for these attributes and ξ_j is the unobserved (to the reseacher) product characteristic.

of the different products. In some applications an unknown ε_{ijt+1} might be appropriate. In which case the option value of waiting as well as the value of the anticipated purchases need to be adjusted to reflect preference uncertainty.

7 Alternative Approaches

There are alternative approaches used in the literature to recover preferences, besides structural estimation. A common method often used to compute long run price responses is to include lagged prices as a way to control for dynamic effects.⁹ Alternatively, lagged quantities, can be motivated through either a partial adjustment model or as approximating the missing inventory variable. The long run effect, in either case, is computed by tracing the effect of a price change if lagged quantity is included, or by summing the coefficients of the lagged prices. We focus on including lagged prices, which is more common.

In principle, including lagged prices as controls seems like an attractive, flexible and model-free way of recovering long run responses. Indeed, in some cases the lagged controls may help approximate a long run effect, but this is not generally the case. In spirit, one could claim that including lagged prices is similar to our model, where we advocate defining events based on lagged sales. However, there is a difference between our approach and this alternative. It is evident from equation 3 that under our model simply including lagged prices is not enough. The lagged prices need to be interacted with the state. This is clearly evident in the cross price effect. Our model advocates inclusion of the effective cross price, which is a function of lagged prices but the function varies by state. Simply including a lagged structure, potentially with many lags, would not recover the actual long run cross price effects. We explore the performance of this method below.

Another common approach to deal with demand anticipation is to aggregate observations over time, for example, from weekly to monthly data. There are cases where aggregation might help solve the problem but they rely on strong assumptions (and may wipe out the variation in prices). Consider the example in Section 3.3. Aggregation can solve the problem under some special conditions. Suppose that we have data from an additional period, $t = 0$, in which the prices and quantities are like period 1. If we aggregate the data from the first

⁹See, for example, van Heerde et al (2000).

two periods ($t = 0, 1$) and the last two periods ($t = 2, 3$) and define prices as revenue divided by quantity, then estimation based on the aggregate data would recover long run effects.

The success of this approach in recovering long run responses relies crucially on several assumptions, like lack of heterogeneity in storage. We provide, in the Appendix, an analytic example that shows this.

We now apply these alternative corrections. The results for Coke are presented in Table 6. The first two columns repeat the results from the store fixed-effects regression, and from our model. The next two columns present the long run effect from models that include 1 and 4 lags, respectively. The results are not very promising. Both lagged prices models impact the own price elasticity in the "right" direction but the magnitude is smaller than our correction. The results do not look good for the cross price effect. The first model does not change the cross-price effect by much. The second, with more lags, does but estimates a negative cross-price elasticity.

The last two columns present the results from aggregating over time: into bi-weekly periods and then by month. In both cases the own price elasticities move in the right direction, in the case of monthly aggregation yielding numbers quite close to our estimates. However, both models estimate a negative cross price elasticity (although in one case the estimate is not significantly different from zero). Overall, these alternative methods do not seem to yield very sensible results.

Table 6: Demand for Coke – Alternative Corrections

	FE	Our	lag 1 p	lag 4 p	agg bi-w	agg month
	(1)	(2)	(3)	(4)	(5)	(6)
P_{Coke}	-1428.2	-938.9	-1131.3	-1239.3	-1130.2	-911.2
	(11.1)	(11.2)	(15.1)	(22.9)	(9.8)	(14.5)
P_{Pepsi}	66.5	150.1	63.3	-252.5	-31.5	-9.0
	(11.9)	(11.4)	(17.1)	(25.6)	(11.1)	(16.0)
ω		0.53				
		(0.01)				

Note: All estimates are from least squares regressions. The dependent variable is the quantity of Coke sold at a store in a week. All regressions include store store fixed effects and teh price of the store brand. The regression in column (1) in a fixed effects regression. The results in column (2) are the model estimates from Table 4. The regressions in columns (3) and (4) allow for 1 and 4 lags

of prices. The regressions in columns (5) and (6) aggregate the data to a bi-weekly and monthly level (and use unit prices). Standard errors are reported in pharenthesis.

8 Concluding Comments

We offer a simple model to account for demand dynamics due to consumer inventory behavior. The model can be estimated using store level data. An application to demand for Coke and Pepsi yields reasonable estimates. At the same time, corrections based on alternative methods, like aggregation or control for lagged variables, do not perform well.

The base results rely on many assumptions, most of which can be relaxed. As we showed we can allow for heterogeneity in preferences, more flexible demand systems, and rational expectations. We can also let the fraction of consumers who store vary with price. Of course, some of these extensions increase the complexity of the model and defeat our goal of delivering a simple model.

We use the simplicity of the model to derive markups implied by dynamic pricing, rather than plugging demand estimates into static first order conditions. The standard static approach underestimates market power for two reasons. First, demand elasticities biases (both own and cross) imply lower markups. Second, the static first order conditions imply lower mark-ups than the dynamic ones.

9 References

- Berry, Steven, James Levinsohn, and A. Pakes (1995), “Automobile Prices in Market Equilibrium,” *Econometrica*, 63, 841-890.
- Berry, Steven, James Levinsohn and Ariel Pakes (1999) “Voluntary Export Restraints on Automobiles: Evaluating a Strategic Trade Policy,” *American Economic Review*, 89(3), 400–430.
- Blattberg, R. and S. Neslin (1990), *Sales Promotions*, Prentice Hall.
- 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.

- Conlisk J., E. Gerstner, and J. Sobel "Cyclic Pricing by a Durable Goods Monopolist," *Quarterly Journal of Economics*, 1984.
- Erdem, T., M. Keane and S. Imai (2003), "Consumer Price and Promotion Expectations: Capturing Consumer Brand and Quantity Choice Dynamics under Price Uncertainty," *Quantitative Marketing and Economics*, 1, 5-64.
- Goldberg, Penny. "Product Differentiation and Oligopoly in International Markets: The Case of the U.S. Automobile Industry," *Econometrica*, Jul. 1995, pp. 891-951.
- van Heerde, H., Leeflang, P., and Wittink, D. "The Estimation of Pre- and Postpromotion Dips with Store-Level Scanner Data." *Journal of Marketing Research*, Vol. 37 (2000), pp. 383-395.
- Hendel, Igal and Aviv Nevo "Sales and Consumer Inventory," Manuscript, *RandJournal of Economics*, 2006a.
- Hendel, Igal and Aviv Nevo "Measuring the Implications of Sales and Consumer Inventory Behavior," *Econometrica*, 2006b.
- Hong P., P. McAfee, and A. Nayyar "Equilibrium Price Dispersion with Consumer Inventories", Manuscript, University of Texas, Austin.
- Jeuland, Abel P. and Chakravarthi Narasimhan. "Dealing-Temporary Price Cuts-By Seller as a Buyer Discrimination Mechanism" *Journal of Business*, Vol. 58, No. 3. (Jul., 1985), pp. 295-308.
- Pesendorfer, M. (2002), "Retail Sales. A Study of Pricing Behavior in Supermarkets," *Journal of Business*, 75(1), 33-66.
- Salop S. and J. E. Stiglitz. "The Theory of Sales: A Simple Model of Equilibrium Price Dispersion with Identical Agents" *The American Economic Review*, Vol. 72, No. 5. (Dec., 1982), pp. 1121-1130.
- Sobel, Joel. "The Timing of Sales," *Review of Economic Studies*, 1984.
- Varian Hal "A Model of Sales" *The American Economic Review*, Vol. 70, No. 4. (Sep., 1980), pp. 651-659.

10 Appendix

10.1 Purchases when $T = 2$

The predicted purchases when $T = 2$ (assuming a single product) are given by:

$$x(p^t) = \begin{cases} \omega q(p^t) & SNN \\ q(p^t) & NNN \\ \omega q(p^t) & \text{if } NSN \text{ or } SSN \\ \omega + 3(1 - \omega)q(p^t) & NNS \\ \omega + 2(1 - \omega)q(p^t) & SNS \\ q(p^t) & NSS \text{ or } SSS \end{cases} \quad (5)$$

First, notice there are 8 states, some of them involve similar predicted purchases. In contrast to equation 3 where demand is affected by lagged prices, when $T = 2$ demand depends on whether there was a sale two periods ago. Second, notice how (some) events are split. Event NN needs to be split into SNN and NNN , because a storer who purchased two periods ago on sale does not buy today at a regular price, while she would buy if two periods earlier there was no sale, namely, in event NNN . Predicted purchases in events SS and NS are not affected by $t - 2$ events, thus they require no modification from equation 3. Purchases differ between SNS and NNS because in SNS current consumption comes out of storage.

10.2 Estimating equations

We choose the parameters to minimize the sum of squares of the difference between observed purchase and those predicted by the model. The data consists of a panel of quantities and prices in different stores. Since purchases are scaled differently in different states in order to account for store fixed effects we need to transform the predicted purchases as follows. Let j denote the store.

$$x_{ijt} = f_t \left(\frac{1}{T} \sum_{\tau=1}^T \left(\frac{x_{j\tau}}{f_\tau} \right) \right) + \beta(p_{ijt} - \bar{p}_{i,t}) + \gamma(pe_{-ijt} - \bar{pe}_{-i,t})$$

where f_t is the factor by which demand is scaled up in period t , $\bar{p}_{i,t}$ is the within store average, and pe is the effective cross price (as defined in the text). Note, that the effective price is a function of ω . For the base model

$$f_t = \begin{cases} 1 & NN \text{ or } SS \\ \omega & \text{if } SN \\ 2 - \omega & NS \end{cases}$$

10.3 Example where aggregation fails

Consider the following example where aggregation fails. Suppose there are two types of consumers. Type A consumers can store for one period, type B cannot store. Assume four time periods with $p_1^2 < p_1^0 = p_1^1 = p_1^3$ (while $p_2^0 = p_2^1 = p_2^2 = p_2^3$). Purchases are given by

$$x_1^A = \begin{bmatrix} q_1^1 \\ q_1^1 \\ 2(q_1^1 - \beta\Delta p_1) \\ 0 \end{bmatrix}, \quad x_2^A = \begin{bmatrix} q_2^1 \\ q_2^1 \\ q_2^1 + \gamma\Delta p_1 \\ q_2^1 + \gamma\Delta p_1 \end{bmatrix}, \quad x_1^B = \begin{bmatrix} q_1^1 \\ q_1^1 \\ q_1^1 - \beta\Delta p_1 \\ q_1^1 \end{bmatrix}, \quad x_2^B = \begin{bmatrix} q_2^1 \\ q_2^1 \\ q_2^1 + \gamma\Delta p_1 \\ q_2^1 \end{bmatrix}$$

Assuming one type of each consumer, and aggregating over periods, aggregate purchases will be

$$x_1 = \begin{bmatrix} 4q_1^1 \\ 4q_1^1 - 3\beta\Delta p_1 \end{bmatrix}, \quad x_2 = \begin{bmatrix} 4q_2^1 \\ 4q_2^1 + 3\gamma\Delta p_1 \end{bmatrix}$$

If we used "consumer-week" weighted prices, i.e. p_1^1 and $\frac{1}{4}p_1^1 + \frac{3}{4}p_1^2 = p_1^1 + \frac{3}{4}\Delta p_1$, we would recover the long run effects. Note, that to figure out the right price we need a model and we need to know the fraction of each type. However, if we use unit value, i.e. revenue divided by quantity, prices will be p_1^1 and $(1 - \lambda)p_1^1 + \lambda p_1^2 = p_1^1 + (1 - \lambda)\Delta p_1$, where $\lambda = 3(q_1^1 - \beta\Delta p_1)/(4q_1^1 - 3\beta\Delta p_1) > 3/4$.

Because of aggregation the unit-value price will generate a price that is too low in the second period. Using this price (even if we know the true slopes of demand) will yield estimates that are biased towards zero. For own price effects this is the right direction, giving the impression that the problem of demand dynamics has been attenuated. For cross price effects this is the wrong direction.