Demand for nondurable goods: a shortcut to estimating long-run price elasticities*

Helena Perrone†

Universitat Pompeu Fabra and Barcelona GSE

September 16, 2016

Abstract

When consumers stockpile, static demand models overestimate long-term price responses. This paper presents a dynamic model of demand with consumer inventories and proposes a shortcut to estimate the long-run price elasticities without having to solve the dynamic program. Using French data on food purchases, I find elasticities consistent with those that result from the full-blown estimations found in the literature.

JEL: D12, L1, M3

Keywords: inventories, consumer behavior, dynamic demand, long-run price elasticities, scanner data.

1 Introduction

Standard demand models assume that consumers purchase choices are static.[1] However, there is substantial empirical evidence that in a number of products, consumers' stockpile.[2] When this is the case, ignoring dynamics may lead to upward-biased estimates of long-run demand price elasticities

*I would like to thank Larbi Alaoui, Steven Berry, Christian Brownlees, Jan Eeckhout, Juanjo Ganuza, Christian Michel, Aviv Nevo, Vincent Requillart, and Michael Waterson for helpful comments. I would also like to thank anonymous referees whose comments and suggestions greatly improved the paper. I am especially grateful to Pierre Dubois for his encouragement, guidance and support throughout this work. Research support by CAPES and INRA is kindly acknowledged. All errors are mine.

†Contact information: helena.perrone@upf.edu

[1] See, for example, Berry, Levinsohn and Pakes (1995) and Nevo (2001), among many others.

[2] Hendel and Nevo (2006a), and Boizot, Robin, and Visser (2001).

because a static model captures not only the effect of prices on consumption but also the short-term variation in inventories (Hendel and Nevo, 2006a, 2006b).

To deal with these issues we need to estimate a dynamic model, but this can be costly because etimation of a dynamic demand model usually requires solving a dynamic programming problem, which is a computer-time consuming numerical exercise. However, because for most policy purposes, for example, competition policy, we require long-run elasticities, it is important to develop techniques to estimate long-run effects without completely solving the dynamic programming problem. This paper contributes to the literature by proposing a simple procedure to estimate long-run own-price elasticities that does not require solving the value function. My method, although simple, is flexible with respect to consumer heterogeneity, price processes, and consumers' future price expectations. Note that the model in its current format does not address brand differentiation, hence, I can estimate long-run own-price elasticities but not cross-price elasticities.

In the model, consumers maximize their discounted utility flow by choosing, in each period, how much to purchase, how much to consume, and how much to leave as inventory of a single homogeneous product. Future prices are uncertain. In the model, when prices are high relative to expected future prices, consumers only purchase to cover current consumption, not to stockpile. Consumers purchase and store at discounted prices, using stocks to avoid having to purchase at high prices. If storage were costless and there was no discounting, they would buy their lifetime consumption when the price is at its lowest. However, storage is costly. Thus, consumers hold limited stocks that they supplement when prices are low and deplete when prices are high, purchasing at non-sale prices only if there is not enough of the product in storage to cover current consumption. Therefore, in periods of high prices, the purchase decision depends on how much of the product consumers already have in stock and on consumption at current prices.

I propose to use these properties of the model to simplify estimation. In particular, the key to the simplicity of the method is twofold. First, I estimate parameters when choices are not affected by future expected variables that require solving the value function. Second, I only require current observable variables and the beginning-of-the-period inventory level, which I am able to construct from the panel data. Note that restricting the estimation sample does not create a selection problem under the standard assumption that the marginal utility of income is independent of the price level.

The empirical implementation requires an assumption concerning how consumers form future price expectations, a utility function specification, and an assumption regarding consumption out of stock. However, the methodology is quite flexible and allows for various alternative assumptions. In the application, I consider a CARA utility specification and 3 alternative price expectation formation hypothesis. The first assumes that consumers always expect prices to return to their regular (mean) level, while the second and third price expectation formation hypotheses assume different first-order Markov processes.

With respect to consumption out of stock (consumption in periods when the consumer does not purchase the product but still consumes it), I assume that it depends on the expected replacement costs, which I proxy by the average price paid for each unit in stock. I empirically check the sensitivity of the results to this assumption by estimating the model in subsamples with increasingly shorter interpurchase periods. The shorter the interpurchase periods, the less I have to assume the out-of-stock level of consumption. Furthermore, I also estimate parameters in a subsample that includes only purchases made at high prices following periods of purchases at high prices. This means that in this subsample, inventories are always zero (at the beginning and end of each period) and consumption equals purchases in each period. Parameter identification therefore does not require the use of the out-of-stock consumption assumption because, according to the model, consumers never consume from stock in this subsample.

The model is estimated using French home scan data on food products over three years (1999, 2000, 2001). In the data, each observation is a purchase occasion and I observe exactly what was purchased, when it was purchased and the price paid. The empirical analysis is performed considering five different product categories: butter, coffee, milk, pasta, and yogurt. Note that in contrast to the US, the most common type of milk in France is shelf stable ultra-pasteurized (UHT) milk. It represents more than 97% of French milk consumption.[3] It can be stored on shelves for a long time while closed and for approximately one week in a refrigerated area once opened. This is the type of milk I consider in the sample, not fresh milk.

Consistent with previous results in the literature, the estimation results show that the price elasticities yielded by a static model that ignores inventories consistently overestimate long-run responses. The bias ranges from 20 up to more than 100 percent. The magnitudes of my long-run price elasticity estimates are also consistent with previous literature estimates obtained by solving the complete dynamic programming model, which is further evidence supporting my method.[4] The

[3] source: Institut Professionel du Lait de Consommation, www.iplc.fr.

[4] See, for example, Hendel and Nevo (2006b and 2012), Sun, Neslin and Srinivasan (2003), Sun (2005), and Hartmann and Nair (2009). Note that these papers focus on different products and, as expected, numbers are similar

results are robust to the assumption on consumption out of stock, to the estimation method, and to the price expectation hypothesis (although the results are more sensitive to the different price expectation hypotheses when the subsamples used for estimation are smaller).

Papers that estimate long-run price elasticities by solving the dynamic program of the stockpiling consumer are Hendel and Nevo (2006b), Erdem, Imai, and Keane (2003), and Sun (2005). My model is very similar to that of Hendel and Nevo (2006a), but their aim is to derive and test the model's implications rather than estimate the long-run price elasticities. In terms of objectives, this paper is most similar to Hendel and Nevo (2013), who develop a simple model of demand with inventories that identifies long-run price elasticities without solving the value function. Their paper differs from mine in a number of important ways. My model is more flexible in terms of the price distribution and expectation, stockpiling technology, and consumer heterogeneity. Their methodology, however, can estimate cross-price effects, is applicable to store-level data, and does not require a large number of observations in the time dimension. For a more thorough comparison between the two methods, please see Section 5.

This paper is organized as follows. The next section introduces the model of consumer choice with stockpiling and derives the main theoretical result, which says that consumers only stockpile at low prices. Section 3 discusses econometrics, including the consumer's purchase decision and estimable equation, identification, and data description. Section 4 presents the results and sensitivity analysis. Section 5 compares my method with other methods in the literature (specifically, Hendel and Nevo, 2006a and 2013). Section 6 concludes.

2 The Model

This section presents a simple dynamic model of demand with inventories in which consumers maximize their discounted utility flow by deciding, in each period, the purchase quantity and how much of this quantity is to be consumed or left as inventory. These decisions depend on how much of the product consumers already have in stock, current prices, and future expected prices. The main result of the model shows that households consume everything they purchase when they purchase at a price that is higher than their expected future price. That is, consumers do not hold as inventory products purchased at prices in the top end of the support of the price distribution. This result is the basis for the identification strategy developed in Section 3.

but not identical.

2.1 The setup

The per period net utility of consumer i from consuming c_{it} is $u(c_{it})$, where u is an increasing and concave function of consumption in period t, c_{it}. I assume that $u'(0) = \infty$, which is sufficient for positive consumption in every period. In each period, consumer i must decide how much to purchase of a certain good (q_{it}) and how much to consume. What is not consumed is retained as inventory[5]. The law of motion of inventories is:

$$y_{it} = q_{it} - c_{it} + y_{it-1} \tag{1}$$

where y_{it} is the end-of-period t level of inventories.

I assume that consumers visit the store at an exogenously given frequency that is determined by their overall shopping needs rather than the need to buy one specific product. That is, I assume that the product is a small part of a longer shopping list or bundle and that it does not individually trigger store visits.

The problem of consumer i in any period τ is:

$$\max_{\{c_{it}, q_{it}, y_{it}\}} E_\tau \left\{ \sum_{t=\tau}^{\infty} \delta^t \left[u(c_{it}) - \alpha_i p_t q_{it} - \Phi(y_{it}) \right] \right\} \tag{2}$$

$$s.t. \quad y_{it} = q_{it} - c_{it} + y_{it-1} \quad (\lambda_{it})$$

$$q_{it} \geqslant 0 \quad (\psi_{it})$$

$$y_{it} \geqslant 0 \quad (\mu_{it})$$

where α_i is the marginal utility of income, and λ_{it}, ψ_{it}, and μ_{it} are the Lagrange multipliers of each constraint. The function $\Phi(y_{it})$ represents the cost of storing inventory, which is an increasing and convex function of inventories. Assume that $\Phi(y_{it}) = \frac{\phi_i}{2} y_{it}^2$.

At time t, consumers know their current inventory y_{it-1} and current prices but are uncertain about future prices.

Note that at this point, there is no need to impose any structure on how prices evolve.

[5] I assume there is no depreciation. This means that products are non-perishable or that their expiration date is never binding.

2.2 Consumer behavior

The Lagrangian of the problem of the consumer is:

$$\pounds = \max_{\{c_{it}, q_{it}, y_{it}\}} E_\tau \left\{ \sum_{t=\tau}^{\infty} \delta^t \left[u\left(c_{it}\right) - \alpha_i p_t q_{it} - \phi_i y_{it}^2 + \lambda_{it}\left(q_{it} - c_{it} + y_{it-1} - y_{it}\right) + \psi_{it} q_{it} + \mu_{it} y_{it} \right] \right\} \tag{3}$$

The first-order conditions with respect to consumption in period t, purchases in period t, and end-of-the-period inventory in period t are, respectively:

$$\frac{\partial \pounds}{\partial c_{it}} = 0 \Rightarrow u'\left(c_{it}\right) = \lambda_{it} \tag{4}$$

$$\frac{\partial \pounds}{\partial q_{it}} = 0 \Rightarrow \alpha_i p_t = \lambda_{it} + \psi_{it} \tag{5}$$

$$\frac{\partial \pounds}{\partial y_{it}} = 0 \Rightarrow y_{it} = \frac{\delta E_t\left(\lambda_{it+1}\right)}{\phi_i} - \frac{\lambda_{it}}{\phi_i} + \frac{\mu_{it}}{\phi_i} \tag{6}$$

Manipulation of the first-order conditions yields the main result:

Proposition 1 *In periods with purchases ($q_{it} > 0$), if the price is higher than the discounted expected price for the following period ($p_t > \delta E_t\left(p_{t+1}\right)$), the utility-maximizing end-of-the-period inventory level is equal to zero ($y_{it} = 0$). Furthermore, in this case, the purchased quantity in the next period is positive ($q_{it+1} > 0$). However, if $p_t \leq \delta E_t\left(p_{t+1}\right)$, then $y_{it} > 0$ in periods with and without purchases, and the quantity purchased in the next period can be either positive or equal to zero.*

Proof. Consider $p_t > \delta E_t\left(p_{t+1}\right)$ and $q_{it} > 0$ ($\psi_{it} = 0$). Assume that $y_{it} > 0$ ($\mu_{it} = 0$). Then, $y_{it} = \frac{\alpha_i(\delta E_t(p_{t+1}) - p_t)}{\phi_i} - \frac{E_t \psi_{it+1}}{\phi_i}$, implying that $y_{it} < 0$, which is a contradiction. Thus, $y_{it} = 0$ in this case. Moreover, $q_{it+1} > 0$ because consumption is positive in every period. Now, consider $p_t \leq \delta E_t\left(p_{t+1}\right)$. Assume that $y_{it} = 0$ ($\mu_{it} > 0$). Then, $y_{it} = \frac{\alpha_i(\delta E_t(p_{t+1}) - p_t)}{\phi_i} + \frac{\psi_{it}}{\phi_i} + \frac{\mu_{it}}{\phi_i} > 0$, which is a contradiction. Hence, $y_{it} > 0$. The purchases in the next period are equal to $q_{it+1} = y_{it+1} - x_{it+1} + c_{it+1}$, which is equal to zero if $y_{it+1} \geqslant y_{it} - c_{it+1}$ and greater than zero otherwise. ■

Proposition 1 is quite intuitive. Stockpiling is a means of intertemporal substitution. To save, consumers advance purchases if they believe that tomorrow's prices will be higher. If they believe that prices will be lower in the future, they only buy to cover current consumption and never stock up at high prices.

Equation 3 is nearly identical to Hendel and Nevo (2006a). They show that the target inventory is a function of prices only. In my case, the target inventory depends also on the marginal cost of holding inventories (and shadow prices). Proposition 1 shows that the target inventory equals zero for any $p_t > \delta E_t(p_{t+1})$.

Proposition 1 is central to the identification strategy developed in the following section. It enables separating periods in which consumers purchase only to cover consumption from periods in which consumers also purchase to stockpile. The purchase decision in periods in which consumers do not stockpile is simpler and depends only on the beginning-of-the-period inventory level and current prices. As I show next, utility parameters can be identified for these periods without having to solve for the full dynamic program. Furthermore, Proposition 1 is also essential in the construction of the beginning-of-the-period inventory level. I use it to bypass the initial conditions problem: for each household, I consider the initial period to be that following the first period of purchases at high prices (thus, the initial inventory for each household is zero, according to Proposition 1).

3 Econometrics

3.1 The empirical purchase decision

This section derives the purchase decision equation in periods of high prices. To reconstruct the beginning-of-the-period inventory level and write the estimable equation, I make an assumption regarding consumption in periods without purchases, the utility function and how consumers form future price expectations. I focus on periods with high prices because, in these, periods consumers purchase only to cover consumption, thereby greatly simplifying the purchase decision and parameter identification.

Proposition 1 implies that conditional on purchasing at a price higher than the future expected price, the probability of stockpiling is zero. Furthermore, if $y_{it} = 0$, then $q_{it} = c_{it} - y_{it-1}$, where $c_{it} = h(\alpha_i p_t)$ and $h = u'^{-1}$. Therefore, conditional on purchases and high prices:

$$E_t(q_{it} \mid q_{it} > 0, p_t > \delta E_t(p_{t+1})) = E_t(h(\alpha_i p_t) - y_{it-1} \mid q_{it} > 0, p_t > \delta E_t(p_{t+1})) \tag{7}$$

Assume that due to some noise in data collection, in each t, the econometrician only observes $q^*_{it} = q_{it} + v_{it}$, where v_{it} is a zero mean error term distributed in $(-q_{it}, q_{it})$.[6] The support of the

[6] The empirical application uses homescan data recorded by the panelists themselves. Hence, errors in the data may be introduced due to recording mistakes by the panelist or misunderstanding of homescan instructions. Einav,

distribution of v_{it} guarantees that there is no misclassification, i.e., $q_{it} > 0 \Leftrightarrow q^*_{it} > 0$, where $q^*_{it} = q_{it} + v_{it}$ is the observed quantity. Substituting q^*_{it} into (7), it follows that:

Assume that in each t, the econometrician observes q_{it} with an error v_{it}, where v_{it} has zero mean and is distributed in $(-q_{it}, q_{it})$. The support of the distribution of v_{it} guarantees that there is no misclassification, i.e., $q_{it} > 0 \Leftrightarrow q^*_{it} \equiv q_{it} + v_{it} > 0$. Substituting q^*_{it} into (7), it follows that:

$$E_t\left(q^*_{it} \mid q_{it} > 0, p_t > \delta E_t\left(p_{t+1}\right)\right) = E_t\left(h\left(\alpha_i p_t\right) - y_{it-1} + v_{it} \mid q^*_{it} > 0, p_t > \delta E_t\left(p_{t+1}\right)\right) \quad (8)$$

3.1.1 The beginning-of-the-period inventory level

The law of motion of inventories implies that:

$$\begin{aligned} y_{it-1} &= y_{i0} + \sum_{n=1}^{t-1} q_{in} - \sum_{n=1}^{t-1} c_{in} \\ &= y_{i0} + \sum_{n=1}^{t-1} q^*_{in} - \sum_{n=1}^{t-1} c_{in} - \sum_{n=1}^{t-1} v_{in} \end{aligned}$$

Proposition 1 states that following a period with purchases at high prices, the beginning-of-the-period stocks are zero. Therefore, I consider household i's initial period to be that immediately following the first observed purchase at a high price. In that way, $y_{t_1(i)-1} = 0$ and:

$$y_{it-1} = \sum_{n=t_1(i)}^{t-1} \left(q^*_{in} - c_{in} - v_{in}\right) \quad (9)$$

where $\sum_{n=t_1(i)}^{t-1} q^*_{in}$ is observable because q^*_{it} is observable for all t.

Concerning consumption, its utility maximizing level in periods with purchases is $c_{it} = h\left(\alpha_i p_t\right)$. In periods without purchas, $c_{it} = y_{it-1} - y_{it}$, which cannot be used in (9) to calculate beginning-of-the-period inventories. To avoid having to solve the value function, I add structure. In particular, impose:

(C) In periods without purchases, consumption is equal to consumption at the average price paid per units in storage, i.e., $c_{it} = h\left(\alpha_i p^i_s\right)$.

Leibtag and Nevo (2010) compare homescan data with information from cash registers of a certain retailer and find discrepancies that seem like standard recording errors (by the panelists or the data collection company).

Assumption (C) is a simplification. Consumption out of stock should depend on the expected replacement costs, which is a function of future expected prices. I proxy this function by the average price paid for the product.[7]

The simplification works well when consumers expect prices to vary for short periods of time around a regular price.[8] It does not work well if consumers expect prices to trend upwards (or downwards). In this case, consumption from storage should be lower (higher) than the "regular" consumption.

In Section 4.2, a series of robustness exercises shows that the qualitative results are not sensitive to (C).

3.1.2 Utility Specification

Assume that $u(c_{it}) = -(1/\rho)\exp[-\rho c_{it}]$, where ρ is a positive parameter.[9] Hence, when $q_{it} > 0$:

$$c_{it} = h(\alpha_i p_t) = -\frac{1}{\rho}\ln\alpha_i - \frac{1}{\rho}\ln p_t \tag{10}$$

Let $T_{i1}^{t-1} \in \{t_1(i), ..., t-1\}$ be the set of periods from $t_1(i)$ (household i's initial period) up to $t-1$ in which consumer i purchased the product, and $|T_{i0}^{t-1}|$, the number of periods in which i did not purchase. Then:

$$\sum_{n=t_1(i)}^{t-1} c_{in} = -\frac{1}{\rho}\left(T_i^{t-1}\ln(\alpha_i) + \sum_{n\in T_{i1}^{t-1}}\ln p_n + |T_{i0}^{t-1}|\ln p_s^i\right) \tag{11}$$

---

[7] Assumption (C) is also consistent with the framework in Hendel and Nevo (2013). When consumers can only stockpile for a pre-specified number of periods, and the first-order Markovian prices can only take two values (high and low), Hendel and Nevo show that the relevant price for out-of-stock consumption is the effective price, i.e., the minimum price actually paid for the units in the pantry. Here, there are no restrictions on the number of periods that a product can be stored. Hence, one cannot know how much each unit remaining in storage actually cost. The average price conditional on purchases, which is the empirical definition of my average price, could then be interpreted as a proxy for the effective price.

[8] This is the usual price pattern observed in supermarket data, for example.

[9] Instead of CARA, Hendel and Nevo (2006b) consider a logarithm function ($u(c_{it}) = \ln(c_{it})$). Here, the log utility would constrain elasticities to be equal to 1. Examples of alternative utility functions can be found in Erdem et al. (2003), where utility is linear in consumption, and in Sun (2005), where $u(c_{it}) = \chi\left(c_{it} - \gamma c_{it}^2\right)$. Identification with a quadratic utility function is straightforward. Identification is also possible assuming other utility functions, for example CRRA functions. However, under a CARA or a quadratic utility function, parameters can be estimated by OLS, whereas parameter estimation under a CRRA utility function would require non-linear least squares.

where $|T_i^t|$ is the total number of periods from $t_1(i)$ up to t ($|T_i^t| = |T_{i1}^{t-1}| + |T_{i0}^{t-1}| + 1 = t - t_1(i)$, where $|T_{i1}^{t-1}|$ is the number of periods in T_{i1}^{t-1}).

3.1.3 Price Expectation

Estimation requires an assumption regarding how consumers form expectations of future prices. The model allows for a wide range of price expectation formation hypotheses. Here, I consider three alternatives. The simplest assumes that consumers always expect prices to return to their regular level, which is equal to the mean price of the good. That is: (E0) $\delta E_t(p_{t+1}) = p_r$, where p_r is the regular or mean price[10] Under this assumption, current price realizations do not affect expectations.

I also estimate two alternative first-order Markov price processes and assume that expected prices are equal to the predicted price:[11]

$$\text{(E1)}\ p_{it} = a + bp_{it-1}$$
$$\text{(E2)}\ \ p_{it} = (a_h + b_h p_{it-1})h + (a_l + b_l p_{it-1})l + (a_r + b_r p_{it-1})r$$

where h, l, and r indicate that p_{t-1} is either higher, lower or equal to the regular price, respectively.

More sophisticated price expectation hypotheses can be used. For example, one could consider a variation of (E1) and (E2) whereby the parameters are updated every period using new price information. This would be consistent with Bayesian updating, and predicted prices would converge to actual prices when the price distribution changes over time.

3.1.4 Estimable equation

Substituting (9), (10) and (11) into the purchased quantity equation (8) and moving $\sum_{n=t_1(i)}^{t-1} q_{in}^*$ to the left-hand side yields the estimable equation:

[10] As the time period considered is brief (one week), I let the discount rate be equal to 1. Thus, $E_t(p_{t+1}) = \delta^{-1}p_r \simeq p_r$, where p_r is the regular price.

[11] Higher order processes do not alter the results in a relevant manner.

$$E_t\left(Q_{it}^* \mid q_{it} > 0, p_t > \delta E_t\left(p_{t+1}\right)\right) = -\frac{\ln\left(\alpha_i\right)}{\rho}|T_i^{t-1}| - \frac{1}{\rho}\left(\sum_{n\in T_{i1}^t} \ln p_n + |T_{i0}^{t-1}| \ln p_s^i\right) + \tag{12}$$
$$E_t\left(\sum_{n=t_1(i)}^{t} v_{in} \mid q_{it} > 0, p_t > \delta E_t\left(p_{t+1}\right)\right)$$

where $Q_{it}^* = \sum_{n=t_1(i)}^{t} q_{in}^*$ and $\delta E_t\left(p_{t+1}\right)$ is given by one of the price expectation hypotheses described above (E0, E1, or E2).

3.2 Identification and Estimation

3.2.1 Identification

The observable variables in the estimable equation (12) are the total number of periods $|T_i^{t-1}|$, the number of periods without purchases $|T_{i0}^{t-1}|$, the sum of current and past prices $\sum_{n\in T_{i1}^t} \ln p_n$, and the average price p_s^i. The parameters to be estimated are α_i and ρ. Errors are not independent because v_{it} is distributed on $]-q_{it}, q_{it}[$. However, $E\left(\sum_{n=t_1(i)}^{t} v_{in} \mid q_{it} > 0, p_t > \delta E_t\left(p_{t+1}\right)\right) = E\left(\sum_{n=t_1(i)}^{t} v_{in}\right) = 0$, which combined with variation of prices over time and of the number of periods with and without purchases ensures the model parameters can be consistently estimated using OLS.

Note that in the purchase equation, the only unobservable variables that vary over time are the measurement errors. To control for potentially unobserved temporal shocks affecting purchase decisions, the estimated equations include time period fixed effects (season and year). These fixed effects capture part of the effect of unobserved variables that could affect consumers' decisions and are not explicitly accounted for in the model.

However, there may be remaining sources of endogeneity that I do not control for. For example, if the marginal utility of income (α_i) varies over time, then the error term is correlated with the total number of periods ($|T_i^{t-1}|$). Another example of endogeneity arises if there are utility differences across brands. In this case, the error term captures the brand differences and is therefore correlated with prices. I believe that this is not a serious issue here because the product categories considered in the empirical application are fairly homogenous (as discussed in Section 3.3). Otherwise, a possible solution would be to include interactions of brand fixed effects and period of purchase in the estimated equations.

In the appendix, I show that the variance of estimated measurement errors v_{it} is reasonably low. The opposite would indicate that the error term is capturing important, missing time-varying effects.

3.2.2 Estimation methods

To allow for consumer-specific marginal utility of income (α_i), I interact $|T_i^{t-1}|$ in (12) with the household indicator. However, as the panel is unbalanced, for some households the longitudinal dimension can be short (although households are observed at least 5 times) and the estimation of the household-specific fixed slope can be noisy. To address this issue, I also consider a linear mixed effect model (LME), in which α_i is treated as a Gaussian random slope coefficient rather than a fixed parameter. The fitted $\alpha_i's$ in the LME approach are equal to the household-specific fixed effects shrunk to the ensemble average. The degree of shrinkage is governed by the variance of the random effect, which is also estimated from the data. The shrinkage effect diminishes the variability of the estimates. However, it requires additional assumptions concerning the distribution of the random effects, as well as independence from the overall error term. The LME is estimated by maximum likelihood, and the random coefficients are estimated (predicted) using the best linear predictor. The standard treatment of random effects is provided in Raudenbush and Bryk (2001).

3.2.3 The price series

I estimate the purchase decision under the three price expectation formation hypotheses discussed in 3.1. Price expectations are household specific.[12] Household i's regular price p_r^i is the average price paid over the 3-year period.[13] In periods when i does not make a purchase, I input the price paid that period in i's favorite store, averaged across households living in the same region as i. When, in a certain period, there are no sampled purchases made in i's favorite store and region, I use the price in i's favorite store averaged across all regions. In the rare instances in which these

[12] Household-specific price expectations partly control for the relevant consumption set of the household. For example, if consumer i never shops at store A (which is too far from home), then the prices in store A should not affect her future price expectations or her quantity decisions.

[13] Although the model features a single variety of each product, there are different sizes and brands in the data, and prices may vary across these different varieties. To construct the average price paid over the three-year period, I consider the average across different brands and sizes purchased by the consumer over the period. When I consider the average price paid for the product by other households (in a certain store and region), I also compute the average price across different brands and sizes.

data are unavailable, I use the average price paid by other household in i's region.[14] Household i's favorite store is defined as the store at which i purchases more frequently. I use this same price input procedure when constructing expected prices under (E1) and (E2).

The regular price p_r^i differs from the average price p_s^i used to calculate consumption from stock. The latter is the average price conditional on purchases made by the household and approximates the average price paid per unit in storage.

3.3 Data

The database is a survey of households distributed across all regions of France. I use information on three years: 1999, 2000, and 2001. Each household is given a scanner with which to register every food product purchased. For each product purchased, there is information on quantities, price, date of purchase, and store, among other characteristics.

I study five product categories: butter, coffee, milk (not fresh but shelf-stable ultra pasteurized, which is the most common in France), pasta, and yogurt. The choice of product categories is made according to three criteria. First, I consider products that are consumed on a regular basis because consumption in the model is positive every period. Second, the products differ with respect to storage costs. For example, storage costs are higher for butter and yogurt because they need to be stocked in a refrigerated area, whereas this is not the case for coffee or pasta. Third, the products differ in terms of storability, or the length of time that a product can be stored. For instance, pasta can be stored for a longer period than coffee, which can be stored for longer than yogurt. Moreover, I account for potential measurement errors arising from the fact that I use a broad definition of product. I treat each category as a single product, capturing the fact that different brands are substitutes. If, however, the consumption of one of the brands is, for a certain household, independent of the consumption of another brand, then by treating the two brands as substitutes, I introduce measurement error in the construction of the beginning-of-the-period inventory level. I therefore consider categories that are relatively homogeneous, increasing the probability that different products are substitutes. An exception is, perhaps, yogurt because it is not clear that all households regard plain and fruit yogurt, for instance, as substitutes.

[14] Note that applying my method to a dataset that also includes store prices further simplifies the procedure because, in this case, prices are directly observed every period, unconditional on purchases. Examples of publicly available datasets that include both household and store level data are the Stanford Basket Data and the IRI Marketing Data Base.

In the data, one observation is one purchase made by the household. For each product category, I consider a subsample of households that purchased an item in that category at least once every six months during the three-year period. As mentioned in the previous section, the subsample used in the empirical exercise is further restricted to only include, for each household, the series of purchases made after the first purchase at a high price (relative to the expected future price). Moreover, as the identification strategy is based on the purchase equations for periods with high prices, I only include households that purchased at least five times at a high price.[15]

Table 1 reports descriptive statistics for the sample used for each product category, including number of households, average quantity per purchase in kilograms/liters, average price paid in euros per kilogram/liter, average interpurchase duration in days, and number of observations (purchases).

Table 1: Summary statistics

Products	number of households	quantity per purchase (kg)	price (€/kg)	days between purchases	Obs
Butter	2858	0.4913	4.896	14.19	161669
		(0.3353)	(1.040)	(15.25)	
Coffee	2353	0.5647	7.183	18.89	111924
		(0.3952)	(2.423)	(19.08)	
Milk	3382	6.386	0.6520	13.89	178614
		(5.490)	(0.2972)	(14.31)	
Pasta	3865	1.879	1.935	10.88	278649
		(1.305)	(0.8028)	(13.26)	
Yogurt	2881	1.063	1.742	17.96	125056
		(0.8854)	(0.9236)	(20.56)	

Notes: (i) quantity, price, and days between purchases are averaged across purchase occasions, standard deviation in parenthesis; (ii) for milk and yogurt, quantity and price are per liter and euros per liter, respectively.

Table 2 shows, for each product, the frequency of purchases made at prices above the expected future prices, as well as the fraction of total purchases made at high prices. These are the purchase occasions that are used to estimate the model parameters. There are three alternative definitions

[15] This may slightly bias the sample toward consumers with higher storage costs, implying that the differences between the estimated static and dynamic price elasticities could actually be even larger than what I find.

of high prices depending on the three price expectation hypotheses discussed above. Columns 3-4 refer to (E0) which assumes that future prices revert to their regular level. Columns 5-6 refer to (E1), in which expected prices are equal to the predicted prices in the simple Markov process $p_{it} = a + bp_{t-1}$, and columns 7-8 refer to (E2), in which prices are equal to the predicted prices in a similar process but in which parameters can vary according to whether the previous price was higher, lower or equal to the regular price.

Table 2: Fraction and frequency of purchases made at a high price considering three alternative price expectation hypotheses

	All	$p_{it} > Ep^{(0)}_{it+1}$		$p_{it} > Ep^{(1)}_{it+1}$		$p_{it} > Ep^{(2)}_{it+1}$	
Product	Freq	Freq	%	Freq	%	Freq	%
Butter	146,079	61,066	41.80	75,403	51.62	71,198	48.74
Coffee	97,884	56,539	57.76	64,671	66.07	58,014	59.27
Milk	137,261	60,576	44.13	57,341	41.78	57,602	41.97
Pasta	255,521	102,623	40.16	113,651	44.48	117,692	46.06
Yogurt	109,097	46,338	42.47	48,157	44.14	53,461	49.00

Note: $Ep^{(k)}_{it+1}$ is the price expectation under assumption (Ek), $k = \{0, 1, 2\}$.

Purchases at high prices under (E0) represent approximately $40\% - 50\%$ of the total sample of purchases. This percentage increases slightly for each product when I consider price expectation hypotheses (E1) and (E2). Note that although the subsample used for estimation represents approximately half of the original sample, it nevertheless includes a very large number of observations (from 20,000 up to more than 100,000 observations, depending on the product and the price expectation hypothesis). Note further that a fraction of high-price purchases close to 1 would imply stockpiling occurred on rare occasions, challenging the need to consider a dynamic model.

4 Results

4.1 Long-run and short-run price elasticities

If consumers stockpile, then in each period, they can purchase more or less than what they consume. Hence, due to inventories, purchases vary more than consumption. The short-run price elasticity captures these short-term purchase responses to prices. It coincides with the price elasticity in a static demand model that, by not considering that consumers can hold stocks, assumes that purchases and consumption are the same in every period. In contrast, the long-run price elasticity measures the effect of prices on long-run consumption. This elasticity can be measured using the estimated parameters of the dynamic model, which allows for stockpiling.

The expression for the long-run price elasticities is:

$$\epsilon_{it} = p_{it} \frac{h'(\alpha_i p_{it})}{h(\alpha_i p_{it})} \tag{13}$$

where $h = u'^{-1}$.

When $u(c_{it}) = -(1/\rho)\exp[-\rho c_{it}]$, (13) becomes:

$$\epsilon_{it} = \frac{1}{\ln \alpha_i + \ln p_t} \tag{14}$$

The static model I consider is a version of my demand model in which inventories (including beginning-of-the-period stocks) are restricted to zero every period. In every period, consumers decide how much to purchase, depending solely on current shocks and prices.[16] The expression for the short-run price elasticity is nearly identical to (14). The difference is that α_i is replaced by the marginal utility of income in the static model. Estimated long- and short-run price elasticities are obtained by plugging the estimated coefficients of the dynamic and static models into the price elasticity expression.

Table 3 presents demand price elasticities at the average price, averaged over households, for each product. The table also reports the ratio of short- to long-run price elasticities at the average price, averaged across households. I compute elasticities using estimated parameters from a fixed coefficients (OLS) and a random coefficients model (LME).[17] In each case, ε_0^{lr} refers to the long-run price elasticities obtained when considering (E0), ε_1^{lr} refers to the long-run price elasticities obtained when considering (E1), and ε_2^{lr} refers to long-run price elasticities when considering (E2). The term ε^{sr} refers to the short-run price elasticities. All means are calculated while excluding values in the extremes of the distribution (price elasticities below the 1% quantile and above the

[16] Therefore, $q_{it} = c_{it} = u'^{-1}(\alpha_i p_t)$ for all t. Assuming that $u(c_{it}) = -(1/\rho)\exp[-\rho c_{it}]$, the estimable equation is $q_{it} = \ln \alpha_i + \ln p_t + v_{it}$.

[17] Table 6 in the Appendix reports the average estimated marginal utility of income (α_i) and Table 7 presents summary statistics for the estimated measurement errors v_{it} for both the OLS and the LME models, as well as the standard deviation of $v_i = (v_{i1}, v_{i2,...})$.

99% quantile). Moreover, for each product and to facilitate comparison, ε_0^{lr}, ε_1^{lr}, ε_2^{lr} are averaged over the same sample of households and periods.

Table 3: Estimated long-run price elasticities at the average price and ratio of short- to long-run price elasticities considering different price expectation hypotheses

		Long-run price elasticities			Ratio short-run to long-run price elasticities		
Products	estimation method	ϵ^{lr0}	ϵ^{lr1}	ϵ^{lr2}	$\epsilon^{sr0}/\epsilon^{lr0}$	$\epsilon^{sr1}/\epsilon^{lr1}$	$\epsilon^{lr2}/\epsilon^{lr2}$
Butter	fixed effect	-1.488	-1.376	-1.354	2.083	2.193	2.197
	random coefficient	-1.312	-1.183	-1.383	2.371	2.463	2.119
Coffee	fixed effect	-1.457	-0.998	-1.012	3.181	4.229	4.198
	random coefficient	-1.333	-1.032	-1.078	2.742	3.439	3.280
Milk	fixed effect	-2.243	-2.035	-2.201	1.955	2.749	2.644
	random coefficient	-2.404	-1.900	-2.085	1.952	2.716	2.575
Pasta	fixed effect	-1.143	-1.056	-1.133	3.406	3.085	2.887
	random coefficient	-1.038	-0.983	-1.088	3.362	3.073	2.862
Yogurt	fixed effect	-2.239	-1.815	-1.922	1.920	1.951	1.942
	random coefficient	-1.879	-1.732	-1.828	1.899	1.972	1.837

Notes: (i) averages across households computed using elasticities in between the $1st$ and the $99th$ percentiles evaluated at the average price;

(ii) ϵ^{lrk} is the long-run price elasticity under price expectation hypothesis (Ek), k=0,1,2, sr denotes short run.

The ratios of short- to long-run price elasticities are always greater than 1. The upward bias of the short-run price elasticities goes from 80 up to more than 200% depending on the product, estimation method and price expectation hypothesis. Note that static and dynamic price-elasticities estimates differ for two reasons. First, the static model is misspecified because variables measuring inventories and expectations are missing. As a result, the static alpha estimates are biased and inconsistent. Second, even with the right controls, the static model measures short-run responses to price variation. The dynamic model, on the other hand, separates short-run inventory responses to price changes from long-run consumption responses.

The long-run price elasticities obtained using the random coefficients are generally lower than those resulting from OLS estimation. Furthermore, the long-run price elasticities under price expectation (E0) tend to be higher than those under (E1) and (E2), although the differences are not substantial. The differences across long-run price elasticities across the different price expectation hypotheses are less important when estimates are obtained using random coefficients. This is likely related to the fact that the random coefficients model shrinks estimates around the mean, reducing the variance.

The long-run price elasticity estimates above are comparable to Hendel and Nevo (2006b). They find long-run price elasticities all above 1, like I do. Their numbers are slightly higher, which is expected because they measure brand own-price elasticities whereas I measure category own-price elasticities, which tend to be lower.

4.2 Sensitivity Analysis

To reconstruct the beginning-of-the-period inventory level, I assume (C), i.e., that in periods without purchases, households consume at the average price level. I study the sensitivity of the empirical results to this assumption by considering subsamples with increasingly shorter interpurchase periods. I also estimate long-run price elasticities using a sample of subsequent purchases at high prices. In this subsample, purchase decisions do not involve inventories because both beginning- and end-of-the-period inventory levels are zero according to the model. Both exercises are tests of the sensitivity of the results to the assumption of consumption out of stock because: (i) the shorter the interpurchase period, the less frequently one needs to use assumption (C), and (ii) when initial and final stocks are zero, there is never consumption out of stock (and thus (C) is never used).

Note that the subsamples used in both exercises are restricted and possibly suffer from selection (for example, it is possible that there is an overrepresentation of households with a high cost of holding inventories, for whom it is difficult to avoid purchases at high prices). For this reason, one should expect to find differences in magnitude between the mean long-run price elasticities found using the whole sample of purchases at high prices and those reported below. However, it is important to determine whether the result that static price elasticities severely overestimate long-run price elasticities is still maintained and is thus robust to the assumption of consumption out of stock.

4.2.1 Shorter interpurchase periods

Table 4 reports the estimated long-run price elasticities at the average price for subsamples considering increasingly shorter interpurchase periods. I vary the maximum interpurchase period from one to four weeks and present estimated long-run price elasticities considering the four interpurchase lengths (1 period, 2 periods, 3 periods, and 4 periods). The estimates are obtained using fixed and random coefficients and under price expectation hypotheses (E0), (E1), and (E2). Table 8 in the Appendix reports the number of observations in each of the subsamples used in the estimation, while Tables 9 and 10 provide the marginal utility of income (α_i) estimates and the ratios of short to long-run price elasticities, respectively.

Given the price expectation hypothesis, varying each product's interpurchase period from 1 to 4 seems to have a limited impact on average long-run price elasticities (a possible exception is yogurt). However, the long-run price elasticities for a certain product and interpurchase period vary across the price expectation hypotheses. This suggests that, compared to those using the complete sample, estimates using subsamples with increasingly shorter interpurchase periods are more sensitive to the different price expectation hypotheses (see Table 3 in which price-elasticity differences across alternative price expectation hypotheses are less important). This is probably because the subsamples with a constrained number of interpurchase periods are smaller (which can be easily seen by comparing the sample sizes in Tables 2 and 8), which likely implies that there is less overlap between the subsamples of purchases at high prices associated with (E0), (E1), and (E2).

Although differences in price elasticities are small given a price expectation hypothesis, shortening interpurchase periods tends to slightly decrease mean magnitudes. This is even clearer when comparing mean long-run price elasticities in Table 4 with those calculated using parameter estimates from the whole sample (Table 3). There are two possible explanations for this. First, the shorter the interpurchase period, the less likely the subsample is to include periods with stockpiling that were misclassified as periods of purchases only to cover consumption (in those periods, price elasticities also capture variation in inventories and hence are generally higher). Second, subsamples with shorter interpurchase periods likely overrepresent households with high costs of holding inventories, which may be less price-elastic than the overall population, for example.

Finally, note that as long-run price elasticities tend to decrease with the number of interpurchase periods, the difference between short- and long-run price elasticities becomes even larger when

Table 4: Estimated long-run price elasticities at the average price with shorter interpurchase periods

Products	interpurchase period	Fixed Coefficients ϵ^{lr0}	Fixed Coefficients ϵ^{lr1}	Fixed Coefficients ϵ^{lr2}	Random Coefficients ϵ^{lr0}	Random Coefficients ϵ^{lr1}	Random Coefficients ϵ^{lr2}
Butter	1 period	-1.084	-1.129	-1.221	-0.993	-1.031	-1.070
	2 periods	-1.284	-1.072	-1.055	-1.102	-1.185	-1.301
	3 periods	-1.362	-1.078	-1.069	-1.160	-1.235	-1.352
	4 periods	-1.411	-1.070	-1.046	-1.209	-1.289	-1.391
Coffee	1 period	-1.386	-0.620	-0.521	-1.052	-0.353	-0.116
	2 periods	-1.393	-0.547	-0.519	-1.181	-0.518	-0.382
	3 periods	-1.422	-0.541	-0.479	-1.250	-0.588	-0.495
	4 periods	-1.394	-0.516	-0.504	-1.249	-0.609	-0.527
Milk	1 period	-2.218	-0.670	-0.568	-2.193	-1.145	-1.124
	2 periods	-2.224	-0.658	-0.600	-2.384	-1.219	-1.160
	3 periods	-2.183	-0.644	-0.595	-2.375	-1.251	-1.194
	4 periods	-2.287	-0.665	-0.609	-2.412	-1.295	-1.226
Pasta	1 period	-0.969	-1.027	-0.914	-0.888	-0.861	-0.842
	2 periods	-1.045	-1.058	-0.969	-0.952	-0.931	-0.930
	3 periods	-1.084	-1.049	-1.017	-0.989	-0.968	-0.975
	4 periods	-1.107	-1.038	-1.013	-2.365	-0.985	-1.002
Yogurt	1 period	-1.379	-0.917	-1.016	-1.511	-1.334	-1.195
	2 periods	-1.803	-0.869	-0.890	-1.604	-1.209	-1.126
	3 periods	-2.019	-0.721	-0.648	-1.677	-1.441	-1.244
	4 periods	-2.116	-0.811	-0.677	-1.734	-1.491	-1.270

Notes: (i) averages across households over elasticities in between the 1st
and the 99*th* percentiles evaluated at the average price;
(ii) ϵ^{lrk} is the long-run price elasticity under
denotes for short run; (iii) 1, 2, 3, 4
price expectation hypothesis (Ek), k=0,1,2, *sr* periods refer to
to the number of interpurchase periods.

parameters are estimated using the constrained subsamples (see Table 10 in the Appendix).

4.2.2 Subsequent purchases at high prices

The second sensitivity exercise relies on an alternative identification strategy based on periods in which the purchase decisions are free of dynamic considerations. In particular, when a consumer purchases a product at a high price for the second consecutive time, then Proposition 1 implies that current purchases equal current consumption and the problem reduces to a static one. More formally, let T_i' be the set of consecutive periods such that $p_{t-1} > \delta E_{t-1} p_t$, $p_t > \delta E_t p_{(t+1)}$, $q_{it-1} > 0$ and $q_{it} > 0$. Then, for all $t-1$ and t belonging to T_i', $y_{it-1} = y_{it} = 0$, and therefore, $q_{it} = c_{it}$. In this case, the equation to be estimated is:

$$q_{it'}^* = h\left(\alpha_i p_{t'}\right) + v_{it'} \tag{15}$$

where $t' \in T_i'$. When the utility function is CARA, as considered before, (15) becomes:

$$q_{it'}^* = -\frac{1}{\rho} \ln \alpha_i - \frac{1}{\rho} \ln p_{t'} + v_{it'} \tag{16}$$

Table 5 presents the estimated long-run price elasticities at the average price and the ratio of short- to long-run price elasticities when considering the subsample of subsequent purchases at high prices, and Table 11 in the Appendix presents the number of observations per product and price expectation hypothesis used in the estimation. Also in the Appendix, Table 12 reports the estimated $\alpha_i' s$, the marginal utility of income, for the subsample of subsequent purchases at high prices. The average long-run price elasticities considering E0 and E1 are generally similar for each product, independent of the estimation method (fixed effects or random coefficients). However, average price elasticities under E2 are consistently lower. Again, this seems to imply that price elasticity measures are sensitive to the price expectation hypothesis, especially when parameters are estimated using a subsample with shorter interpurchase periods. Moreover, mean elasticities in Table 5 are generally smaller than those in Table 3. As mentioned above, discrepancies between the estimates in the two tables should be expected because the latter uses a much smaller sample and a more selected group of households. The same intuitive arguments used to explain why price elasticities in Table 4 are smaller than the price elasticities in Table 3 are valid here. Accordingly, the upward bias of the short-run price elasticities is even more severe here.

Table 5: Estimated long-run price elasticities at the average price and ratio of short- to long-run price elasticities for subsequent purchases at high prices

		Long-run price elasticities			Ratio short-run to long-run price elasticities		
Products	estimation method	ϵ^{lr0}	ϵ^{lr1}	ϵ^{lr2}	$\epsilon^{sr0}/\epsilon^{lr0}$	$\epsilon^{sr1}/\epsilon^{lr1}$	$\epsilon^{lr2}/\epsilon^{lr2}$
Butter	fixed effect	-1.089	-1.143	-0.785	1.903	1.817	2.486
	random coef	-1.049	-0.992	-0.641	1.978	2.118	3.090
Coffee	fixed effect	-0.655	-0.820	-0.576	3.991	2.944	5.074
	random coef	-0.452	-0.667	-0.393	6.198	4.848	6.320
Milk	fixed effect	-0.810	-0.713	-0.632	4.868	6.528	5.675
	random coef	-0.852	-0.637	-0.588	4.161	6.442	7.467
Pasta	fixed effect	-1.145	-1.008	-0.748	2.669	2.550	3.776
	random coef	-1.069	-0.966	-0.650	2.576	2.629	3.540
Yogurt	fixed effect	-0.798	-0.829	-0.780	4.420	4.679	5.253
	random coef	-0.722	-0.766	-0.603	3.802	3.802	3.584

Notes: (i) averages over elasticities between the $1st$ and the $99th$ percentiles;
(ii) ϵ^{lrk} is the long-run price elasticity under price expectation hypothesis (Ek), k=0,1,2, sr denotes short-run.

5 Comparison with other Methods

Hendel and Nevo (2006a) derive and empirically test the implications of a model of demand with stockpiling. Although their empirical exercise is unlike mine, their theoretical model is similar. One difference is that their model includes a random preference shock. That is, per period utility is a function not only of consumption but also of an additive random shock, $u(c_{it} + v_{it})$. In my setting, the main result (Proposition 1) would still hold if there were a shock to utility. Further, I could derive the same estimable equations. However, if I considered v_{it} to be a preference shock to consumption, then in each period, the decision to purchase a positive quantity would depend on that period's preference shock. Therefore, in (12), the random component (which includes past shocks) would be correlated with the number of periods without purchases ($|T_{i0}^{t-1}|$), creating an endogeneity problem.

Hendel and Nevo (2013) develop a simple method to estimate demand with inventories that does not require solving the consumers' value function. They allow for product differentiation, and their estimates can be obtained using store-level data. The simplicity of their model relies on assuming that consumers are able to store for a pre-specified number of periods. Storage is free, and there is a constant proportion of consumers who never store. Consumers who stockpile only do so at sale prices, never at regular prices. The model is applied to weekly store-level data on purchases of 2-liter bottles of cola. The empirical application introduces consumer heterogeneity by allowing price sensitivity to vary between storers and non-storers. They consider two alternative price expectation hypotheses: perfect foresight and rational expectations.

Compared to their model, mine is more flexible in terms of consumer heterogeneity, as price elasticities and inventory costs in my model can be consumer specific. However, my model does not allow for multiple products, whereas they consider two products within the same product category. Allowing for multiple products does not change the main result that consumers never stockpile at high prices. However, it would require further assumptions concerning how consumers choose among different brands. One possibility would be to assume, as in Hendel and Nevo (2006b), that brand differentiation exists only at the moment of purchase (that is, products become homogeneous once they enter the home). This assumption implies that there is no need to track the different brands in storage, only of the total quantity stored. Then, one can estimate the discrete brand choice based on brand characteristics and current prices and plug the results from this first stage into the dynamic quantity decision, replacing prices with hedonic prices.

6 Conclusion

Estimating long-run price elasticities usually requires solving a complete dynamic programming problem. I propose a shortcut that greatly simplifies estimation. I consider a model in which consumers stockpile and prices are uncertain and show that consumers never stockpile when they purchase at a price higher than the expected future price. I focus on the purchase decision at periods of high prices and make an assumption on out of stock consumption. In this way, I am able to estimate the utility function parameters without having to solve the value function.

The empirical analysis uses a comprehensive dataset on household food purchases. The results are consistent with long-run price elasticties obtained with the full-blown estimation. I find that price elasticities resulting from a static demand model significantly overestimate the long-run price elasticities. I show that the results are robust to the assumption of consumption from stock and the different estimation methods. Price elasticities are also robust to the different price expectation hypotheses when the sample with an unconstrained number of interpurchase periods is used to estimate the parameters. However, price elasticities become more sensitive to the price expectation hypothesis considered as the number of interpurchase periods included in the estimation sample is reduced. This last result highlights the importance of appropriately characterizing how consumers form future price expectations. One of the advantages of the method proposed here is that it allows for a large number of alternative price expectation hypotheses, enabling the comparison of results across them.

7 References

Berry, Steven, James Levinsohn, and Ariel Pakes (1995). "Automobile Prices in Market Equilibrium", *Econometrica*, 63(4): 841-90.

Boizot, Christine, Jean-Marc Robin, and Michael Visser (2001). "The Demand for Food Products: An Analysis of Interpurchase Times and Purchased Quantities," *Economic Journal*, 111(470): 391-419.

Einav, Liran, Ephraim Leibtag, and Aviv Nevo (2010). "Recording Discrepancies in Nielsen Homescan Data: Are They Present and Do They Matter?" Quantitative Marketing and Economics 8 (2): 207–39.

Erdem, Tülin, Susumu Imai, and Michael Keane (2003). "Brand and Quantity Choice Dynamics under Price Uncertainty", *Quantitative Marketing and Economics*, 1: 5-64.

Hartmann, Wesley R.and Harikesh S. Nair (2009) "Retail Competition and the Dynamics of Demand for Tied Goods", *Marketing Science*, 29 (2): 366-386.

Hendel, Igal and Aviv Nevo (2006a) "Sales and consumer Inventory", *Rand Journal of Economics*, 37 (3): 543-561..

Hendel, Igal and Aviv Nevo (2006b) "Measuring the Implications of Sales and consumer Inventory Behavior", *Econometrica*, 74 (6): 1637-73.

Hendel, Igal and Aviv Nevo (2013) "Intertemporal Price Discrimination in Storable Goods Market", *American Economic Review*, 103(7): 2722-51.

McFadden, Daniel (1980) "Econometric Models for Probabilistic Choice Among Products", *Journal of Business*, 53(3), Part 2: Interfaces between Marketing and Economics: 13-29.

McFadden, Daniel (1984) "Econometric Analysis of Qualitative Response Models", in Z. Griliches and M. Intilligator, eds., Handbook of Econometrics, Vol. II, Amsterdam: North-Holland, 1396-1456.Nevo, Aviv (2001) "Measuring Market Power in the Ready-to-Eat Cereal Industry." *Econometrica*, 69 (2): 307-42.

Raudenbush, Stephen W. and Anthony S.Bryk (2001) *Hierarchical Linear Models: Applications and Data Analysis Methods (Advanced Quantitative Techniques in the Social Sciences)*, London: Sage Publications.

Sun, Baohong (2005) "Promotion Effect on Endogenous Consumption." *Marketing Science*, 24 (3): 420-43.

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

8 Appendix

8.1 Estimated marginal utility of income and measurement errors of main purchase equation specification

Table 6 reports the estimated marginal utility of income (α_i) averaged across households and for each price expectation hypothesis and estimation method using the full sample of purchases at high prices.

Table 6: Estimated marginal utility of income averaged across households

		Alpha estimates		
Products	estimation method	$E0$	$E1$	$E2$
Butter	fixed effect	0.085	0.073	0.082
	random coef	0.079	0.073	0.082
Coffee	fixed effect	0.059	0.044	0.046
	random coef	0.055	0.045	0.046
Milk	fixed effect	0.900	0.729	0.754
	random coef	0.882	0.733	0.751
Pasta	fixed effect	0.169	0.154	0.170
	random coef	0.156	0.156	0.170
Yogurt	fixed effect	0.311	0.277	0.291
	random coef	0.283	0.276	0.291

Notes: (i) averages across households computed using elasticities between the $1st$ and the $99th$ percentiles;
(ii) ϵ^{lrk} is the long-run price elasticity under price expectation hypothesis (Ek), k=0,1,2, sr denotes short-run.

Table 7 presents summary statistics for the estimated measurement errors v_{it}. The table also includes the standard deviations of $v_i = (v_{i1}, v_{i2,...})$ for each i, $sd_i(vt)$. Note that if the purchase equation were estimated for every period t, one could recover v_{it} for all t ($v_{it} = \sum_{n=t_1(i)}^{t} v_{in} - \sum_{n=t_1(i)}^{t-1} v_{in}$, for every t and i). However, I only have estimates of $\sum_{n=t_1(i)}^{t} v_{in}$ for periods of purchases at high prices, which means that v_{it} is identified only for purchase periods with high prices following another purchase period with high prices (in this case, I have estimates of both

$\sum_{n=t_1(i)}^{t} v_{in}$ and $\sum_{n=t_1(i)}^{t-1} v_{in}$). I calculate the variance of the measurement errors for this selected subsample. For conciseness, I report only results for the specification using fixed coefficients and price expectation (E0). Values using the other model specifications are quite similar.

Table 7: Estimated variance of the measurement errors

Product	Variable	Obs	Mean	Standard Deviation	Min	Max
Butter	v_{it}	40432	-0.005	0.375	-14.595	3.362
	$sd_t(v_i)$		0.322	0.208	0.000	3.428
Coffee	v_{it}	44912	0.002	0.433	-9.156	5.863
	$sd_t(v_i)$		0.366	0.237	0.000	2.016
Milk	v_{it}	44690	0.382	4.746	-123.412	97.807
	$sd_t(v_i)$		3.783	2.852	0.000	58.292
Pasta	v_{it}	59327	-0.014	1.343	-27.696	9.506
	$sd_t(v_i)$		1.193	0.684	0.000	8.515
Yogurt	v_{it}	27028	0.002	0.756	-11.639	10.325
	$sd_t(v_i)$		0.667	0.443	0.000	4.913

Notes: (i) summary statistics of estimated measurement errors; (ii) v_{it} are the estimated measurement errors recovered directly from estimated shocks, for the subsample of the second consecutive purchase at a high price; (iii) $sd_t(v_i)$ is the standard deviation of the distribution of household $i's$ measurement errors; (iv) estimates under price expectation hypothesis (E0), and fixed coefficients.

8.2 Sensitivity analysis: sample sizes, estimated marginal utility of income, and ratio of short to long-run price-elasticities

Table 8 shows the number of observations of the subsamples of purchases at a high price with a constrained number of interpurchase periods, for each product and price expectation hypothesis (E0, E1, and E2). For each price expectation hypothesis and product, there are 4 subsamples depending on the maximum number of periods between purchases, which varies from 1 to 4 periods. Table 9 reports the marginal utility of income averaged across households for each price expectation hypothesis using samples with increasingly shorter interpurchase periods, while Table 10 reports the ratio of short- to long-run price elasticities for the same subsamples. Note that this ratio is greater than 1 for all products.

Table 8: Number of observations in samples with shorter interpurchase periods

Products	Sample Size			
		$E0$	$E1$	$E2$
Butter	1 period	2700	17880	13361
	2 periods	41594	26748	21308
	3 periods	48616	30779	25446
	4 periods	52543	32971	28009
Coffee	1 period	13873	12364	10706
	2 periods	23388	20358	17837
	3 periods	28671	24510	21825
	4 periods	31931	26913	24286
Milk	1 period	2663	9922	7259
	2 periods	40570	14970	11518
	3 periods	47657	17496	14036
	4 periods	51620	18929	15654
Pasta	1 period	59864	38365	27013
	2 periods	81721	51425	38950
	3 periods	90375	56246	44966
	4 periods	94414	58542	48658
Yogurt	1 period	15127	7944	4651
	2 periods	24950	13006	8009
	3 periods	30964	16179	10609
	4 periods	34872	8175	121534

Table 11 displays the number of observations in the subsamples that include only subsequent purchases at high prices. For each product, that are 3 subsamples depending on the price expectation hypothesis (E0, E1, and E2). Table 12 reports the marginal utility of income averaged across households for each price expectation hypothesis using the subsample of subsequent purchases at

Table 9: Estimated marginal utility of income (alpha) with shorter interpurchase periods

Products	interpurchase period	Fixed Coefficients α^0	Fixed Coefficients α^1	Fixed Coefficients α^2	Random Coefficients α^0	Random Coefficients α^1	Random Coefficients α^2
Butter	1 period	0.067	0.079	0.075	0.064	0.068	0.071
	2 periods	0.077	0.084	0.074	0.069	0.077	0.084
	3 periods	0.080	0.076	0.076	0.072	0.079	0.087
	4 periods	0.082	0.075	0.074	0.074	0.082	0.088
Coffee	1 period	0.057	0.027	0.020	0.047	0.009	0.000
	2 periods	0.057	0.022	0.020	0.051	0.020	0.011
	3 periods	0.058	0.021	0.017	0.053	0.024	0.018
	4 periods	0.057	0.020	0.019	0.053	0.025	0.020
Milk	1 period	0.843	0.311	0.243	0.814	0.552	0.555
	2 periods	0.889	0.308	0.269	0.859	0.584	0.573
	3 periods	0.897	0.300	0.266	0.865	0.598	0.587
	4 periods	0.900	0.314	0.276	0.870	0.620	0.604
Pasta	1 period	0.147	0.177	0.157	0.138	0.138	0.137
	2 periods	0.157	0.183	0.168	0.146	0.150	0.152
	3 periods	0.162	0.182	0.177	0.150	0.155	0.159
	4 periods	0.164	0.180	0.177	0.152	0.158	0.165
Yogurt	1 period	0.405	0.177	0.196	0.251	0.224	0.214
	2 periods	0.340	0.166	0.170	0.259	0.239	0.224
	3 periods	0.326	0.133	0.112	0.266	0.251	0.230
	4 periods	0.312	0.154	0.121	0.270	0.256	0.235

Notes: (i) averages across households over estimates between the $1st$ and the $99th$ percentiles; (ii) α^{lrk} is the marginal utility of income under price expectation hypothesis (Ek), k=0,1,2;
(iii) 1, 2, 3, 4 periods refer to the number of interpurchase periods.

Table 10: Ratio of short-run to long-run price elasticities at the average price for shorter interpurchase periods

Products	interpurchase period	Fixed Coefficients $\epsilon^{sr}/\epsilon^{lr0}$	Fixed Coefficients $\epsilon^{sr}/\epsilon^{lr1}$	Fixed Coefficients $\epsilon^{sr}/\epsilon^{lr2}$	Random Coefficients $\epsilon^{sr}/\epsilon^{lr0}$	Random Coefficients $\epsilon^{sr}/\epsilon^{lr1}$	Random Coefficients $\epsilon^{sr}/\epsilon^{lr2}$
Butter	1 period	2.773	1.929	2.035	2.895	2.522	2.071
	2 periods	2.357	1.794	2.072	2.706	2.089	1.681
	3 periods	2.220	2.047	2.023	2.606	2.107	1.714
	4 periods	2.140	2.041	2.066	2.512	2.050	1.704
Pasta	1 period	3.877	2.247	2.490	3.735	3.522	3.625
	2 periods	3.586	2.137	2.336	3.585	3.174	3.189
	3 periods	3.490	2.154	2.212	3.513	3.113	3.203
	4 periods	3.432	2.181	2.204	3.484	3.064	3.056
Coffee	1 period	3.387	4.496	5.262	3.084	11.023	28.957
	2 periods	3.403	5.049	5.561	2.831	7.951	10.453
	3 periods	3.338	5.308	6.168	2.731	6.413	6.493
	4 periods	3.380	5.291	5.609	2.766	6.104	6.213
Milk	1 period	2.270	5.530	6.818	2.236	5.012	4.646
	2 periods	2.027	5.492	6.036	2.037	4.112	3.876
	3 periods	1.968	5.914	6.087	1.993	3.708	3.622
	4 periods	1.954	5.460	5.963	1.971	3.460	3.449
Yogurt	1 period	1.458	3.382	3.110	2.212	1.610	2.627
	periods	1.803	3.554	3.470	2.134	1.929	2.258
	3 periods	1.880	4.215	4.866	2.082	2.064	2.522
	4 periods	1.920	3.762	4.439	2.037	1.961	2.283

Notes: (i) considers price elasticities between the $1st$ and the $99th$ percentiles; (ii) $\epsilon^{sr}/\epsilon^{lrk}$ is the ratio of the short-run to long-run price elasticity under price expectation hypothesis (Ek), k=0,1,2; (iii) 1, 2, 3, 4 periods refer to the number of interpurchase periods.

high prices.

Table 11: Number of observations in sample of subsequent purchases at high prices per price expectation hypothesis

Products	Sample size		
	$E0$	$E1$	$E2$
Butter	17189	27784	16701
Coffee	9183	21133	13417
Milk	19032	19753	11727
Pasta	34200	49838	33586
Yogurt	8268	10573	8133

Table 12: Estimated marginal income (alpha) for subsequent purchases at high prices

		Alpha estimates		
Products		$E0$	$E1$	$E2$
Butter	fixed effect	0.078	0.081	0.055
	random coef	0.077	0.070	0.042
Coffee	fixed effect	0.029	0.040	0.024
	random coef	0.016	0.029	0.011
Milk	fixed effect	0.396	0.345	0.298
	random coef	0.429	0.292	0.267
Pasta	fixed effect	0.200	0.178	0.127
	random coef	0.192	0.168	0.106
Yogurt	fixed effect	0.155	0.162	0.149
	random coef	0.139	0.146	0.107

Notes: (i) averages across households computed using elasticities between the 1*st* and the 99*th* percentiles;
(ii) ϵ^{lrk} is the long-run price elasticity under price expectation hypothesis (Ek), k=0,1,2, *sr* denotes short-run.