

The Effect of More Accurate Cost-of-Living Indexes on Welfare Recipients

Emily Yucai Wang

January 28, 2013

Abstract

As an economic indicator, the Consumer Price Index (CPI) is most widely regarded as a measurement of inflation. It also plays an important role in adjusting payments to social welfare programs for millions of American households. As a fixed weight index on the cost-of-living, the CPI captures little intertemporal substitution introduced by temporary price reductions. This can be particularly problematic in the context of storable goods. Recent literature has sought to amend this shortcoming and improve the ability of the index in tracking price movements. Adopting these methods will likely result in a lower cost-of-living estimate since the current fixed weight index assigns too little weight to reduced prices. However, it is unclear whether the improved measurements, based on the behavior of the average household, should be used in adjusting social welfare payments. Households on welfare differ from the average household in that they are more likely to face binding budget and space constraints, both of which inhibit their ability to stockpile and thus lead to higher living expenses. Therefore, switching from a fixed weight to a more accurate price measurement, estimated on samples of average household, could drastically underestimate the true cost-of-living for the welfare recipients. I quantify the welfare impact of this switch by specifying and structurally estimating a dynamic model of consumer preferences for storable goods with flexibly incorporated unobservable persistent heterogeneity. Using the Supplemental Nutritional Assistance Program ("Food Stamps") as an example, I quantify the welfare effect of adopting a more accurate cost-of-living under two scenarios: 1) using only a sample of households on welfare and 2) using a sample of all representative households. I find households in each sample experience a welfare loss of 6% and 15% respectively. I discuss the implications of these results in the paper.

1 Introduction

Being one of the most important economic indices, the consumer price index (CPI) is used not only in measuring inflation but also in adjusting welfare payments for nearly all federal assistance programs. Given the need for a fast turnaround, the CPI has been calculated simply as a weighted average using a sample of monthly prices from a fixed basket of goods and a set of weights from representative households. This leads to shortcomings such as its failure to take consumers' substitution patterns into account. An abundance of research has been done to improve the accuracy of the CPI, including calculating utility based measurements such as a cost-of-living-index (COLI). Along with the increasing availability of scanner data, as well as the development of techniques that allow for the estimation of sophisticated models of consumer behavior, new methods are becoming better at tracking and approximating the price movements of household purchases. As this research is slowly adopted to improve the CPI, the resulting cost-of-living adjustments can have important implications on the wellbeing of households who depend on social welfare. If not carried out carefully, more accurate index measurements may end up decreasing the welfare of the very households who depend on it to adjust their income. This paper seeks to quantify the effects of such measurements on the welfare of poor households in the presence of storable goods with periodic temporary price reductions.

When products are storable, households may take advantage of sales by making extra purchases when prices are low and storing the additional items for later consumption. This type of consumer behavior creates intertemporal substitutions and causes standard price indices to deviate from indices that are derived from households' true cost-of-living. This causes a biased estimate of inflation which affects welfare payment adjustments. To see this, consider a storable product such as soft drinks. Consumers engaging in stockpiling tend to delay their purchases when prices are high and wait for reduced prices in a later period. This behavior is difficult to capture in a standard fixed-based index, where period t prices are weighted by period 0 quantities. Since promotions happen infrequently, these prices will on average receive less weight than non-promotional prices. However, when households anticipate promotions and wait for them, the standard fixed-based index should assign more weight to prices in promotional periods than to non-promotional periods. This implies that the prices accounted for in fixed-based indices are consistently higher than the actual prices paid and hence the measurement is biased upward. This further implies that welfare payments based on such indices are likely to overestimate the actual cost-of-living.

As the methodologies advance, more accurate measurements on the cost-of-living, such as Feenstra and Shapiro (2003) and Osborne (2012), are able to account for intertemporal substitution through stockpiling. With the adoption of these advancements, our ability to capture the price movements of the representative households would also improve. Assuming welfare recipients have similar purchase behavior as representative household, more accurate measurements of food costs will likely imply a decrease in welfare payment. However, poorer households' behavior may not match that of the average representative household. These households tend to live in smaller houses and have more limited transportation options. Both aspects restrict poorer households' abil-

ity to stockpile. Thus, if we use measurements derived from the average representative household in adjusting the welfare payments to the poor, we risk underestimating their cost-of-living.

This paper quantifies possible effects of using more accurate cost-of-living measures calculated using a sample of representative households on the welfare of recipients of federal assistance. I specify and structurally estimate a dynamic model of consumer preferences for storable goods with flexibly incorporated unobservable persistent heterogeneity. I use the 2002 household panel collected by Information Resources Inc. (IRI) on soft drinks to estimate parameters that govern households' consumption and purchase behavior. I then apply the resulting parameters in simulating welfare effects under different methods of constructing cost-of-living indices for a storable product with frequent price reductions.

To be more specific, using the Supplemental Nutrition Assistance Program (SNAP) as an example, I estimate the changes in welfare for households who qualify for these benefits under two counterfactual experiments. First, for households whose annual income is less than the federal poverty line I compare the their utilities under actual observed prices to utilities under monthly averaged prices. The actual observed prices represent index measures that are more accurate in tracking consumer price payments. The monthly average prices represent the standard fixed-base indices. I find that price paths tracked by standard indices leads to a 6% higher welfare than that under prices actually paid by poor households.

Since price indices are rarely computed using samples of poor households, this result can be misleading. As poor households tend to be more limited in their ability to stockpile, they may on average pay higher prices than the average household. Indeed, evidence from the data, the IRI 2002 weekly household soft drink purchases, suggests that households who live below the poverty line on average pay five cents more than others. Accounting for this fact, I again calculate welfare changes but instead use parameter estimates from all households in the sample. Results suggest that price paths tracked by standard indices leads to over 15% higher welfare than prices actually paid by households. This implies that if we adopt a more accurate price index measurement based on purchase behaviors of representative households, and use it to adjust welfare benefits, welfare recipients could be negatively impacted.

Estimating realistic parameters of households' utility functions can be computational challenging. Due to the intertemporal substitution generated by temporary price reductions, traditional static demand analyses misestimate long-run price elasticities. To avoid this, the dynamic demand model specified here explicitly accounts for consumers' inventory. On top of storage concerns, most products in the market are also differentiated. Literature on static demand models has shown that incorporating consumer heterogeneity into differentiated product demand systems is important in obtaining realistic predictions. However, due to the heavy computation burden, similar components have been difficult to implement in dynamic demand models of storable goods. The model specified here allows for heterogeneity as persistent random components in addition to observable attributes. Agents in the model differ not only in their observed characteristics but also in their unobservable product preferences and sensitivities to price and advertising. By allowing for unobservable

persistent consumer heterogeneity, I obtain flexible and realistic substitution patterns.

However, these models are computational intensive since they carry large state spaces, which include not only all current prices and promotional activities of all products and sizes but also the current period inventories and consumption shocks. Allowing for random coefficients in this dynamic demand model further exacerbates the computation burden. Since households' preferences for product characteristics are unknown, they have to be integrated out. That is, the model, the value function, and the choice probabilities need to be re-calculated for all households, heterogeneous tastes and preferences, and parameter tries, and then integrated over the distribution of all preferences. Estimating this model using a fixed point approach would take months, if not years. I follow Wang (2012) and adopt a simulated maximum likelihood approach and adapt importance sampling techniques developed in Akerberg (2003) to significantly reduce the computation time.

There is a large body of literature that seek to improve the CPI. One line of this literature concentrates on the inclusion of new products. These include Boskin et al (1998) and Hausman (1998). More recently, with the availability of scanner data, some researchers are pushing for more realistic models that capture actual prices paid for goods. Two of these are directly relevant for this project. Both seek to incorporate the intertemporal substitution that arises in storable goods into the calculation of CPI and other indices. Osborne (2012) estimates a storage goods demand model on data from canned tuna. He finds that accounting for storage leads to a cost-of-living index that averages 82% over the course of a year. In contrast, a fixed-based Laspeyres index averages about 90% over the course of the same year. Feenstra and Shapiro (2003) also use data on canned tuna but at a store level instead of a household panel. Their result is similar to that of Osborne (2012) with an index averages at 84% over the course of a year.

There are two strands of literature on structural demand models that this paper fits into: static demand models with consumer heterogeneity and dynamic demand models of storable goods. In terms of static demand models, the literature started by Bresnahan (1981) and Berry, Levinsohn, and Pakes (1995) (henceforth BLP) and continued on with Nevo (2000) has modeled static consumer decisions for differentiated products systems with heterogeneous products. These papers and those that have followed show that, in order to obtain realistic predictions for differentiated products, it is important to incorporate consumer heterogeneity in the demand systems. While this paper also incorporates consumer heterogeneity into the demand system, it differs from these papers in two ways. First, the demand model specified here is dynamic in nature. And second, whereas BLP-type models are adapted for aggregated market level data, this model is better suited for the more detailed disaggregated retail data.

This paper also fits into the literature on structural dynamic models of storable goods.¹ In particular, three papers are closely related to my work: Erdem, Imai, and Keane (2003), Hendel

¹There is also a line of literature on dynamic demand models of durable goods. Examples include Gordon (2006), Conlon (2009), and Nair (2007). Some of these papers, like Gowrisankaran & Rysman (2009) incorporate unobservable heterogeneous tastes into dynamic durable goods demand models. However, in most cases, durable goods are purchased once instead of repeatedly. The decisions that consumers make in these models focus on the timing of product replacement instead of stockpiling. Hence these models and methods are not appropriate in studying consumption storable goods.

and Nevo (2006b), and Hartmann & Nair (2010). Hartmann and Nair study the demand for razor blades. While this product is storable, they assume inventory is observable to a limited degree: when a new brand is purchased, old inventories disappear and the new inventory is the observed purchase. In this paper, however, inventory is unobservable and hence has to be estimated. Erdem, Imai, and Keane (2003) estimate a dynamic inventory model of demand for households' ketchup consumption. They build a model that allows for a limited degree of unobserved consumer heterogeneity. The model developed in this paper is more flexible in two respects. First, I allow consumption to depend on current and expected future prices. Second, while Erdem, Imai, and Keane (2003) take consumer heterogeneity into account, households are appropriated in a number of types and are identical within each type. This paper takes a random coefficients approach that more flexibly incorporates persistent consumer heterogeneity: households' preferences for product characteristics and prices/advertising follow distributions whose parameters are estimated within the model.

The papers by Hendel and Nevo (2006 a and b) are closest to my work. Both papers study households' purchases of laundry detergents. Hendel and Nevo (2006a) provides evidence for stockpiling and Hendel and Nevo (2006b) builds a demand model that allows households to stockpile. Their model incorporates heterogeneity as a function of observable household attributes. That is, households' tastes depend on their demographics such as the size of the household, race, and income. However, due to the aforementioned computational complexity, the authors do not allow for persistent random components which capture households' intrinsic preferences for products and prices. This means two households within the same demographic group share the same unobservable persistent preference. While it is important to account for observable household characteristics, it is not obvious that there is a large correlation between households' demographics and their intrinsic preferences. For instance, a poor household may have strong tastes for regular Pepsi while another may instead like diet Coke. I build on Hendel and Nevo's framework but incorporate heterogeneity both as a function of observable household attributes and as a function of unobservable persistent heterogeneity.

The rest of the paper is organized as follows. Section 2 presents industry details and the data. Sections 3 presents the model. Section 4 details the estimation procedure and the empirical results and examines the fit of the model. Section 5 provides a discussion on possible welfare implications of using more accurate price index measures and section 6 concludes.

2 Data

I use weekly scanner data provided by Information Resources Inc. (IRI) from 2002. The data is comprised of two components: a household panel and a store panel. The households in the data set were randomly selected.² For each household in each week, I observe whether any soda was purchased; if so, I observe which product was purchased, where it was purchased, how much was

²IRI randomly selects a sample of households in each of the two suburban areas around the Midwest and New England and requests their participation. The participating households conform to the local demographics fairly well. The data does not suggest a selection bias in the panel.

bought, and the total dollar amount paid. From the store panel, in each week, I observe the price charged, the total quantity sold, and all promotional activities for each product that was sold in the store. Sales from these stores account for over 97% of all purchases of soft drinks observed from the households in the panel. There are also two sets of attributes: one set for the households and one set for the products. For each household, I observe a few basic demographic variables, such as race, income, and the size of the households. For each product, I observe the product name, brand, packaging and volume.

For each observation from both the store and the household panel, unique identifiers on the product, store, and week allows me to link the two sets of data and track prices and promotions for all products available to households on any given shopping trip. That is, I observe not only information on the purchase itself but also information regarding each household’s complete choice set.

Household Panel I group households into two non-mutually-exclusive sets. The first sample includes only households whose aggregate annual income is below the federal poverty line and has at least three household members³. Table (1) below provides some information on the 2002 poverty guidelines along with the number of such households in the sample. The second sample includes all

Table 1: 2002 Household Poverty Guideline

Size of Family	Aggregate Income ¹	Number of HH in Sample
3	\$15,020	70
4	\$18,100	114
5	\$21,180	75
6	\$24,260	58
7	\$27,340	N/A ²
8	\$30,420	N/A ²
Total		317

1. Source: U.S. Department of Health and Human Services

2. IRI sample caps family size at 6. Households with more members are encoded as 6.

households in the data. I will refer to sample 1 as the low income households and sample 2 as all households. Table (2) below reports some statistics associated with each sample.

There are a few patterns to note in the data. First, the average household size in the sample with low income household appears to be much larger than the sample of all households. This is due to the fact the first sample is restricted to households with three members or above as mentioned previously. This also leads to a much higher average annual volume of soda purchased. In fact, converting all volume and size to 2-litre bottles, we see that the sample of poorer households

³In the data there is a large share of young households with one or two members. These households are unlikely to be welfare recipients and may still receive help from their parents. Hence, I limit the sample to households with at least three members.

Table 2: Household Demographics

	Low Income	All Households
<hr/> Demographics <hr/>		
avg Income (K)	\$15K	\$49K
avg household size	4.25	2.73
number of observations	317	1,546
<hr/> <hr/>		
Soda Purchase <hr/>		
avg annual vol purchased (2L)	97.20	69.68
avg annual dollars spend	\$143.86	\$97.55
avg price paid, (2L)	\$1.48	\$1.40
<hr/> <hr/>		
Brand Purchase Shares <hr/>		
Coke market share	46.81%	54.54%
Pepsi market share	50.64%	43.77%
Store Brand market share	2.56%	1.69%

purchase 27.52 more bottle of soda over the course of a year.

A more interesting observation is that poor households on average pay a slightly higher price for soda than the average household. This is true despite the fact that they purchase more private label items. There are two possible explanations for this. First, households with strong brand preferences substitute less to other cheaper brands and hence pay a higher price. Second, the poorer households may be more constrained in terms of storage spaces and transportation options and thus are more limited in their ability to take advantage of price reduction and intertemporally substitute.

Table (2) also presents a breakdown of the products households purchased in terms of the shares of Cola drinks (Coke, Pepsi, and store brand). It is clear that low income households purchase much less Coke than the average. Part of this could be explained by the fact that Coke prices are on average slightly higher than Pepsi and definitely much higher than Generic. Since poorer households are more price sensitive, those who do not strongly prefer Coke are more likely to purchase cheaper products.

Store Panel In the store panel, I observe weekly prices and advertising information for all items sold.⁴ For each store in each week, there are over 250 different products offered on average, a combination of all brands, regular/diet, packaging, volume, and flavors. In the estimation of the model, prices and promotions of all products become part of the state variables. Carrying prices and promotions of 250 products is computationally infeasible in the dynamic programming problem. Hence, I restrict the set of products used in estimation to include all cola products. I allow for all regular and diet versions of Coca-Cola, Pepsi Cola, and Private Label Cola drinks in different

⁴I do not directly observe the price of each product. Instead, I observe the weekly total revenue and the total units sold for every UPC (product). I take the average and use it as a proxy for the product price. In terms of promotional activities, for each product in every week, I observe whether it is on display, on feature, and/or has price discounts.

variations. Cola products are unique in flavor and thus are naturally distinct from other soft drinks. Alternatives such as lemon-lime soda or root beer are not close substitutes.

Table 3: Leading Brands by Market Share

Rank	Brand	2002 Market Share	2003 Market Share	2004 Market Share
1	Coke Classic (Coke)*	19.3	18.6	17.9
2	Pepsi-Cola (Pepsi)*	12.6	11.9	11.5
3	Diet Coke (Coke)*	9.0	9.4	9.7
4	Mt. Dew (Pepsi)	6.4	6.3	6.3
5	Diet Pepsi (Pepsi)*	5.5	5.8	6.1
6	Sprite (Coke)	6.2	5.9	5.7
7	Dr Pepper (Cadbury)	5.9	5.7	5.6
8	Private Label Cola*	2.1	2.0	1.7
9	CF Diet Coke (Coke)*	1.7	1.7	1.7
10	Sierra Mist (Pepsi)	0.7	1.4	1.4

This study includes items marked with * and other cola products.

Furthermore, as table (3) above shows, 6 of the 10 leading brands are cola related products. In 2002, cola products took 50.2% of the total market share of all soft drinks. Similar pictures are apparent in 2003 and 2004, where 49.4% and 48.6% of the total market respectively were cola drinks. While the total market shares are declining slightly over the years, none of the other products come close to compete.

Prices Table (4) presents the percentage of times I observe different prices offered in the same week between Coke and Pepsi conditioned on the type of product as well as different prices between regular and diet sodas of the same brand.

Table 4: Price Variation

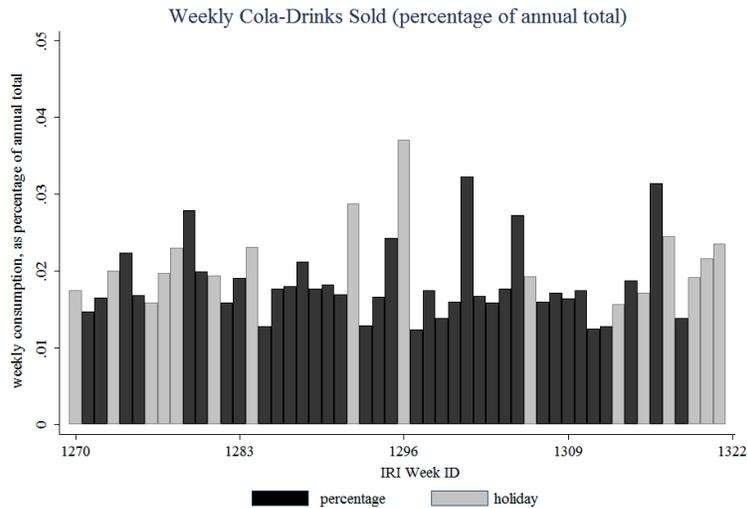
Price Difference	Coke vs Pepsi	Diet vs Regular
	Share	Share
no difference	19.63%	72.39%
(\$0 \$0.1)	9.26%	13.27%
[\$0.10 \$0.25)	36.63%	7.56%
[\$0.25 \$0.50)	24.95%	3.96%
[\$0.50 \$0.75)	8.83%	0.94%
[\$0.75 \$1.00)	0.62%	0.94%
[\$1.00 \$2.00)	0.08%	0.94%

note: price differences are in units of 2 liters.

From table (4), we see that most of the time Coke and Pepsi offer different prices. In fact the same price is offered at less than 20% of the time. Most of the price differences observed are between 10 to 25 cents. There are fewer price variations for products within the same brand. Regular and diet drinks of a given brand are priced differently about 30% of the time. Most of these price

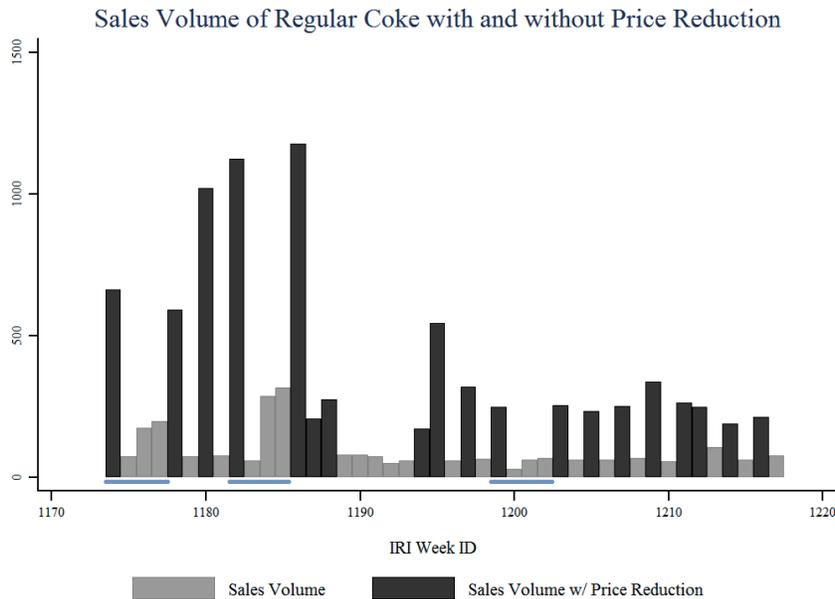
differences are small, mostly just below 10 cents. However, large differences do occur. We see that about 8% of the time regular and diet sodas of the same brands have price differences between 10 to 25 cents. These price variations are used to identify households' taste parameters.

Seasonality One concern of soft drinks is that their consumption could be influenced by holidays such as Christmas or Super Bowl Sunday. In the data, I observe a very small effect. In the following graph I show weekly cola product sales by volume, as a percentage of the annual total observed in stores. The gray bars indicate the presence of a holiday, where holidays are defined very liberally. For example, Super Bowl Sunday is counted as a holiday.



We see that there are increases in the volumes sold around some holidays. The dominant ones are around 4th of July and Thanksgiving. But this effect is strong. In this paper, I do not model holiday effects.

Stockpiling Behavior In this subsection I show evidence to support consumers' stockpiling behavior and argue that an inventory model is necessary for modeling the consumption of soda. The main competing model is one in which consumption is directly influenced by price and no storage occurs. If this were true, then we should expect to see the following pattern in the data: sales volume increases when price reductions are available but decreases and remains largely constant when there is no sale. This, however, is not the case. In the following graph we see that sales volume drops dramatically immediately after each sale ends; then it slowly grows again. This can be explained by stockpiling: households fill their inventories during price reductions. Hence, they do not need to purchase much soda immediately afterwards. As time passes households deplete their inventories. Therefore, more and more households have to restock.



As an example, I use the sale of regular Coke from a representative store in 2002. This graph shows how sales volume evolves over the duration of a year and how it is influenced by price reductions. Each bar shows the sale volume of a week. Black bars indicate price reductions while the gray ones are used for weeks without sales.

It is clear that the demand for regular Coke dramatically increases when there is a price reduction. Moreover, we see that the sales volume decreases drastically right after the price reductions; it picks up again in the following weeks. (These occurrences are marked by the horizontal lines beneath the bars.) These dips in purchases are consistent with households stocking up on soda when there are price reductions, reducing the need for buying soda immediately afterwards. However, as their stocks run low they make more purchases again. Static models do not account for this effect. They assume all purchased units are immediately consumed and hence are inherently misspecified.

To further distinguish between static and dynamic behavior I analyze how past prices influence current purchase size choices. To be more precise, I regress the quantity purchased on prices and the time since the last sale ("duration"). Table 5 below shows the results from this simple regression: in a static model, duration should not have any effect; purchase decisions only depend on current

Table 5: Regression of Quantity Purchased

	Coefficient
duration to last sale	10.8128* (3.6670)
price of product	738.137** (14.1617)
constant term	-1630.061** (35.1733)

Note: * at 95% significance level and ** at 99% significance level.

prices and marginal utilities of consumption. However, if the inventory model is correct, purchase quantities should be greater if more time passed since the last sale. The reason for this is that households are able to stockpile items not needed for immediate consumption. But when they go through a long time of high prices their inventories are depleted so that it takes a larger purchase to restock. Indeed, coefficients in both regressions are positive and statistically significant. This shows that purchases are highly correlated with past prices, justifying a dynamic specification.

Persistent Heterogeneous Preferences Lastly, I present evidence for persistent consumer heterogeneity. The following table shows the probability of observing brand purchases conditional on the observed last purchases:

Table 6: Transitional Probability of Purchase

Prob(brand in per. 2 brand in per.1)	
Prob(coke coke)	84.81%
Prob(pepsi coke)	15.19%
Prob(coke pepsi)	26.66%
Prob(pepsi pepsi)	73.34%

If households do not have intrinsic brand preferences, then the product choice in the second period should be independent to that of the first. That is the probability of buying product x in period $t + 1$ conditional on buying the same product in period t should be roughly the same as that conditional on buying product y in period t . However, this is not the case, as seen table (6). The probability of purchasing Coke in period $t + 1$ conditional on purchasing Coke in period t is 84.81%. The probability of purchasing Coke in period $t + 1$ conditional on purchasing Pepsi in period t , however, is only 26.66%. This suggests that households have persistent brand preferences.

3 Model

In line with Wang (2012), this is a dynamic discrete choice model with random coefficients. It flexibly incorporates observable consumer characteristics as well as unobservable persistent heterogeneity in preferences. Households in the model are forward looking. They form expectations over future prices and may choose to build inventories to ward off future price increases. In each period, they maximize the expected sum of their discounted future utilities by making three decisions: they choose how much they want to consume, how much they want to purchase, and if they chose to buy anything at all, which product they want to buy. Since soft drinks are storable, households are not forced to consume everything they purchase. They choose the utility maximizing consumption level taking prices, promotions, and available inventories into account. All leftover products become part of the storage for future consumption. Households face two sources of uncertainty in this model:

random and potentially nonlinear future prices (along with promotional activities) and future utility shocks.

The decision period is one week. In each period, the timing of events goes as follows: first, households observe their inventory, obtain information on prices and promotional activities of all products in the market, and receive privately drawn utility shocks associated with consumption and product choices. With information on the current prices and advertising activities, they form expectations over future prices and promotions. During their supermarket visits, they choose the level of consumption and the amount to purchase. Should they make any purchase, households will base their product choices on their preferences, prices, available promotions, and the previously drawn shocks. At the end of each period, households incur costs associated with their leftover inventories.

Broadly speaking, heterogeneity is present in the model in two ways. First, households are differentiated according to their observable attributes, such as their income, race, etc. Second, they differ by their unobservable persistent tastes for product characteristics, such as brand, as well as their sensitivities for prices/advertising and their abilities to stockpile. As shown in the following subsection, this specification of the unobservable tastes allows for a full correlation structure between brand preferences. It provides unrestricted substitution patterns across all products and hence presents a very flexible way of incorporating household taste heterogeneity.

3.1 Model Setup

Household h 's problem can be represented by the following infinite horizon maximization problem:

$$V^h(\phi_1) = \max_{\{c^h(\phi_t), s^h(\phi_t), i^h(\phi_t)\}} \sum_{t=1}^{\infty} \delta^{t-1} E \left[u_{t si}^h | \phi_1 \right] \quad (1)$$

where $V^h(\phi_1)$ denotes household h 's present value of future utility for a given state ϕ_t at time $t = 1$. The state space for period t , ϕ_t , consists of current prices and promotions, the beginning-of-period inventory, a shock to utility from consumption, and an idiosyncratic shock to utility for each product. In each period, households maximize their objectives with respect to the level of consumption c^h , the volume of purchase s^h , and the product choice i^h . I use the standard notation δ to denote the discount factor and the small letter u for the flow utility, where household h 's flow utility for period t is defined as:

$$u_{t si}^h = U \left(c_t^h, v_t^h | \Theta^h \right) - C \left(inv_{t+1}^h | \Theta^h \right) + U \left(p_{t si}, a_{t si}, i_{t bd}^h, \varepsilon_{t si}^h | \Theta^h \right). \quad (2)$$

It is comprised of three components: the utility of consumption $U(c_t^h, v_t^h | \Theta^h)$, the cost of inventory $C(inv_{t+1}^h | \Theta^h)$, and the utility to purchase $U(p_{t si}, a_{t si}, i_{t bd}^h, \varepsilon_{t si}^h | \Theta^h)$. The set of household specific taste parameters is denoted by Θ^h .

The utility of consumption is a function of c_t^h , the amount household h consumes in period t , and v_t^h , a stochastic shock to utility that changes the utility from consumption. The consumption

is equal to the total of all i products consumed by household h in period t . That is, $c_t^h = \sum_i c_{it}^h$. The shock v_t^h households receive in each period introduces randomness in households' need to consume and allows their consumption to depend on factors that are unobserved by the researcher. It influences the utility in the following way. If household h draws a high realization of v_t^h , then its marginal utility of consumption would be smaller than if a low realization were drawn. Therefore, a high realization decreases households' need to consume and hence makes demand more elastic. I make the following assumption regarding v_t^h 's distribution:

Assumption 1. *The consumption shocks, v_t^h , are distributed independently and identically across households and over time.*

In principle, serial correlation in v_t^h can be accommodated, but at a significant increase in computational burden. I follow previous literature in assuming that the utility of consumption takes the following form:

$$U(c_t^h, v_t^h | \Theta^h) = \alpha^h \log(c_t^h + v_t^h) \quad (3)$$

where the consumption level and the stochastic shock enter additively into the utility. Unlike previous literature, I allow the marginal utility of consumption, α^h , to be household specific.⁵

Households' cost of inventory follows a quadratic functional form:

$$C(inv_t^h | \Theta^h) = \beta_1^h inv_t^h + \beta_2^h (inv_t^h)^2 \quad (4)$$

where β_1^h and β_2^h are household specific linear and quadratic cost terms. These parameters represent households' abilities to store leftover products: households with higher marginal costs of inventory find storage more costly and hence might not be able to take advantage of sales as readily as households with lower β_1^h and β_2^h . Following simple accounting rules, the end-of-period inventory, inv_{t+1}^h , is equal to the beginning-of-period inventory, inv_t^h , plus the total current purchase, s_t^h , minus the total consumption, c_t^h . That is, $inv_{t+1}^h = inv_t^h + s_t^h - c_t^h$. Since households cannot consume more than what is available in their storage or make negative purchases, the total inventory, inv_t^h , and purchase, s_t^h , are restricted to be non-negative.

The last piece of the flow utility is associated with household h 's purchase in period t , which is a function of prices, promotional activities, product preferences, and a random shock to the household's choice. It can be represented as:

$$U(p_{t si}, a_{t si}, i_{t bd}^h, \varepsilon_{t bs}^h | \Theta^h) = -\gamma_1^h p_{t si} + \gamma_2^h a_{t si} + \sum_b \mathbb{1}(b_t^h) \zeta_b^h s_t^h + \varepsilon_{t si}^h, \quad (5)$$

where $-\gamma_1^h p_{t si}$ represents household h 's disutility of price. It is equal to the negative of household h 's marginal utility of income, γ_1^h , multiplied by the price, $p_{t si}$, for product i of volume s in time t . The

⁵Technically, the marginal utility of consumption is $\alpha^h / (c_t^h + v_t^h)$, which depends on both the consumption level and the consumption shock. For simplicity, I call α^h the marginal utility of consumption.

second component $\gamma_2^h a_{t si}$ captures the effect of advertising on demand. Variable $a_{t si}$ is an indicator. It denotes the presence of any features or displays. Households in the model are forward-looking. They form expectations over the uncertain future prices and promotions. For simplicity, I assume the transitions of these prices and promotions follow the assumption below:

Assumption 2. *Prices and advertising activities follow an exogenous first order Markov process.*

This assumption is often employed in the empirical literature to reduce the state space in the dynamic programming problem. It is a simplifying assumption on households' expectations for future prices and promotions and implies that consumers use only the current period prices and promotions in predicting future prices and promotion. This assumption seems to be a reasonable approximation of the formation of consumers' price expectations. Higher order processes would require households to recall more prices from previous visits. The main concern might be seasonal price fluctuations, when the probability of advertising increases. However, as described in the data section, there is no clear evidence of significant seasonal effects in the household panel. The third component, $\mathbb{1}(b_t^h) \zeta_b^h s_t^h$, represents household h 's utility for consuming s units of brand b in time t . It is equal to household h 's intrinsic preferences for brand b , ζ_b^h , multiplied by the volume purchased, s_t^h , in period t . The last component, $\varepsilon_{t si}^h$, denotes an idiosyncratic shock to household h 's product choice, which follows the commonly used distribution in discrete choice problems:

Assumption 3. *The product choice shocks, $\varepsilon_{t si}^h$, follow a type I extreme value distribution and are independent and identically distributed across household, period, and purchase.*

The idiosyncratic shock absorbs the apparent randomness in consumers' purchases that are based on factors unobserved by the econometrician. This assumption significantly reduces the computational burden. It could be relaxed; however, since households' product preferences are explicitly estimated and the volumes of purchases are allowed to be correlated across periods, assuming the shock is i.i.d. across households, time, and purchase is reasonable.

In line with Hendel and Nevo (2006b) and Wang (2012), product differentiation in this model takes place at the moment of purchase. While households have intrinsic preferences for different brands and diet drinks, these preferences and the differences between products affect households' behaviors exclusively at the store but do not give different utilities at the moment of consumption. Since consumptions and inventories are not observed, this specification avoids putting extra structure on consumption rates of different items in the inventory. Furthermore, this assumption significantly decreases the state space. Instead of carrying the whole vector of inventories for each brand, only the total quantity in stock is tracked. Differentiation at purchase, represented by $\mathbb{1}(b_t^h) \zeta_b^h s_t^h$ captures the expected value of the future differences in utility from consumption. More specifically, the term $\zeta_b^h s_t^h$ captures the expected utility from future consumption of s units of brand b at the time of purchase.

4 Estimation

Solving the infinite horizon maximization problem faced by each household in each period can be computationally demanding and time consuming, especially in the presence of a large state space. Several key factors of the model further complicate this process. First, current period inventories are unknown since neither consumption levels nor initial inventories are observed. To calculate the current inventory, both the current consumption and the initial inventory have to be estimated. Second, since all parameters are random coefficients, the likelihood has to be integrated over the joint distribution of these parameters. The necessity of having to solve the consumers' dynamic programming problem at every guess of the parameters and for each possible consumption, inventory level, and product choice presents a significant computation burden. Using traditional maximum likelihood methods implies that the estimation for a set of parameters would take months to compute. I nest a simulated maximum likelihood approach into the importance sampling method introduced by Akerberg (2009) to reduce the computation time. I start the discussion of the estimation by providing an overview and then move on to a more detailed description of the technical details.

4.1 Overview

Recall from the data section that we observe only the purchases by households over time. In each period, we know whether household h made any purchases. If a purchase was made, then we observe which product was purchased and how many units were bought. However, no information is available on households' consumptions, inventories, or their intrinsic tastes. The goal is to maximize the probability of observing this sequence of purchases with respect to the parameters that govern the distribution of the random coefficients, the initial inventory, as well as the distribution of household type.

Assume for now that we are in a world where the econometrician observes not only the purchases but also both the households' preferences and their initial inventories. Then, for a given set of parameters, the probability of observing a sequence of purchase decisions, pur_t^h , across all households and over time, $(pur_1^1, \dots, pur_T^1, \dots, pur_1^H, \dots, pur_T^H)$, as a function of the state variables, $(\phi_t^1, \dots, \phi_t^H)$, is given by the following integral:

$$L = \int_{\Theta^h} \Pr(pur_1^1 \dots pur_T^H | \phi_t^h, inv_0^h, \Theta^h) dF(\Theta^h | \mu_\Theta, \Sigma_\Theta)$$

where $\phi_t^h = \left\{ p_t, a_t, v_t^h, inv_t^h(pur_1^h \dots pur_{t-1}^h, v_1^h \dots v_{t-1}^h, inv_0^h) \right\}$

The state ϕ_t^h includes the vector of prices and promotions, p_t and a_t ; the current period consumption shock, v_t^h ; and the beginning of the period inventory, inv_t^h . The inventory is a function of previous decisions and states. It depends on last period's consumption and purchase decisions, which in turn depends on the consumption shocks and the idiosyncratic purchase shocks of that period as well as

all earlier periods' inventories.

Assuming households are independent from each other, the probability of observing a sequence of purchases from household h is then independent from that of household h' . Hence, the probability of observing the entire purchase sequence of all households can be re-written as:

$$L = \prod_{h=1}^H \int_{\Theta^h} \Pr \left(pur_1^h, \dots, pur_T^h | \phi_t^h, inv_0^h, \Theta^h \right) dF \left(\Theta^h | \mu_{\Theta}, \Sigma_{\Theta} \right) \quad (6)$$

Furthermore, recall that purchases can be decomposed into product purchases, i_t^h , and volume purchases, s_t^h . That is $\Pr \left(pur_1^h, \dots, pur_T^h | \phi_t^h, inv_0^h, \Theta^h \right) = \Pr \left(s_1^h \dots s_T^h, i_1^h \dots i_T^h | \phi_t^h, inv_0^h, \Theta^h \right)$. Since the product choice shocks are i.i.d., households' brand purchases are independent over time. Therefore rewriting the probability of observing the sequence of purchases once again, we arrive at:

$$L = \prod_{h=1}^H \int_{\Theta^h} \left[\prod_{t=1}^T \Pr \left(i_t^h | s_t^h, \phi_t^h, \Theta^h \right) \right] \Pr \left(s_1^h \dots s_T^h | \phi_t^h, inv_0^h, \Theta^h \right) dF \left(\Theta^h | \mu_{\Theta}, \Sigma_{\Theta} \right) \quad (7)$$

From this equation, it is straight forward to see that household's two purchase decisions, product and size, are treated differently. The product choices, i_t^h , are uncorrelated across periods, hence the decision making process is static. That is, conditional on households' intrinsic preferences, their product choices in period t are independent from those in period $t + 1$. By assumption 3, product shocks are distributed according to a Type I extreme value distribution. Therefore, the probability of seeing any particular purchase follows the standard logit formula:

$$\Pr \left(i_t^h | s_t^h, \phi_t^h, \tau^h, \Theta^h \right) = \frac{\exp \left[\left(-\gamma_1^h p_{tsbd} + \gamma_2^h a_{tsbd} + \sum_b \mathbb{1}(b_t^h) \zeta_b^h s_t^h \right) \right]}{\sum_b \sum_d \exp \left[\left(-\gamma_1^h p_{tsbd} + \gamma_2^h a_{tsbd} + \sum_b \mathbb{1}(b_t^h) \zeta_b^h s_t^h \right) \right]} \quad (8)$$

The volume purchase decisions, s_t^h , on the other hand, depend on current consumption as well as past consumption levels, purchase decisions, and initial inventories. Consumptions and volume purchases are determined by households' optimal policies and are correlated over time. Hence, the probability of observing any sequence of volumes purchased depends on that household's optimal policy rule, which is a dynamic programming problem and cannot be estimated separately for each period.

Recall, households' consumption shocks, v_t^h , and initial inventories, inv_0^h , are not actually observed. Therefore, to estimate the model, we need to integrate equation (7) over the distribution of consumption shocks and initial inventories. Mathematically, this can be expressed as:

$$L = \prod_{h=1}^H \left\{ \int_{\Theta^h} \left[\int_{v} \int_{inv_0} \left(\prod_{t=1}^T \Pr \left(i_t^h | s_t^h, \phi_t^h, inv_0^h, \Theta^h \right) \Pr \left(s_1^h \dots s_T^h | \phi_t^h, inv_0^h, \Theta^h \right) \right) dF \left(inv_0^h \right) dF \left(v_1^h \dots v_T^h \right) \right] dF \left(\Theta^h | \mu_{\Theta}, \Sigma_{\Theta} \right) \right\} \quad (9)$$

Calculating the above likelihood function is a 2-step process. First, solve the dynamic programming problem and find the optimal policy rule that govern households' optimal consumption and purchase choices. Second, compute households' choice probabilities. Both of steps are carried out for each household, consumption shock, initial inventory, and parameter try. This becomes a very time consuming process. Part of the computation burden comes from the unobserved initial inventory; part from the random coefficients whose parameters govern the distribution of households' persistent preferences; and part from carrying a large state space. I discuss these complicating factors in detail below.

Households' initial inventories are unknown to the econometrician and hence need to be estimated. Assume, for now, that we observe households' inventory at the beginning of period one. Then once the optimal consumption and purchase choices are estimated, we can calculate the end of period inventories for the first period. Roll this process forward and we can obtain all inventories across all periods. Estimating the entire sequence of inventories boils down to an initial conditions problem.

One way of estimating the initial inventory is to start at an arbitrary level and use the first few weeks of observations to generate the distribution of inventories. To obtain a good estimate using this method, the panel has to be relatively long. For example, Hendel and Nevo (2006b) take this approach. Their panel spans 2 years and they use the first 11 weeks to recover estimates for the initial inventories. In this paper, however, the length of panel is one year. Hence, I adopt a different approach. I use the model and observed households' purchase behaviors to explicitly estimate a distribution over the initial inventory. I assume initial inventories follow the distribution specified below:

Assumption 4. *The initial inventories inv_0^h are distributed log-normal with mean μ_{inv} and standard deviation σ_{inv} .*

Households' persistent preferences are also unobserved. These heterogeneous preferences are modeled using random coefficients which follow distributions whose parameters are estimated in the model. I assume the distribution follows the assumption below:

Assumption 5. *The random coefficients $\Theta^h = \{\alpha^h, \beta_1^h, \beta_2^h, \gamma_1^h, \gamma_2^h, \zeta_C^h, \zeta_P^h\}$ are distributed jointly log-normal with mean μ_Θ and variance-covariance matrix Σ_Θ .*

These random parameters capture households' persistent heterogeneous tastes. Since the econometrician does not observe them, the joint distribution of these preferences is estimated. These parameters are naturally bounded on either side of the real line. For example, the marginal utilities of consumption are non-negative. Hence, the log normal distributions are fitting. I further assume that the off diagonal elements of the variance-covariance matrix are zero. This restriction can be lifted at the cost of an increased computation burden.

Solving the dynamic programming problem and calculating the choice probabilities for one household at an initial inventory level for a given set of parameter try is not time consuming, per

se. This process takes roughly 10 seconds once.⁶ However, if we adopted the standard maximum likelihood estimation approach, this computation process is repeated for each household, initial inventory guess, and parameter try.⁷ The repetition quickly increases the computation time to become unmanageable. I adopt a simulated maximum likelihood approach and build on the importance sampling method proposed by Akerberg (2009) to significantly cut down computation time. In estimating this model, allowing for parallel computing on 8 cores, the entire estimation process takes about 3 days.

The reduction in computation time predominantly comes from the fact that importance sampling allows computing the value functions and the choice probabilities to be performed outside the maximization routine. This speeds up the calculation in two ways: 1) it allows for easy parallel computing of all dynamic programming problems across all different household types and parameter draws; this further allows the set of choice probabilities for all households and all parameter tries to be computed simultaneously. 2) since the value function calculations are not nested within each parameter search, the search routine costs very little time. I fully describe the estimation procedure below.

4.2 Estimation Procedure

The estimation procedure can be roughly broken into five steps.

Step 1 (Initial Parameter Distribution) Importance sampling methods require specifications of initial distributions over parameters, which are updated during the convergence process. The initial distributions can be chosen arbitrarily at the discretion of the econometrician as long as they have a fairly large range that well covers the true parameter value.⁸ To start the estimation process, I pre-define an initial joint distribution of the random coefficients, $g\left(\tilde{\Theta}^h \mid \mu_{\Theta}^0, \Sigma_{\Theta}^0\right)$, and a distribution over the initial inventory, $g\left(\tilde{inv}_0 \mid \mu_{inv_0}^0, \sigma_{inv_0}^0\right)$, where the parameters that govern these distributions, $(\mu_{\Theta}^0, \Sigma_{\Theta}^0)$ and $(\mu_{inv_0}^0, \sigma_{inv_0}^0)$, are chosen arbitrarily. Then I draw N sets of parameter values and initial inventories according to the pre-defined distributions. I denote the randomly drawn parameters by $\tilde{\Theta}^h$ and the randomly drawn initial inventory by \tilde{inv}_0 .

Step 2 (Product Choice) With the randomly drawn parameter tries, we can calculate the dynamic programming problem and compute the choice probabilities associated with each household's observed sequence of purchases. Recall, household's purchase decisions are decomposed into a product choice, a static component, and a volume choice, a serially correlated component. Since product choices are static, conditional on the quantity purchased, they follow the standard logit formula. For each household, household type, and parameter draw, the probability of observing any product choice is:

⁶Using a standard PC with 8 cores and 2.80GHz.

⁷Recall that there are roughly 800 households in the data and 9 possible household types. Even at a convergence criterion of 1e-3, the maximization routine generally takes at least 3000 iterations. That is, the calculation is repeated at least 21 million times and thus takes 216 million seconds, which is about 6 years.

⁸It's important to choose distributions that have a sufficiently large variance. Otherwise, certain parameters may be excluded from being drawn. This would lead to misestimation if the true parameters lay outside the range of the initial parameter draws. For a discussion of how to choose initial distributions, refer to Akerberg (2009).

$$\Pr \left(i_t^h | \phi, s_t^h, \tilde{\Theta}_{sim}^h \right) = \frac{\exp \left[\left(-\tilde{\gamma}^h p_{tsbd} + \tilde{\gamma}^h a_{tsbd} + \sum_b \mathbb{1}(b_t^h) \tilde{\zeta}_b^h s_t^h \right) \right]}{\sum_b \sum_d \exp \left[\left(-\tilde{\gamma}^h p_{tsbd} + \tilde{\gamma}^h a_{tsbd} + \sum_b \mathbb{1}(b_t^h) \tilde{\zeta}_b^h s_t^h \right) \right]}$$

Step 3 (Dynamic Programming Problem) Households' decisions on how much to purchase depend on their current level of inventory, which in turn depends on past consumption and purchase decisions. Hence these volume choices are correlated over time. The Bellman equation equivalent to households' infinite horizon maximization problem is:

$$\begin{aligned} & V \left(\phi_1^h \right) \\ = & \max_{\{c, s, i\}} \left\{ \left[\tilde{\alpha}^h \log \left(c + v^h \right) - \left(\tilde{\beta}_1^h inv' + \tilde{\beta}_2^h inv'^2 \right) - \tilde{\gamma}_1^h p_{si} + \tilde{\gamma}_2^h a_{si} + \sum_b \mathbb{1}(b^h) \tilde{\zeta}_b^h s^h + \varepsilon_{si}^h \right] \right. \\ & \left. + \delta E \left[V \left((\phi^h)' \right) | \phi_1^h, c, s, i \right] \right\}. \end{aligned} \quad (10)$$

The statespace associated with this model is large. It encompasses not only all prices and promotional activities of all products and sizes but also the consumption shocks and initial inventories. To be able to estimate the I approximate the expected future utilities by using the sieve value function iteration method developed in Arcidiacono, Bayer, Bugni, and James (2012).⁹ To most accurately approximate the dynamic programming problem, I use all interactions.

With the computed optimal consumption and purchase decisions, the probability of observing any sequence of volume purchases, for a given household type and parameter draw, is:

$$\begin{aligned} & \Pr \left(s_1^h \cdots s_T^h | \phi_t^h, inv_0^h, \Theta^h \right) \\ = & \prod_t \Pr \left(s_t^h | p_t, v_t^h, inv_t^h \left(s_1^h \cdots s_{t-1}^h, v_1^h \cdots v_{t-1}^h, inv_0^h \right), \Theta^h \right) \\ & \text{where } \Pr \left(s_t^h | p_t, v_t^h, inv_t^h \left(s_1^h \cdots s_{t-1}^h, v_1^h \cdots v_{t-1}^h, inv_0^h \right), \Theta^h \right) \\ & = \frac{\max_{c_t^h} \left\{ u_t^h + \delta E \left[V \left((\phi^h)' \right) | \phi^h, c_t^h, s_t^h \right] \right\}}{\sum_{\tilde{s}} \left\{ \max_{c_t^h} \left\{ u_t^h + \delta E \left[V \left((\phi^h)' \right) | \phi^h, c_t^h, \tilde{s}_t^h \right] \right\} \right\}} \end{aligned} \quad (11)$$

After combining both components of households' purchase decisions, calculating households' choice probabilities is fairly straight forward.

Step 4 (Choice Probability) For each subset of household and parameter draws, the choice

⁹In comparing the approximation with the exact solution obtained from policy function iteration using a simpler model on a smaller state space, I find the EV obtained by sieve value function iteration closely approximates the exactly calculated solution.

probabilities conditional on household h 's type and its beginning-of-period inventory is given by:

$$\begin{aligned} & \tilde{L}_{sim}^h \left(pur_1^h \cdots pur_T^h | \phi_t^h, inv_0^h, \tau^h, \Theta^h \right) \\ &= \left[\prod_t \Pr \left(i_t^h | \phi_t^h, s_t^h, \tilde{\Theta}_{sim}^h \right) \right] \Pr \left(s_1^h \cdots s_T^h | \phi_t^h, inv_0^h, \tilde{\Theta}_{sim}^h \right) \end{aligned} \quad (12)$$

This conditional choice probability is equal to the probability of observing a sequence of product choices conditional on a sequence of observed volume purchases multiplied by the probability of observing that particular sequence of volume purchase. To remove some of the conditioning arguments, integrate the above over initial inventories and consumption shocks for each parameter draw:

$$\begin{aligned} & \tilde{L}_{sim}^h \left(pur_1^h \cdots pur_T^h | \phi_t^h, \tilde{\Theta}_{sim}^h \right) \\ &= \int_v \int_{inv_0} \tilde{L}_{sim}^h \left(pur_1^h \cdots pur_T^h | \phi_t^h, inv_0^h, \Theta^h \right) dF \left(inv_0^h \right) dF \left(v_1^h \cdots v_T^h \right) \end{aligned} \quad (13)$$

Therefore, each household has a set of choice probabilities, one for each parameter draw.¹⁰ The objective of the calculations is to find the parameters that maximize the total likelihood.

Step 5 (Log Likelihood Function) To compute the total likelihood function, randomly select a subset of the parameter draws along with their corresponding choice probabilities for each household. The simulated maximum likelihood function is a weighted average of all the randomly selected choice probabilities across all households. That is,

$$\tilde{L} = \prod_{h=1}^H \left\{ \frac{1}{\#\text{draws}} \sum_{\#\text{draws}} \left[\tilde{L}_{sim}^h \left(pur_1^h \cdots pur_T^h | \phi_t^h, \tilde{\Theta}_{sim}^h \right) \frac{h \left(\tilde{\Theta}_{sim}^h | \mu_{\Theta}, \Sigma_{\Theta} \right)}{g \left(\tilde{\Theta}_{sim}^h | \mu_{\Theta}^0, \Sigma_{\Theta}^0 \right)} \right] \right\} \quad (14)$$

where $h \left(\tilde{\Theta}_{sim}^h | \mu_{\Theta}, \Sigma_{\Theta} \right)$ denotes the estimated distribution over random coefficients and over initial inventories. The component $h \left(\tilde{\Theta}_{sim}^h | \mu_{\Theta}, \Sigma_{\Theta} \right) / g \left(\tilde{\Theta}_{sim}^h | \mu_{\Theta}^0, \Sigma_{\Theta}^0 \right)$ can be seen as a weight on the likelihood. Parameters are allocated more weight if they explain households' decisions well. Note, since all the choice probabilities are calculated prior to maximization, $\tilde{L}_{sim}^h \left(pur_1^h \cdots pur_T^h | \phi_t^h, \tilde{\Theta}_{sim}^h \right)$ is nothing more than a large matrix of values. Moreover, since $g \left(\tilde{\Theta}_{sim}^h | \mu_{\Theta}^0, \Sigma_{\Theta}^0 \right)$ is pre-defined, the maximization routine updates only the weight $h \left(\tilde{\Theta}_{sim}^h | \mu_{\Theta}, \Sigma_{\Theta} \right)$.

4.3 Parameter Estimates

As described in the introduction, the estimation routine is performed on two samples. The first is a sample of 317 households. These are households whose annual incomes are below the federal poverty line. The second sample includes all 1,546 households who have met the IRI reporting criterion.

¹⁰Recall, N is the number of parameter draws specified by the econometrician. In the estimation, I use $N = 10,000$.

However, due to the size of the state space and its associated computation burden, I include only households who have made purchases of no more than 8.5 litres¹¹ per shopping trip. Packaging of all types, i.e. bottles or cans, are included in the sample. This limitation excludes around 3% of all households from the original sample but does not limit households' inventory buildup. Households may go shopping on two consecutive days or may choose to consume little. Over time, this leads to substantial inventories.

Table (7) presents the mean and the standard deviation that govern the distribution of all parameters. for each parameter. Recall that all parameters in the model are random coefficients. Estimates for different income brackets are listed in separate columns.

Table 7: Parameter Estimates

	Low Income		All Households	
	mean	std. dev.	mean	std. dev.
Marginal Utility of Consumption	1.8354 (0.1928)	1.1732 (0.1043)	1.6512 (0.2116)	1.4128 (0.0911)
Cost of Inventory, Linear Term	1.5e-3 (3.9e-4)	8.1e-4 (1.1e-4)	1.3e-3 (3.2e-4)	1.2e-3 (1.1e-4)
Cost of Inventory, Quadratic Term	1.6e-3 (3.2e-4)	5.3e-3 (1.2e-4)	1.2e-3 (3.1e-4)	5.1e-3 (1.2e-4)
Marginal Utility of Income	1.3159 (0.2154)	1.3304 (0.0954)	1.1345 (0.1937)	1.5065 (0.0649)
Utility of Promotions	1.2084 (0.5103)	1.1456 (0.0955)	1.1002 (0.4692)	1.2546 (0.0812)
Coke Preference Parameter	2.0695 (1.0295)	1.6398 (0.4638)	2.4109 (1.1932)	2.3396 (0.5236)
Pepsi Preference Parameter	2.0954 (1.1456)	1.7194 (0.3216)	2.1058 (1.2584)	1.7850 (0.6664)
Initial Inventory, (Litre)	2.5212 (2.8749)	1.5046 (0.4997)	3.7524 (2.5676)	1.9524 (0.6793)

Note: The estimates are presented on top and standard errors in parentheses.

The parameter estimates roughly conform to our intuition. For instance, comparing poor households to the entire sample, the average poor household has a slightly higher marginal utility of income and hence value money more than the average household in the sample. This result is consistent with decreasing marginal utility of income: as households' income per capita increases, money becomes less important for them. The same pattern holds for the utility of promotions: low income households have a larger affinity for promotional activities than the average. Poor households also have a slightly higher marginal utility of consumption. This implies that holding the cost of inventory constant, these households would consume more soda than the average household. The cost of inventory has an opposite story. Poor households have a higher cost of inventory. It seems that it is more difficult for poor households to stock up than for the average household.

In terms of persistent brand preferences, there is little difference between Coke and Pepsi for low income households. While the data suggests that these households tend to purchase Coke a bit

¹¹This is roughly equivalent to either 4 units of the standard 2-litre bottles or a pack of 24 cans.

less than Pepsi, this could be a result of Pepsi having lower prices most of the time. The average family in the sample, on the other hand, prefers Coke more to Pepsi and purchases much more Coke. The mean of the initial inventory distribution is increasing in income. Poor households, on average store a little over a bottle of cola at home, whereas the average household stores just under two bottles. There are several factors that influence households' inventory behavior. First, the marginal utility of income: a higher marginal utility of income leads to a higher sensitivity to price changes. This implies that households with higher marginal utilities of income would want to buy more when a sale is offered. Second, the cost to inventory: The higher the cost to inventory, the less the households would store. In general, the estimated cost to inventory of an average household is much lower than that of a poor household. Hence, while poor households want to make larger purchases when prices are low, they also have to balance the extra savings against the inconvenience of finding extra storage space.

4.4 Model Fit

To test the fit of the model, I compare the predicted market shares to the actual market shares observed in the data across different products. Taking the observed price paths as given, I simulate households according to the distribution of the parameter estimates. I compare the resulting market shares of different brands with the corresponding ones from the data. Overall, they match fairly well. For instance, the observed Coke market share is 54.54% for all households. In the simulation, the market share is predicted to be about 52.08%, which is slightly underestimated. I take the fact that they closely align as a positive signal.

Table 8: Observed and Predicted Market Shares

	Low Income		All Households	
	Observed	Predicted	Observed	Predicted
Coca Cola	46.81%	48.01%	54.54%	52.08%
Pepsi Cola	50.63%	50.77%	43.77%	46.65%
Store Brand	2.52%	1.22%	1.77%	1.27%

4.5 Identification, Heterogeneity, and Standard Errors

4.5.1 Identification

I now briefly and informally discuss the identification of the model. The identification of the static parameters follows the standard arguments (See Bajari, Fox, Kim, and Ryan (2012) for a detailed discussion on the identification in random coefficient logit models). Parameters governing households' sensitivities to prices and advertising are identified from variations over time in prices and promotional activities. As price reductions are not perfectly correlated with products on feature or on display, the effects can be separately identified. The unobservable persistent household taste parameters are identified off households' purchase behaviors as prices and promotions change over

time. Consider the case where we observe a household consistently purchasing Coke even during most sales of Peps. For households such as this, we can conclude that it has a strong preference for Coke.

Similar to Hendel and Nevo (2006b) the identification of the dynamic parameters is more subtle. The complication arises from the fact that neither the consumption nor the inventory is observed. If both were observed, then the identification would follow standard arguments. (See Aguirregabiria (2005), Magnac and Thesmar (2002), and Rust (1996)). To recover the utility of consumption and the cost of inventory, I rely on the relationship between current purchases, previous purchases, and prices. Given the process of prices and advertising, preferences, storage costs, and initial inventories determine consumer purchases. For a given storage cost function, the utility of consumption determines the level of demand. That is, households with a high utility of consumption will purchase more on average. Hence, if we observe a household making large purchases frequently, we may conclude that the household has a high utility of consumption. For a given preference, storage costs determine interpurchase duration and the extent to which consumers can take advantage of price reductions. Here consider the case of no storage, or very high storage cost. Such households cannot exploit the price reduction opportunities and stock up on their preferred products. Their interpurchase durations are dependent only on current prices but not past prices. Hendel and Nevo (2006b) provide a more in depth discussion of the identification of dynamic parameters.

4.5.2 Heterogeneity

Heterogeneity in product preferences or marginal impacts of prices and promotions can lead to interesting dynamics. For instance, consider a household with a high unobservable preference for Coke. This household is likely not to react to sales on other products and instead chooses to wait for regular Coke to go on sale. Suppose the inventory of this household runs low, it may purchase small amounts of an alternative while waiting for the price of Coke to decrease. Previous models are unable to accommodate this feature due to the computation intensive nature of the problem.

Accounting for unobservable persistent preferences such as the example above is computationally demanding. Indeed, as pointed out by Hendel and Nevo (2006b), in allowing random effects in product preferences or marginal impacts of prices, the taste parameters are solutions to dynamic programming problems instead of being static parameters. That is, to compute the choice probabilities associated with the model unconditional on the unobservable preferences implies integrating over the distribution of the preferences. The probability of product choice conditional on size purchased and conditional on the preferences of consumers will still be independent of the various dynamic components. However, to compute the probability of the product choice conditional on size (but unconditional on persistent preferences) we would have to integrate over the distribution of preferences conditional on the size purchased. Solving for this distribution requires solving the dynamic programming program. In this paper, I estimate a model similar to that in Wang (2012), which allows for a flexible form of the unobservable persistent preferences.

4.5.3 Standard Errors

Standard errors were calculated by random subsampling without replacement. I follow Politis and Romano (1994) and randomly drew subsamples of 100 households 500 times.

5 Policy Study

Currently, the CPI is calculated using weighted average monthly prices and hence does not capture intertemporal substitution. In the presence of storable products with frequent and periodic price reductions this measurement is biased upward. To be more specific, households may purchase goods when they are on sale and store them for later use. Failing to account for this behavior the CPI would consistently overestimate the prices paid for such products. Along with the increasing availability of household and store scanner data, as well as the development of techniques that allow estimation of sophisticated dynamic models of consumer behavior, more and more research is being conducted to better track and approximate the price movements of household purchases. These methods are able to accommodate intertemporal substitution and capture the actual prices paid.

If these changes in the methodology were adopted by social programs to approximate the cost-of-living, they could have profound implications on the wellbeing of households who depend on social welfare programs such as the Supplemental Nutritional Assistance Program (SNAP). Assuming poor households have similar purchase behaviors as any representative household, more accurate measurements of food costs will likely imply a decrease in welfare payment. In this section, I quantify the effects of a more accurate measurement of the cost-of-living-index on poor households' welfare. I calculate poor households' welfare in two situations: one with the actual observed prices and one with monthly average prices. The monthly averages reflect the current CPI method of determining price levels. The difference between the utilities achieved under the actual observed prices and the monthly average prices is the welfare loss poor households would experience should social welfare programs adopt a cost-of-living measurement that takes intertemporal substitution into account.

Using the estimated parameters from the sample of households who earn below the federal poverty line, I first simulate the purchase and consumption behaviors of a set of 1,000 households using sequences of observed prices from the 2002 store panel. I further calculate and record the total utility each household achieves in the year. This provides a distribution of utilities households would experience under real prices. Then using the same 1,000 households, I alter the price sequences and instead use a monthly average price for each product. To be more specific, following the practice used in sampling prices for the CPI, for each product I take the average of all prices observed across all the stores in a market in the first week of a month and use this average price in simulating households' consumption and purchase behavior for the month. As previously, I calculate and record the total utility each household achieves. I compare the average of these two distributions of utility. I find that on average, poor households achieve a 6% higher welfare under constant price estimates. That is to say, if social welfare programs switch from an average price based measurement such as CPI to a more accurate estimate of cost-of-living, using a sample of only households who qualify

for social welfare programs, then households on welfare would experience a 6% welfare loss on average. This result comes from the fact that poor households reduce their cost-of-living through intertemporal substitution, which the current method of calculating the CPI does not take into account. Unfortunately, this result can be misleading, since price indices are usually estimated not using only poor households but all representative households.

The CPI does not only serve as inflation adjustment for social welfare, it is also used for many other purposes. And as such the weights used in calculating the CPI are estimated using a sample of randomly selected households. Comparing these households to households who need federal assistance, the average household tends to be less price elastic but also has a smaller cost to storage. They tend to live in larger homes and have better transportation options. These factors allow the average household to take advantage of sales more and hence on average pay less than the poorer households who lack these provisions. Indeed, from Table(9) below we see that the average price an average family pays is above that paid by poorer households.

Table 9: Average Prices Paid

	Low Income	All Households
Average Price Paid, (2L)	\$1.48 (\$0.53)	\$1.40 (\$0.31)

Note: Difference significant at 95% level.

This is all the more surprising since the average family tends to have a higher persistent brand preference for Coke, which is usually more expensive. It can be explained by poorer households' limited ability to stock up and hence end up paying higher prices. From the estimated parameters, we see the initial inventory distribution has a higher mean for the sample with all households than for the sample with only poorer households. Given that poor households seem to benefit less than the average from stockpiling, this means that poor households' welfare could be reduced further should an accurate measure of the cost-of-living be employed on a representative sample of all households.

To estimate this effect, I repeat the above exercise. Using estimated parameters, this time from the estimation on the entire sample, I again simulate the consumption and purchase behaviors of 1,000 households under two different sequences of prices. The first is a set of randomly drawn price sequences from observed store data and the second a monthly average across stores for each product. I find that the average utility of a representative household, under actual pricing, is over 15% that of monthly average prices. This would suggest that should we use cost-of-living measures that account for intertemporal substitution on a sample of all households welfare recipients could be substantially negatively impacted.

6 Concluding Remarks

Previous literature has shown that the current way of calculating the CPI is flawed by not taking households' substitutions into account. With the ubiquity of scanner data and powerful computers this could be remedied easily. I develop and estimate a dynamic demand model for storable goods to study how the welfare recipients would be affected if a more accurate index accounting for households' ability to stockpile were introduced.

I find that such a change would hurt poor households in two ways: First, the current fixed weight CPI does not take inventories and households' ability to intertemporally substitute into account; this means that it assigns too little weight to low prices and thus overestimates the cost of living. Therefore, introducing a more accurate cost-of-living measurement would result in reduced welfare payments.

Second, since the CPI is calculated using a sample representative for the whole population, welfare recipients would suffer an even greater loss. The average household can exploit intertemporal prices fluctuations more heavily than poor households due to larger homes, better transportation options, and the lack of binding budget constraints. Thus, calculating the CPI off a sample of all households could severely underestimate prices paid by poor households and thus lead to even larger payment reductions.

Together, these findings show that we have to tread carefully when adjusting inflation measures. As long as the CPI has multiple uses there is a trade-off between accuracy and preserving social benefits. The most logical conclusion seems to be to introduce a separate cost-of-living measure for welfare recipients.

References

- [1] Akerberg, D. A.(2009). A new use of importance sampling to reduce computational burden in simulation estimation. *QME: Quantitative Marketing and Economics*, 7(4), pp 343-376.
- [2] Boskin, M., E. Dulberger, R. Gordon, Z Griliches, and D. Jorgenson (1998) “Association Consumer Prices, the Consumer Price Index, and the Cost of Living”, *The Journal of Economic Perspectives*, Winter, Vol. 12, No. 1, pp. 3-26.
- [3] Bajari, P., J. Fox, K. Kim, and S. Ryan (2012) “The Random Coefficients Logit Model is Identified”, *Journal of Econometrics*, February, Volume 166 (2), pp. 204-212.
- [4] Bajari, P. and L., Benkard (2005): “Demand Estimation with Heterogeneous Consumers and Unobserved Product Characteristics: A Hedonic Approach,” *Journal of Political Economy*, Vol. 113, No. 6.
- [5] Benitez-Silva, H., G. Hall, G. Hittsch, G. Pauletto, an J. Rust (2000): “A Comparison of Discrete and Parametric Approximation Methods for Continuous-State Dynamic Programming Problems” *mimeo*, Yale University.
- [6] Berry, S., J. Levinsohn, and A. Pakes (1995): “Automobile Prices in Market Equilibrium,” *Econometrica*, 63, 841-890.
- [7] Bresnahan, T.F. (1981): “Departures from Marginal Cost Pricing in the American Automobile Industry: Estimates for 1997-1978,” *Journal of Industrial Economics*, 17, 201–277.
- [8] Conlon, C (2010): “A Dynamic Model of Costs and Margins in the LCD TV Industry,” *mimeo*.
- [9] Erdem, T., S. Imai, and M. Keane (2003): “Consumer Price and Promotion Expectations: Capturing Consumer Brand and Quality Choice Dynamics under Price Uncertainty,” *Quantitative Marketing and Economics*, 1, 5-64.
- [10] Feenstra, R. and M. Shapiro (2003): “High-Frequency Substitution and the Measurement of Price Indexes,” Univ. of Chicago and NBER. in Robert Feenstra and Matthew Shapiro (eds.), *Scanner Data and Price Indexes*.
- [11] Goettler, R. and B. Gordon (2011): “Does AMD Spur Intel to Innovate More?,” *Journal of Political Economy*, Vol. 119, No. 6, pp. 1141-1200.
- [12] Gowrisankaran, G. and Rysman (2011): “Dynamics of Consumer Demand for New Durable Goods,” *mimeo*.
- [13] Hartmann, W (2006): “Intertemporal effects of consumption and their implications for demand elasticity estimates,” *Quantitative Marketing and Economics*, Vol 4, No. 4, 325-349.
- [14] Hausman, J (1998): “New Products and Price Indices,” *NBER*, Retrieved on Nov 14, 2012, from http://www.nber.org/reporter/fall98/hausman_fall98.html.

- [15] Hartmann, W. and H. Nair (2010) : “Retail Competition and the Dynamics of Consumer Demand for Tied Goods,” *Marketing Science*, Volume 29 Issue 2, March 2010.
- [16] Hendel, I. and A. Nevo (2006a): “Sales and Consumer Inventory,” *Rand Journal of Economics*, Vol. 37, No.3, Autumn 2006 pp.543-561.
- [17] Hendel, I. and A. Nevo (2006b): “Measuring the Implications of Sales and Consumer Inventory Behavior,” *Econometrica*, Vol. 74, No.6, 1637-1673.
- [18] Nevo (2000): “Mergers with Differentiated Products: The Case of the Ready-to-Eat Cereal Industry,” *Rand Journal of Economics*, Vol. 31, No.3, Autumn 2000 pp.395-421.
- [19] Osborne, M (2012): “Estimation of a Cost-of-Living Index for a Storable Product Using A Dynamic Structural Model,” *mimeo*
- [20] Powell L.M. et al, “Food Store Availability and Neighborhood Characteristics in the United States," *Prev Med.* 2007; 44(3):189–95.
- [21] Rust, J. (1987): “Optimal Replacement of GMC Bus Engines: An Empirical Model of Harold Zurche,” *Econometrica*, Vol. 55, 999-1033.
- [22] Wang (2012): “Estimating the Distributional Impact of Taxes on Storable Goods: A Dynamic Demand Model with Random Coefficients,” *mimeo*