

# Online Appendix for

## *Approximating the Cost-of-Living Index for a Storable Good*

Matthew Osborne

### **A Appendix: Additional Evidence of Stockpiling Behavior**

The time-series patterns of price and quantity shown in Figure 1 of the main paper are consistent with consumer stockpiling of canned tuna and canned soup: when product prices fall for a short time, quantity spikes to a level much higher than the baseline level that is observed at regular prices. In this section I present some additional evidence that patterns of consumer purchases are consistent with stockpiling behavior. To do this I perform a regression analysis following Hendel and Nevo (2006) using individual household purchases. I regress the total quantity a household purchases during a given week on a measure of the household's inventory, the price, a dummy variable for whether a sale occurs, the interaction of that dummy variable with price, and feature and display variables. Household inventory and consumption are both unobserved, so I compute a proxy for inventory under the assumption that a household's consumption rate is constant over time. I estimate a household's consumption rate by dividing the total quantity consumed by the number of weeks the household is observed. One can construct an estimate of the household's inventory during a given week as the sum of quantity purchased prior to that time, minus total consumption, plus initial inventories. Initial inventories are also unobserved, so a household fixed effect is included to control for them. Note that in the Nielsen data, household store visits where no purchases were made were only tracked from the middle of 1986 onwards. Due to this limitation the regressions only use the final year of data for the canned tuna category, and the final 2 years of data for canned soup.<sup>1</sup> Regression results for both categories are shown in Table A1.

In both categories, the impact of inventory is negative and significant, indicating that when a household's inventory increases, the quantity purchased decreases, consistent with stockpiling behavior. Turning to the three price interaction variables, in the canned tuna category less is purchased for higher prices, and when a product goes on sale, more is purchased. Furthermore, when a product is on sale the quantity purchased is higher, and the

---

<sup>1</sup>The Nielsen data tracks household purchases for the entire 3 years of the data, but it is not known if a household visits a store but makes no purchase prior to the middle of 1986.

interaction of the sale dummy variable is negative and significant, suggesting that consumers are more sensitive to lower prices when a product is on promotion. These facts are also consistent with stockpiling behavior. The dummy variable for feature is also significant. In the canned soup category, consumers appear to be relatively insensitive to changes in the regular price (this may be a result of the fact that the regular price in soup does not change a lot). However, they are extremely sensitive to sales. It may be the case that since promotions are very infrequent in the soup category, when consumers observe them they stockpile a lot. Additionally, in the soup category consumers appear to be more sensitive to advertising, as the display variable is insignificant.

Table A1: Evidence for Inventory Behavior: Household-Level Regression of Quantity on Inventory

Regressor	Canned Tuna		Canned Soup	
	Estimate	Std Err	Estimate	Std Err
Inventory	-0.124**	(0.004)	-0.035**	(0.001)
Price	-47.31**	(9.82)	7	(7.25)
Sale	4.64**	(1.32)	45.58**	(3.59)
Price $\times$ Sale	-40.83**	(14.2)	-112.48**	(9.18)
Display	0.04	(0.19)	0.74**	(0.27)
Feature	1.03**	(0.16)	1.36**	(0.3)

Notes: Both regressions include household, store, brand and variety (in the case of tuna, whether the tuna is packed in oil or water, and for soup, the flavor of the soup) fixed effects. Household quantity purchased and inventory are measured in ounces and price is measured in cents per ounce. A single star indicates significance at the 10 percent level while two stars indicate significance at the 5 percent level.

## B Appendix: Evidence Related to Nonlinear Pricing

This appendix describes evidence in the data that suggests that nonlinear pricing does not have a strong presence in either product category. It is important to verify this as the presence of nonlinear pricing would complicate estimation of the structural model. In some product categories retailers will use quantity discounts, where one can buy 2 or 3 units of a product for a lower per-unit price than a single can. There are a number of ways to test for the presence of nonlinear pricing in the data. One arises from the Nielsen store-level files, which include a variable called *price multiple*. The data documentation states that this variable captures the number of units to be purchased at the regular retail price. It is

possible that if this number is above 1, it may indicate nonlinear pricing. In the tuna data, this variable is above 1 for 69 store-week-UPC (Universal Product Code) combinations, out of 1,597 total store-week-UPC observations, indicating nonlinear pricing is not likely to be a major factor in the canned tuna data. The term UPC in these appendices corresponds to a variety of a product from the main paper. A drawback to using the store data for this type of inference is that the store data's coverage is incomplete: for canned tuna, in the household data 61 unique UPC's are observed being purchased but price information is available for only 30 of them. Additionally, for the 30 UPC's that are tracked in the store-level data, no information is recorded for many weeks. The purchase data covers 123 weeks and 19 stores, which means a dataset covering all possible store-week-upc combinations would contain 70,110 observations.

Two alternative ways to check for nonlinear pricing come from using the purchase data rather than the store data. When a purchase is made the purchase data records the store, the week, the day of week, and the UPC purchased, as well as the total price and number of units purchased. I aggregate the purchase data across households, and compute the mean price given the store, UPC, day of week, week, and number of units purchased. One way to check for quantity discounts is to compute the remainder of the price of an UPC after division by the number of units purchased. If the remainder is positive, then the price is not divisible by the number of units, which could occur on a 3-for-1-dollar deal, for example. In the tuna data nonzero remainders are observed in only 0.4 percent of upc-week-day-store combinations. Another way is to isolate store-week-day observations where some households purchased 1 unit of a UPC and other purchased 2 units (or 1 and 3 units, etc), and to examine how often the price per unit differs, and if it does, by how much. The results of such analysis are shown in Table A2. The first panel shows the results for canned tuna. An entry in the table shows the percentage of observations where different per unit prices are observed when different consumers purchase the number of units indicated in the row and in the columns for the same UPC in the same store during the same day and week. For example, in only 0.4 percent of observations where some consumers purchase 1 unit and some consumers purchase 2 units are different prices observed. For observations where a different price exists, the average magnitude of the difference, in cents, is shown in brackets. For canned tuna, almost all the percentages are below 1, and when differences occur, they are very small, on the order of 1 to 4 cents per can (compared to an average price per can of around 60 cents). Thus, we can conclude that nonlinear pricing is not a significant phenomenon in canned tuna.

There is somewhat more evidence of nonlinear pricing in canned soup, but it still does

not appear to be a widespread phenomenon. In the store data, the multiple unit variable is above 1 for about 27 percent of observations in that data, but the coverage of the store data is especially poor. Only 4,744 observations are recorded, but the purchase data covers 138 weeks and 315 UPCs. Turning to the information in Table A2, the percentages of instances where different per unit prices are observed are quite small. The instances where the percentages are larger tend to be cases where consumers purchase 3 and 5 units, 4 and 5 units, or 5 and 6 units. Relative to the sample size there are not many observations where these combinations of units are observed: they are between about 100 and 350 observations, whereas there are about 12,000 week-UPC-store-day combinations where 1 and 2 units are observed together. Additionally, when differences occur they are not large. Taken together, the extent of nonlinear pricing seems small enough that I feel comfortable abstracting away from it in the empirical model.

Table A2: Evidence Describing the Extent of Nonlinear Pricing

Canned Tuna						
	1 Unit(s)	2 Units(2)	3 Unit(s)	4 Unit(s)	5 Unit(s)	6 Unit(s)
1 Unit(s)	-	0.4 (3.6)	0.9 (1.6)	0.4 (1.9)	-	0.9 (4)
2 Unit(s)	-	-	0.9 (1.2)	0.5 (1)	0.6 (0.9)	0.3 (1)
3 Unit(s)	-	-	-	0.5 (1.2)	0.4 (10)	1.4 (0.4)
4 Unit(s)	-	-	-	-	-	0.2 (26.3)
Canned Soup						
	1 Unit(s)	2 Units(2)	3 Unit(s)	4 Unit(s)	5 Unit(s)	6 Unit(s)
1 Unit(s)	-	0.4 (1.5)	2.1 (0.2)	0.8 (0.6)	4.8 (1.6)	5.5 (0.4)
2 Unit(s)	-	-	2.3 (0.6)	1 (0.4)	7 (1)	5.7 (0.3)
3 Unit(s)	-	-	-	1.6 (0)	8.3 (1.5)	6 (0.1)
4 Unit(s)	-	-	-	-	10.2 (1.8)	4.8 (0)
5 Unit(s)	-	-	-	-	-	31.6 (0.1)

Note: An entry in each table shows the percentage of observations where different per unit prices are observed when different consumers purchase the number of units indicated in the row and in the columns for the same UPC in the same store during the same day and week. For example, in only 0.4 percent of observations where some consumers purchase 1 unit and some purchase 2 are different prices observed. For observations where a difference exists, the average magnitude of the difference, in cents, is shown in brackets.

## C Appendix: Formulation of the Model with Multiple Variety Purchase

This section describes the more general case of my IVS modifications when consumers are allowed to purchase more than one variety of a product in a single purchase occasion (such as purchasing two flavors of soup at once). To motivate my modifications to IVS, I frame the consumer problem as follows: when the quantity choice decision is made, the consumer also decides what number of varieties to purchase, and the proportions in which to purchase them. For example, the consumer might decide to purchase 4 cans of soup, 2 different varieties, in a 3-to-1 ratio. The inclusive value can then be formed over all variety combinations that would lead to a 3-to-1 ratio of purchases. To see how this works, denote  $M_{it}$  as the number of varieties chosen, and  $\mathbf{R}_{it}(M_{it})$  as a vector of ratios of length  $M_{it}$ , in order from largest to smallest. In the example given we would have  $M_{it} = 2$  and  $\mathbf{R}_{it}(M_{it}) = (3/4, 1/4)$ . Denote the set of allowable variety-ratio pairs as  $\Psi$ , where an element is  $\boldsymbol{\psi}_{it} = (M_{it}, \mathbf{R}_{it}(M_{it}))$ . I assume that rather than tracking individual prices, consumers track the inclusive values for each feasible variety, ratio pair:

$$\Omega_{it}(\boldsymbol{\psi}_{it}) = \log \left( \sum_{j_1=1}^J \sum_{j_2=j_1+1}^J \cdots \sum_{j_{M_{it}}=j_{M_{it}-1}+1}^J \exp \left( \sum_{k=1}^{M_{it}} R_{ikt} [\gamma_{i,j_k} - \alpha_i N_{UP} p_{i,j_k,t} + \boldsymbol{\phi}'_i \mathbf{Y}_{i,j_k,t}] \right) \right). \quad (\text{A1})$$

I present the vector  $\mathbf{R}_{it}(M_{it})$  in order from largest to smallest as a convention, because the identities of the varieties purchased do not affect utility outside of the inclusive value. The same reasoning justifies the sums in equation (A1) being “triangular” sums. Note that the dimensionality of the price state space will now depend on how many combinations of  $(M_{it}, \mathbf{R}_{it}(M_{it}))$  are included in the model, rather than the number of products. In the data one would typically observe most purchases falling into a small number of categories, and could restrict choices to be in these categories. In the soup data almost 70 percent of purchases are either of 1 UPC, or of 2 UPCs in either a 1:1, 2:1, or 3:1 ratio. In this case, the dimensionality of the continuous part of the state space is 4, which is much more manageable than 12.

I modify Assumption 1 to incorporate multiple variety choice as follows:

**Assumption 2.** *Nested Logit Error Distribution*

*The choice specific error term  $\epsilon_{qit}$  can be decomposed into two errors as follows:*

$$\epsilon_{qit} = \xi_{X(q),it} + \frac{X(q)}{N_U} \epsilon_{j(q),it},$$

where  $q$  indexes a possible choice (i.e. a quantity- $\psi_{it}$  pair),  $X(q)$  indexes a quantity choice associated with  $q$ ,  $j(q)$  indexes a variety-ratio ( $\psi_{it}$ ) choice, the distribution of  $\epsilon_{j(q),it}$  is Type 1 extreme value and the distribution of  $\xi_{X(q),it}$  has a distribution of the form denoted as  $C(\frac{X(q)}{N_U})$  from Cardell (1997).

The value function corresponding to equation 7 from the main paper can be similarly expressed as:

$$V(\Omega_i, I) = \log \left( \sum_{x_{it}=0}^{N_U} \sum_{\psi \in \Psi} \exp \left( \frac{x_{it}}{N_U} \Omega_i(\psi) - scI_{it}^2 - FC \mathbf{1}\{X_{it} > 0\} \right. \right. \quad (\text{A2}) \\ \left. \left. + \max_{c \in C(X_{it}, I_{it})} [u(c; \beta) + \delta E_{\Omega'|\Omega} V(\Omega'_i, I'(c))] \right) \right).$$

## D Appendix: Simulation Study Details

This section provides some details on the simulation performed to help understand the identification of key model parameters in Section IV, as well as the simulation study described at the end of Section VII.C of the main paper. The simulation approach begins by simulating the purchases and consumption decisions of 10,000 consumers for 200 periods. Turning to the distribution of prices, the low values of price are chosen to be 0.20, 0.25, 0.30, 0.35, 0.40, 1.00, 1.09 and 1.18. The 9 remaining high prices are evenly spaced between 1.27 and 1.99. If the price is low in a given period, the probability it stays in the low region is 1/8 (each low price occurs with equal probability within this region). 1.27, 1.36 and 1.45 each occur with probability 0.25 conditional on a low price. The rest of the probability is evenly allocated to the remaining high prices. In the high region the probability of transitioning to any of the first 5 low prices is 6 percent (so the probability of a sale given no sale in the current period is 30 percent). The probability the price stays the same in the high region is 40 percent and the remaining probability is equally allocated to the other prices. When running the simulation, I assume that in period 1 all consumers start with 0 inventory. To ensure that the moments I calculate are not affected by the period 1 inventory I will remove the first 50 periods from the data set and only compute the moments on the last 150 periods. I chose 50

periods because it appears that the aggregate distribution of inventories becomes stationary after this time.

When constructing the purchase hazard in rows 2 and 3 of Table 2, I simulate consumer purchases under the Markov price series I described until period 100. In period 100 I assume the product goes on promotion to the lowest price, and in the periods afterwards I hold prices fixed at the lowest of the high prices (1.27). The probability shown in the first row of the table is the simulated probability ratio of purchasing 1 unit at the lowest price relative to 2 units at the lowest price.

Turning to the simulation study in Section VII.C, the basic setup of the study is similar to the identification study, except that prices only take on two values, and I allow for two types of individuals who vary in their storage costs. The simulation is run for 250 periods and the first 50 periods are removed to reduce dependence on initial inventories. The base period used for the indexes is therefore period 50. The dynamic index is computed using the same procedure as the empirical model. The FS and CK indexes are constructed using a window of 50 periods. There are a few small differences in how I construct these two indexes relative to the empirical exercise. First, for the FS index, rather than using modal prices, and shares at those prices as base prices and shares I follow Caves, Christensen and Diewert (1982) and use the alternative formula

$$I_t^{FS} = \exp \left( \sum_{k=1+T^{FS}*(m_t-1)}^{T^{FS}*m_t} \sum_{i=1}^N \frac{1}{2} (s_{i0,k} + s_{ik}) \ln \left( \frac{p_{ik}}{p_{i0,k}} \right) \right), \quad (\text{A3})$$

where  $s_{i0,k}$  corresponds to the share in periods 1 to 50, and  $p_{i0,k}$  is the corresponding price. I do this because I found that the FS formulation with modal prices tended to be heavily biased towards zero, and did not vary over time, since modal prices were high and modal shares received a large weight in equation (A3). In the CK index, the Best Price index is constructed in a slightly different way from the empirical work. I do not use a rolling minimum price, but rather I simply use the low value of the price when constructing the index. I found that using a rolling minimum price generated extremely wide variation in the Best Price index, making the CK index perform poorly. Finally, in the simulation study I also allow prices to vary across individuals, as is the case in the empirical exercise. I simulate ten different price series from the price Markov process, and assign each individual to a price series. When constructing indexes, each price series is treated analogously to how a store is treated in the empirical exercise.

# E Appendix: Estimation Details

## 1 Choice Probability Details

This section expands upon the discussion related to the construction of choice probabilities described in Section V of the main paper, and relates it to the Metropolis-Hastings steps outlined in online Appendix Section E.4. The vector of individual-specific utility coefficients being estimated is  $(\gamma_{i1}, \dots, \gamma_{iJ}, \alpha_i, \phi, \beta_i, sc_i, FC)$ . Denote the vector of continuous population-varying parameters as  $\tilde{\theta}_i = (\tilde{\gamma}_i, \ln(\alpha_i), \ln(\beta_i))$  where  $\tilde{\gamma}_i$  is the vector of population-varying product coefficients. I assume that  $\tilde{\theta}_i \sim N(\theta^1, \theta^2)$ , where the matrix  $\theta^2$  is diagonal. The storage cost parameter  $sc_i$  takes on two possible values, 0 with probability  $\pi^{sc}$ , or  $sc_1 > 0$  with probability  $1 - \pi^{sc}$ , corresponding to low and high storage cost types. The vector of population-fixed parameters is denoted  $\bar{\theta} = (\bar{\gamma}_i, \phi, sc_1, FC)$ , where  $\bar{\gamma}$  are the product coefficients that are fixed across the population.

In the Gibbs steps, it is necessary to compute the probability of a consumer's sequence of observed choices conditional on a set of draws on initial inventories,  $I_{i0}$ , inclusive values  $\Omega_{it}$ , and consumer utility coefficients  $\theta_i$ . An individual's optimal consumption in period  $t$ , conditional on her purchase decision,  $\mathbf{x}_{it}$ , and the states  $(\Omega_{i,t+1}, I_{i,t+1}(c), S_i)$  can be computed as

$$c_{it}^* = \arg \max_{c \in C(X_{it}, I_{it})} \left\{ u(c; \beta) + E_{\Omega_{i,t+1} | \Omega_{it}} V(\Omega_{i,t+1}, I_{i,t+1}(c), S_i; \tilde{\theta}_i, \bar{\theta}) \right\}. \quad (\text{A4})$$

Note that in the formulation above I treat the storage cost type as a state that does not change to improve the accuracy of taking draws on the storage cost type for each consumer. When I take these draws it is necessary to compute a consumer's sequence of purchase probabilities given they are either high or low. Having an estimate of the value function saved at each state improves accuracy but does not increase the size of the state space very much. An approximation to  $E_{\Omega_{i,t+1} | \Omega_{it}} V(\Omega_{i,t+1}, I_{i,t+1}(c_{it}^*), S_i; \tilde{\theta}_i, \bar{\theta})$  is computed using a nearest neighbor algorithm which is described in Section 2. Inventories are unobserved in period  $t$ , but they can be computed conditional on the initial inventories in period 0 (the method used to compute period 0 inventories is described at the end of online Appendix E.4.). Denote the dynamic utility (which does not depend on individual product choice) from the optimal consumption in period  $t$  as



$$\nu^d(X_{it}, \boldsymbol{\psi}_{it}(\mathbf{x}_{it})) = \max_{c \in C(X_{it}, I_{it})} \left\{ u(c; \beta) + E_{\Omega_{i,t+1} | \Omega_{it}} V(\Omega_{i,t+1}, I_{i,t+1}(c), S_i; \tilde{\boldsymbol{\theta}}_i, \bar{\boldsymbol{\theta}}) \right\} \\ - \frac{X}{N_U} \Omega_{it}(\boldsymbol{\psi}_{it}(\mathbf{x}_{it})) - sc_i I_{jit}^2 - FC \mathbf{1} \left\{ \sum_{j=1}^J x_{jit} > 0 \right\}$$

and denote the static utility (derived from a particular product choice and proportion) as

$$\nu^s(\mathbf{x}_{it}) = \sum_{j=1}^J \frac{x_{ijt}}{X_{it}} (\gamma_{ij} - \alpha_i N_U p_{ijt} + \phi'_i Y_{ijt}),$$

where  $\boldsymbol{\psi}_{it}(\mathbf{x}_{it})$  corresponds to a pair  $(M_{it}, \mathbf{R}_{it}(M_{it}))$  of number of varieties and variety ratios, arising from a particular choice of  $\mathbf{x}_{it}$  as described in online Appendix C.

Recall that the error term is distributed using the nested logit distribution as described in Assumption 2. Index the observed  $\mathbf{x}_{it}$  in period  $t$  as  $\mathbf{x}_{it}^{obs}$ , and the corresponding total quantity and variety-number, proportion pair as  $X_{it}^{obs}, \boldsymbol{\psi}_{it}^{obs}$ . Denote the possible choices of quantities and proportions as  $q_1 = 1, \dots, Q_1$  and the possible combinations of products a consumer could choose given the proportion choice as  $\mathbf{x} \in Q_2(X_{it}, \boldsymbol{\psi}_{it})$ . The nested logit form of the model implies the probability of a particular sequence of  $\mathbf{x}_i^{obs} = \{\mathbf{x}_{it}^{obs}\}_{t=1}^T$  being chosen by individual  $i$  will be

$$Pr_i(\mathbf{x}_i^{obs}, \tilde{\boldsymbol{\theta}}_i, \bar{\boldsymbol{\theta}}, S_i) = \prod_{t=1}^T \frac{\exp(\nu^d(X_{it}^{obs}, \boldsymbol{\psi}_{it}^{obs}))}{\sum_{q_1=1}^{Q_1} \exp(\nu^d(X_{it}^{q_1}, \boldsymbol{\psi}_{it}^{q_1}))} \frac{\exp(\nu^s(\mathbf{x}_{it}^{obs}))}{\sum_{\mathbf{x} \in Q_2(X_{it}^{obs}, \boldsymbol{\psi}_{it}^{obs})} \exp(\nu^s(\mathbf{x}))}. \quad (\text{A5})$$

Note that in the equation above the number of observations per individual,  $T$ , is constant across individuals, but there are some weeks in which no store is visited. During such weeks the choice likelihood is simply 1. When a consumer does not visit the store she will still consume the optimal amount out of her inventory. The probability in equation (A5) still depends on the individual  $i$  outside of the parameters and  $S_i$  due to the fact that different individuals have different probabilities of visiting a store each week, and this probability is incorporated into the individual's dynamic programming problem.

## 2 Approximating the Expected Value Function

One of the most computationally intensive parts of the estimation procedure is computing an estimate of the value function, which is necessary for constructing the choice probability

in equation (A5), to compute simulated inventories, and to update the value function. Here, I follow the Imai, Jain and Ching (2009) procedure of starting with a guess of the value function, and computing one update to the value function at each Gibbs iteration. In each update, it is necessary to compute an estimate of the value function at the current Gibbs draw. This is done by averaging over value functions saved in previous Gibbs draws where the parameters are “close” to the current draw. I use a nearest neighbor approach to select close parameter draws (Norets, 2009).

At step  $g$  of the Gibbs sampler, there will be  $N(g)$  saved draws on  $\tilde{\theta}_i$  for every household and  $N(g)$  saved draws on  $\bar{\theta}$ . There will also be  $N(g)$  saved value function estimates for each state space point at each of those draws.<sup>2</sup> Denote the draw on  $\bar{\theta}$  (or  $\tilde{\theta}_i$ ) from iteration  $k$  as  $\bar{\theta}^k$  (or  $\tilde{\theta}_i^k$ ),  $k = g - N(g), \dots, g - 1$ . Denote the vector composed of  $\tilde{\theta}_i^k$  and  $\bar{\theta}^k$  as  $\Gamma_i^k = (\tilde{\theta}_i^k, \bar{\theta}^k)$ .

For each  $i$  and  $k$  there is an associated saved value function,  $V_{i,k}(\Omega, I, S)$ . There will be a saved value function for all values of  $I$  and  $S$ , since they are discrete, but not for all possible  $\Omega$  values, since  $\Omega$  is continuous. The value functions are computed on a random grid of  $N_p$  inclusive values drawn from an importance distribution  $h(\cdot)$ , and these inclusive value draws are also saved. Denote the saved draws as  $\{\Omega_{g,s}\}_{s=1, g=N(g)}^{N_p, g-1}$ . Suppose that an estimate of the expected value function at some parameter vector  $\Gamma$ , an inclusive value  $\Omega$ , an inventory  $I$  and storage cost type  $S$  is needed. This estimate is constructed in two steps. First, for household  $i$ , I find the  $m = 1, \dots, \tilde{N}$  closest  $\Gamma_i^k$ s to  $\Gamma$ . Indexing these  $\Gamma_i^k$ s as  $\tilde{l}^m$ , the algorithm to find them can be written

$$\begin{aligned} \tilde{l}^1 &= \min_{l \in \{1, \dots, N(g)\}} \{\|\Gamma - \Gamma_i^l\|\} \\ \tilde{l}^2 &= \min_{l \in \{1, \dots, N(g)\} \setminus \tilde{l}^1} \{\|\Gamma - \Gamma_i^l\|\} \\ &\vdots \\ \tilde{l}^{\tilde{N}} &= \min_{l \in \{1, \dots, N(g)\} \setminus \{\tilde{l}^1, \dots, \tilde{l}^{\tilde{N}-1}\}} \{\|\Gamma - \Gamma_i^l\|\}, \end{aligned}$$

where the  $\|\|\|$  function indicates the Euclidean norm and  $\setminus$  indicates set subtraction. The estimated value function will average over the  $\tilde{l}^m$ 's, as well as the saved inclusive values  $\Omega_{g,s}$  at each Gibbs iteration associated with  $\tilde{l}^m$ . I denote these associated Gibbs iterations as  $g(\tilde{l}^m)$ . Note that I need to average over the saved  $\Omega_{g,s}$ 's, because when I compute the expected value function at a current inclusive value  $\Omega$ , I must integrate over the distribution

---

<sup>2</sup>When starting the algorithm, I begin with an estimate of the value function that is zero in all states.

of possible inclusive values for the next period. To average over the saved  $\Omega_{g,s}$ 's, I use importance sampling. To do this I compute the probability that  $\Omega_{g,s}$  occurs in period  $t + 1$  given the period  $t$  inclusive value  $\Omega$  using the transition density at the current parameter draw,  $P(\Omega_{g,s}|\Omega, \tilde{\theta}_i, \bar{\theta})$ , and use  $P(\Omega_{g,s}|\Omega, \tilde{\theta}_i, \bar{\theta})/h(\Omega_{g,s})$  as the importance weight. The estimated expected value function is then

$$\hat{E}V_i(\Omega, I, S; \Gamma) = \frac{1}{\tilde{N}} \sum_{m=1}^{\tilde{N}} \frac{\sum_{s=1}^{N_p} V_{i, \tilde{I}^m}(\Omega_{g(\tilde{I}^m), s}, I, S) \frac{P(\Omega_{g,s}|\Omega, \tilde{\theta}_i, \bar{\theta})}{h(\Omega_{g(\tilde{I}^m), s})}}{\sum_{s=1}^{N_p} \frac{P(\Omega_{g,s}|\Omega, \tilde{\theta}_i, \bar{\theta})}{h(\Omega_{g(\tilde{I}^m), s})}}. \quad (\text{A6})$$

Additional implementation details regarding the value function approximation are described in online Appendix 6.

### 3 Updating the Value Function

After all the parameters are drawn, the last part of each Gibbs iteration is to update the value function  $V_{i,k}(\Omega_{g,s}, I, S)$  at every  $I, S$ , and  $\Omega_{g,s}$  combination for each individual in the data. Note that in every iteration, I only perform a single value function update, as outlined in Imai, Jain and Ching (2009). Over the course of running the Gibbs sampler the value function will converge. At the end of iteration  $g$ , a new parameter vector  $(\tilde{\theta}_i^g, \bar{\theta}^g)$  will be obtained for each individual  $i$ . First, given some  $\Gamma = (\tilde{\theta}_i^g, \bar{\theta}^g)$  I compute  $\hat{E}V_i(\Omega_{g,s}, I, S; \Gamma)$  for all possible  $I, S$  values and every individual  $i$ . Then, I compute optimal consumption at all possible  $X, \psi$  choices, which I will denote  $c^*(X, \psi)$ . Net of the error term, utility at this point will be

$$\hat{v}(X, \psi) = \frac{X}{N_U} \Omega_i(X, \psi) + u(c^*; \beta) - sc_i I^2 + \delta \hat{E}V_i(\Omega_{g,s}, I', S; \Gamma) - FC \mathbf{1}\left\{ \sum_{j=1}^J x_j^q > 0 \right\},$$

where  $I'$  denotes period  $t + 1$  inventory. The updated value function will then be

$$V_{i,l}(\Omega_{g,s}, I, S) = \log \left( \sum_{q_1=1}^{Q_1} \exp(\hat{v}(X^{q_1}, \psi^{q_1})) \right).$$

### 4 Gibbs Steps and Initial Inventories

This section describes the details of the Gibbs steps used to estimate the dynamic structural model.

First I draw  $\tilde{\theta}_i$  for each consumer  $i$ , conditional on  $\bar{\theta}$ ,  $\{I_{i0}\}_{i=1}^N$  and  $S_i$ . The posterior distribution of  $\tilde{\theta}_i$  does not have a closed-form solution, so I use the Metropolis-Hastings (MH) algorithm to draw a new  $\tilde{\theta}_i$ . Denoting the previous draw on  $\tilde{\theta}_i$  as  $\tilde{\theta}_i^0$ , my procedure is to draw a candidate  $\ln(\tilde{\theta}_i^1) \sim N(\tilde{\theta}_i, \rho_1^2 \mathbf{W})$ , and to accept the new draw with probability

$$\frac{Pr_i(\tilde{\theta}_i^1, \bar{\theta}, S_i)}{Pr_i(\tilde{\theta}_i^0, \bar{\theta}, S_i)}.$$

The parameter  $\rho_1$  is adjusted every period so that about 30 percent of the candidate  $\tilde{\theta}_i^1$  draws are accepted every period (see Train 2003). Once I have drawn  $\{\tilde{\theta}_i\}_{i=1}^N$ , I draw out  $\theta^1$  and  $\theta^2$  conditional on the logarithm of the draws. I assume that the prior on  $\theta^1$  is normal with a prior variance of  $100\mathbf{I}$ , so the prior are proper but relatively uninformative. This prior generates a normal posterior distribution for  $\theta^1$ , which I can draw from using standard methods. Similarly, I put an inverse gamma prior on each of the diagonal elements of  $\theta^2$  with scale parameter 3 and shape parameter 0.5, which results in an inverse gamma posterior for the elements of  $\theta^2$ .<sup>3</sup>

Next, I draw  $\bar{\theta}$  given the drawn  $\{\tilde{\theta}_i\}_{i=1}^N$ , the previous draw's  $I_{i0}\}_{i=1}^N$  and  $S_i$ . I also use a random walk MH for this step, with a parameter  $\rho_2$  on the variance that periodically updates to keep the acceptance rate at 30 percent. Given a previous draw  $\bar{\theta}^0$ , I draw a candidate  $\bar{\theta}^1$  where

$$\bar{\theta}^1 \sim N(\bar{\theta}^0, \rho_2^2 \mathbf{I}).$$

The candidate draw is accepted with probability

$$\left( \prod_{i=1}^I \frac{Pr_i(\tilde{\theta}_i, \bar{\theta}^1, S_i)}{Pr_i(\tilde{\theta}_i, \bar{\theta}^0, S_i)} \right) \frac{k(\theta^1)}{k(\theta^0)},$$

where  $k$  is the prior distribution on  $\bar{\theta}$ . I assume that  $k$  is normal, with a variance matrix of  $100\mathbf{I}$  so that priors are close to uninformative. The prior means are set to the average parameter estimates from a similar model presented in Osborne (2010).

Last, we need to draw the  $S_i$ . This parameter is drawn for each consumer from a binomial distribution. The probability of being a low-cost type I denote as  $\pi^{sc}$ . For each

---

<sup>3</sup>Most of the mass of this prior distribution is between 0 and 0.5, which would be too tight if I was not exponentiating the utility parameters. However, my utility parameters are lognormal, which means a relatively small variance in the underlying normal can lead to a large variance in the resulting lognormal. For example, the exponential of a normal distribution with mean 1 and variance 0.5 has a variance of about 13.7.

individual, the posterior probability of  $S_i = 0$  is

$$\frac{Pr_i(\tilde{\theta}_i^1, \bar{\theta}, 0)\pi^{sc}}{Pr_i(\tilde{\theta}_i^0, \bar{\theta}, 0)\pi^{sc} + Pr_i(\tilde{\theta}_i^0, \bar{\theta}, 1)(1 - \pi^{sc})}$$

$\pi^{sc}$  is drawn from a Beta distribution, given the draws on  $S_i$  for each household.

My choice probabilities above depend on inventories prior to period 1,  $I_{i0}$ , which are unobserved. I integrate them out using simulation. My approach is to simulate  $I_{i0}$  conditional on  $\tilde{\theta}_i$ ,  $\bar{\theta}$  and  $S_i$  for each consumer. Estimation is performed on the final year of the data for both product categories and the data before that, which has incomplete information about store visits, is used for simulation (for canned tuna, store visits are unobserved when purchases are not made. In canned soup, store visits are unobserved when purchases are not made for the first 25 weeks of the data). I proceed by first drawing a sequence of store visits for each household where a visit is unobserved from the empirical density of observed store visits for that particular household.<sup>4</sup> I then assume that at the beginning of the data (73 weeks prior to the beginning of the estimation data for canned tuna, and 87 weeks for canned soup) initial inventories are zero. I then compute optimal consumption in each period, conditional on the parameter draws, prices and quantities chosen in each period, which gives us end of period inventories.  $I_{i0}$  is then taken to be the level of inventories computed at the end of the 73rd week. A new draw on the initial inventories is taken in each Gibbs step.

## 5 Implementation Details of Inclusive Value Transition Process

This section describes the construction of the spline function used for estimating the inclusive value transition process and some related implementation details. For a particular component of  $\Omega_{it}$ , to compute the basis I first compute the minimum and maximum values of  $\Omega_{it}(\boldsymbol{\psi})$  for each value of  $\boldsymbol{\psi}$  and each individual  $i$ , and choose a number of equidistant interpolation nodes  $n_d$ . Denoting the minimum and maximum  $\Omega(\boldsymbol{\psi})$  as  $a_{i,\boldsymbol{\psi}}$  and  $b_{i,\boldsymbol{\psi}}$  respectively, the nodes are defined as  $\Omega_i^k(\boldsymbol{\psi}) = a_{i,\boldsymbol{\psi}} + kh_i$ ,  $h_i = (b_{i,\boldsymbol{\psi}} - a_{i,\boldsymbol{\psi}})/n_d$ ,  $k = 0, \dots, n_d$ .

---

<sup>4</sup>A slightly different procedure was presented in Osborne (2010) where the  $I_{i0}$ 's were assigned a distribution and draws were taken on the  $I_{i0}$  in a separate Gibbs step to rationalize the final year of choice data. The procedure used in Osborne (2010) produces a distribution of initial inventories that is zero for more than 99 percent of households. Although the approach of Osborne (2010) is more nonparametric, my approach produces a distribution of initial inventories which more closely mimics the distribution of inventories I compute at the very end of the sample, which I feel is more realistic. Since the consumers in the sample are fairly regular purchasers of canned tuna, it seems unlikely that almost all of them would have no tuna in inventory at the beginning of the sample. The distribution of initial inventories may be difficult to identify nonparametrically.

For a particular value of  $\Omega(\boldsymbol{\psi})$  the spline basis is the vector of functions

$$u_{ik}(\Omega) = \Phi \left( \frac{\Omega - a_{i,\boldsymbol{\psi}}}{h_i} - (k - 2) \right), \quad k = 1, \dots, n_d + 3,$$

where

$$\Phi(t) = \begin{cases} (2 - |t|)^3 & 1 \leq |t| \leq 2 \\ 4 - 6|t|^2 + 3|t|^3 & |t| \leq 1 \\ 0, & \text{elsewhere} \end{cases}$$

At a particular dimension the regression will be  $\hat{\Omega}(\boldsymbol{\psi})$  is to use the predicted value of the regression

$$\Omega_{it}(\boldsymbol{\psi}) = \sum_{\boldsymbol{\psi} \in \Psi} \sum_{k=1}^{n_d+3} \kappa_{i,\boldsymbol{\psi},k} u_{ik}(\Omega_{it}(\boldsymbol{\psi})) + \eta_{it,\boldsymbol{\psi}}, \quad (\text{A7})$$

where the  $\kappa_{i,\boldsymbol{\psi},k}$  coefficients are estimated using OLS on an individual-by-individual basis.

Sometimes when running the regression (A7) the coefficients may not be stable or the variance of the error may not be stable, especially if some components of the regression are highly collinear. In canned tuna, there is only a single dimension to be included in the spline. However, for some individuals it isn't possible to estimate a regression with  $n_d > 1$  due to the aforementioned collinearity problem. To address this I check the stability of the regression for each individual, and reduce the dimensionality of the spline if they are unstable. I check stability in two ways. First, if some coefficients are unreasonably large (over 100) then that suggests a problem. Second, I compute transition probabilities from all inclusive values to a grid of inclusive values that covers the set of possible inclusive values. If all transition probabilities are zero for some observations then, as I will discuss below there will be numerical problems in constructing the expected value function. This occurs as a result of overfitting, which reduces the error variance to a very small number. If either case occurs then reduce the dimension by 1, and try again. I stop if the number of dimensions is 1 (a linear regression is run). For canned soup there are 4 dimensions to the inclusive value process. I have found that the process for this product category is not stable unless a linear functional form is chosen. Moreover, the inclusive values for proportions where more than 1 product is consumed are highly collinear with each other. Thus, when I run the regression for the case where the proportion is that only a single product is consumed, I just include the inclusive value for (1, 0) and (0.5, 0.5) on the right-hand side. For the other proportions I include the lagged inclusive value for that proportion, as well as (1, 0). I also assume a diagonal variance matrix for the error term.

As a final note, in the data there are occasionally weeks where a consumer does not visit a store. Since during these periods it is not clear what prices and marketing variables the consumer observes, some assumption must be made. I construct the inclusive value using the prior period’s price and assume the marketing variables feature and display are set to 0. This assumption seems reasonable as it relies on the idea the consumer would remember the most recent price. My results are unlikely to be sensitive to this assumption as well, since consumers visit the store in about 85 percent of weeks. Additionally, in periods prior to the estimation period when I draw out store visits, I account for not visiting the store by constructing a household-specific store visit probability, and draw out whether the household goes to a store from this probability.

## 6 Value Function Approximation Details

When constructing the value function approximations, I save  $\tilde{N} = 30$  previous value function draws, choosing the closest previous value function draw (Norets 2009 finds that it is sufficient to average over a single draw as typically the closest draw to the current candidate is usually the draw from the previous iteration, due to serial correlation in the Markov Chain). 100 inclusive value draws are taken in each iteration and used to construct the expected value function approximation. There is also a minor difference between the way I construct the inclusive value grid in practice and the way it is presented in this section. I have found that specifying an importance distribution over the space of possible inclusive values can lead to instability in the estimation routine. The reason for this is sometimes a parameter draw will be taken that is far from any of the previous draws, and the transition probabilities to past inclusive values used to construct the denominator of equation (A6) will always be zero, leading to numerical problems when constructing the value function. To construct the inclusive value grid, I first draw underlying price and marketing variables from their empirical distribution (which corresponds to the importance distribution), and then use those to construct an inclusive value at the current parameter draw. The price/advertising vectors are saved rather than the inclusive values themselves. Theoretically, this approach should work properly as long as the probability a particular inclusive value is drawn corresponds to the probability a particular price/advertising vector is drawn. This would be violated if at a parameter draw, there were two different price/advertising vectors corresponding to a single inclusive value. As long as the price coefficients are not zero, this should only happen on sets of measure zero, meaning drawing a grid of prices and advertising rather than inclusive values themselves should be innocuous.

## F Appendix: Construction of Estimation Data

This section discusses the construction of the estimation data for the two product categories. The estimation data is smaller than the entire data for a few reasons: 1) there are many individuals who appear to stop using their swipe cards for some period of time. 2) Some individuals purchase smaller brands, and are dropped from the estimation as a result. 3) some individuals purchase unit/size combinations that are not modeled as part of the choice set. For both product categories, for a household to be included they must make at least 25 store visits during the estimation period (the final year of the data), and make at least 5 purchases over the course of the period.

In canned tuna, a Universal Product Code (UPC) identifies a product and records features such as the package size, type of tuna (regular or albacore), whether the tuna is packed in oil or water, and whether the product contains multiple units in a single package. In terms of package size by far the most popular size is the 6 ounce can, which comprises over 91 percent of purchases in the entire sample. 98.8 percent of purchases are single cans rather than multipacks, 96 percent of purchases are chunk rather than solid, there are no purchases of albacore, and the share of tuna canned in oil is 26 percent (with the rest being in water). Since the bulk of purchases correspond to individually packaged 6-ounce cans, the analysis of this paper will focus on purchases of cans for that particular size of the product, for the top brands.

In canned tuna, I include the top 6 UPCs, purchases of 10 cans or less, and weeks where only a single product is purchased, or more than 2 stores are visited within the same week. By far the majority of purchases are in this category. Less than 0.2 percent of purchases are of more than 10 cans. The top 6 UPCs comprise over 80 percent of purchases. Over 93 percent of weeks are comprised of purchases of a single UPC in a single store. Any household who makes a purchase during the estimation period outside of these constraints is removed from the sample. The final household level data set is comprised of 557 households. The prices of each UPC are mostly taken from daily price files provided by Nielsen. There are a few UPCs in the soup category (corresponding to private labels) where no price information is provided. I construct weekly prices for these UPCs from household-level purchase data using a procedure similar to what Nielsen uses to construct its price files, which is described in the data documentation. Finally, in some weeks where no purchase is made, a consumer will visit multiple stores. It is necessary to decide in these cases what store the household used to consider prices for canned tuna. The data collects the total amount spent on groceries by each household in each store for every week. I assume that the prices observed during such



weeks corresponds to the prices at the store where the most money was spent. The idea here is that the household did most of its grocery shopping at that store, and it was most likely to consider purchasing tuna from that store.

The Nielsen soup data records purchases of both standard-size canned soup and other types of soup that are in dry packages (powdered soup) or larger cans of beef/chicken/vegetable broth. I do not include these latter two categories in the analysis since they are mostly used for cooking and are not likely to be substitutes for standard-size cans of soup.

For the canned soup data more households must be dropped from the data due to the complexity of the choice set. The complexity arises due to the fact that consumers often purchase different UPCs in a single occasion. Most commonly, consumers purchase either a single UPC, or 2 UPCs in either 1:1, 1:2, or 1:3 ratios. These combinations are included in the choice set and correspond to about 67 percent of all purchases. In terms of UPCs, there are many more UPCs in canned soup than canned tuna. I include all purchases of either the Campbell's brand or the private label brand. Soup is available in many different flavors. The most popular is cream of mushroom with a 17 percent purchase share, followed by tomato with a 13 percent share, followed by chicken noodle with a 12 percent share. The top ten flavors comprise about 75 percent of canned soup purchases. I model the purchases of the 8 most popular Campbell's flavors and aggregate the rest of the Campbell's flavors in a composite brand. I also model the top 2 flavors of the private label separately and aggregate the rest of the private label flavors into a composite product. Last, in the soup data consumers tend to purchase more cans at once so I allow people to purchase up to 20 cans in a week. Only about 0.2 percent of purchase occasions contain more than 20 cans in a week. After cutting out all households that ever make a purchase outside the choice set, I am left with 257 households in the sample. As a final note, in the soup data there are many purchases of powdered soup packets, and cans of soup that are 2 or 3 times the size of a standard can (these tend to be purchases of chicken, beef or vegetable stock which are used in cooking). I consider these to be in a separate product category from standard cans of soup, and so I drop these purchases but do not drop households that purchase these products.

## **G Appendix: Formulas for Simpler Indexes**

This section presents formulas for some of the simpler indexes. The index formulas I present below correspond to those used for the week 1 base period. During the time that my data

was collected (in 1988), the price index for food products used by the BLS was a Laspeyres index:

$$I_t^{Laspeyres} = \sum_i w_{i0} \frac{p_{it}}{p_{i0}}.$$

In the above equation,  $i$  indexes a brand-store combination.  $w_{i0}$  is the base period expenditure share for product  $i$ , and  $p_{i0}$  is the modal price in the previous year.<sup>5</sup>

The Törnqvist index, in contrast to the Laspeyres, allows weights to vary over time, and allows substitution across products:

$$I_t^{Törnqvist} = \exp \left\{ \sum_i 0.5(w_{i0} + w_{it}) \ln \left[ \frac{p_{it}}{p_{i0}} \right] \right\}.$$

The weights in the base period 0 are the quantities sold for the first week of the data. The Törnqvist index will approximate cost of living changes well when consumers behave statically and their expenditure functions are approximately translog.

The third fixed-base index I use is the Fisher index, which also aims to correct substitution bias. This index is a geometric average of the Laspeyres and the Paasche index. The Paasche index is defined as:

$$I_t^{Paasche} = \sum_i w_{it} \frac{p_{it}}{p_{i0}},$$

where  $w_{it}$  is the period  $t$  expenditure share, and the Fisher is defined as

$$I_t^{Fisher} = \sqrt{I_t^{Laspeyres} I_t^{Paasche}}$$

## H Appendix: Weighting of the Best Price Index

This section discusses some additional details related to the construction of the Best Price index. As discussed in the body of the paper, the Best Price index does not always capture the movements in the cost of living of low-storage-cost consumers who stockpile the most. Two possible reasons could be a) the time window for the index is not correct, and/or b) the minimum prices, which are averaged across stores, are not weighted properly. The time window should cover the planning horizon of households - if it is too narrow or too wide it will not correctly capture cost-of-living changes. Weighting can matter if low-storage-cost

---

<sup>5</sup>In 1999, the BLS started using a geometric price index to compute the CPI.

consumers shop at different stores than the general population, or if such consumers tend to be more likely to purchase in stores where promotions are occurring. To check this in Tables A3 and A4 I compute the best price indexes for different weights and different time windows. The first three rows of each table replicate the corresponding rows of Table 4 from the main paper. In the base case, the best price index is computed using the purchase shares across all stores, constructed over the entire sample period. The next row shows the best price index where the weights are the overall store shares for low-cost types only. The final row shows the index when weights are allowed to vary over time (the weights used are the weekly shares of each store). Using the purchase shares of low-cost types, and allowing them to vary over time, does bring the best price index slightly closer to the low-storage cost COLI (except for the second period of the tuna data). However, the impact of changing weights is small. The next rows show the same calculations for a bimonthly window. Changing the window brings the best price index closer to the low-cost type cost of living, and has a larger effect than changing the weights. For soup, the same patterns generally hold, although changing weights has a larger effect. Overall, as with canned tuna, the best fit seems to occur with a bimonthly index and with weights that vary over time, and reflect the purchase shares of low-cost types. This analysis suggests that if the measurement objective is to capture the cost of living changes of low-cost types, a wider window with time-varying weights reflecting low-cost type purchases is the best approach. In the body of the paper, I keep the weights fixed over time because to construct time-varying weights it is necessary to have time-varying quantity data. Such data is often not available to statistical agencies. Also, I only use population-level weights, rather than the weights for low-cost types only, since identifying such consumers without a structural model may be difficult.

Table A3: Sensitivity of the Best Price index to time window and weights (tuna)

Index	Weeks 1-72	Weeks 73-102
Dynamic Index (Monthly)	2.73	6.46
Dynamic (Low Cost, Monthly)	4.89	5.84
Best Price (Monthly, Fixed Weights, All Households)	-0.22	8.91
Best Price (Monthly, Fixed Weights, Low Cost)	0.34	9.7
Best Price (Monthly, Varying Weights, Low Cost)	0.5	8.86
Dynamic Index (Bimonthly)	0.92	4.64
Dynamic (Low Cost, Bimonthly)	3.38	4.32
Best Price (Bimonthly, Fixed Weights, All Households)	-1.3	5.14
Best Price (Bimonthly, Fixed Weights, Low Cost)	-0.9	5.68
Best Price (Bimonthly, Varying Weights, Low Cost)	-0.3	6.08

Table A4: Sensitivity of the Best Price index to time window and weights (soup)

Index	Weeks 1-34	Weeks 35-102
Dynamic Index (Monthly)	0.07	-4.28
Dynamic (Low Cost, Monthly)	0.66	-7.55
Best Price (Monthly, Fixed Weights, All Households)	0.67	-4.32
Best Price (Monthly, Fixed Weights, Low Cost)	0.94	-4.28
Best Price (Monthly, Varying Weights, Low Cost)	-0.54	-6.42
Dynamic Index (Bimonthly)	-0.26	-4.61
Dynamic (Low Cost, Bimonthly)	0.11	-8.1
Best Price (Bimonthly, Fixed Weights, All Households)	-2.5	-9.49
Best Price (Bimonthly, Fixed Weights, Low Cost)	-2.19	-9.77
Best Price (Bimonthly, Varying Weights, Low Cost)	-0.43	-7.74

# I Appendix Tables

Table A5: Purchase Shares by Product Category

Tuna		Soup	
Brand	Share	Brand	Share
StarKist	50.7	Campbell's	84.8
Chicken of the Sea	41.6	Store Brand	12.5
Store Brand	4.7	Swanson	1.2
3 Diamonds	1.9	-	-

Note: This table shows the purchase share of major brands of canned tuna and soup across the entire household panel for Sioux Falls. Market shares are only shown for brands with purchase share above 1 percent in both markets.

Table A6: Sale/Nonsale Transition Probabilities for Canned Tuna Brands

	Chicken of the Sea		StarKist	
	Sale	Nonsale	Sale	Nonsale
Sale	0.58	0.42	0.52	0.48
Nonsale	0.15	0.85	0.13	0.87

Note: The data used to compute this table is derived from the store level Nielsen files. The first column of the first row shows the probability that any Chicken of the Sea UPC goes on promotion in a given week conditional on being on promotion the prior week.

Table A7: Probability of Sale/Nonsale Conditional on Competitor Being on Sale/Nonsale for Canned Tuna Brands

	Chicken of the Sea		StarKist	
	Sale	Nonsale	Sale	Nonsale
Sale	0.38	0.62	0.22	0.78
Nonsale	0.26	0.74	0.23	0.77

Note: The data used to compute this table is derived from the store level Nielsen files. The first column of the first row shows the probability that any Chicken of the Sea UPC goes on promotion in a given week conditional on StarKist being on promotion the prior week. Similarly the first row, third column shows the probability that StarKist goes on sale in a given week conditional on Chicken of the Sea being on sale the prior week.

Table A8: Sale/Nonsale Transition Probabilities for Canned Soup Brands

	Campbell's		Private Label	
	Sale	Nonsale	Sale	Nonsale
Sale	0.45	0.55	0.13	0.87
Nonsale	0.02	0.98	0.01	0.99

Note: The data used to compute this table is derived from the store-level Nielsen files. The first column of the first row shows the probability that any Campbell's UPC goes on promotion in a given week conditional on being on promotion the prior week.

Table A9: Probability of Sale/Nonsale Conditional on Competitor Being on Sale/Nonsale for Canned Soup Brands

	Campbell's		Private Label	
	Sale	Nonsale	Sale	Nonsale
Sale	0.36	0.64	0.07	0.93
Nonsale	0.28	0.72	0.07	0.93

Note: The data used to compute this table is derived from the store-level Nielsen files. The first column of the first row shows the probability that any Campbell's UPC goes on promotion in a given week conditional on the private label being on promotion the prior week. Similarly the first row, third column shows the probability that the private label goes on sale in a given week conditional on Campbell's being on sale the prior week.

Table A10: Estimates of Utility Coefficients (Canned Tuna)

Coefficient	Mean	95% Conf. Int.	Std. Dev.	95% Conf. Int.
COS Water ( $\gamma$ )	0.704	(0.616, 0.791)	1.092	(0.996, 1.194)
SK Water ( $\gamma$ )	0.988	(0.894, 1.081)	0.878	(0.791, 0.975)
SK Oil ( $\gamma$ )	-0.163	(-0.299, -0.036)	-	-
Private Label Water ( $\gamma$ )	-1.202	(-1.529, -0.925)	-	-
3 Diamonds Water ( $\gamma$ )	-0.554	(-0.841, -0.266)	-	-
Consumption ( $\beta$ )	1.515	(1.366, 1.711)	0.499	(0.403, 0.637)
Storage Cost	0.177	(0.159, 0.196)	-	-
Price ( $\alpha_i$ )	1.233	(1.121, 1.344)	0.578	(0.486, 0.699)
Feature	1.682	(1.533, 1.823)	-	-
Display	0.823	(0.615, 1.008)	-	-
Fixed Cost of Purchase	2.705	(2.591, 2.834)	-	-
Prob of Low Cost Type	0.101	(0.074, 0.132)	-	-

Note: This table shows the estimates of the population average of the individual-specific draws on the utility coefficients (column 1), as well as the estimates of the population standard deviation of those coefficients (column 3). Columns 2 and 4 show the 95% confidence bounds on the population means and standard deviations. Dashes in the standard deviation columns indicate a coefficient is fixed across the population. The means and confidence intervals are taken over 20,000 Gibbs draws. 25,000 Gibbs draws are taken in total, and the first 5,000 are excluded to remove the influence of starting values. Abbreviations: SK is StarKist, COS is Chicken of the Sea.

Table A11: Estimates of Utility Coefficients (Canned Soup)

Coefficient	Mean	95% Conf. Int.	Standard Deviation	95% Conf. Int.
Campbell's Tomato ( $\gamma$ )	-2.716	(-2.923, -2.506)	-	-
Campbell's Chicken Noodle ( $\gamma$ )	-2.276	(-2.641, -1.952)	1.908	(1.583, 2.273)
Campbell's Cream of Chicken ( $\gamma$ )	-0.91	(-1.455, -0.344)	2.721	(2.139, 3.289)
Campbell's Cream of Celery ( $\gamma$ )	-2.493	(-3.741, -1.661)	2.593	(1.848, 3.578)
Campbell's Bean and Bacon ( $\gamma$ )	-4.077	(-5.988, -2.044)	3.678	(2.184, 5.037)
Campbell's Vegetable Beef ( $\gamma$ )	1.358	(0.752, 1.82)	-	-
Private Label Cream of Mushroom ( $\gamma$ )	-1.924	(-2.307, -1.542)	-	-
Private Label Tomato ( $\gamma$ )	-3.532	(-3.866, -3.217)	-	-
Campbell's Vegetable ( $\gamma$ )	-2.264	(-2.857, -1.791)	-	-
Campbell's Composite ( $\gamma$ )	6.342	(5.958, 6.748)	-	-
Private Label Composite ( $\gamma$ )	-2.007	(-2.572, -1.487)	3.231	(2.67, 3.837)
Consumption ( $\beta$ )	0.843	(0.732, 0.965)	0.52	(0.464, 0.583)
Storage Cost	0.133	(0.122, 0.145)	-	-
Price ( $\alpha_i$ )	1.183	(1.113, 1.256)	0.544	(0.45, 0.66)
Feature	-0.012	(-0.652, 0.538)	-	-
Display	2.007	(1.066, 3.198)	-	-
Fixed Cost of Purchase	2.859	(2.749, 2.985)	-	-
Prob of Low Cost Type	0.046	(0.021, 0.08)	-	-

Note: This table shows the estimates of the population average of the individual-specific draws on the utility coefficients (column 1), as well as the estimates of the population standard deviation of those coefficients (column 3). Columns 2 and 4 show the 95% confidence bounds on the population means and standard deviations. Dashes in the standard deviation columns indicate a coefficient is fixed across the population. The means and confidence intervals are taken over 20,000 Gibbs draws. 25,000 Gibbs draws are taken in total, and the first 5,000 are excluded to remove the influence of starting values.



Table A12: Consumption and Price Coefficients for Each Storage Cost Type

Coefficient	Canned Tuna		Canned Soup	
	Low	High	Low	High
Consumption	1.306	1.537	0.684	0.85
Price	1.816	1.172	1.87	1.152

Note: The table shows the averages of individual-specific draws for each storage cost type, for the final 20,000 of 25,000 Gibbs draws.

Table A13: Regression of Storage Cost Type Indicator on Demographic Variables

Coefficient	Canned Tuna	Canned Soup
Intercept	0.012 (0.052 )	0.018 (0.038 )
Income (20-50%)	0.028 (0.042 )	0.02 (0.036 )
Income (50-75%)	0.02 (0.044 )	0.033 (0.038 )
Income (75-100%)	-0.015 (0.044 )	-0.003 (0.039 )
Size > 2	0.025 (0.027 )	0 (0.031 )
Condo	-0.012 (0.298 )	-0.017 (0.124 )
Single Family	0.061 (0.055 )	0.021 (0.04 )
Multiple Family	0.009 (0.081 )	-0.025 (0.069 )
Mobile Home	0.138 (0.103 )	-0.03 (0.077 )
Other	0.475** (0.214 )	-0.04 (0.125 )
$R^2$	0.018	0.012

Note: An observation is a household. One star indicates significance at the 10% level, and two at the 5% level.

Table A14: Yearly Indexes

Product	Tuna			Soup		
	Average	Average	Change	Average	Average	Change
Weeks	1 - 51	52 - 102		1 - 51	52 - 102	
Laspeyres	-0.69	6.08	6.76	-0.06	-1.36	-1.3
FS	-1.76	3.74	5.5	-0.29	-1.62	-1.34
CK	-2.59	4.06	6.65	-1.85	-4.55	-2.7
Dynamic Model	-0.33	3.48	3.81	-1.63	-4.79	-3.16

Table A15: Regressions of Predicted Compensating Variations on Simulated Inventory

Storage Cost Type	Regressor	Canned Tuna	Canned Soup
High Cost Types	Intercept	0.176** (0.004)	0.475** (0.072)
	Simulated Inventory	0.903** (0.037)	0.539** (0.072)
	$R^2$	0.855	0.4812
Low Cost Types	Intercept	0.087** (0.005)	0.45** (0.103)
	Simulated Inventory	0.083** (0.003)	0.16 (0.104)
	$R^2$	0.9127	0.1954

Notes: The dependent variable in the regressions is aggregate average compensating variation, and the independent variable is average weekly inventory. An observation is a week (102 weeks are included). Standard errors are in parentheses. One star indicates significance at the 10% level and two at the 5% level.

Table A16: Summary of Regressions of Dynamic Index on Laspeyres and Best Price Indexes

Regressor	Canned Tuna	Canned Soup
Laspeyres	0.691** (0.086)	0.691** (0.045)
Best Price	0.292** (0.084)	0.313** (0.039)
$R^2$	0.9959	0.9991

Notes: Standard errors in parentheses. One star indicates significance at the 10% level, and two at the 5% level.

Table A17: Price Series in Simulation Study

Parameters:	
Parameter	Value
$\alpha$ (Price coefficient)	1
$\gamma$ (Consumption utility)	1
$SC$ (Storage cost quadratic for high-cost types)	-0.01
$FC$ (Fixed cost)	-0.25
Fraction of zero storage cost types	0.25

Price Series for Simulation 1 (Analog to Tuna):

Before			After		
$p_{t-1} / p_t$	0.5	2	$p_{t-1} / p_t$	0.5	1.5
0.5	0.1	0.9	0.5	0.1	0.9
2	0.3	0.7	1.5	0.3	0.7

Price Series for Simulation 2 (Analog to Soup):

Before			After		
$p_{t-1} / p_t$	0.5	10	$p_{t-1} / p_t$	0.4	10.1
0.5	0.01	0.99	0.4	0.01	0.99
10	0.1	0.9	10.1	0.04	0.96

## J Appendix Figures

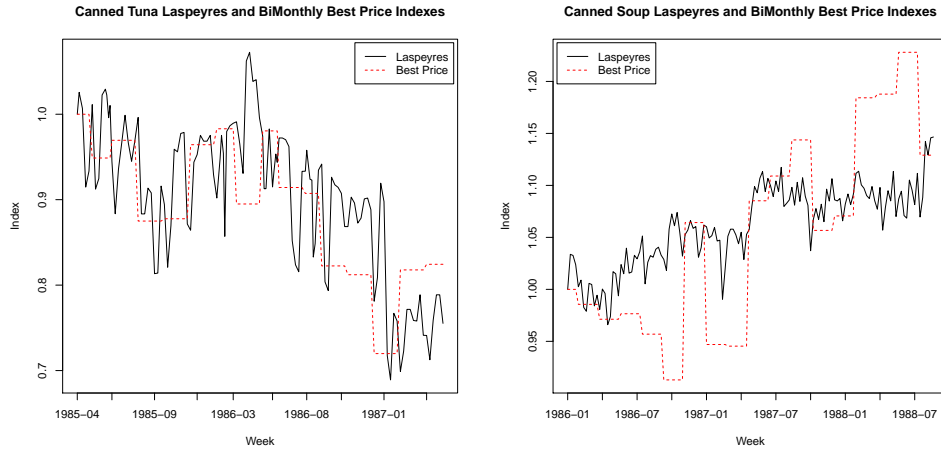


Figure A1: Laspeyres and Best Price (Minimum Bimonthly Price) Indexes




Figure A2: Timing of Purchase and Consumption within a Period

Notes: time series plots of the model parameters are shown in Appendix Figures A3 through A6. For population-varying parameters, the plots show draws of transformed hyperparameters. For utility parameters that are normally distributed the draws will show mean and variance. If the parameter is lognormal then the draws show the mean and variance of the lognormal distributions, ie, for a draw with mean parameter  $b$  and variance  $w$ , I would show the standard deviation of the lognormal, which is  $\sqrt{(\exp(w) - 1) \exp(2b + w)}$ .

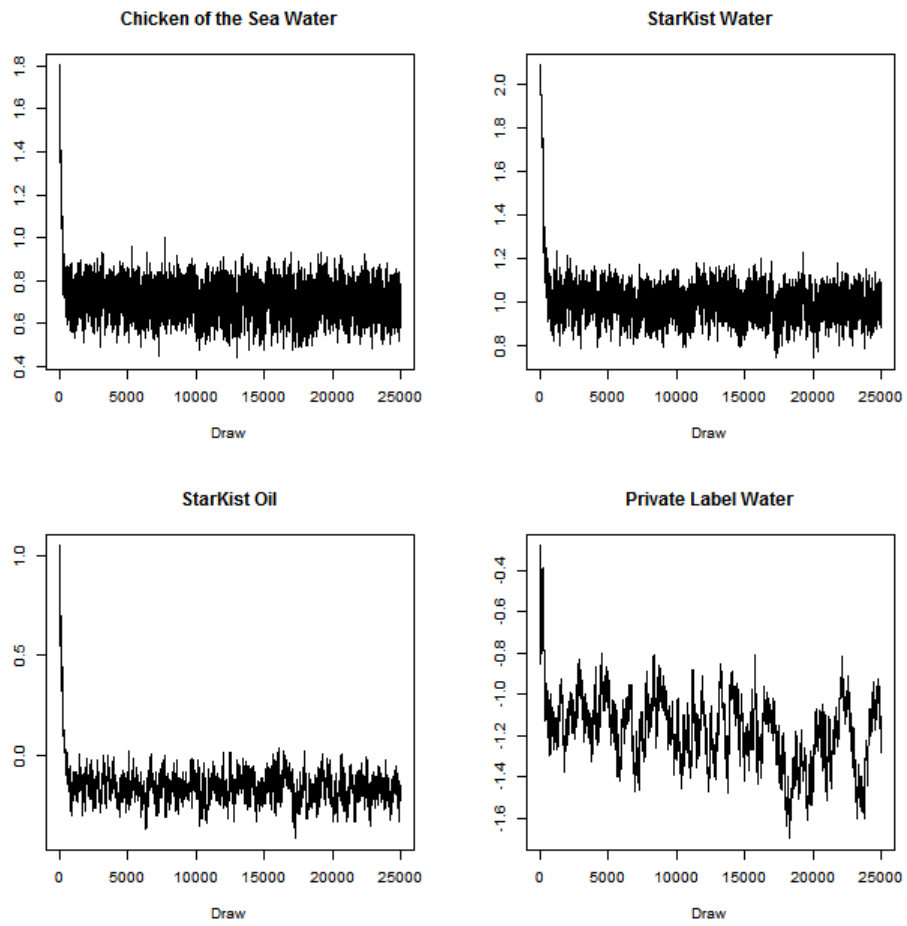


Figure A3: Plots of MCMC Draws for Means of Selected Product (UPC) Coefficients, Canned Tuna

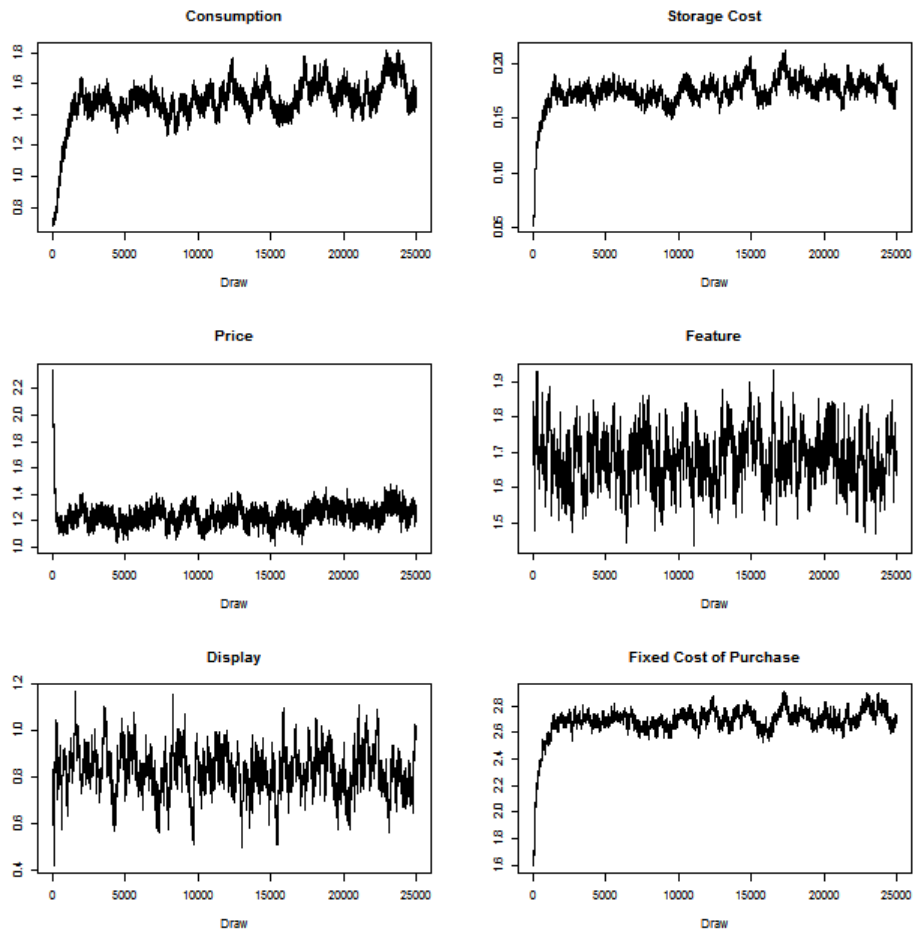


Figure A4: Plots of MCMC Draws for Means of Utility Coefficients, Canned Tuna

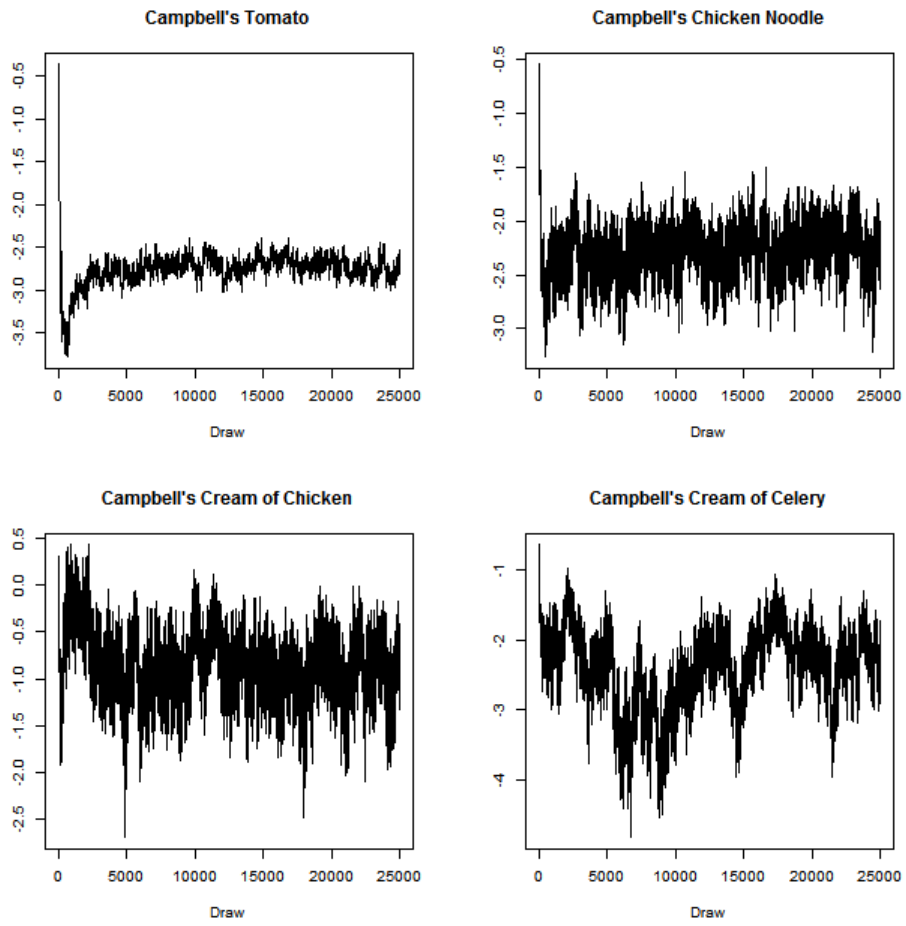


Figure A5: Plots of MCMC Draws for Means of Selected Product (SKU) Coefficients, Canned Soup

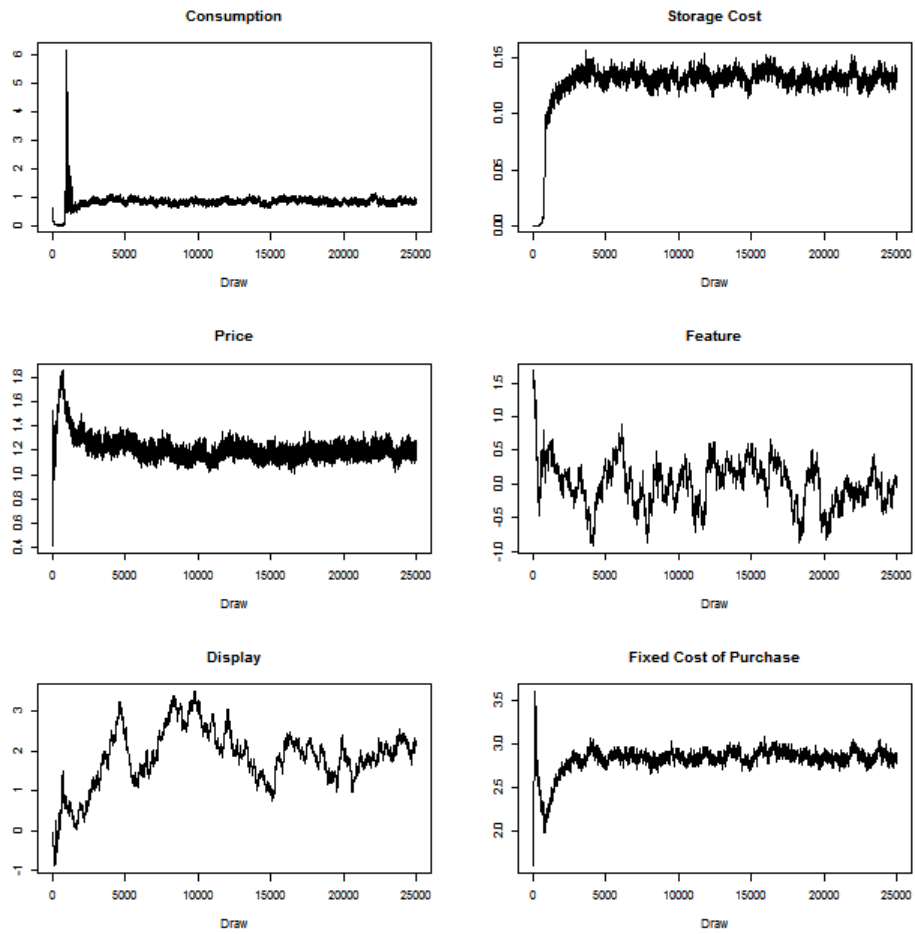


Figure A6: Plots of MCMC Draws for Means of Utility Coefficients, Canned Soup






Figure A7: Plots of Log-Likelihoods across Iterations




Figure A8: Plot of FS Index and Dynamic Index, with Local Linear Trends, for Canned Tuna

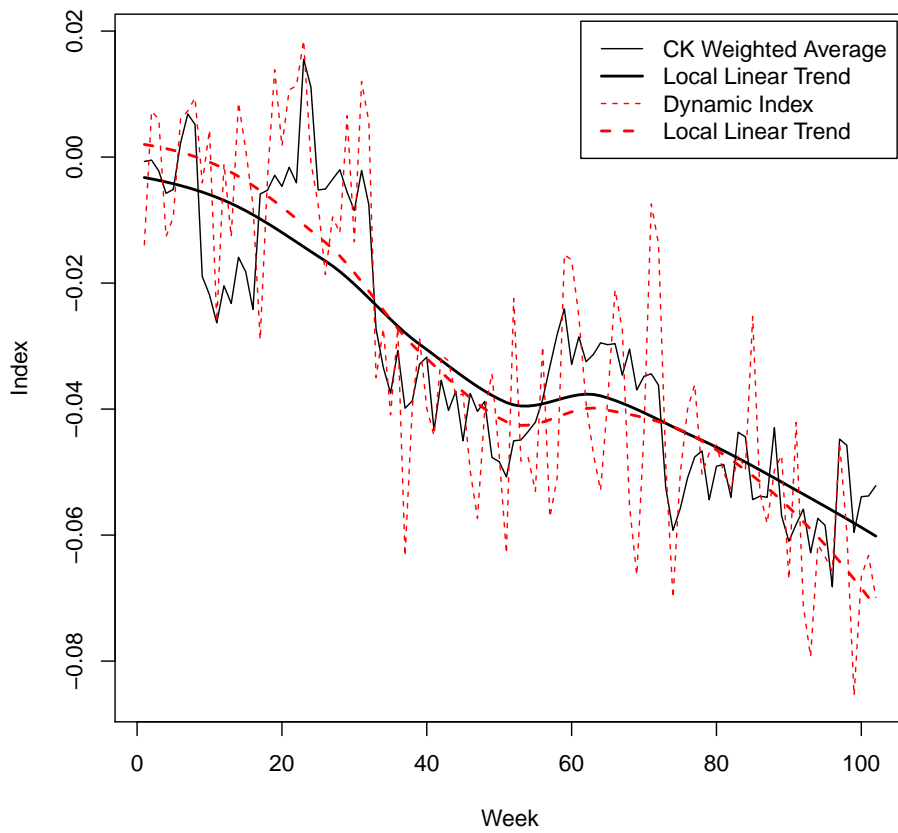


Figure A9: Plot of CK Weighted Average Index and Dynamic Index, with Local Linear Trends, for Canned Soup

## References

- Cardell, N. S.** 1997. “Variance Components Structures for the Extreme-Value and Logistic Distributions with Application to Models of Heterogeneity.” *Econometric Theory*, 13(2): 185–213.
- Caves, D., L. Christensen, and W. E. Diewert.** 1982. “Multilateral Comparisons of Output, Input and Productivity Using Superlative Index Numbers.” *Economic Journal*, 92: 73–86.
- Hendel, I., and A. Nevo.** 2006. “Sales and Consumer Inventory.” *The RAND Journal of Economics*, 37: 543–561.
- Imai, S., N. Jain, and A. Ching.** 2009. “Bayesian Estimation of Dynamic Discrete Choice Models.” *Econometrica*, 77(6): 1865–1899.
- Norets, A.** 2009. “Inference in Dynamic Discrete Choice Models with Serially Correlated Unobserved State Variables.” *Econometrica*, 77: 1665–1682.
- Osborne, M.** 2010. “Frequency Versus Depth: How Changing the Temporal Process for Promotions Impacts Demand for a Storable Good.” Working Paper.
- Train, K.** 2003. *Discrete Choice Methods with Simulation*. Cambridge University Press, New York.