Best Prices

Judith A. Chevalier
Yale School of Management and NBER

and

Anil K Kashyap
University of Chicago Booth School of Business,
Federal Reserve Bank of Chicago and NBER

This draft: December, 2010

Abstract: We explore the role of strategic price-discrimination by retailers for price determination and inflation dynamics. We model two types of customers, "loyals" who buy only one brand and do not strategically time purchases, and "shoppers" who seek out low-priced products both across brands and across time. Shoppers always pay the lowest price available, the "best price. Retailers in this setting optimally choose long periods of constant regular prices punctuated by frequent temporary sales. Supermarket scanner data confirm the model's predictions: the average price paid is closely approximated by a weighted average of the fixed weight average list price and the "best price". In contrast to standard menu cost models, our model implies that sales are an essential part of the price plan and the number and frequency of sales may be an important mechanism for adjustment to shocks. We conclude that our "best price" construct provides a tractable input for constructing price series.

The views expressed here are our own and not necessarily those of any institutions with which we are affiliated. Kashyap thanks the Chicago Booth Initiative on Global Markets for research support. We thank Cecilia Gamba, Aaron Jones, and Ashish Shenoy for outstanding research assistance. All mistakes are our own.

1 Introduction

We explore the importance of retailer price discrimination strategies for macroeconomic descriptions of price dynamics. The critical observation underpinning our analysis is that some consumers are active shoppers who chase discounts, substitute across products in a narrowly defined product category, and potentially use storage to maintain smooth consumption whilst concentrating their purchases in sale weeks. Other customers are passive, and retailers will employ strategies to charge these two groups different prices. Due to the actions of these strategic consumers, we find that weighted average prices paid differ substantially from posted prices. Moreover, accounting for the price discrimination motive on the retailer's side helps resolve several pricing "puzzles" that arise in the recent macroeconomics literature and provides several new facts about properly measured prices.

Our paper is related to the recent literature on pricing dynamics. Macroeconomists have struggled with the question of whether and how to incorporate intermittent price discounts, or "sales", into models of price setting and into price index construction. Since episodic sales occur quite frequently in some sectors, price series that ignore sale prices display infrequent price changes, while price series that contain sales prices display very frequent price changes. Nakumura and Steinsson (2008) estimate a median price duration of approximately 5 months including sales, and approximately 8 to 11 months excluding them. Klenow and Kryvtsov (2008) find that considering only regular (e.g., non-sale) prices raises the estimated median price duration in their dataset from 3.7 months to 7.2 months. In general, many researchers have documented that the price series for a given product from retailers such as grocery stores exhibit long periods of a constant "regular" price punctuated by occasional sales.

The question of how to treat sales in constructing price series is closely related to the question of whether price changes due to sales have implications for theories of monetary neutrality. The data from grocery stores and other retailers are fundamentally difficult to reconcile with the workhouse macroeconomic models of price setting. In particular, the coexistence of frequent weekly price changes due to sales with extended periods of constant "regular" prices cannot be reconciled readily with either flexible price models or menu cost models. This has led Kehoe and Midrigan (2010) and Eichenbaum, Jaimovich, and Rebelo (forthcoming) to modify the standard menu cost model to allow the firm to have a low cost of producing temporary (or sale) price changes but a high cost of producing permanent (or reference) price changes. Kehoe and Midrigan accomplish this by having two menu costs—one for a temporary price change and one for a permanent price change. Eichenbaum, Jaimovich, and Rebelo accomplish this by allowing the firm to have a "price plan" which consists of two prices—firms can freely change between the two prices, but bear a fixed cost of altering the price plan. While both of these approaches do a much better job of matching the data than a standard menu cost model, neither approach provides a microfoundation for why a firm might have a low cost of shifting between two prices, but a higher cost of changing a "regular" price.

The "modified menu cost approach" has the implication that temporary price changes can be ignored by macroeconomists in considering monetary neutrality. The central argument for ignoring sales when examining macroeconomic questions stems in part from the (unproven) belief that sales motivated by price discrimination are orthogonal to price responses to macroeconomic processes. As Mackowiak and Smets (2008) note in their recent survey, "It is possible that variation in "reference" prices captures most of the variation in prices that matters for macroeconomics, that is, most of the variation in prices reflecting the response to macro shocks. Deviations from reference prices tend to be transient, whereas macro price shocks tend to be persistent." Eichenbaum, Jaimovich, and Rebelo note, "a lot of high-frequency volatility in prices and quantities has little to do with monetary policy and is perhaps best ignored by macroeconomists." Similarly, in Kehoe and Midrigan's model, since sale-related prices are transient, they do not allow firms much flexibility to respond to persistent nominal shocks.

Our goal in this paper is to provide theory and preliminary evidence suggesting that sales are a crucial phenomenon, even for macroeconomists. Our theoretical model is similar in construction to Varian (1980), Sobel (1984), or Pesendorfer (2002). The retailer price discriminates between two types of consumers-those who have high valuations for particular products, and consumers who have lower valuations, are willing to store goods, and do not have a particular affinity for a given brand. Thus, our model contains consumer heterogeneity in preferences and shopping behavior as its primitives. We show that such a model produces frequent temporary sales and long periods of constant "regular" prices.

Previous examinations of pricing data have diagnosed menu costs in part because a standard Dixit-Stiglitz (1977) model with flexible prices predicts constant markups; the prediction of constant markups is not supported by the data. In contrast, in our model, there are a cost and demand shocks for which adjustment of the regular price is not the optimal response—despite the absence of any menu costs. We show that adjustment of the frequency and the depth of sales is the optimal to response to cost shocks and to certain types of demand shocks. In this framework, even with perfect price flexibility and with unchanging costs and demand, optimal markups are not held constant across items or across time.

Finally, we show that our model produces a natural means of price measurement. Average price paid (or the variable weight price index) is a weighted average of the prices paid by the different types of consumers. In particular, we show that the average price paid is an average of the fixed weight price index and the "best price", the lowest price for any good in the choice set in a given week.

We conduct an empirical examination of data for particular supermarket products collected over parts of seven years at Dominick's Finer Foods (DFF), a supermarket chain in the Chicago area. We supplement our analysis with data for a later six-year period, using a dataset provided by Symphony IRI. The IRI dataset covers stores in 47 markets around the country and we (randomly) selected one store in each of the nine census regions for evaluation. Prices for individual products at DFF and at the IRI stores display the now-familiar pattern of very infrequent regular price changes combined with frequent temporary sales. But once the data are analyzed through the lens of our model, four regularities predicted by the model emerge.

First, the model suggests that average prices paid for goods within a bundle of close substitutes are the relevant "price" for consumers and retailers. Because retailers time sales strategically, the model suggests, and the data show, that sales for close substitute products tend not to occur in the same weeks. Consumers, we show, "chase" sales, and thus, effective prices paid are substantially less than regular prices, and even measurably less than average prices posted. For example, for our peanut butter data, in our reference store in Charlotte North Carolina, the actual price paid in a quarter ranges from 1.0% to 17.6% less than the quarterly "reference" (e.g., regular) price and from 0.3% to 15.3% less than the fixed weighted average list price.

Second, we introduce the concept of a "best" price, defined as the lowest price charged for any good in the narrow product category during a short multi-week time window. The model predicts that best prices should be the relevant prices for sales-chasing consumers. We show that the actual price paid tracks the "best" price and is well-approximated by an average of the "best price" and the fixed weight index. The data match the structural form of our model.

Third, the data confirm the model's key prediction of strong spillovers in quantities purchased due to price changes for close substitutes. So, for example, when the price of Minute Maid orange juice is reduced for a temporary discount, sales of Tropicana orange juice plunge (provided the Tropicana price does not change). In previous literature, large quantity variation along with constant prices has lead to a diagnosis of substantial demand variability. However, we show that much of this apparent demand volatility derives from choices made by the retailer in setting prices for substitute products. We show that total ounces sold of all products within a narrow product category are much less volatile than total ounces sold of any individual product. When examining total ounces sold for the categories we examine, there is not tremendous unexplained week to week variation in demand.

Finally, the model suggests that varying the intensity and depth of sales can substitute for adjusting regular prices. We examine the extent to which the number and frequency of sales can change through time, even over periods during which the "regular" price is kept constant. We identify quarters where the regular price posted for a product or set of products is constant, but changes in the frequency and intensity of temporary discounts for products in the product category have led to substantial differences across quarters in the average price paid.

Overall, we conclude that a reconsideration of the motivation for sales leads to a redirection of the literature. Care needs to be taken in drawing inferences about demand variability using data on individual items. Because consumers chase sales across products within a category and across time and because retailers can alter the frequency and depth of these discounts, our results suggest that sales are important, even for macroeconomists. A more robust approach requires looking at bundles of items over time. The price path that accounts for cross- product and intertemporal substitution looks different from the price path for an individual item. This perspective suggests that the effect of inflation (or monetary policy) on price setting will depend critically on how consumers update reservation prices for individual goods. We also provide some preliminary thoughts on how to construct summary price series in sectors characterized by price discriminating retailers.

Our paper proceeds as follows. Section 2 provides the simple model of a price-discriminating retailer and highlights empirical predictions. Section 3 describes the data. Section 4 demonstrates the empirical results highlighted above. Section 5 discusses how to use our model to measure and summarize price series. Section 6 concludes.

2 A Model of Price Discriminating Retailers and Heterogeneous Consumers

We begin by presenting a simple model that is similar in spirit to, and borrows significantly from, Varian (1980), Sobel (1984), or Pesendorfer (2002). The baseline version of our model takes consumer heterogeneity as its primitive. The firm knows about this heterogeneity and accounts for it in price setting. In this model, the firm bears no menu cost of changing prices. Nonetheless, we will show that the firm will iterate between a small number of prices, even in the face of some cost changes and some types of demand changes. The "regular" price will change infrequently but "sales" will be utilized. The model suggests that it is possible that the retailer will respond to a nominal shock by changing the frequency of sales while holding the "regular" price fixed.

2.1 Model Assumptions

Consider a single retailer selling two substitute differentiated products, A and B. We will focus on a single retailer for simplicity. However, we note that it would be fairly straightforward to embed our model into a model of two retailers competing in geographic space. In such a model (see, for example, Lal and Matutes (1994), Pesendorfer (2002) and Hosken and Reiffen (2007)), consumer reservation prices would be determined by the price that would trigger consumer travel to another store. Thus, for tractability, we focus on a single retailer, but a monopoly assumption is not necessary.

Assume that all customers have unit demand in each period but are differentiated in their preferences for the two substitute goods. A share $\alpha/2$ of the customers value product A at V^H and product B at V^L, where $V^H > V^L$. We call these consumers the high A types. For convenience, we consider the symmetric case where a share $\alpha/2$ of the customers, the high B types, value product B at V^H and product A at V^L. The remaining share of consumers $(1-\alpha)$, the "bargain hunters", value both products at V^L. We normalize the total number of consumers to be 1 and consider N shopping periods.

The seller has a constant returns to scale technology of producing A and B and the marginal cost of producing either is c.

Each period, customers arrive at the retailer to shop. Consider the choices for the high A types (which will be symmetric for the high B types). If the price is less than or equal to their reservation value ($P^A < V^H$), then they buy their preferred good, A. If $P^A > V^H$ and the price for product B is less than or equal to their reservation price for B ($P^B < V^L$), the high A type customers substitute to good B. If $P^A > V^H$ and $P^B > V^L$, then the high A types make no purchase. For simplicity we also assume that, if the high types do not buy, their demand for the period is extinguished so that next period there are no implications of them having been out of the market.[1]

Next consider the choices made by the bargain hunters. If $P^A < V^L$ and/or $P^B < V^L$, they will buy whichever product is cheaper.[2] If $P^A > V^L$ and $P^B > V^L$, the bargain hunters do not buy. However, we capture the idea of shoppers being willing to engage in intertemporal storage by assuming that their demand partially accumulates to successive periods, deteriorating at rate ρ. Thus, for example, if they made a purchase in period $t-1$, but $P^A = P^B = V^H$ in period t so that no purchases are made, then their total demand entering period $t+1$ will be $(1-\alpha)+(1-\alpha)\rho$. Similarly, if a good was available at a price of V^L in period $t-1$, but $P^A = P^B = V^H$ in periods t, $t+1$, ..., $t+(k-1)$, total demand from the bargain hunters in period k will equal:

$$(1-\alpha)(1+\rho+\rho^2+\cdots+\rho^k) = (1-\alpha)\frac{(1-\rho^k)}{1-\rho} \quad \text{if } 0 < \rho < 1.$$

Note then, that we are assuming that the high types are inactive shoppers —they do not wait for and/or stock up during bargains, while the low types do. In this sense, our model reflects well the empirical facts described in Aguiar and Hurst (2007), in which they document that some consumers in a local area pay systematically lower prices for the same goods as other consumers. That is, some consumers are strategic in bargain-hunting, and others are not.[3]

Total profits for the retailer depend on the total amount of A and B sold. The retailer has three basic choices: (i) the retailer can charge high prices and service only the high types, foregoing any potential margins to

[1]In our model, as long as costs are less than V^H, the firm will always set price so that the high types purchase. We could make an assumption about high type demand accumulating, but it wouldn't have any important implications for the model.

[2]We will see that in equilibrium it will not be profit maximizing to put both goods on sale in the same week.

[3]As in Pesendorfer (2002), we combine "bargain-hunting" behavior with low willingness to pay. We could provide a more detailed model with more types—brand loyals who are willing to intertemporally substitute purchases, brand loyals who do not intertemporally substitute, non-loyals who are willing to intertemporally substitute and non-loyals who do not intertemporally substitute. We think most of the interesting implications are evident with these 4 types collapsed into the two extremes.

be earned on the low types. (ii) the retailer can charge low prices and serve both types of customers, thus foregoing the extra willingness to pay that could have been extracted from the high types. Or (iii) the retailer can strategically iterate between high and low prices in an attempt to capture some of the potential demand from the bargain-hunters while exploiting some of the extra willingness to pay of the high types.

2.2 Model Results

Depending on the parameters, strategy (i), (ii), or (iii), described above can be optimal. We will explain when each is optimal, with particular attention to parameter values under which (iii) is optimal, since the pricing behavior associated with (iii) is roughly consistent with our empirical observation of occasional sales at supermarkets. We characterize the retailer's behavior in several steps.

Proposition 1: As long as $V^L > c$, the optimal price in any period is either V^L or V^H.

Sketch of Proof: Choosing a price between V^H and V^L would reduce margins on the high types but would produce no offsetting demand increase for the bargain hunters. Thus, the optimal price is either V^H or V^L.

Proposition 2: It is never optimal to charge V^L for both good A and good B in the same period.

Sketch of Proof: Charging V^L for the second good leads to a loss of margins on the high types that prefer that good, but produces no offsetting demand increase for the bargain hunters. The bargain hunters' demand is satisfied by charging V^L for *either* good. This delivers one of the prediction that we will test empirically, namely that discount periods or "sales" for products within a product category are not synchronized.

Thus, we have shown that the retailer will charge V^H for at least one good in every period, and may charge V^L for one good in some periods. Below, we show the conditions under which the retailer will charge V^L for one good in every period, conditions under which the retailer will never charge V^L for either good, and conditions under which the retailer will charge V^L for one good in some periods but not all periods (intermittent sales). In order to derive these results, we first provide two straightforward intermediate results.

Proposition 3: If the retailer chooses to hold one and only one sale during the N periods, then the optimal time to hold it will be in the Nth (final) period.

Sketch of Proof: If the retailer charges V^H for both products during

$N-1$ periods, and V^L for one product in 1 period, profits will be:

$$(N-1)\alpha(V^H - c) + \frac{\alpha}{2}(V^L - c) + \frac{\alpha}{2}(V^H - c) + \frac{(1-\rho^k)}{1-\rho}(1-\alpha)(V^L - c) \quad (1)$$

where k represents the period in which the sale is held. The first term of (1) is the profits from the high types in all of the non-sales periods, the second and third term represent the profits from the high types in the sale period (if the sale is on A, the high A types pay V^L, but the high B types still pay V^H). The fourth term of (1) represents the profits from the bargain hunters. Note that the first three terms are invariant to the timing of the single sale. The fourth term is maximized when $k = N$ (as long as $\rho > 0$).

Proposition 4; For a retailer who chooses to hold j sales (i.e. charge V^L in j periods), the optimal strategy is to hold evenly spaced sales every k periods, where $k = N/j$.

Sketch of Proof: The prior proposition covered the case of $j = 1$ and $k = N$. Note that the logic underlying Proposition 3 for the entire period carries forward straightforwardly to sub-periods, so within each sub-period it makes sense to delay the sale as long as possible. Hence sales will be equally spaced.

With the results of Propositions 3 and 4 in hand, we can characterize the remaining decision about when it pays to have any sales at all. We consider a retailer who holds a sale every kth period. If $k = 1$, the retailer holds a sale every period. If $k = N$, the retailer holds the minimum positive number of sales—one sale at the end to "sweep up" the low demanders. If $k > N$, the retailer never holds sales. We will focus on interior solutions where $1 < k < N$.

We consider the profits of a retailer who charges $P^A = P^B = V^H$ every period except during a "sale" and holds a sale at $P = V^L$ for one or the other good every kth period. Assuming no discounting, total profits for this retailer over all N periods are:

$$N\frac{k-1}{k}\alpha(V^H - c) + \frac{N}{k}\frac{\alpha}{2}(V^L - c) + \frac{N}{k}\frac{\alpha}{2}(V^H - c) + \frac{N}{k}\frac{(1-\rho^k)}{1-\rho}(1-\alpha)(V^L - c) \quad (2)$$

The four terms in (2) are very intuitive. The first piece represents the profits from selling to the high types only, which will occur during all the non-sale periods. The second term is the profits from the high types during the periods where they are able to buy their preferred good on sale; during these sale periods the other high type still pays V^H so that explains the third term. The last term is the profits from the bargain hunters. Note that, the larger ρ is, the less the bargain hunters' demand has depreciated by the time the sale is held.

Proposition 5: The retailer will find it optimal to hold some sales if:

$$\frac{\alpha}{2}(V^H - V^L) < (1-\alpha)\frac{1-\rho^N}{1-\rho}(V^L - c) \Leftrightarrow$$

$$V^H < V^L + 2\frac{1-\alpha}{\alpha}\frac{1-\rho^N}{1-\rho}(V^L - c) \tag{3}$$

Sketch of Proof: The left hand side of first expression shows the loss from allowing the high types to pay less than they are willing to pay by offering a single sale. The right hand side of the first expression gives the profits from selling to the bargain hunters in an optimally timed single sale.

The second expression in (3) helps build intuition. Essentially, holding at least one sale will be optimal as long as V^L is large "enough" relative to c and relative to V^H and if there are enough bargain hunters. The total demand of the bargain hunters depends on both their share in the population and the extent to which their unmet demand cumulates. An increase in the share of bargain hunters (or increase the persistence of their demand) will lead to more sales.

Note also that, if the stock of consumers depreciates completely ($\rho = 0$) from period to period, then the condition degenerates to the condition for profits from charging the low price being higher than profits from charging the high price (period by period):

$$\frac{\alpha}{2}(V^H - c) < \left(1 - \frac{\alpha}{2}\right)(V^L - c)$$

Assume that the condition in Proposition 5 holds so that the retailer will hold at least 1 sale. By construction, for given values of α, ρ, V^H, V^L, and c, the seller will then choose k to maximize (2). That is, the seller will deterministically hold a sale every kth period. We are assuming that k and N are positive integers. However, we can examine comparative statics involving the optimal choice of k by treating these functions as continuous and maximizing profits with respect to k. We are considering the cases here where N is large.

With profits given by (2), the first order condition for k is:

$$k = \left(\frac{1}{\log \rho}\right) \times \left(1 + W\left(\frac{\alpha(V^H - c) - (2-\alpha)(V^L - c) - \alpha\rho(V^H - V^L)}{2e(1-\alpha)(V^L - c)}\right)\right) \tag{4}$$

Where $W(\cdot)$ is the Product Log or Lambert function. The Lambert function provides the solution for z such that $W = ze^z$. It is used extensively to model time delay problems in operations research—a family of problems akin to the timing of sales problem studied here. For our purposes, it is useful to note that W is a monotonically increasing nonlinear function. Since k is the number of periods between sales, a larger k means less frequent sales.

While analytic solutions are unavailable, by fixing some parameters we can present a few interesting and illustrative comparative statics numerically. For example, consider the case where $V^H = 4$, $c = 1$, $\rho = 0.9$, $\alpha = 0.4$. Figure 1 shows the optimal k for V^L varying from 0 to 4.

Not surprisingly k is decreasing in V^L: the higher the valuation of the low valuation types, the shorter the time between sales; serving the low valuation types is attractive. Note that k is not defined if V^L is too low, because at some point it ceases to make sense to have sales. Also, note that as $V^L \to V^H = 4$, the optimal k falls below 1. That is, as V^L gets large enough, it makes sense to charge V^L for one product every period.

Alternatively consider the effect of c on the optimal k. Fix $V^H = 4$, $\rho = 0.9$, $\alpha = 0.4$ as before, and fix $V^L = 2$. Figure 2 shows that, as would be expected k is increasing in c: the greater is marginal cost, the lower the profits from serving the bargain hunters and the less frequently one would want to hold sales. As $c \to V^L = 2$, $k \to \infty$.

Given the structure of our model, the firm is always charging V^H when there is not a sale. So changes in marginal cost that are unaccompanied by changes in buyer willingness to pay will not result in a change in the non-sale price.

Lastly, consider changes in α, the share of customers with high willingness to pay for one of the goods. As before, fix $V^H = 4$, $V^L = 2$, $c = 1$, and $\rho = 0.9$. Figure 3 shows the optimal k as α ranges from 0 to 1. As α increases, i.e. the share of bargain hunters falls, sales become less attractive. If $\alpha = 1$, optimal k is infinite, because sales are never attractive.

From the retailer's point of view the relevant "price plan" is the full sequence of high and low prices that prevail over the cycle of $2k$ periods. Note that even with unchanging cost and demand parameters, for many parameter values, the firm changes prices from period to period as it optimally iterates between capturing rents from high types and capturing the demand of the "bargain hunter" low types. Fixing tastes and technology in this set up, the main choice variable of the retailer is k. P^A and P^B are choices in each period but are set by the willingness to pay of the consumer types.

It is useful to compare the outcomes of this model to the models proposed in Kehoe and Midrigan (2010) and Eichenbaum, Jaimovich, and Rebelo (forthcoming). Both of those models can produce a result of a firm charging a fixed regular price and sometimes charging a sale price. However, in both of these papers, the decision to charge a sale price is driven by some change in the cost or demand environment. In our model, a sale would occur every k weeks with no change in the cost or demand environment.

It is also useful to enumerate the circumstances in our model that would lead to a change in the regular price. The regular price is held constant if there are shocks to any of the following parameters: V^L, c, ρ, and α. If any of those parameters change, the retailer's optimal response is to alter the frequency and depth of sales.

One way to combine the comparative statics described above is to simulate the model allowing for changes in preferences and costs, under the assumption that the seller expects each change to be permanent. In the following figure we fix the persistence of unmet demand, but allow the share of the bargain hunters, the valuations of the bargain hunters and high types, and marginal cost to each change once per quarter. In this experiment the taste/cost shocks occur every thirteen weeks.

A resulting price path from this simulation for good A is plotted in Figure 4. Having allowed for a full set of shocks, the model now predicts changes in prices, the frequency of sales and the size of the discount that occurs during a sale. While our model is highly stylized, the price pattern from our simulation in Figure 4 does look extremely similar to a price series for substitute grocery products.

Finally, it is helpful to use the model to think about quantities sold. In the simplest case where demand does not deteriorate at all from period to period but just stockpiles ($\rho = 1$), total units purchased over all periods equals N, independent of the number and frequency of sales. This follows literally from the assumption of unit demand per period. But even so, the number of units sold in each individual period, however, is a function of the timing of sales; bargain hunters move all of their demand into sale periods.

The only distortion in total units purchased over all N periods relative to the scenario in which $P^A = P^B = V^L$ in all periods stems from the deterioration in bargain hunter demand while waiting for sales. Thus, for general ρ, total units sold equal $= N\alpha + \frac{N}{k}\frac{(1-\rho^k)}{1-\rho}(1-\alpha)$.

Note that, in any individual period, sales of an individual good could be as low as $\alpha/2$ and as large as $\alpha + k(1-\alpha)$, if the good is on sale and deterioration of "bargain hunter" demand is negligible. Note that this volatility in demand across products and periods occurs despite the static environment that we modeled; quantity sold varies across periods and across goods despite the fact that the demand and supply primitives are constant through time. This implies that, while individual good sales are volatile period by period, by summing across goods A and B the total amount sold will be invariant across k-period cycles. Building on this prediction, we look whether the volatility in quantities sold are smaller for a collection of close substitutes over a full cycle of weeks that include at least one discount, than in individual weeks or for individual goods.

The preceding results are also helpful for thinking about effective prices paid. Summing over all periods and the products A and B, V^H will be the price charged in at least half of the product-periods. In cases in which there are interior solutions for k (as depicted above), V^H will be the price charged in more than half of all product-periods. Specifically, if A and B are discounted equally frequently, the share of periods in which a given product is "on sale" is $1/2k$. Thus, applying the algorithms proposed by either Eichenbaum, Jaimovich and Rebelo (forthcoming) or Kehoe and Midrigan (2008) would identify V^H as the regular/reference price for both goods. A "fixed weight" index measuring average price over the entire period would measure an average price of:

$$\frac{1}{2k}V^L + \frac{2k-1}{2k}V^H \tag{5}$$

This price (and certainly the "regular" price) does not capture average price from the retailer's perspective or from the perspective of the bargaining-hunting customer. Our effective price paid puts much more weight on the low price, because it reflects the strategic shift of the bargain-hunters into the low priced product (when one is on sale) and the stockpiling of bargain hunter demand (when nothing is on sale).

This is easiest to see by first neglecting/assuming away the deterioration in bargain hunter demand. With no demand deterioration, average revenue per unit equals:

$$\frac{\alpha + 2k(1-\alpha)}{2k}V^L + \frac{\alpha(2k-1)}{2k}V^H \tag{6}$$

Of course, we know that lower prices expand demand, but in this case, that factor is magnified by the presence of multiple products that bargain hunters view as substitutes and because bargain hunters stockpile demand. It will be very helpful to note the relationship between (5) and (6). Recall that (5) is the fixed weight index and that (6) is the average revenue per unit or effective price paid. Note also that the lowest or "best price" achieved over the k period sales cycle is V^L. The average price paid in (6) can be rearranged to be equal to:

$$\alpha\left(\frac{1}{2k}V^L + \frac{2k-1}{2k}V^H\right) + (1-\alpha)V^L \tag{7}$$

That is, the effective price paid in (6) is equal a weighted average of the fixed weight price index in (5) and the "best price", where the share of loyals and of bargain hunters in the population are the weights.

The average revenue per unit (or effective price paid per unit) is more complicated when demand deterioration is taken into account. The weighted average price then becomes:

$$\frac{\left(\frac{\alpha}{2k} + \frac{(1-\alpha)(1-\rho^k)}{k(1-\rho)}\right)V^L + \frac{\alpha(2k-1)}{2k}V^H}{\alpha + \frac{(1-\alpha)}{k}\left(\frac{1-\rho^k}{1-\rho}\right)} \tag{8}$$

The average revenue in (8) approaches that in (7), or equivalently (6), as ρ approaches 1. For to $\rho < 1$, the weighted average price in (8) is slightly higher than in (6) because demand deterioration destroys some of the demand when sale prices are not offered. Thus, the average revenue per unit, or effective price paid, is a weighted average of the price paid by the high types, and the price paid by the low types.

Note also from (8) that, if ρ is small, then as α approaches 0, the weighted average price will approach the "best price"—the lowest price posted for any of the substitute products within the k period planning cycle. Note also that if α is fairly small and if ρ is small and V^H remains constant for a long period of time, then a time series of the weighted average price will resemble a fixed increment over the time series of the "best price". If α is very large, then the weighted average price will more closely resemble the "regular" or "reference" price. For small to modest values of ρ and intermediate values of α, the price paid resembles a weighted average of the "best" price and the fixed price index as illustrate in (7). We examine these possibilities in the next section.

The model extends straightforwardly to more substitute goods. One can imagine more complex demand relationships that could be exploited by the retailer. For example, in our model, each product is associated with a cadre of brand-loyal consumers. However, it is possible that the retailing landscape includes some goods for which some consumers are brand loyal and other goods for which no consumers are brand loyal (possibly private label goods). As mentioned above, there may be some consumers who inter-temporally substitute actively, but are highly brand loyal. There may be other consumers who are not brand loyal but do not inter-temporally substitute. In these more complex cases, the formulation and intuition in (6) is particularly helpful. Specifically, we can think of the prices paid for a set of closely related goods over a shopping cycle as being characterized by the weighted average of the prices paid by non-bargain hunter loyal shoppers and the prices paid by bargain hunters. As shown above, the prices paid by the loyal shoppers are essentially a fixed weight average of the prices posted. However, the prices paid by bargain hunters more closely resemble the lowest price charged for any substitute good over a reasonable shopping horizon. Below, we will show empirically that the prices paid in our data resemble a hybrid of a fixed weight index and the "best price". We use these insights to speculate about pricing implications in more complicated environments than those described by the model and we provide a preliminary discussion of approaches to measuring and summarizing prices in such environments.

Summing up, the model comfortably explains the familiar price pattern observed for individual goods of a regular price with intermittent sale prices. In addition, it makes the following testable predictions. First, sales across items should be staggered. Second, demand will be more stable for bundles of close substitutes across several weeks (that include at least one sale) than for individual weeks or for individual items over that same period. Third, the actual price paid will be much lower than a regular or normal price if bargain hunters are important. Fourth, prices paid should generally be a combination of a conventional fixed weight price index and the best available price within the group of close substitutes over the course of several weeks. Finally, it is possible that the frequency of sales (or the size of discounts) will be a meaningful margin of adjustment in a pricing strategy. More generally, the updating rules that consumers use in revising their reservation prices will play a critical role in determining prices.

3 The Data

The data for two of the categories we analyze are taken from Dominick's Finer Foods (DFF). DFF is a leading supermarket chain in the Chicago metropolitan area; they have approximately 90 stores and a market share of approximately 20%. Dominick's provided the University of Chicago Booth School of Business with weekly store-level scanner data by universal product code (UPC) including: unit sales, retail price, cents of profit per dollar sold, and a deal code indicating shelf-tag price reductions (bonus buys) or in-store coupon. The DFF relationship began in 1989 (with week 1) and our dataset ends in 1997 (with week 399).[4] To explore the predictions of the model we focus on data for frozen concentrated orange juice (frozen OJ in what follows) and oatmeal products.

DFF has four types of stores that vary in the average level of prices. These average price differences are based on the wealth of the surrounding neighborhoods and the amount of local competition. For the results below, we focus on a single store, although we have prepared robustness checks for groups of stores in all of the price tiers. The store we use is located in the northwest part of the city of Chicago (near the boundary with Skokie) and has prices that are in the medium Dominick's pricing tier.

We also use data from a different time period and location using Symphony IRI's "IRI Marketing Data Set" (Bronnenberg, Kruger and Mala (2008)). This dataset allows us to study pricing at other stores outside Chicago from a more recent period. The IRI dataset contains information on retailers in 47 market areas. For our benchmark calcuations, we use data from "Chain 35" which has 110 stores between Pennsylvania and South Carolina, with the largest concentrations in Raleigh/Durham, Washington DC, and Charlotte. As part of its efforts to preserve the anonymity of chains, IRI assigns different chain numbers to the regional divisions of large retailers. Thus, Chain 35 is possibly part of a retailer with more than 110 stores.

We focus on creamy peanut butter. Stores within chain 35 appear to charge widely varying prices in the peanut butter category. We chose a store in Charlotte NC, store 250517, with medium prices and no missing data for all 313 weeks. Chain 35 overall has about a 33% share of the

[4]There are some missing data. For weeks 219, 232–233, 266–269, 282–283, 358–361, 370–71, 388, and 394 all categories are missing; the particular store we use had data missing for weeks 33 to 36 and 50; most stores have incomplete information in some weeks. In addition for frozen concentrated orange juice week 148 is missing and oatmeal is missing weeks 1–90 and week 151. More information about the Dominick's database is available at Kilts Marketing Center homepage at the Chicago Booth School of Business.

Charlotte market. Below, we also consider 8 other stores from the other census regions to explore regional variation; however, the Charlotte store is representative in most respects of our findings for other stores throughout the country.

In choosing the categories to analyze, we sought to highlight several characteristics that the model suggests should be important. First, we wanted categories where it would be possible to identify sets of goods that were close substitutes for each other. This leads us to choose relatively simple items with only one or two dominant characteristics. For instance, frozen OJ is all pretty similar, so that the primary choice we had to make was to exclude other juices (e.g. tangerine) that Dominick's includes in its frozen juice category. Hence, by focusing on the 6 top selling brands of frozen concentrate we capture over 80 percent of market share in the category.

In all of what follows we always group UPCs of a single brand for which the prices are perfectly or nearly perfectly correlated. For example, Minute Maid sells four different types of 12 ounce frozen concentrated orange: regular orange juice, country style, pulp free and added calcium. The prices of each type move in lock step, with the cross-correlations between the prices uniformly above 0.98. The pulp free juice was not available at the start of the sample and appears in the store we study about 35 weeks into the sample.[5] While for some purposes one might want to distinguish between regular orange juice and these other varieties, given the near perfect co-movement in prices (and in particular the perfectly co-incident sales) it would be practically impossible to estimate an elasticity of substitution between them since the relative prices do not vary.

Because package sizes need not be the same, we also convert prices to price per ounce to facilitate comparisons; for orange juice, the house brand of frozen OJ is sold in both a 12 ounce and six ounce can, while for the other 4 brands (Minute Maid, Tropicana, Florida Gold, Citrus Hill), the vast majority of sales derive from 12 ounce cans.

For oatmeal we concentrate on the two sizes of regular Quaker Oats (18 ounce and 42 ounce packages). They account for roughly 35% of all hot oatmeal sold. The remainder of the purchases is largely concentrated amongst instant oatmeal. It was our judgment that the substitutability across these types of cereal would be low; we will offer some direct evidence of this below.

[5] Chevalier, Kashyap and Rossi (2001) note that in some categories UPCs are discontinued only to have the same product appear with a new UPC. Hence, splicing series by hand is the only sure way to capture all the same sales of these types of similar items.

For the Charlotte store, we include the 3 top selling national brands of creamy peanut butter: Jif, Peter Pan and Skippy. Each is sold in an 18 ounce jar and collectively these three 18oz products account for roughly a quarter of all peanut butter sales in the Charlotte store, with much of the rest of the category being either much different jar sizes, or being chunky peanut butter or peanut butter blends such as peanut and honey.

Because the model emphasizes the possibility of sales, we also wanted categories that differ with respect to perishability and the degree to which stock-piling is possible. We expect that many consumers would buy frozen OJ every week or two, whereas oatmeal would be purchased much less frequently, especially in the summer and be more likely to be stored. Peanut butter would likely be purchased less often than frozen OJ but would not have purchases as seasonally concentrated as oatmeal.

Representative prices for particular UPCs of frozen OJ and oatmeal are shown in Figures 5 and 6. The acquisition cost for the item is also included as a point of reference. We see three salient facts in these two pictures. First, both of the items show the familiar pattern of regular prices that change occasionally, mixed together with intermittent sales. Second, the frequency of sales varies across the categories, in the expected fashion, with sales being much less common for oatmeal than frozen OJ. This suggests that these two categories will allow us to explore different aspects of the model's predictions. The frequency of the frozen OJ sales will highlight the role of high frequency cross-brand willingness to switch, while oatmeal discounts will only matter if consumers are willing to engage in storage.

Third, DFF seems to be consciously making choices about prices so that prices do not simply mirror a pass-through of acquisition cost changes. While we recognize that the acquisition costs data are imperfect measures of marginal cost, there are times when sales take place with no movements in acquisition costs and other cases where the regular acquisition cost changes and the regular price does not. It seems very unlikely that these patterns are being caused by the problems with the acquisition cost calculations.

Turning to the Charlotte store, prices for 18 ounce Peter Pan creamy peanut butter are shown in Figure 7. The price pattern is similar to the ones from the DFF dataset (and to what we find for the other cities in the IRI sample). Over the six year period, most of the price variation reflects switching between two different regular prices and one sale price. Peanut butter is discounted more often than oatmeal but less often than frozen OJ. Store 250517 sometimes runs sales for two weeks, whereas most Dominick's sales last only a week.[6]

These characteristics of the data have been responsible for several of the debates mentioned in the introduction. In particular, although prices change frequently, "regular" prices change infrequently. These features of the data are difficult to reconcile with either a standard flexible price model or a standard menu cost model. As we discussed above, one solution that the literature has offered is to simply ignore the sale prices, focusing on the regular prices. However, as our model suggests, these sale periods are recurring and appear to reflect some strategy beyond passing through changes in acquisition costs. Thus, we explore below whether focusing only on changes in regular prices is appropriate and consider alternative approaches to price measurement.

To characterize the high frequency price variation apparent in these categories, we require definitions for "sale prices" and "regular prices". In order to provide comparability to the literature, we will consider different definitions of "regular" and "sale" prices. Eichenbaum, Jaimovich, and Rebelo (forthcoming), focus on "reference" prices and departures from "reference" prices. A reference price is defined quite simply as the modal price for an item in a given quarter. We will examine the behavior of reference prices, as well as prices below the reference price, and prices above the reference price. As Eichenbaum, Jaimovich, and Rebelo note, however, while reference prices may provide a reasonable measure of the "regular" price, a reference price methodology does not necessarily cleanly identify sale prices. For example, if the regular price is reduced toward the end of the quarter, so that the new regular price is not the modal price for the quarter, we would not want to characterize the new price as a sale price. Thus, while we will examine the behavior of quantities purchased when prices are at their reference price or below their reference price, we do not use reference price departures as our primary measure of sales.

Kehoe and Midrigan (2010), report a different sales identification methodology, which we also calculate. Kehoe and Midrigan propose measuring a sale as a price cut which is reversed within five weeks. We adopt a similar definition. However, we note that the data contain very small apparent price changes; there are cases where the price in a week appears to be less than a cent or two lower than the price in the previous week. As in most scanner datasets, the price series is actually constructed by dividing total revenues

[6]For Dominick's, the weekly sales circular and sales price cycle coincides with the data collection week. This is not necessarily true at all, or even any, of the stores in the IRI dataset.

by total unit sales. There may be product scanning input errors, situations in which a consumer uses a cents off coupon, situations with store coupons, etc, where tiny shifts in measured prices occur but do not reflect real changes in posted prices. We thus set a tolerance for a price change—requiring the price change to be "large" enough to be considered either a sale price or a change in the regular price. We set this tolerance at 2 cents per item.

4 Results

The simplest prediction of our model is that sales should not be coincident for branded close substitutes. Specifically, our model shows that price discrimination between brand-loyal and non-loyal consumers can be exploited by holding sales on only one branded product. Figures 8 to 10 show the price series' for the top two selling items in each of the categories. Simple eyeball tests suggest that the sales are staggered. We investigate this issue more systematically below.

4.1 The Staggering of Sales

We use a simple methodology to examine the extent to which sales are staggered. In Table 1, we compare the observed frequency of simultaneous sales to that which would be expected if the sales of the individual UPCs were randomly timed. For example, for the three peanut butters that we study, at any point in time there can be between zero and three of the UPCs on sale. Using our data, we calculate the share of weeks for which no product is on sale, a total of one of the products is on sale, a total of two of the products are on sale and all three of the products are on sale. Using the unconditional probability of a sale for each of the three products, we compute the predicted probability that 0, 1, 2, and 3 products would be on sale if the sale/no sale decision was independent across products. We can compare the predicted coincidence of sales to the actual coincidence of sales in the data.

We present these results for the three peanut butter products and for the two Quaker oatmeal products. For frozen OJ, we present a comparison of the three branded products and these same three plus the two DFF house brands; we exclude the Citrus Hill juice that dropped out of the sample halfway through. For frozen OJ and peanut butter the probability of exactly one item in the category being on sale is higher than would be predicted if the sale probabilities were independent. Conversely, the probabilities of more than one item being on sale are lower than would be predicted under the assumption of independent sales.

For oatmeal, unlike the other two categories, both products are on sale together three times during the 291 weeks, which is about as frequently as would be expected under the independence hypothesis. This is the first of many indications that oatmeal is very different from our other two categories. That is precisely why we selected it for study. Oatmeal is purchased infrequently and is purchased seasonally. It is also essentially a monopoly category-there are no brands with any significant market share competing with Quaker in our data.

We make two observations from the pattern of oatmeal price discounts. First, the retailer does not appear to be actively price discriminating among consumers who are loyal to one size of oatmeal over the other. Second, for the 34 sales we record over 291 weeks, only one takes place during the months from April to August. Thus, price discounts for oatmeal are concentrated during the period when demand is highest. While the retailer does not appear to be price discriminating across individuals loyal to the two sizes, the retailer may be engaged in discriminating between consumers who can wait for sales and consumers who cannot.

A comparison of the three-product orange juice results and the five-product orange juice results is interesting. The five-product results include the two unbranded products. While sales are generally staggered, it is clear that overlapping sales are more common in the five-product universe than in the three-product universe. Indeed, of the cases where 2 or more of the juices are on sale, 52% have the Dominick's private label12 ounce being on sale, and another 24% involve the private label 6 ounce product. Only in 44% of the cases of multiple items on sale do we see the duplication due to a coincidence of discounts of branded items.[7] Thus, sales are much less likely to be coincident among the branded goods than for branded and unbranded goods. This pattern accords with our expectations. The marketing literature generally finds that branded goods and the unbranded goods are not seen as close substitutes by many consumers (Blattberg and Wisniewski (1989)).[8] Furthermore, since the "regular" price of the unbranded goods is low, consumers may not need the inducement of a sale on the unbranded

[7]These percentages need not sum to 1 because we can have cases where the two private label items are on sale together, etc.

[8]See also Gicheva, Hastings, and Villas-Boas (2007) for some evidence that the willingness to substitute towards unbranded goods may be time-varying. Although they find the propensity to look for sale prices may vary even more then the willingness to switch to store brands.

goods to get the bargain-hunters to purchase. We provide additional information on unbranded goods in our analysis of peanut butter pricing in the penultimate section of the paper.

We conclude that the retailers strategically time sales. We have found that, in each of our three product categories, for the leading products that we examine, there is at least one sale in 12% of the weeks for oatmeal, 50% of the weeks for peanut butter, and 88% of the weeks for frozen orange juice (73% of the weeks for national brand frozen orange juice). These observations hint at a possibility that we will explore below; an alert consumer who is not loyal to a particular brand and who is willing to store goods for a short period of time need almost never pay the "regular" or reference price. Next, we explore the extent to which consumers are and are not, in fact, paying the "regular" prices.

4.2 Purchase Responses to Sales

Having concluded that sales are strategically timed by the retailer, we next explore their implications for consumers. Specifically, we examine the effect of sales on prices paid by consumers and quantities purchased. Tables 2 and 3 show, for each of our UPCs, data on percent of weeks during the sample and the percent of total units sold with discounted prices. Panel A defines the sales using our variant of the Kehoe–Midrigan algorithm. Panel B repeats the calculations when we use the EJR algorithm to select regular prices and define sales prices to be those prices which are below the reference price. It is unsurprising that quantity sold increases substantially when a product experiences a price reduction. However, the combination of frequent, staggered discounts along with consumers who readily switch brands and time purchases means that a substantial fraction of all of the units sold are sold at prices below the "regular" price.[9] For the two DFF categories, we find that the share of total ounces purchased on sale is approximately three times the number of product-weeks during which sales were held. Consumers at the North Carolina Chain 35 store preliminarily

[9]The large percentages of transactions that take place at sales prices are not surprising. Kehoe and Midrigan mention this observation as one of their observations about the Dominick's data. Bronnenberg, Kruger and Mala (2008)'s IRI data set covers 30 categories of goods at over 100 grocery chains in 50 different geographical markets. Their Table 2 shows the fraction of products that are sold on any deal and the mean percentage is 36.8%; more than 30% are sold on deal in 25 of the 30 categories they study. Griffith et al. (2009) also find that about 29.5 percent of total food expenditure from a large sample of British households is on sale items. Hence, the findings for our three categories are very typical of what happens in grocery stores.

seem less responsive to sales; for the peanut butter category we find that the share of total ounces purchased on sale is approximately twice the number of product-weeks during which sales were held.

But there is some interesting heterogeneity in the fraction of weeks on sale and the responses. First, the six ounce store brand of orange juice goes on sale infrequently relative to the other juices—about ten percent of total weeks, yet it still garners nearly seventy percent of its sales in those weeks. Loosely speaking it appears that shoppers mostly prefer the larger package sizes and but are quite willing to shift towards the smaller house brand when it is being discounted.[10]

Second, we chose to examine oatmeal because it is more conducive to storage than frozen OJ. Increased durability should lead the retailer to optimally reduce the frequency of sales; we observe that sales are much less frequent for oatmeal than for frozen OJ. We also see substantial hoarding in sales of the 18 ounce package, with quantities purchased more than tripling when it is put on discount. The purchase response by buyers of the large oatmeal packages to a discount is less pronounced. This might be due to the difficulty of storing the packages; notice also that the average price per ounce of the large package is about 25 percent cheaper than the smaller packages. From Figure 9 we can see that early in the sample there are some cases where even when the smaller package is on sale, it is not cheaper on a per ounce basis than the larger package; see Griffith et al. (2009) for additional more direct evidence of the heterogeneity in households willingness to buy in bulk to save money. The percentage declines in the price when the large package is discounted are not as great as for discounts on the smaller package.

The peanut butter data illustrate that the pattern of sales and regular prices can differ significantly across brands. Discounts are more frequent for Peter Pan (22 percent of weeks) versus Jif (14.4 percent of weeks). Unsurprisingly, then, a large fraction of the units of Peter Pan peanut butter are sold on sale. This raises an important fact; ignoring sales in examining pricing behavior can distort inferences not only across categories, but even across brands within the category.

Finally, notice that using our variant of the Kehoe and Midrigan algorithm, we find fewer sales, in general, than when using the Eichenbaum,

[10] The 6oz juice discounts are concentrated in the latter half of the sample. Dominick's appeared not to use the 6oz product as a promotional vehicle during the first half. Around this time, Dominick's was very enthusiastic about running "fifty cent" sales promotions in which all of the items featured prominently in their sales circular were on special for fifty cents. The sale price of the 6oz juice appears to fit in with the fifty cent sale strategy.

Jaimovich, and Rebelo reference price methodology. This stems mechanically from the fact that the methodology in EJR is not designed to measure "regular" price changes within the quarter. For example, if a price is reduced two weeks before the end of the quarter, and the new price continues throughout the next quarter, our methodology codes that change as a shift in the regular price, not a sale. For EJR, that same price sequence is coded as a price that deviates from the reference price for two weeks, and equals a new reference price at the beginning of the next quarter.

4.3 Actual Prices Paid

To quantify the impact of strategic shopping on prices paid and facilitate comparisons of results with previous studies, we convert the weekly data into quarterly data. This allows us to compare the weighted average prices paid to the reference prices identified by EJR. We calculate the "price paid" by looking over the quarter and using transactions as the weighting mechanism, rather than time. That is, we calculate the average price for every unit purchased during the quarter.

The results for frozen OJ are shown in Figure 11. The departure between reference prices and actual prices paid is striking. In virtually every quarter the actual price paid by consumers is lower than the reference price of any of the individual product. This difference is present even during 1991 when the average (reference) price per ounce of the branded products is nearly identical, so that the gap between the effective price paid and the reference prices requires active timing of purchases.

Figure 12 shows reference prices and actual prices paid for Oatmeal. In this case the effective price paid usually lies between the prices for the two goods. This is not surprising given the information in Figure 9 and Table 4. We know from Figure 9 that the per-ounce price of large package is almost always much cheaper than the smaller package, so a shopper that cared only about price would almost always buy the larger package. Nonetheless, purchases of the smaller package account a substantial portion of total category sales and purchases more than triple when discounts are offered. So it appears that shoppers are timing their purchases to exploit sales but do not view price as the sole consideration in making the purchase decision.

Figure 13 shows the reference prices and actual prices paid for Peanut Butter. As for frozen OJ, the actual price is consistently lower than the reference price. In this case, we see that between 2003 and 2005, the reference prices for Skippy and Peter Pan are constant, while the reference price for Jif is always between the two. Nonetheless, the average price paid of over this period changes noticeably. Obviously this is only possible because of the strategic behavior of the shoppers.

Figures 11 to 13 suggest that prices paid only weakly relate to reference prices. The second column of Table 4 formalizes this by showing the R^2 from a regression of the quarterly average price paid measure on each of the individual item reference prices and a constant. Unsurprisingly, given the figures, the R-squared values range from 0.17 for oatmeal to 0.78 for frozen OJ; keep in mind that for frozen OJ there are 31 quarters and 7 explanatory variables.

The third column of Table 4 presents a regression of prices paid on each of the constituent item list prices. By construction, the average price paid for a product category is a *time-varying* weighted average of list prices. Loosely speaking, the R^2 of this regression is higher the closer the prices paid are to the fixed weight benchmark. If the market shares of the constituent products were time invariant, the R^2 would be 1. The more volatile the week to week market shares of the constituent products, the lower the explanatory power that this regression will have. The difference between 1 and the R^2 for the specification using list prices and 1 shows the importance of the changing weights that arise from active substitution by shoppers. For peanut butter and oatmeal the R^2 is around 60 percent, while it rises to 87 percent for frozen OJ (where again the number of degrees of freedom relative to the number of right hand side variables is low). Overall, we interpret Table 4 as saying that both inter-temporal substitution and inter-brand substitution is important.

4.4 Best Prices

Yet another way to gauge the importance of bargain hunting is to compare the actual price paid to the best possible price that could be obtained by a consumer who is willing to undertake storage and views all brands as perfect substitutes. To compute the "best" price, we consider a hypothetical shopper in a product category who concentrates all purchases over a certain interval into whichever good is cheapest. The best price reflects the limiting case in our model, the case where every shopper is a bargain hunter and has no deterioration in demand from waiting across periods to buy.

We are forced to make an assumption about the horizon over which bargain hunters can be expected to hunt for sales and stockpile. We use the information about how frequently sales are held to infer this. Empirically, we know that some discount occurs in roughly 93% of the weeks for frozen OJ, about half of the weeks for peanut butter and about 12% of the weeks for oatmeal. Thus, we posit that the purchasing horizon is shortest for frozen OJ and longest for oatmeal. For frozen OJ we assume that optimization occurs over 3 week intervals and that the shoppers have perfect foresight about coming sales. So, for each week we compute the "best price" as equal to the lowest price (per ounce) in the category that week, the week before and the week after. Thus, we construct a best price over a three week window for every week. The average best price for the quarter is the average of the 13 weekly best prices. For oatmeal the storability leads us to allow for a 7 week window so that the hypothetical shopper is scanning three weeks forwards and backwards at each week. We use a 5 week window for peanut butter, reflecting the fact that sales for peanut butter are more common than oatmeal but not as frequent as for orange juice. For each category, for the horizon chosen, a consumer can almost always find a sale if s/he is willing to search weekly over the horizon.

We compare the effective price paid to the "best price" series as well as the two fixed weight price series discussed above. One is the fixed weight average of the reference price for each constituent UPC and the other is the fixed weight average of the list prices. The fixed weights are all computed based on the constituent product's share of total ounces over the first quarter of our sample. The quarterly prices to which the fixed weights are applied are constructed by equal weighting the weekly prices for the UPCs in the quarter.

Figure 14 shows the resulting series for frozen OJ. The effective price paid tracks the best price remarkably well; the correlation is 0.92. Recall that in the model, if there were a constant fraction of shoppers who were loyal to one brand, then these people's prices paid would track the list prices for that brand. If there were groups loyal to each brand plus bargain hunters, the average price paid of the loyals would equal a standard fixed weight index and the average price paid of the bargain hunters would equal the "best" price. In this case, the inactive shoppers would lead the average price paid to be a relatively constant amount above the best price. This description seems to describe well Figure 14.

Figure 15 shows the analogous data for oatmeal. As we saw in Figure 12, sales are much less common for oatmeal than the other two categories, yet the price paid still closely tracks the best price; the correlation is 0.95. This tracking indicates that shoppers are timing their purchases to take considerable advantage of the sales when they do occur.

Figure 16 shows the four series for creamy peanut butter. Once again the price paid mirrors the best price; the correlation is 0.81. In this case the gap between the average list price and the price paid is lower than in the other two categories. This is because there are a substantial number of consumers who choose Jif despite the fact that Peter Pan is almost always cheaper. This accounts for the earlier finding that fraction of units bought on sale is only twice the fraction of weeks where sales occur in the peanut butter category in Chain 35. In the two categories at Dominick's, we find that purchases are more responsiveness to discounts.

Recall from Table 1 that sales happen about 88% of the weeks for frozen OJ, about half the time for creamy peanut butter, and only about 12% of the time for oatmeal. Yet, price paid is strongly correlated with best price for all three categories. For frozen OJ and creamy peanut butter, the tracking indicates either a willingness to actively switch across brands every week or to hoard substantially every few weeks when the preferred brand comes up for a sale. For oatmeal the only explanation for the tight association is inter-temporal storage, whereby people bulk up purchases during the sale periods.

4.5 Demand Variability

Our model posits that the high percentages of purchases on sale reflect strategic cross-product and cross-time substitution by bargain hunters. If so, then the quantities sold for an individual product and the quantities sold in a given week will fluctuate as consumers substitute across time and across products. However, this does not imply that demand is volatile, nor that consumption is volatile. Thus, we expect that the total quantity purchased within the product category will be more stable than the purchases of individual UPCs. Furthermore, if consumption is unaffected by sales and consumers are merely stockpiling, then the surge in quantities purchased associated with a price discount will largely be "borrowed" from adjacent weeks. We explore these predictions in Table 5 and Table 6.

If the demand for the goods in the category were independent, then the variance of the sum of the whole category's ounces sold would equal the sum of the variances of the individual products. If the demand for goods in the category were instead hit by significant common demand shocks, then the variance of the sum of the category's ounces would exceed the sum of the variances of the individual products. Finally, if there is a negative correlation in demand across the products due to price discounts and cross-product substitution, the variance of total sales will be lower than for the individual items. Of course, the data in the table only show the net of these effects—it is possible that there are both positively correlated demand shocks and negative correlation induced by substitution for the same set of goods.

Table 5 shows the weekly variance of ounces sold for each UPC. It also shows the weekly variance of total ounces sold in the 2-good oatmeal group, the 6-good orange juice group, and the 3-good peanut butter group. Table 5 shows that frozen OJ sales from week to week are about half as volatile as the sum of the variance of the individual UPCs. Similarly, the total weekly creamy peanut butter is about 20 percent less volatile than predicted based on the fluctuations of the three underlying brands. In contrast, for oatmeal there is no difference between the variance of weekly sales and the sum of the variances of the two different sized packages.

The results suggest that, for orange juice and peanut butter, negative correlation in the purchases of individual products (presumably due to sales) is important and is not offset by large correlated demand shocks. For oatmeal, this appears not to be true. This is unsurprising for two reasons. First, sales are infrequent in oatmeal—the number of weeks in which there is a surge in demand for one product and a decline in demand for the other due to the sale is small relative to the total number of weeks. Second, there is a substantial offsetting factor that works in the opposite direction. Demand for oatmeal is seasonal; there are substantially more sales—of both package sizes—in the winter than in the summer. That means that the two products face positively correlated demand shocks; this would tend to lead the variance of the sum of all product sales to be greater than the sum of the variances.

Next, we examine the issue of inter-temporal substitution. Table 6 shows a very simple regression. We regress total sales in the product grouping on two indicator variables—a dummy that equals one if any of the items in the category is on sale this week and a dummy that equals one if any of the items in the category was on sale last week. As controls, we include quarterly dummies (to account for seasonality) and a time trend to allow for any secular shifts. If consumers are forward looking and timing purchases to exploit deals then we should see sales surge during the weeks of sales and drop off the next week. The effect should be weaker for frozen OJ since there is a sale in over 80 percent of the weeks, but the pattern should be easy to detect in the other two categories where sales are less prevalent.

The results are as expected. Frozen OJ ounces sold are very similar across the four quarters of the year, and sales are elevated in all weeks where a sale is offered in a given week. If, as occasionally happens, there is nothing is on sale, then purchases are lower, but there is no measurable drop off in demand if there was a sale in the previous week. Note that, for orange juice, there are only three incidents in the entire 376 week sample in which two weeks in a row go by with no sale taking place for any of our juices.

Oatmeal purchase patterns are very different. First, there is a very pronounced seasonal cycle with amounts bought being much higher in the winter quarters than in the warmer quarters. Second, when a sale occurs, quantities purchased increase by even more than for frozen OJ. But, in contrast to frozen OJ there is a measurable "payback" the week following the sale, whereby ounces sold the subsequent week drop by about a third of the amount of the jump associated with the sale. This suggests important hoarding by shoppers.

The creamy peanut butter spending shows an intermediate pattern between these two polar cases. Peanut butter purchases increase when a sale is offered on any of the three major brands, but the response is less than for frozen OJ or oatmeal. Note that this is consistent with our earlier findings that the total share of orange juice and oatmeal sold on discount was roughly three times the share of weeks in which discounting occurred; for peanut butter at Chain 35, this ratio was roughly two. For peanut butter, we see that, in the week following the sale, there is a payback as sales retreat by about one quarter of the amount of the surge.

We conclude that the variability in quantity sold from week to week or for an individual UPC paints a misleading picture of the inherent volatility of demand. Once we account for retailer-induced substitution across time and brands, a significant fraction of the variance disappears, as predicted by the model. This provides an important caveat for anyone engaged in the exercise of calibrating a pricing model using data for each UPC in isolation. In particular, by examining each UPC in isolation, one would be tempted to conclude that a product is experiencing significant variation in demand for which the retailer is not responding by changing prices. Using these data to calibrate a menu cost model, one would diagnose significant menu costs in order to fit the observation of significant demand fluctuations unanswered by a price response.

4.6 Time-varying impact of sales

Our model and simulation in Figure 4 suggest that retailers may adjust the frequency of sales in response to shifts in the share of consumers who are brand loyal or even in response to cost shocks. Testing this hypothesis directly is challenging in our setting because we do not have direct observation on most of the key driving factors in the model, namely reservation prices, the shares of bargain hunters, or the deterioration in bargain hunter demand. So we pursue a more modest goal. We examine quarterly data for our products and show that, across quarters with the same regular or reference price, there is substantial variation in the effective price paid.

We start with peanut butter, for which we have no acquisition cost data. Note that, for peanut butter, we have two long periods with no changes in the regular price of any of our products. Specifically, the average reference price is constant for the first five quarters of the sample, and then is constant again (at a lower level) for four quarters (11 to 14). During the first five quarters, the average quarterly price paid ranges from 4.5% below the reference price to 17.6% below the reference price. During the second long stretch, the quarterly price paid ranges from 4.3% below the reference price to 12.5% below the reference price. Thus, while the reference price is constant, the effective price paid varies substantially.

For orange juice, we see fairly frequent regular price changes. Nonetheless, we find that the effective price paid varies considerably relatively to the reference price. The mean discount of the price paid relative to the reference price is 22.0%, with a standard deviation of 6.0%. Interestingly, there isn't much evidence that the sales are being used to smooth out cost shocks—based on the limited cost data that we have available. The estimated margins using the reference price averages 72.1% across quarters, with a standard deviation of 10.6%. The actual margins earned in the category average 34.0%, with a standard deviation of 10.3%.

For oatmeal, prices remain constant for a long period of time. There are eight consecutive quarters with a constant reference price. However, the quarterly discount of the price paid relative to the reference price ranges from −3.0% to 20.3%. Again, this volatility in the price paid largely stems from the retailer adjusting margins—margins are more volatile using price paid than using reference prices.

These large margin shifts may be surprising. However, even our model and our analysis heretofore understates the strategic scope for margin shifting by profit maximizing retailers. In some sense, even the categories that we have identified are too narrow to incorporate all of the demand spillovers over which the retailer may be optimizing. In Chevalier, Kashyap, and Rossi (2003), we provide evidence that retailers may be strategically shifting margins between seasonal categories of products and non-seasonal categories of products. Consistent with that, in this paper, we noted that oatmeal prices are only discounted in the winter.

Our analysis is preliminary and limited by the data availability; however, it is clear from our analysis that sales are an important strategic tool for the retailer. We think it is irresponsible to rule out a priori the hypothesis that the frequency and depth of sales can play a role in monetary transmission.

5 Model Fit and Price Measurement Proposal

The recent debate in macroeconomics has sometimes been framed as a debate between a "keeping sales in" versus "taking sales out" approach to price series construction. We have shown that even episodic sales for a few items can reduce the average price paid in a product category in a calendar quarter by twenty to thirty percent. Thus, we suggest that sales are important. Furthermore, our model and empirical results highlight the importance of the multiproduct nature of the retailer's price-setting decision and its inter-temporal nature. We would argue that, for a multiproduct retailer like a grocery store, the price series for a single UPC is simply not an object that should be of any interest to macroeconomists. Like the retailer, the economist has to be concerned with the pricing plan over time and over close substitute products.

When considering close substitute products, we have repeatedly referred to the average price paid or variable weight index as the relevant construct for a product category. For a narrow product category, changes in this index reflect changes in the prices that consumers pay per unit for a fairly homogeneous product. This presents a conundrum, however, because variable weight price index has two important downsides for price index construction. First, one has to be quite cautious about identifying an appropriate set of close substitutes.[11] If two few substitutes are included then the category will exhibit spurious volatility when the omitted good goes on sale. In this case, the average price paid will also be too high since it will miss the substitution into the discounted product. Alternatively if too many items are thrown into a category, the degree of substitution will appear low and average prices paid could become less informative. For example, grouping orange and apple juices together would likely lead to a much noisier prices paid series than if both categories were properly modeled.

Secondly, and crucially, however, the high frequency quantity data necessary to construct a variable weight price index is frequently unavailable. In contrast, construction of neither the "reference" price series nor the fixed weight average price series requires high frequency data on quantities. (At some point, some quantity data has to be collected to construct weights, but that data collection can be infrequent.) Thus, for example, the Bureau of Labor Statistics (BLS) employs the fixed expenditure share geometric

[11] See Nevo and Hatzitaskos (2005) for additional discussion of this issue.

mean index within item strata. An important feature of this index is that it requires no quantity data.

Fortunately, our model and empirical results suggest an alternative methodology. Note that our model, in Equation (7) provides a very tractable formulation for estimating a price series without high frequency data on quantities. In (7), we show that the variable weight index or price paid is approximated by a weighted average of the fixed weight price index and the "best price" over the k-period sales cycle. Note that the best price is simply the lowest price (per ounce) of any product in the product category over some number of weeks. Finding the "best price" does not require quantity data.

Thus, an approximation for the price paid (variable weight index) can be constructed that takes a fixed weight average of the "best price" and the normal fixed weight price index. All that remains then is to find the weights. The weights are the share of "loyals" in the marketplace and the share of non-loyals in the marketplace.[12] Of course, one would like to have some kind of data to estimate these shares and ultimately, to estimate those shares one needs some kind of quantity data. We posit that assuming any share of loyals between 0 and 1 will produce a price series that is more informative than a price series constructed just using the regular price (which actually implies a share of loyals *greater* than one). However, even with limited quantity data one could produce weights for the best price and the fixed weight index. One might be willing to assume that the loyal share is the same across product categories within a store, or the same across stores within a city. More research needs to be done on this point. For now, we conclude by approximating the weights for the fixed price index and the best price index in our sample, where we actually have quantity data and thus the variable weight price index.

For each of our three product categories, we conduct a simple regression of the variable weight price index ("price paid") on the fixed weight index and the "best price" series with no constant. Note that if (7) is a good approximation for our data, the weights on the two price series should add up to one and the constant will be zero. Note also that these conditions are not hard-wired to hold. Consumer preferences and therefore market shares could drift away from the fixed weights of the fixed weight indices; some demand model could hold where best price is not particularly relevant.

[12] Note that, in our model, we assumed that the share of the loyals attached to each of the two brands is equal. That assumption is not necessary and plays no role in the analysis of this section.

Results are shown in Table 7 for each of the three product groups. Table 7 also shows the p value for a test of the hypothesis that the coefficients for the fixed weight index and best price sum to one and also shows the p value of a test of the joint hypothesis that the constant equals zero and the best price coefficients sum to one. In all cases, the estimated value of the constant term is not significantly different from zero at standard confidence levels. We also cannot reject that the coefficients for best price and average list price sum to one. (For orange juice, however, we do reject the joint hypothesis of both the constant equalling zero and the coefficients summing to one). The coefficient estimates are themselves informative. For oatmeal and frozen OJ, the data prefer approximately a 60% weight on the best price and a 40% weight on the fixed weight index. For the peanut butter at Chain 35, the data choose approximately equal weights on the best price index and the fixed weight index. For that product and store, the data suggest that more consumers are brand-loyal. This is consistent with our findings above that the impact of sales is smaller for peanut butter at Chain 35 than for either of the products at DFF.

In order to explore the this more fully, we also examine data on peanut butter for stores in eight additional cities. We select the city randomly, choosing one from each of the eight Census Regions. In each city, we examine the largest chain. We discard chains where regional brands dominate Skippy, Jif, and/or Peter Pan in sales in order to provide results maximally comparable to the results for Chain 35 in Charlotte. We repeat the analysis in Table 7 for the 18oz peanut butters in each of these cities. The results, shown in Table 8 broadly support the hypothesis that best prices are important. For all eight of the cities (in addition to Charlotte), we cannot reject that the sum of the coefficients for average list price and for best price sum to one. The coefficient for best price range from a low of 0.132 in Knoxville to 1.07 in New York. In New York, puzzlingly, the coefficient for average list price is actually negative, although insignificantly different from zero. The hypothesis that the constant term equals zero cannot be rejected at the 5 percent level in any of the 8 cities. However, the joint hypothesis, that the price coefficients sum to one and that the constant is zero is rejected at the 5 percent level in New York and Houston (and Charlotte, as previously noted). These results suggest that our model fits reasonable well across a variety of settings, but also reveals interesting variation across cities to be explored. Differences across cities may or may not, for example, be correlated across different categories of products. Rich scanner datasets such as the IRI dataset will allow expansive exploration of this issue.

These findings may bear on the large literature debating how best to construct consumer price indices. Griffith et al. (2009) have a helpful description of the many ways in which substitution (across brands or over time) can thwart the construction of cost of living measures. Readers familiar with government price index construction methodologies may be interested in the question of how our methodology compares to the BLS's methodology. The BLS constructs a fixed expenditure share geometric mean index within item strata. This methodology does allow for a limited amount of cross-item substitution. However, this substitution differs substantially from what we propose here. The BLS methodology effectively assumes a cross-price elasticity of demand of -1 between items within the strata (with a strata corresponding to a category like peanut butter, oatmeal, or frozen orange juice). Our analysis thus far has focused on more aggressive item substitution for a subset of consumers. Our analysis also falls outside usual cross price elasticity frameworks in that our measures emphasize ordinal rather than cardinal price relationships.

We evaluate the relationship between our methodology and the BLS's methodology in Table 9. Specifically, in Table 9, we demonstrate the results of the following experiment: we consider whether, if a BLS-type geometric mean index is used in place of the fixed weight index, does the Best Price still have significant explanatory power for the average price paid? That is, we regress average price paid on the geometric index and our best price measure. In all cases, the best price measure still has significant explanatory power for the average price paid. Indeed, the coefficients for best price when using the geometric weight index in the regression are nearly identical to the coefficients for best price when using the fixed weight index in the regression. We obtain similar results for peanut butter in the eight supplemental IRI cities that we examined in Table 8; regression results using the geometric mean index are very similar to those show in Table 8 using the fixed weight index.

We note that our analysis potentially has implications for the measurement of inflation. As discussed above, the ratio of any list price measure to the best price is by no means constant, even at the quarterly frequency. Thus, using list price measures in calculating inflation misses time or location varying discounts that are economically important in magnitude.

6 Conclusion

We provide a simple model of consumer heterogeneity and show how that heterogeneity motivates temporary price discounts by retailers. The simulated price path for this model looks remarkably like the empirical path of prices observed for many retailers. Margins that vary dramatically over time—even when consumer preferences are stable—are a natural outcome of our model. Because some consumers are strategic in seeking out sales, the share of all goods in our sample that are sold at sale prices is two to three times as large as the number of product-weeks in which sales occur. For the goods we study, the effective price paid falls far below the posted prices and indeed, closely tracks the "best price"—a price which is calculated simply as the minimum price for any good in the set of close substitutes over a short interval.

We further show that measuring quantities sold for a single UPC will erroneously lead to the impression that underlying demand is volatile. Indeed, much of that volatility derives from the deliberate price-setting behavior of the retailer for other products in the category. We note that the extent to which the price paid differs from the regular price varies substantially from quarter to quarter—focusing on the "regular" price masks considerable differences across quarters. Lastly, we show that our model can be exploited as a structural model of prices paid. We show that, even without high frequency quantity data, a variable weight index can be approximated using the "best price" concept.

Clearly, the importance of strategic consumer responses to temporary sales is of paramount importance in some sectors, and of more limited importance in others. However, as Varian notes in his 1999 Handbook of Industrial Organization survey of price discrimination, sellers almost always *want* to engage in price discrimination and price discrimination schemes involve substantial computational costs. Both the consolidation of the retailing sector over the last decades and the rapid decline in IT costs suggest that data-driven price discrimination schemes are likely to become more, rather than less important in the future. Thus, if macroeconomists are to successfully model price-setting, confronting price discrimination appears to be an inevitable challenge.

References

Aguiar, Mark and Erik Hurst, “Life-Cycle Prices and Production,” *The American Economic Review*, December 2007, 97(5), pp. 1533–1559.

Bils, Mark and Peter J. Klenow, “Some Evidence on the Importance of Sticky Prices”, *Journal of Political Economy*, October 2004, 112(5), pp. 947–985.

Blattberg, Robert C. and Kenneth J. Wisniewski, “Price-induced patterns of competition”, *Marketing Science*, Autumn 1989, 8(4), pp. 291–309.

Bronnenberg, Bart. J., Michael W. Kruger and Carl F. Mela, “The IRI marketing data set,” *Marketing Science*, July–August 2008, 27(4), pp. 745–748.

Chevalier, Judith A., Anil K. Kashyap and Peter Rossi, “Why Don’t Prices Rise During Periods of Peak Demand? Evidence from Scanner Data”, *American Economic Review*, March 2003, 93(1), pp. 15–37.

Dixit, Avinash K. and Joseph E. Stiglitz, “Monopolistic Competition and Optimum Product Diversity,” *American Economic Review*, 1977, 67(3), pp. 297–308.

Eichenbaum, Martin, Nir Jaimovich, and Sergio T. Rebelo, “Reference Prices and Nominal Rigidites,” forthcoming, *American Economic Review*.

Gicheva, Dora, Justine Hastings, and Sofia B. Villas-Boas, “Revisiting the Income Effect: Gasoline Prices and Grocery Purchases,” 2007, NBER Working Paper 13614.

Rachel Griffith, Andrew Leicester, Aviv Nevo, and Ephraim Leibtag “Consumer Shopping Behavior: How Much Do Consumers Save?” *Journal of Economics Perspectives* 23(2), Spring 2009, pp. 99–120.

Daniel Hosken and David Reiffen, “Pricing Behavior of Multiproduct Retailers,” *The B.E. Journal of Theoretical Economics*, 2007, Vol. 7: Iss. 1 (Topics), Article 39. Available at: `http://www.bepress.com/bejte/vol7/iss1/art39`

Kehoe, Patrick J. and Virgiliu Midrigan, "Sales and the Real Effects of Monetary Policy", 2008, Working Papers 652, Federal Reserve Bank of Minneapolis.

Klenow, Peter J. and Oleksiy Kryvtsov, "State Dependent of Time-Dependent Pricing: Does it Matter for Recent U.S. Inflation?" *The Quarterly Journal of Economics*, August 2008, 123(3), pp. 863–904.

Lal, Rajiv and Carmen Matutes. "Retail Pricing and Advertising Strategies," *Journal of Business*, July 1994, 67(3), pp. 345–70.

Mackowiak, Bartosz Adam and Frank Smets, "On Implications of Micro Price Data for Macro Models", 2008, CEPR Discussion Papers 6961.

Nakamura, Emi and Jn Steinsson, "Five Facts about Prices: A Reevaluation of Menu Cost Models," *Quarterly Journal of Economics*, November 2008, 123(4), pp. 1415–1464, November.

Nevo, Avi and Konstantinos Hatzitaskos, "Why Does the Average Price of Tuna Fall During Lent?", NBER Working Paper W11572, August 2005.

Pesendorfer, Martin, "Retail Sales: A Study of Pricing Behavior in Supermarkets", *Journal of Business*, January 2002, 75(1), pp. 33–66.

Sobel, Joel, "The Timing of Sales", *Review of Economic Studies*, 1984, L1, pp. 353–368.

Varian, Hal R. "A Model of Sales," *American Economic Review*, 1980, 70(4), pp. 651–9.

Varian, Hal R. "Price Discrimination", in R. Schmalensee & R. Willig (ed.), *Handbook of Industrial Organization*, edition 1, volume 1, 1999, chapter 10, pp. 597–654 Elsevier.

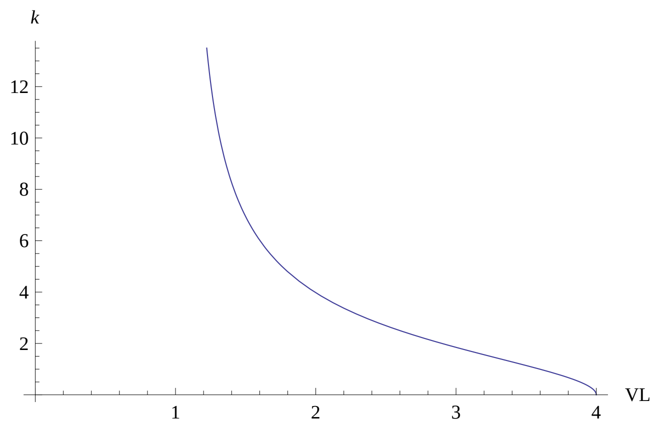

Figure 1: The time between sales as a function of the valuation of bargain hunters. Note: $V^H = 4$, $\alpha = 0.4$, $\rho = 0.9$, $c = 1$

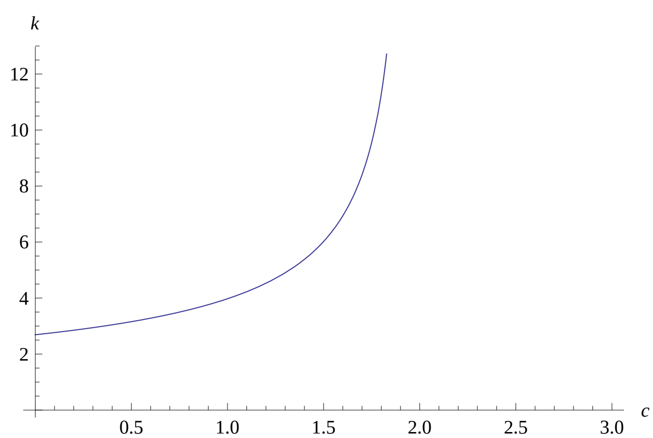

Figure 2: The time between sales as a function of marginal cost. Note: $V^H = 4$, $V^L = 2$,$\alpha = 0.4$, $\rho = 0.9$

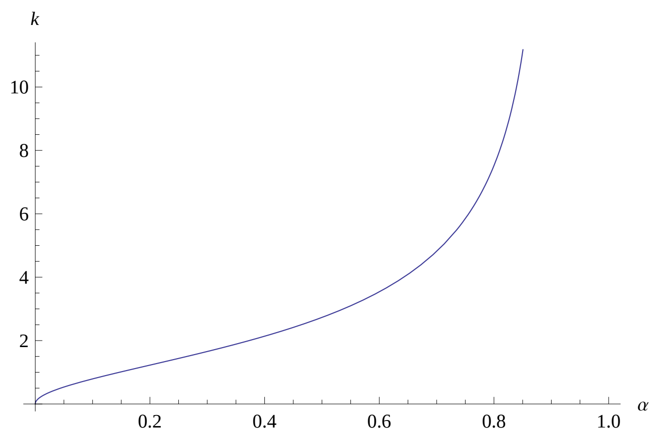

Figure 3: The time between sales as a function of the share of loyals (non-bargain hunters). Note: $V^H = 4$, $V^L = 2$, $\rho = 0.9$, $c = 1$

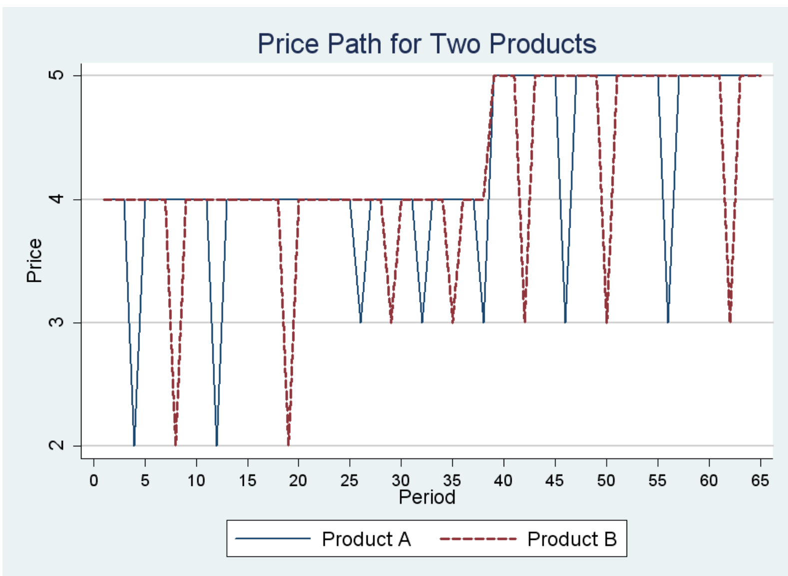

Figure 4: Simulated Price from Quarterly Changes in Demand and Cost. Note: Simulation shows initial values of $V^H = 4$, $V^L = 2$, $\rho = 0.9$, $c = 1$, $\alpha = 0.4$. Then, demonstrates the effect of demand and cost shocks. Week 13: Cost shock from 1 to 1.5. Week 26: V^L increases from 2 to 3. Week 39: V^H increases from 4 to 5. Week 52: α increases from 0.4 to 0.6.

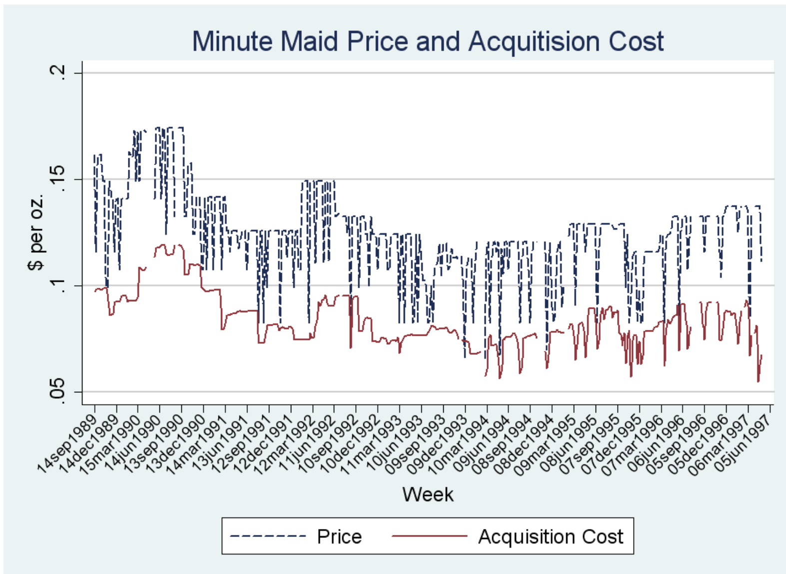

Figure 5: Minute Maid 12 ounce Frozen OJ Price and Acquisition Cost

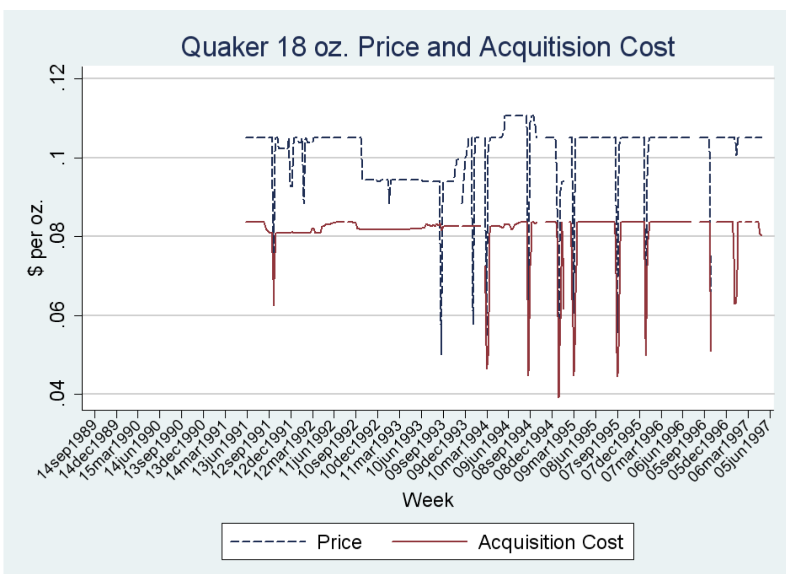

Figure 6: Quaker 18 ounce Oatmeal Price and Acquisition Cost

Figure 7: Peter Pan 18 ounce jar of Creamy Peanut Butter

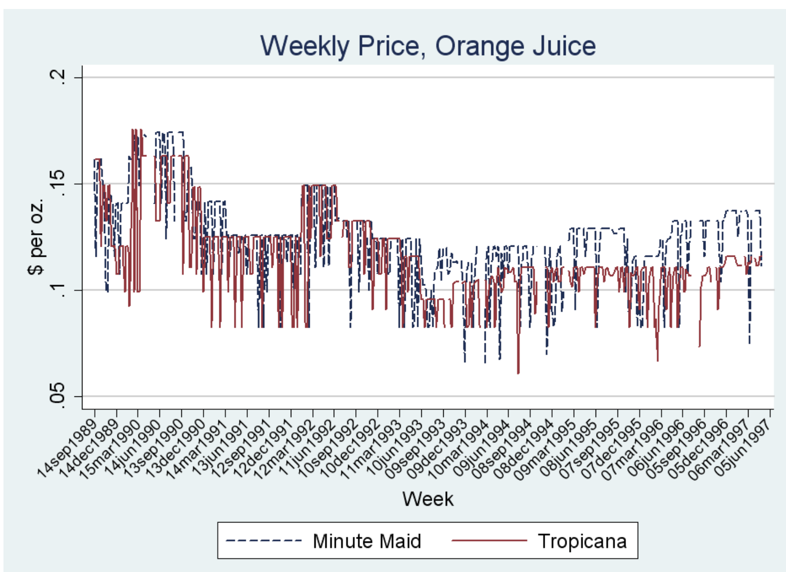

Figure 8: Prices for Two Top Selling Frozen Orange Juices

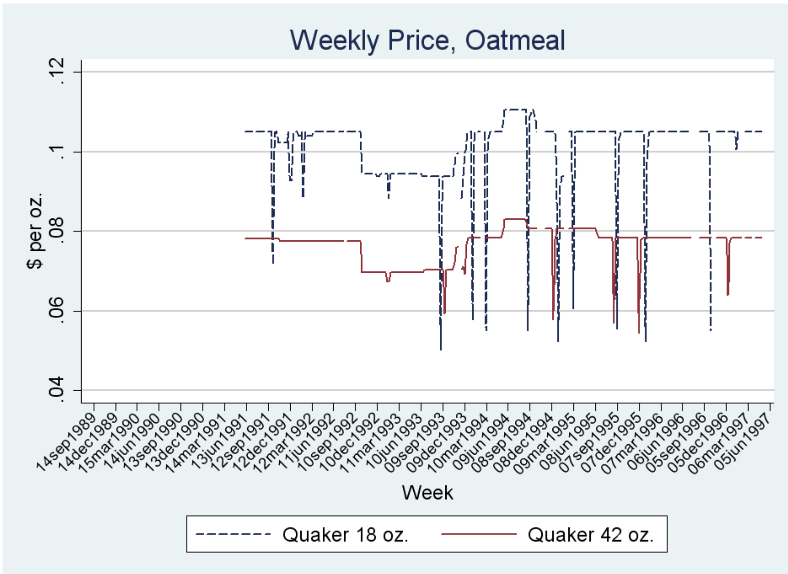

Figure 9: Prices for Two Top Selling Oatmeals

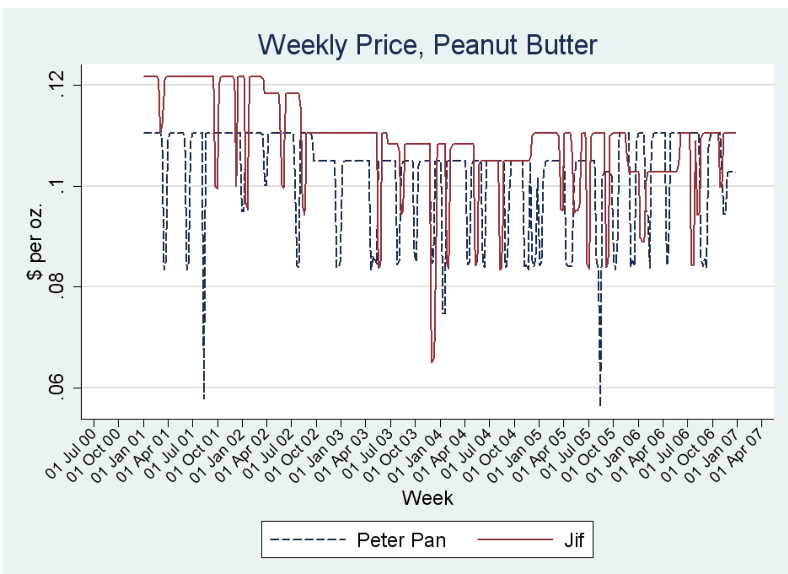

Figure 10: Prices for Two Top Peanut Butters

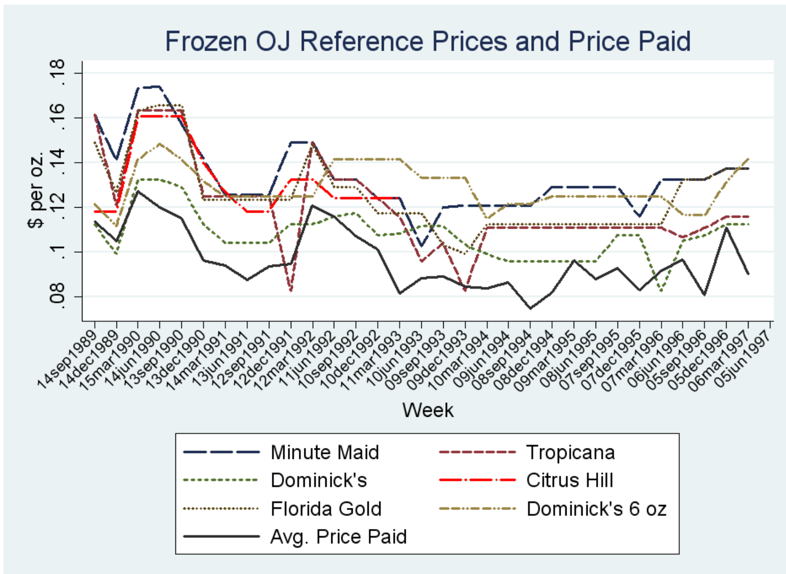

Figure 11: Quarterly Frozen OJ Reference Prices and Actual Prices Paid

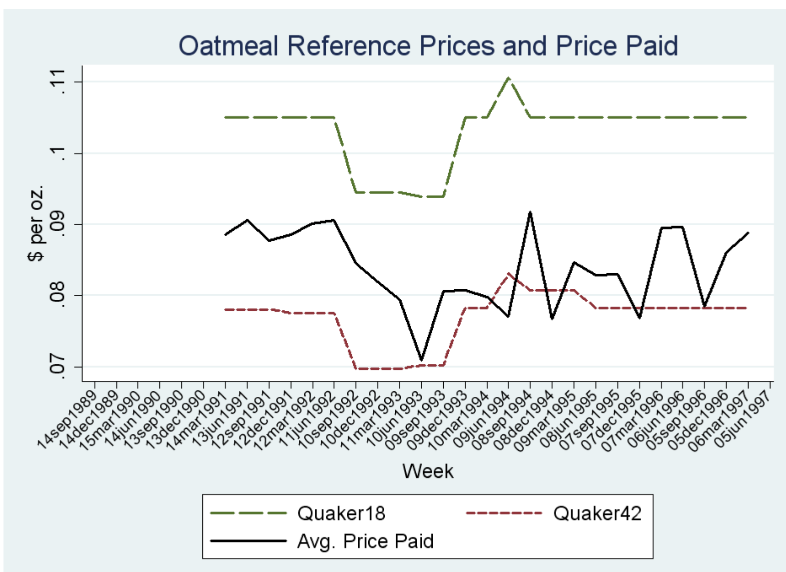

Figure 12: Quarterly Oatmeal Reference Prices and Actual Prices Paid

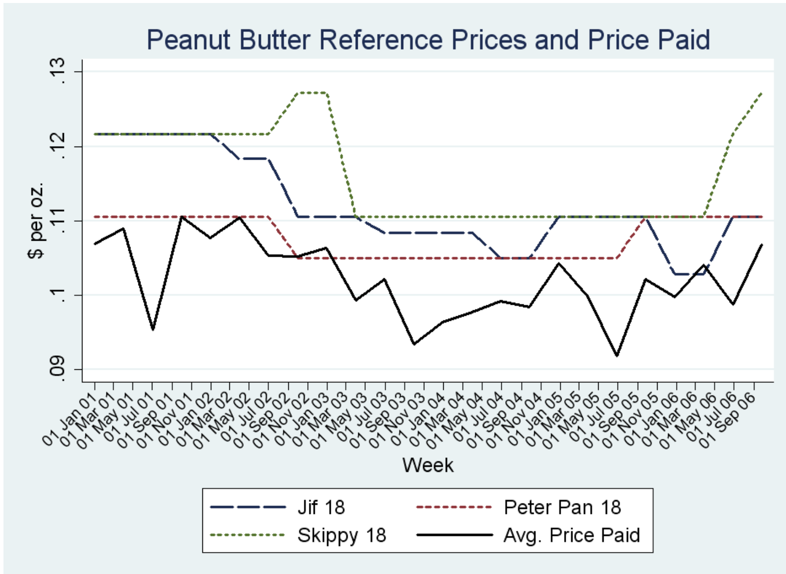

Figure 13: Quarterly Creamy Peanut Butter Reference Prices and Actual Prices Paid

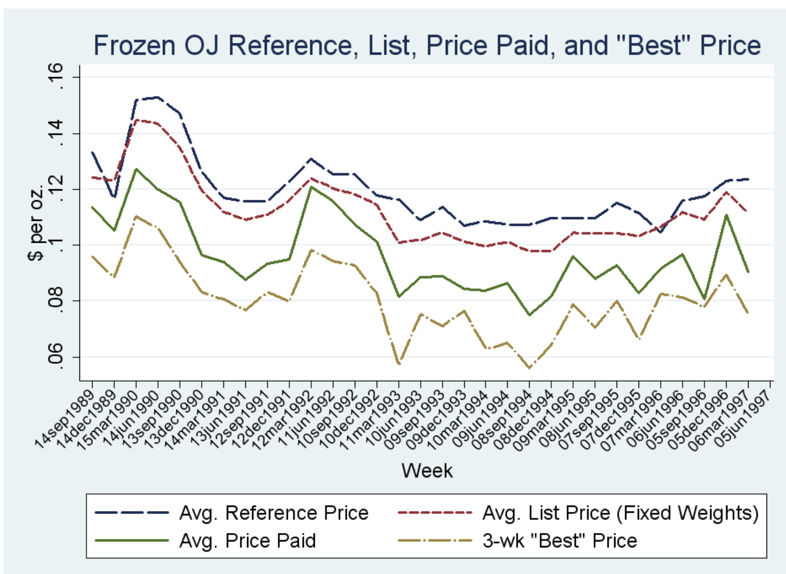

Figure 14: Quarterly Frozen OJ Reference Prices, Average List Prices, Effective Prices Paid, and Best Prices.

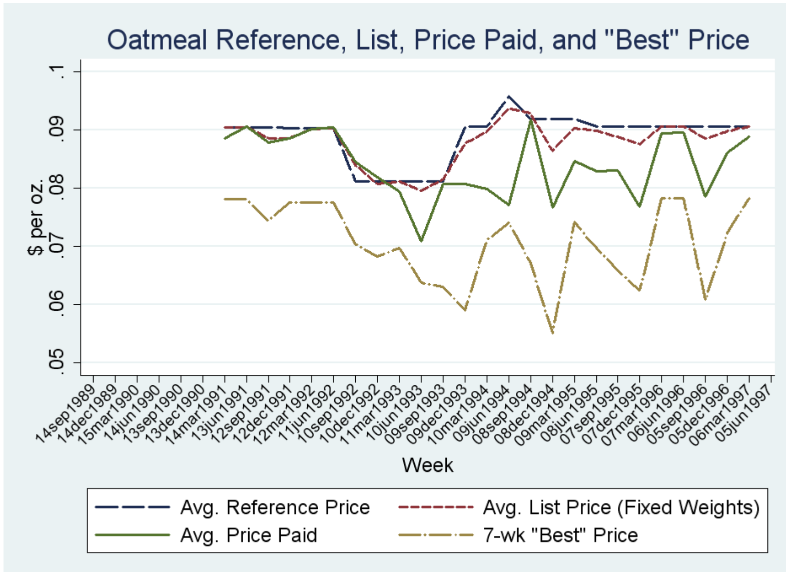

Figure 15: Quarterly Oatmeal Reference Prices, Average List Prices, Actual Price Paid, and “Best” Prices

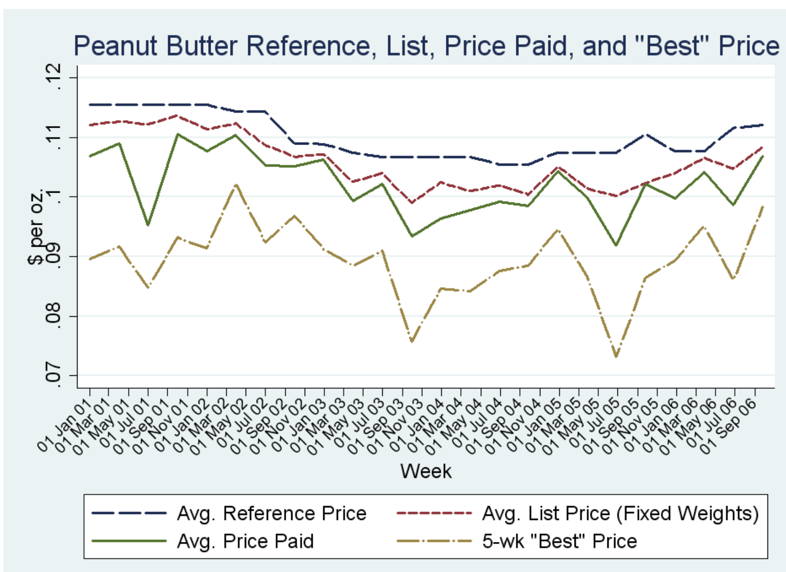

Figure 16: Quarterly Peanut Butter Reference Prices, Average List Price, Actual Prices Paid and “Best” Prices

Table 1: Actual and Expected Distribution of Sales

# of UPCs on sale	Expected Percentage	Actual Percentage
All Frozen OJ		
0	21.66%	12.50%
1	39.58%	52.13%
2	27.90%	29.52%
3	9.34%	5.59%
4	1.44%	0.00%
5	0.08%	0.27%
Note: Contains the 5 frozen OJ products that we examine that are in the data sample for the entire time period. Unconditional probabilities of a sale are Minute Maid 12 ounce = 0.3005, Tropicana 12 ounce = 0.3227, Florida Gold 12 ounce = 0.2553, Dominick's (Heritage House) 12 ounce = 0.3200, Dominick's (Heritage House) 6 ounce = 0.0970		
Branded Orange Juice Products		
0	35.28%	28.99%
1	44.06%	55.32%
2	18.18%	14.63%
3	2.48%	1.06%
Note: Contains the same UPCs as above, excluding the Dominick's private label brand.		
Oatmeal Products		
0	87.64%	88.32%
1	12.01%	10.65%
2	0.35%	1.03%
Note: Unconditional probabilities of a sale for 18 ounce is 0.086 and for 42 ounce is 0.041		
Creamy Peanut Butter, 18 oz.		
0	54.57%	49.52%
1	36.79%	46.65%
2	8.07%	3.83%
3	0.58%	0.00%
Note: Unconditional probabilities of a sale are Jif= 0.144, Peter Pan= 0.224, Skippy= 0.179		

Table 2: Ounces Sold and Weeks at Various Prices

	Ounces sold		Weeks	
Frozen OJ	Regular price	Sale price	Regular price	Sale price
Minute Maid 12	29.69%	70.31%	69.95%	30.05%
Tropicana 12	22.06%	77.94%	67.82%	32.18%
Dominick's 12	29.65%	70.35%	68.09%	31.91%
Citrus Hill 12	24.43%	75.57%	79.66%	20.34%
Florida Gold 12	17.98%	82.02%	74.47%	25.53%
Dominick's 6	35.35%	64.65%	90.43%	9.57%
TOTAL	26.50%	73.50%	74.62%	25.38%
Oatmeal				
	Regular price	Sale price	Regular price	Sale price
Quaker 18	69.65%	30.35%	91.41%	8.59%
Quaker 42	91.21%	8.79%	95.88%	4.12%
TOTAL	80.18%	19.82%	93.64%	6.36%
Peanut butter				
	Regular price	Sale price	Regular price	Sale price
Jif 18	71.01%	28.99%	85.62%	14.38%
Peter Pan 18	60.51%	39.49%	77.64%	22.36%
Skippy 18	55.89%	44.11%	82.08%	17.92%
TOTAL	65.01%	34.99%	81.78%	18.22%

Table 3: Share Sold at Various Prices

	Ounces Sold			Weeks		
	Ref Price	Sale Price	Above Ref Price	Ref Price	Sale Price	Above Ref Price
Frozen OJ						
Minute Maid 12	24.66%	72.72%	2.62%	56.12%	39.36%	4.52%
Tropicana 12	26.22%	70.45%	3.33%	56.38%	36.70%	6.91%
Dominick's 12	25.88%	70.42%	3.71%	54.26%	37.23%	8.51%
Citrus Hill 12	21.67%	76.27%	2.05%	69.49%	23.73%	6.78%
Florida Gold 12	15.06%	83.46%	1.48%	60.90%	32.18%	6.91%
Dominick's 6	28.35%	68.54%	3.11%	77.39%	15.16%	7.45%
TOTAL	24.32%	72.70%	2.98%	61.74%	31.40%	6.85%
Oatmeal						
Quaker 18	61.41%	35.60%	2.99%	82.13%	13.06%	4.81%
Quaker 42	86.34%	9.35%	4.31%	90.72%	4.81%	4.47%
TOTAL	73.58%	22.78%	3.63%	86.43%	8.93%	4.64%
Peanut Butter						
Jif 18oz	64.11%	33.95%	1.93%	77.96%	20.13%	1.92%
Peter Pan 18 oz	56.84%	41.82%	1.34%	73.48%	24.92%	1.60%
Skippy 18 oz	46.97%	51.91%	1.11%	76.22%	21.50%	2.28%
TOTAL	59.42%	38.99%	1.59%	75.88%	22.19%	1.93%

Reference Price is exactly at EJR Reference Price; sales defined as in EJR

Table 4: Explanatory Power of List Prices and Reference Prices for Prices Actually Paid

	R^2 for Reference Price Regression [number of obs] {number of regressors}	R^2 for List Price Regression [number of obs] {number of regressors}
Frozen OJ	0.75 [31] {7}	0.87 [31] {7}
Oatmeal	0.17 [25] {2}	0.60 [25] {2}
Creamy Peanut Butter	0.45 [24] {3}	0.59 [24] {3}

Note: The table reports the R^2 from a regression of quarterly price paid on quarterly reference prices and a constant in column 2 and quarterly prices paid on quarterly average list prices and a constant in column 3. Because the Citrus Hill brand drops out part way through the sample, the frozen OJ regressions also include a dummy variable for the weeks where that brand is missing and the Citrus Hill prices are set to 0 for those weeks. The intercept is not counted in the number of regressors reported above.

Table 5: Volatility of Ounces Sold

	Variance Over the Sample	
Frozen OJ		
Minute Maid 12	10,846,714	
Tropicana 12	18,696,622	
Dominick's 12	20,655,582	
Citrus Hill 12	5,730,370	
Florida Gold 12	5,070,080	
Dominick's 6	2,173,754	
Sum of UPC Variances	63,173,120	
Weekly Category Variance	33,835,204	
Note: Category Variance/Sum of UPC Variances		53.6%
Oatmeal		
Quaker 18	5,345,582	
Quaker 42	650,746	
Sum of UPC Variances	5,996,331	
Weekly Category Variance	6,032,957	
Note: Category Variance/Sum of UPC Variances		100.6%
Creamy Peanut Butter		
Jif	134,239	
Peter Pan	299,379	
Skippy	17,936	
Sum of UPC Variances	451,555	
Weekly Category Variance	369,632	
Note: Category Variance/Sum of UPC Variances		81.9%

Table 6: Regressions of quantity sold on sale, sale last week, time trend, and seasonal dummies.

Dependent variable: log(weekly total ounces sold)	Frozen Orange Juice	Oatmeal	Peanut Butter
Dummy: Any item on sale	0.600	0.708	0.397
	(0.097)	(0.066)	(0.031)
Dummy: Any item on sale	-0.017	-0.276	-0.097
last week	(0.097)	(0.066)	(0.031)
Quarter 1	9.051	8.245	7.071
	(0.143)	(0.055)	(0.038)
Quarter 2	8.829	7.919	7.065
	(0.144)	(0.052)	(0.038)
Quarter 3	8.855	7.896	7.183
	(0.147)	(0.049)	(0.038)
Quarter 4	8.794	8.294	7.080
	(0.145)	(0.052)	(0.041)
Time Trend	-0.002	-0.0004	-0.0001
	(0.0002)	(0.0002)	(0.0002)
Num Obs	375	290	312
R-squared	0.997	0.998	0.999

Notes: Standard errors in parentheses.

Table 7: Regressions of Price Paid on the fixed weight index (Avg List Price), "Best price", and a constant.

	Frozen Orange Juice	Oatmeal	Peanut Butter (Charlotte)
Avg List Price	0.360	0.404	0.450
	(0.16)	(0.23)	(0.13)
Best Price	0.666	0.482	0.504
	(0.15)	(0.12)	(0.09)
Constant	0.002	0.014	0.009
	(0.009)	(0.018)	(0.011)
R-squared	0.900	0.596	0.826
p value for test that weights sum to 1	0.79	0.57	0.67
p value for test that weights sum to 1 AND constant = 0	0.70	0.15	0.03

Notes: Standard errors in parentheses.

Table 8: Structural Estimates of Orange Juice Price Paid, 9 cities

	Constant	Average List Price	Best Price	p-value that sum of coefficients=1	p-value of joint test: coefficients sum to 1 and constant =0
Charlotte	0.009	0.45	0.504	0.67	0.03
	(0.011)	(0.132)	(0.094)		
Los Angeles	-0.014	0.942	0.163	0.28	0.20
	(0.015)	(0.131)	(0.065)		
West Texas	-0.011	0.376	0.808	0.48	0.31
	(0.039)	(0.433)	(0.314)		
Saint Louis	-0.017	0.556	0.603	0.29	0.53
	(0.015)	(0.168)	(0.199)		
Chicago	-0.014	0.623	0.517	0.47	0.26
	(0.025)	(0.179)	(0.109)		
New York	0.022	-0.115	1.07	0.86	0.03
	(0.033)	(0.355)	(0.192)		
Hartford	0.004	0.003	1.023	0.95	0.49
	(0.053)	(0.417)	(0.194)		
Houston	0.025	0.352	0.478	0.12	0.00
	(0.012)	(0.129)	(0.051)		
Knoxville	0.001	0.864	0.132	0.92	0.63
	(0.004)	(0.034)	(0.028)		

Notes: Standard errors in parentheses.

Table 9: Regressions of Price Paid on the geometric index price, Best price, and a constant.

	Frozen Orange Juice	Oatmeal	Peanut Butter
Geo Mean Price	0.390	0.429	0.454
	(0.163)	(0.225)	(0.128)
Best Price	0.634	0.468	0.495
	(0.154)	(0.124)	(0.094)
Constant	0.002	0.013	0.010
	(0.009)	(0.017)	(0.011)
R-squared	0.902	0.603	0.831

Notes: Standard errors in parentheses.