

NBER WORKING PAPER SERIES

BEST PRICES

Judith A. Chevalier  
Anil K Kashyap

Working Paper 16680  
<http://www.nber.org/papers/w16680>

NATIONAL BUREAU OF ECONOMIC RESEARCH  
1050 Massachusetts Avenue  
Cambridge, MA 02138  
January 2011

The views expressed here are our own and not necessarily those of any institutions with which we are affiliated, nor of the National Bureau of Economic Research. 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. We thank numerous seminar participants for helpful comments. We are grateful to the Kilts Marketing Center and SymphonyIRI Group for the data. As a condition of use, SymphonyIRI reviews all papers using their data to check that the data are not described in a misleading fashion. However, all analyses in this paper based on SymphonyIRI Group, Inc. data are the work of Chevalier and Kashyap, not SymphonyIRI Group, Inc.

NBER working papers are circulated for discussion and comment purposes. They have not been peer-reviewed or been subject to the review by the NBER Board of Directors that accompanies official NBER publications.

© 2011 by Judith A. Chevalier and Anil K Kashyap. All rights reserved. Short sections of text, not to exceed two paragraphs, may be quoted without explicit permission provided that full credit, including © notice, is given to the source.

Best Prices

Judith A. Chevalier and Anil K Kashyap

NBER Working Paper No. 16680

January 2011, Revised March 2011, Revised August 2012

JEL No. E3,E31,L11,L16

**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.

Judith A. Chevalier

Yale School of Management

135 Prospect Street

New Haven, CT 06520

and NBER

judith.chevalier@yale.edu

Anil K Kashyap

Booth School of Business

The University of Chicago

5807 S. Woodlawn Avenue

Chicago, IL 60637

and NBER

anil.kashyap@chicagobooth.edu

# 1 Introduction

We explore theoretically and empirically the importance of retailer price discrimination strategies for descriptions of price dynamics. We depart from the usual approach in the macro literature that ignores consumer heterogeneity in the price-setting problem and instead make this heterogeneity central to the analysis. In particular, we posit that 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.

The model that we propose to account for this behavior is a model which does not necessarily speak to nominal price rigidities. But the model does suggest that conscious price discrimination rather than restrictions on the ability of firms to change prices (such as menu costs) is critical for understanding price dynamics. In the context of our framework, the most natural way to understand nominal rigidities would be to look to the consumer side of the pricing problem rather than the firm side. But even without taking a stand on this issue, the evidence we develop does speak to a pair of debates that have emerged amongst macroeconomists looking at micro data.

For example, there is a measurement question of whether (and potentially how) intermittent price discounts, or “sales”, should be incorporated in 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. These measurement issues come to the fore when macroeconomists must decide how to calibrate models that presume that prices are sticky. Our empirical analysis argues strongly for keeping track of sale prices because they play a major role in governing stores’ profits and our model suggests that their depth and/or frequency

could be altered to adjust to shocks.

Likewise, a largely theoretical literature has emerged to look at the micro-foundations for why sales exist. For example, Matejka (2010) builds on the literature of rational inattention (see Sims (1998, 2003), Mankiw and Reis(2002,2006), for example) to explore its implications for price setting. In his model firms do not know the exact cost of the good each is selling and the firms find it expensive to keep track of all the information that would be necessary to perfectly deduce this cost. Thus, firms must decide which cost signals to monitor and because of their finite processing ability will set prices based on the estimated distribution of cost. In order to economize on information, the seller chooses signals that lead to a small number of distinct prices. Sales in this framework corresponding to situations in which a signal suggests that costs were low.

There is a complementary long literature, beginning with Varian (1980) and Sobel (1984) that rationalize price-discrimination strategies that include temporary discounts. Guimaraes and Sheedy (2011) were the first to bring this perspective to the macroeconomic literature on price-setting. Like us, they assume heterogeneous consumers. They have two types of consumers, bargain hunters and loyal shoppers, who are price insensitive; they explore how price discrimination plays out when retailers compete to sell to these customers. The key finding in their analysis is that competition bounds the importance of sales as a mechanism to adjust to shocks.

Interestingly, the theoretical results in Guimaraes and Sheedy (2011) contrast the empirical results of Klenow and Willis (2007), who find that sale prices are just as responsive to macroeconomic conditions as regular prices. That is, Klenow and Willis find that the depth of sales decreases when recent inflation has been high and that sales prices give way to higher regular prices when recent inflation has been high.

Our contribution is largely empirical, but to guide the analysis we begin by introducing a model in the spirit of Varian (1980), Sobel (1984), and Pendorfer (2002) that allows for the possibility that pricing strategies might involve variation in the usage of discount prices. In contrast to the previous literature, we focus explicitly on a retailer controlling the prices of multiple products. The robust implication of this kind of model is that store pricing patterns ought to reflect the presence of different consumers and be strategically coordinated across products. The model can account for frequent temporary sales and long periods of constant regular prices. Furthermore, even with unchanging costs and demand, optimal markups are not held constant across items or across time. In contrast to Guimaraes and Sheedy, our model also implies that changing the frequency and depth of sales is the

optimal response to cost shocks and to certain kinds of demand shocks.

To quantify the importance of consumer heterogeneity we turn to detailed microeconomic data. Some of our data are for particular supermarket products collected over parts of seven years at Dominick’s Finer Foods (DFF), a supermarket chain in the Chicago area. We also analyze data for a later seven-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 upon more careful scrutiny, it appears that purchase patterns associated with the sales clearly reflect the important role played by bargain hunters. In particular, we show three sets of results.

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, actual prices paid are substantially lower than regular prices, and even measurably below average posted prices.

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 price index. The data match the structural form of our model.

Third, the data exhibit 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 led 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. For example, 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 suggesting that while customers actively shift between items when they go on sale, making individual demand quite volatile, the total demand across all the relevant items is much more stable.

We read the evidence as suggesting that the models of price setting that emphasize consumers taste, cognitive processing, or information-gathering deserve more attention. For instance, in our model, occasional sales provide a tool for retailers to effectively charge different consumers different prices and pricing decisions are driven by the different consumer groups' reservation prices. Thus, how consumers update reservation prices for individual goods becomes a critical factor affecting inflation. This mechanism has been much less explored than have those which emphasize firms costs of changing prices.

Similarly, price index construction may need to be revised to account for the importance of bargain hunters. We provide some preliminary thoughts on how this might be done.

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 establishes a number of new facts about pricing and purchase patterns that are consistent with the model. 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 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. We discuss the competitive case more below.

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.<sup>1</sup>

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.<sup>2</sup>

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.<sup>3</sup> 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  $1-\rho$ . Thus, for example, if they made a purchase in period  $t - 1$ ,

---

<sup>1</sup>In our model,  $V^H$  and  $V^L$  are real. However, it is possible that one or both types of consumer have nominal illusion. This would influence the retailer’s optimal response to changes in monetary policy. We leave exploration of this set of issues to future work.

<sup>2</sup>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.

<sup>3</sup>We will see that in equilibrium it will not be profit maximizing to put both goods on sale in the same week.

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  $t, t + 1, \dots, t + (k - 1)$ , total demand from the bargain hunters in period  $k$  will equal:

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

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.<sup>4</sup>

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 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  $V^H$ .

---

<sup>4</sup>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.



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  $N$ th (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  $k$ th 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  $k$ th period. Assuming no discounting, total profits for this retailer over all  $N$  periods are:

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

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  $\rho$  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:

$$\begin{aligned} & \frac{\alpha}{2} (V^H - V^L) < (1 - \alpha) \frac{1 - \rho^N}{1 - \rho} (V^L - c) \Leftrightarrow \\ & V^H < V^L + 2 \frac{1 - \alpha}{\alpha} \frac{1 - \rho^N}{1 - \rho} (V^L - c) \end{aligned} \quad (3)$$

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  $\alpha$ ,  $\rho$ ,  $V^H$ ,  $V^L$ , and  $c$ , the seller will then choose  $k$  to maximize (2). That is, the seller will deterministically hold a sale every  $k$ th 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:

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

Where  $W(\cdot)$  is the Product Log or Lambert function. The Lambert function provides the solution for  $z$  such that  $z = W(z)e^{W(z)}$ . 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. Note that, 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  $V^L \rightarrow V^H = 4$ , 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  $c \rightarrow V^L = 2$ ,  $k \rightarrow \infty$ .

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  $\alpha$ , 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  $\alpha$  ranges from 0 to 1. As  $\alpha$  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 (2011). 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 compare this model to Pesendorfer (2002), the microeconomic model that ours most closely resembles. In Pesendorfer, the sale decision is stochastic, but shocks to cost are an important driver of the decision to hold a sale.

Another important distinction between all of these models and ours is that none of these models explicitly examine a retailer managing a portfolio of close substitute products. Indeed, the fact that cost drivers are important

in these models implies that the prices for close substitute products would tend to be *positively* correlated. In our model, the time series of prices for close substitute products are *negatively* correlated (unless a common shock leads to the change in  $V^L$  and  $V^H$ ).

Guimaraes and Sheedy (2011) share the feature of our model that sales are not driven primarily by cost shocks. Their model, like Sobel (1984), examines two competitors (competing brands or competing retailers) chasing a fixed pool of bargain hunters. They do not examine a single retailer managing a portfolio of close substitute products. In their model, sales will not be negatively correlated across competing brands.

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:  $V^L$ ,  $c$ ,  $\rho$ , and  $\alpha$ . If any of those parameters change, the retailer's optimal response is to alter the frequency or 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  $\rho$ , total units sold equal  $= N\alpha + \frac{N(1 - \rho^k)}{k(1 - \rho)}(1 - \alpha)$ .

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 modelled; 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 actual 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 (2011) or Kehoe and Midrigan (2010) 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. The actual 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 actual 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:

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

That is, the actual 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 loyal and of bargain hunters in the population are the weights.

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

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

The average revenue in (8) approaches that in (7), or equivalently (6), as  $\rho$  approaches 1. For  $\rho < 1$ , 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 actual 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  $\rho$  is small, then as  $\alpha$  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  $\alpha$  is fairly small and if  $\rho$  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  $\alpha$  is very large, then the weighted average price will more closely resemble the “regular” or “reference” price. For small to modest values of  $\rho$  and intermediate values of  $\alpha$ , 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.

Our model is highly stylized and so we discuss ways in which it may be extended. First, 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).

Furthermore, there are numerous ways to contemplate an extension of the model to multiple types. In particular, 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. In a model with more types, the pricing outcomes will be as in (6) but the weights on the different prices will reflect the behavior of the different consumer types.

Also, while there is a long tradition of two-type models in the literature on temporary discounts (including Varian (1980) and Sobel (1984)) extension to a continuum of consumer willingnesses to pay is possible. As in Varian (1980), as long as enough consumers are not adept at moving all of their purchases into sale periods, intertemporal variation in prices will enhance the retailer's ability to price discriminate.

Our model also abstracts from competition among multiple stores. Obviously, a model with perfect competition would limit the opportunity to price discriminate. However, an extension of our model along the lines of Lal and Matutes (1994) should preserve the basic insights of our model. Lal and Matutes (1994) model multiproduct retailers with travel costs between stores. The ability of consumers to shop both stores, at a cost, disciplines the overall rents across products that stores can extract from each consumer. However, the ability to extract different levels of rents from different consumer types and from different products is preserved. Ellison (2005), in a different context, provides an extension of Lal and Matutes in which price discrimination across consumer types emerges in an equilibrium with spatial competition. The easiest way to understand the effects of introducing competition into our model is to view competition as impacting consumer



willingness to pay. Spatial competition opens additional scope for price discrimination as retailers may benefit from price discriminating between consumers who will and won't travel to obtain a price discount.

While our model is stylized, it provides a useful framing for our analysis. Below, we will show empirically that the prices paid in our data resemble a weighted average of a fixed weight index and the "best price", as in (6). 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 characterizing 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

Given our interest in examining both the incidence of sales and consumer responses to them, our project requires a data set that contains both prices and quantities. This eliminates many data sets, most notably those from the Bureau of Labor Statistics that have been the source for many of the most important recent papers on price setting patterns. We instead we rely on two scanner data sets (from Dominick's and IRI Symphony) that include quantity information.

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).<sup>5</sup> 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 Mela (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 the period 2001 to 2007. For our benchmark calculations, we use data from “Chain 35” which has more than 100 stores in the southeastern United States. 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 may or may not have more stores in other regions.

For the IRI dataset, 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 the first 313 weeks.<sup>6</sup>

Below, for robustness, we also examine results for peanut butter for eight other 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 the national brands that dominate in Charlotte. Throughout the paper, we discuss the extent to which the

---

<sup>5</sup>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.

<sup>6</sup>Peter Pan prices are missing for all stores between the week of February 12 2007 and August 13 2007.

Charlotte store is representative of our findings for other stores throughout the country.

In choosing the categories to analyze, we traded off several considerations. The model highlights the importance of accounting for potential substitution among different goods within a category. This leads us to categories where it is possible to identify sets of goods that were close substitutes for each other. Substitutability is easier to determine for 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.

The three simple categories we identified also have the advantage of each featuring a single major cost driver; in section IV, we introduce data on underlying commodity prices for oranges, peanuts, and cereals and measure the extent to which regular prices and the frequency of sales reflect changes in the underlying commodity price.

Focusing on only a few products raises concerns about how the results generalize. In all such studies, there is a tradeoff between comprehensiveness of categories and the care that can be paid to the details of each category. For example, Chevalier, Kashyap and Rossi (2003) document 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. But, the task of splicing the data to capture UPC changes, as well as grouping UPCs within brands as described below, requires a substantial investment in learning about the category and cleaning data. Of course, this could in principle be done on a large scale, but it is very costly. The Bureau of Labor Statistics devotes some portion of its field force to cleaning data to create consistency across categories in construction the CPI (see, for example, Bureau of Labor Statistics, 2011.)

We therefore chose to study three very different categories of goods that are representative of the range of goods found in grocery and drug stores. Bronnenberg et al (2008) show that the propensity of retailers to use sales depends primarily on the perishability and stockpiling ability of a good. Highly perishable items can rarely be stored, but less perishable items might be storable depending mostly on the natural frequency that the good would be used. For example, Bronnenberg et al (2008) note that diapers are not perishable, but the rate at which they are used and the size of an individual diaper means that it is not very easy or practical to stockpile them either.

Conversely, while frozen orange juice is more perishable than diapers, the serving sizes and typical frequency of use makes it an excellent candidate for stockpiling.

Bronnenberg et al (2008) classify categories according to whether they are of high, medium or low regarding perishability and storability. By definition, anything highly perishable cannot be stored, while some items like diapers cannot be stockpiled for other reasons regardless of perishability. This suggests focusing on three types of categories: one like oatmeal where incentives to discount are minimal (due to either perishability or lack of storability), one like peanut butter with moderate incentives to use sales because either the purchase cycle or perishability limits the storability, and one like frozen orange juice which is conducive to frequent discounts because of high stockpilability.

Another important data convention used through out the analysis is that 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. 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.

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. It is important to note that, when we undertake the robustness checks in other cities, we use these brands and sizes although 18 oz is not always the top selling size in those other locations. These cities also vary substantially in the extent to which private label brands are important.<sup>7</sup> While the packages we examine account for roughly a quarter of all peanut butter sales in Charlotte, they account for as little as 15 percent of the peanut butter sales in the other cities.

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 Jif 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 most of the other cities in the IRI sample). 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.<sup>8</sup>

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

---

<sup>7</sup>See Kashyap (2012) for a discussion of the complications that arise in analyzing categories in which the market share of private labels is high.

<sup>8</sup>For Dominick’s, the weekly sales circular and sales price cycle coincides with the data collection week. This is not necessarily true of the stores in the IRI dataset.

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 (2011), 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 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 \$0.002 per ounce.

## 4 Results

One of the simplest predictions of our model is that retailer-driven 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.<sup>9</sup>

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

---

<sup>9</sup>Our finding of a high probability of one and only one sale holds true for peanut butter in the other cities where the three selected brands are the market leaders, but not in the cities where a private label brand has a dominant position in the market.

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 31 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, 50% have the Dominick’s private label 12 ounce being on sale, and another 27% involve the private label 6 ounce product. Only in 45% of the cases of multiple items on sale do we see the duplication due to a coincidence of discounts of branded items.<sup>10</sup> 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)).<sup>11</sup> Furthermore, since the “regular” price of the unbranded goods is low, consumers may not need the inducement of a sale on the unbranded 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

---

<sup>10</sup>These percentages need not sum to 1 because we can have cases where the two private label items are on sale together, etc.

<sup>11</sup>See also Gicheva, Hastings, and Villas-Boas (2010) 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 than the willingness to switch to store brands.



we examine, there is at least one sale in 11% of the weeks for oatmeal, 51% of the weeks for peanut butter in Charlotte, and 88% of the weeks for frozen orange juice (70% of the weeks for national brand frozen orange juice). In seven of the eight comparator cities for peanut butter, there is a sale on at least one of our three products in over half of the weeks. These observations suggest 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. Table 2 defines the sales using our variant of the Kehoe–Midrigan algorithm. Table 3 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.<sup>12</sup> 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. Peanut butter consumers in Charlotte are somewhat 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. Within the peanut butter category, there is some variation across

---

<sup>12</sup>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 Mela (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.

our eight comparator cities in the measured overall responsiveness to sales. The lowest responsiveness is Knoxville, where the share of ounces sold on sale is 23 percent greater than the share of product-weeks with sales. The greatest responsiveness is our data from Hartford, where, over our seven year sample, only 7 percent of product-weeks have sales, but 57 percent of the total ounces sold are purchased at a sale price.

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 sixty five percent of its sales in those weeks. Loosely speaking it appears that shoppers mostly prefer the larger package sizes and are quite willing to shift towards the smaller house brand when it is being discounted.<sup>13</sup>

Second, we chose to examine oatmeal because we expect it to be discounted less frequently than frozen OJ. That is what we find. Interestingly, we do see evidence of 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 contrast also seems reasonable given the difficulty in storing the larger package; notice also that the average price per ounce of the large package is about 25 percent cheaper than the smaller package. 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. For the Charlotte store, 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 algo-

---

<sup>13</sup>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.

rithm, we find fewer sales, in general, than when using the Eichenbaum, 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 Best Prices

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. We calculate two useful measures of pricing. First, we calculate the “price paid” by aggregating purchases over the quarter and using transactions as the weighting mechanism, rather than time. That is, we calculate the average price actually paid by consumers for every ounce purchased during the quarter.

One 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 88% of the weeks for frozen OJ, about half of the weeks for peanut butter in our Charlotte store and about 11% 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 actual 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 11 shows the resulting series for frozen OJ. The actual 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 11.

Figure 12 shows the analogous data for oatmeal. As we saw in Figure 11, 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 13 shows the four series for creamy peanut butter in Charlotte. 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 Charlotte. In the two categories at Dominick’s, we find that purchases are more responsiveness to discounts. Turning to peanut butter in the other cities, our data suggest that the correlation between price paid and best price ranges from a low of 0.63 in Knoxville to a high of 0.93

in Houston and averages 0.82.

Recall from Table 1 that sales happen about 88% of the weeks for frozen OJ, about half the time for creamy peanut butter in Charlotte, and only about 11% 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.

#### 4.4 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 4.

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 4 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 4 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. These results are robust for peanut butter for the eight other cities that we examined (total category sales volatility ranges from 26 percent to 4 percent below what would be predicted based on the volatility of the individual products). 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.

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.

#### **4.5 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 in response to cost shocks. Furthermore, if macroeconomic factors lead to shifts in share of consumers who are “shoppers”, even a fixed

discount strategy will have varying effects over the cycle. Testing these hypotheses 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. In this section, we examine quarterly data for our products, and we examine the extent to which both the regular price and the frequency of sales responds to cost shocks.

It is first useful to note that, while our data show fairly “sticky” regular prices, the actual price paid by consumers can nonetheless vary substantially. For example, for peanut butter in Charlotte, we have two long periods with no changes in the regular price of any of our products. Specifically, the average Kehoe-Midrigan base price is constant for the first four quarters of the sample, and then is constant again (at a lower level) for four quarters (10 to 13). During the first five quarters, the average quarterly price paid ranges from 4.3% below the base price to 18.1% below the base price. During the second long stretch, the quarterly price paid ranges from 4.7% below the regular price to 13.1% below the regular price. Thus, while the regular price is constant, the effective price paid varies substantially.

We further investigate these issues by examining the relationship between the regular price (defined using the Kehoe-Midrigan base price), the frequency of sales, and the world commodity price of the underlying commodity being sold. For peanut butter, we examine the IMF’s world peanut price; for orange juice, we examine the IMF’s world orange price; for oat-meal we examine the IMF’s world cereals price index. The IMF prices are all in metric tons. We convert them to ounces, and for oranges, we adjust for the production relationship between oranges and frozen concentrated orange juice. To give a sense of magnitudes, in our data, the mean price paid by consumers for packaged peanut butter is approximately 11 cents per ounce; the commodity nut price is about 2.4 cents per ounce.

For each product (and in each of the IRI cities), we conduct three very simple regressions to measure the relationship between our observed market data and the underlying commodity price. First, we regress the quarterly average KM base price on the commodity price. Nearly any theory of price adjustment would predict that the regular price of the product would covary positively with the underlying commodity price. We would expect the coefficients in this regression to be between zero and one for orange juice and peanut butter (for oats, since we have a cereals index, rather than the

oat price specifically, the levels are not interpretable).<sup>14</sup>

Next, we measure the share of total weeks in the quarter in which a sale takes place for some product in the category; this is the frequency of sales. We regress the share of weeks in the quarter that have a sale on the commodity price. Our model predicts that the retailer will hold a sale more frequently when costs are lower and less frequently when costs are higher. In our model, adjusting the frequency of sales can be used as a substitute for altering the base price (specifically, in situations in which the cost changes but consumer reservation values do not).

Finally, we regress the price paid by consumers on the commodity price. Again, in our model, the price paid (which incorporates both the base price and sales including consumer responsiveness to sales) should be responsive to the underlying costs. Again, for orange juice and peanut butter, we would expect coefficients between zero and one.

The results are presented in Table 5. The first column of Table 5 shows the regressions of the base or regular price on the world commodity price. Coefficients would, of course, be expected to be positive. They are in nine of the eleven cases. In six of the cases where the relationship has the correct sign, it is statistically different from zero at least at the ten percent confidence level. Coefficient estimates vary significantly. The coefficients for peanut butter in New York and Chicago are “too big”, as they exceed one.

The second column of Table 5 shows the regressions of the share of weeks with a sale in the category regressed against the commodity price. Our hypothesis is that, if sales are used as an adjustment margin, sales should be less frequent when commodity prices are high, leading to a negative coefficient for the commodity price. This is true for nine of the eleven cases. One of the two exceptions is oats; as we have discussed earlier, oatmeal price reductions appear to have a strong seasonal component. The winter price reductions in oatmeal overwhelm any price changes related to commodity price movements. An alternative unreported specification that includes a winter quarter dummy for oats shows the predicted (albeit insignificant) coefficient for the cereals price once controlling for the winter quarter dummy. The results are significant at conventional confidence levels in two of the cases.

The third column of Table 5 shows the regressions of actual price paid by consumers on the commodity price. This price will incorporate changes in

---

<sup>14</sup>Of course, peanuts and oranges are not the only inputs into peanut butter and orange juice— there is packaging, labor, etc. If the prices of those other inputs are systematically positively or negatively correlated with the key commodity input, then the prediction of coefficients between zero and one would not hold.



the base price, changes in discounts, and consumer responsiveness to those discounts. We have argued before that this object really should be the object of interest for macroeconomists rather than the base price in situations in which retailers use sales strategically. Here, all of the specifications have the predicted sign, and in contrast to the regressions using the base price, all of the coefficients for the commodity price are bounded by zero and one. Houston and St. Louis are both interesting cases here; for these cities, the base price appears negatively correlated with the underlying commodity price; using price paid yields results more in line with price adjustment.

This analysis is quite simple and limited by the data availability; however, our analysis suggests that sales may be an important strategic tool for the retailer. Furthermore, incorporating sales to determine the price paid does generate realistic responsiveness to underlying commodity prices in a way that the base price does not. 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. Indeed, our model and evidence conform closely with the finding in Klenow and Willis (2007) that sale prices are as responsive to shocks as regular prices. It also conforms with the finding in Gicheva, Hastings, and Villas-Boas (2010) that consumers become much more responsive to grocery store promotions when gasoline prices are high.

## 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 homo-

geneous 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.<sup>15</sup> If too 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 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.<sup>16</sup> 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

---

<sup>15</sup>See Nevo and Hatzitaskos (2005) for additional discussion of this issue.

<sup>16</sup>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.

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.

Results are shown in Table 6 for each of the three product groups. Table 6 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 the Charlotte store, 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 the Charlotte store than for either of the products at DFF.

We also examine our results for peanut butter in the additional cities. For seven of the eight cities, 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 0.129 in Knoxville to 0.973 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 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 St. Louis, Chicago, 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 8. Specifically, in Table 8, 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 7; regression results using the geometric mean index are very similar to those show in Table 7 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.

A final open issue is how to account for nominal price stickiness. The framework we have pursued ties price determination to consumers reservation values for different goods. These reservation values should be in terms of real prices, not nominal ones. But the data suggests that nominal prices are in fact rigid. There are many reasons why consumers who are routinely making nominal purchases (especially in low and moderate inflation environments) might confuse real and nominal prices. But this topic deserves a much fuller treatment and would naturally bring many considerations about consumer decision-making into the analysis. Given the limited success of pure menu cost models of price setting in explaining the kinds of facts we have documented, and the fact that these theories implicitly have focused on firms difficulties in adjusting prices, this alternative direction seems promising.

## 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.

Bureau of Labor Statistics, *BLS Handbook of Methods*, Chapter 17, pp. 29-31.

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 Rigidities,” 2011, *American Economic Review*.

Ellison, Glenn, “A Model of add-on pricing”, 2005, *The Quarterly Journal of Economics*, 120(2): 585-637.

Gicheva, Dora, Justine Hastings, and Sofia B. Villas-Boas, “Investigating Income Effects in Scanner Data: Do Gasoline Prices Affect Grocery Purchases?” 2010, *American Economic Review* 100(May), pp. 480-484.

Golosov, Mikhail, and Robert E. Lucas Jr., “Menu Costs and Phillips Curves,” 2007, 115(2), pp. 171-199.

Griffith, Rachel, 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.

Guimares, Bernardo, and Kevin D. Sheedy, “Sales and Monetary Policy”  
*American Economic Review* 101(2), pp. 844-76.

Hosken, Daniel 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/vol17/iss1/art39>

Kashyap, Anil K, “Comments on Gagnon, Lopez-Salido, and Vincent”,  
*National Bureau of Economics Research Macroeconomics Annual* 2012, forthcoming.

Kehoe, Patrick J. and Virgiliu Midrigan, “Prices are Sticky After All”,  
2010, Research Department Staff Report Papers 413, 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.

Klenow, Peter J and Jonathan L. Willis, “Sticky Information and Sticky Prices”, *Journal of Monetary Economics*, 2007, Vol 54, 79-99.

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

Mankiw, N. Gregory and Ricardo Reis. “Sticky Information Versus Sticky Prices: A Proposal to Replace the New Keynesian Phillips Curve,” *Quarterly Journal of Economics*, 2002, vol 117(4), 1295- 1328.

Mankiw, N. Gregory and Ricardo Reis. “Inattentive consumers,” *Journal of Monetary Economics*, 2006, vol. 53(8), pp. 1761-1800.

Matejka, Filip, “Rationally Inattentive Seller: Sales and Discrete Pricing”, 2010, CERGE-EI Working paper Series No. 408, SSRN No. 1653308.



Nakamura, Emi and Jon 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.

Sims, Christopher A., “Stickiness”, *Carnegie-Rochester Conference Series on Public Policy* 49(1),1998, pp. 317-356.

Sims, Christopher A., “Implications of Rational Inattention”, *Journal of Monetary Economics* 50(3), 2003, pp. 665-690.

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 loyal (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:  $\alpha$  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: Jif 18 ounce jar of Creamy Peanut Butter, Chain 35




Figure 8: Prices for Two Top Selling Frozen Orange Juices




Figure 9: Prices for Two Top Selling Oatmeals

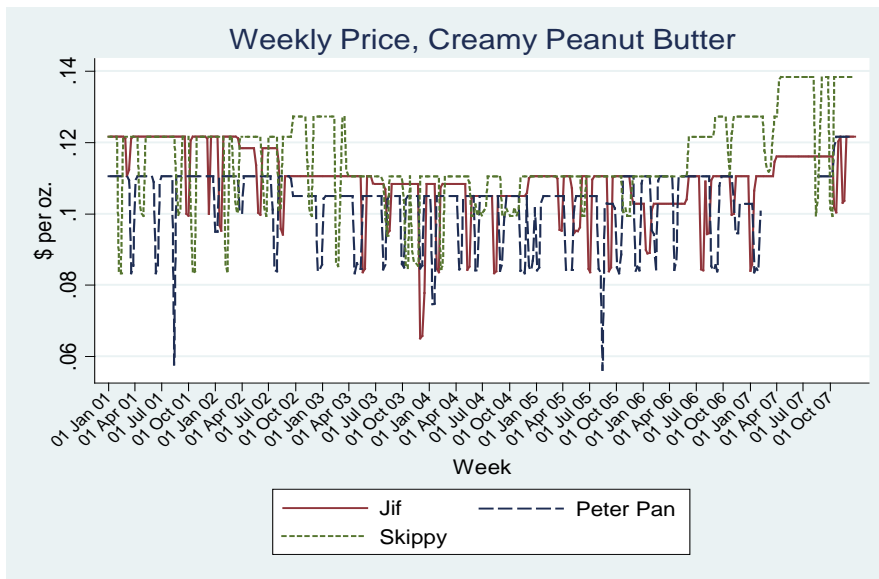


Figure 10: Prices for Three Top Peanut Butters, Charlotte




Figure 11: Quarterly Frozen OJ Reference Prices, Average List Prices, Effective Prices Paid, and Best Prices.




Figure 12: Quarterly Oatmeal Reference Prices, Average List Prices, Actual Price Paid, and "Best" Prices




Figure 13: 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
<i>All Frozen OJ</i>		
0	21.91%	12.50%
1	39.59%	52.66%
2	27.70%	28.72%
3	9.27%	5.85%
4	1.45%	0.00%
5	0.08%	0.27%
<p>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 = 0.295, Tropicana 12 ounce = 0.322, Florida Gold 12 ounce = 0.255, Dominick's (Heritage House) 12 ounce = 0.311, Dominick's (Heritage House) 6 ounce = 0.106</p>		
<i>Branded Orange Juice Products</i>		
0	35.59%	29.52%
1	44.00%	54.79%
2	17.98%	14.63%
3	2.43%	1.06%
<p>Note: Contains the same UPCs as above, excluding the Dominick's private label brand.</p>		
<i>Oatmeal Products</i>		
0	88.61%	89.35%
1	11.09%	9.62%
2	0.30%	1.03%
<p>Note: Unconditional probabilities of a sale for 18 ounce is 0.079 and for 42 ounce is 0.038</p>		
<i>Creamy Peanut Butter, 18 oz.</i>		
0	54.28%	48.80%
1	36.92%	47.29%
2	8.20%	3.92%
3	0.60%	0.00%
<p>Note: Unconditional probabilities of a sale are Jif= 0.151, Peter Pan= 0.223, Skippy= 0.178</p>		

Table 2: Ounces Sold and Weeks at Various Prices

	Ounces sold		Weeks	
	Regular price	Sale price	Regular price	Sale price
<i>Frozen OJ</i>				
Minute Maid 12	31.68%	68.32%	70.48%	29.52%
Tropicana 12	22.06%	77.94%	67.82%	32.18%
Dominick's 12	30.06%	69.94%	68.88%	31.12%
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	35.02%	64.98%	89.36%	10.64%
TOTAL	27.09%	72.91%	74.67%	25.33%
 <i>Oatmeal</i>				
	Regular price	Sale price	Regular price	Sale price
Quaker 18	70.05%	29.95%	92.10%	7.90%
Quaker 42	91.58%	8.42%	96.22%	3.78%
TOTAL	80.56%	19.44%	94.16%	5.84%
 <i>Peanut butter</i>				
	Regular price	Sale price	Regular price	Sale price
Jif 18	73.31%	26.69%	85.75%	14.25%
Peter Pan 18	60.92%	39.08%	78.11%	21.89%
Skippy 18	56.45%	43.55%	81.89%	18.11%
TOTAL	66.72%	33.28%	82.02%	17.98%

Table 3: Share Sold at Various Prices

	Ounces Sold			Weeks		
	Ref Price	Sale Price	Above Ref Price	Ref Price	Sale Price	Above Ref Price
<i>Frozen OJ</i>						
Minute Maid 12	25.34%	72.57%	2.09%	57.18%	38.83%	3.99%
Tropicana 12	23.21%	74.28%	2.51%	57.87%	37.60%	4.53%
Dominick's 12	26.24%	70.11%	3.65%	55.73%	36.27%	8.00%
Citrus Hill 12	21.67%	76.27%	2.05%	62.32%	31.31%	6.03%
Florida Gold 12	15.11%	83.41%	1.48%	61.17%	31.91%	6.91%
Dominick's 6	28.22%	68.61%	3.16%	77.63%	15.90%	6.47%
TOTAL	23.88%	73.48%	2.64%	62.54%	31.41%	6.05%
<i>Oatmeal</i>						
Quaker 18	63.86%	33.30%	2.85%	84.88%	10.65%	4.47%
Quaker 42	86.34%	7.83%	2.91%	93.47%	3.44%	3.09%
TOTAL	76.26%	20.86%	2.88%	89.18%	7.04%	3.78%
<i>Peanut Butter</i>						
Jif 18oz	66.89%	31.58%	1.53%	78.36%	20.00%	1.64%
Peter Pan 18 oz	56.68%	43.32%	0.00%	73.37%	26.63%	0.00%
Skippy 18 oz	48.03%	51.08%	0.89%	76.32%	21.73%	1.95%
TOTAL	61.03%	38.13%	0.84%	76.08%	22.69%	1.22%

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

Table 4: Volatility of Ounces Sold

	Variance Over the Sample	
<i>Frozen OJ</i>		
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%
<i>Oatmeal</i>		
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%
<i>Creamy Peanut Butter- Chain 35</i>		
Jif	149,677	
Peter Pan	287,680	
Skippy	18,806	
Sum of UPC Variances	456,163	
Weekly Category Variance	341,530	
Note: Category Variance/Sum of UPC Variances		74.9%

Table 5: Regressions of price/sale measures on commodity input price

	FW Index of Base Price KM	Share of Weeks with sale in category	Actual Price Paid	Num Obs
Oats Chicago	1.676 (0.937)	10.32 (34.90)	0.975 (1.389)	25
Oranges Chicago	0.352 (0.183)	-2.053 (1.54)	0.346 (0.181)	31
Peanuts Charlotte	0.281 (0.144)	-8.04 (6.88)	0.448 (0.203)	28
Peanuts Chicago	1.45 (0.285)	24.16 (7.36)	0.174 (0.178)	28
Peanuts Hartford	0.545 (0.235)	-4.44 (3.60)	0.563 (0.210)	28
Peanuts Houston	-0.013 (0.100)	-17.42 (7.44)	0.324 (0.218)	28
Peanuts Knoxville	0.106 (0.203)	-5.74 (7.65)	0.142 (0.155)	28
Peanuts Los Angeles	0.176 (0.132)	-10.4 (7.34)	0.521 (0.257)	28
Peanuts New York	1.098 (0.195)	-9.04 (6.42)	0.954 (0.235)	28
Peanuts St. Louis	-0.553 (0.165)	-18.13 (5.30)	0.142 (0.292)	28
Peanuts West TX	0.351 (0.284)	-3.4 (6.06)	0.607 (0.333)	28
Share with expected sign	0.82	0.82	1.00	

Each cell is a regression. Standard errors in Parentheses

Table 6: 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 (0.16)	0.404 (0.23)	0.448 (0.12)
Best Price	0.666 (0.15)	0.482 (0.12)	0.469 (0.08)
Constant	0.002 (0.009)	0.014 (0.018)	0.013 (0.008)
R-squared	0.900	0.596	0.902
$p$ value for test that weights sum to 1	0.79	0.57	0.26
$p$ value for test that weights sum to 1 AND constant = 0	0.70	0.15	0.03

Notes: Standard errors in parentheses.

Table 7: Structural Estimates of Peanut Butter 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	0.013 (0.008)	0.448 (0.119)	0.469 (0.081)	0.26	0.03
Los Angeles	0.012 (0.008)	0.684 (0.281)	0.261 (0.140)	0.77	0.66
West Texas	-0.04 (0.029)	0.682 (0.339)	0.676 (0.266)	0.06	0.09
Saint Louis	-0.028 (0.017)	0.380 (0.186)	0.999 (0.174)	0.01	0.01
Chicago	-0.011 (0.023)	0.390 (0.146)	0.776 (0.115)	0.37	0.01
New York	0.038 (0.020)	-0.177 (0.253)	0.973 (0.196)	0.27	0.07
Hartford	-0.036 (0.034)	0.378 (0.299)	0.962 (0.153)	0.22	0.23
Houston	0.022 (0.011)	0.420 (0.124)	0.422 (0.050)	0.12	0.00
Knoxville	0.00 (0.003)	0.871 (0.031)	0.129 (0.024)	1.00	0.69

Notes: Standard errors in parentheses.

Table 8: 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 (0.163)	0.429 (0.225)	0.461 (0.116)
Best Price	0.634 (0.154)	0.468 (0.124)	0.465 (0.078)
Constant	0.002 (0.009)	0.013 (0.017)	0.012 (0.008)
R-squared	0.902	0.603	0.906

Notes: Standard errors in parentheses.