NBER WORKING PAPER SERIES

PASS-THROUGH IN RETAIL AND WHOLESALE

Emi Nakamura

Working Paper 13965
http://www.nber.org/papers/w13965

NATIONAL BUREAU OF ECONOMIC RESEARCH
1050 Massachusetts Avenue
Cambridge, MA 02138
April 2008

I would like to thank David Bell, Charles Engel, Tim Erickson, Penny Goldberg, John Greenlees, Rebecca Hellerstein, Daniel Hosken, Ephraim Leibtag, Rob McClelland, Alice Nakamura, Aviv Nevo, Ariel Pakes, Jon Steinsson and seminar participants at several venues for useful conversations related to this project. I would like to thank the U.S. Department of Agriculture, Economic Research Service (USDA ERS) for financial support for this project. The views expressed herein are those of the author(s) and do not necessarily reflect the views of the National Bureau of Economic Research.

Pass-Through in Retail and Wholesale
Emi Nakamura
NBER Working Paper No. 13965
April 2008
JEL No. E30,F40,L16,L81

ABSTRACT

This paper studies how prices comove across products, firms and locations to gauge the relative importance of retailer versus manufacturer-level shocks in determining prices. I make use of a large panel data set on prices for a cross-section of retailers in the U.S. I analyze prices at the barcode or "Universal Product Code" (UPC) level for individual stores. I find that only 16% of the variation in prices is common across stores selling an identical product. 65% of the price variation is common to stores within a particular retail chain (but not across retail chains), while 17% is completely idiosyncratic to the store and product. Product categories with frequent temporary "sales" exhibit a disproportionate amount of completely idiosyncratic price variation. My results suggest that most of the observed price variation arises from retail-level rather than manufacturer-level demand and supply shocks. However, the behavior of prices is difficult to relate to observed variation in costs and demand at the retail level. This suggests that retail prices may vary largely as a consequence of dynamic pricing strategies on the part of retailers or manufacturers, rather than static demand and supply shocks.

Emi Nakamura
Columbia Business School
3022 Broadway, Uris Hall 820
New York, NY 10027
and NBER
and Economics Department, Columbia University
enakamura@columbia.edu

International economists have long studied retail prices to investigate the central question of how prices respond to exchange rates (e.g. Engel, 1993). Retail price data have also played a key role in assessing empirical models of pricing in industrial organization and empirical macroeconomics. Yet, theoretical pricing models in these literatures have traditionally focused on manufacturer behavior. Recent empirical work suggests important differences in price dynamics at the retail versus the wholesale level of production (Goldberg and Hellerstein, 2007; Nakamura, 2007; Nakamura and Steinsson, 2007). This evidence suggests that understanding the link between retail and wholesale prices is key to developing pricing models that can fit the retail price data.[1]

This paper studies how prices comove across products, firms and locations to gauge the relative importance of retailer versus manufacturer-level shocks in determining prices. I make use of a large panel data set on prices for a cross-section of retailers in the U.S. I analyze prices at the barcode or "Universal Product Code" (UPC) level for individual stores. I find that only 16% of the variation in prices is common across stores selling an identical product. 65% of the price variation is common to stores within a particular retail chain (but not across retail chains), while 17% is completely idiosyncratic to the store and product.[2] Product categories with frequent temporary "sales" exhibit a disproportionate amount of completely idiosyncratic price variation.

My results suggest that most of the observed price variation arises from retail-level rather than manufacturer-level demand and supply shocks. However, the behavior of prices is difficult to reconcile with a model in which desired prices move due to contemporaneous demand and supply shocks, a common set-up in macroeconomics, international economics and industrial organization. This suggests that retail prices may vary largely as a consequence of dynamic pricing strategies on the part of retailers or manufacturers.[3]

The analysis I present here regarding the importance of price variation at the level of individual retail stores is related to recent work in macroeconomics showing that large "idiosyncratic shocks" are needed to explain retail price fluctuations (Golosov and Lucas, 2007; Klenow and Kryvtsov, 2007). In these models, the "idiosyncratic shocks" driving price dynamics are shocks to manufac-

[1]This work is also closely related to recent papers in international economics attempting to measure and study the theoretical implications of "distribution margins". See for example, Burstein, Neves, and Rebelo (2003).

[2]Retailers are, of course, not necessarily the source of price variation idiosyncratic to particular retail chains since manufacturers may adjust their prices differently to different retailers. I discuss this issue in section 3.

[3]For example, see Varian (1980), Sobel (1984), Aguirregabiria (1999) and Lazear (1986) for models in which the firm's desired price varies endogenously. Kehoe and Midrigan (2007) study an alternative model of sales, in which sales arise due to transitory demand and supply shocks.

turers' productivity. Such productivity shocks would, however, generate substantial comovement across prices for the same good at different retail stores. I show that we observe little such comovement. My results suggest that we must delve deeper for the source of the large observed fluctuations in retail prices.

1 Data

This paper uses a new data set on prices from AC Nielsen. The novel feature of the data set is its large cross-sectional dimension. The data consist of weekly price and quantity series for about 7000 grocery stores across the U.S.[4] These grocery stores are members of 33 major chains, and cover 50 major U.S. cities. The time series coverage is short: the data cover all 12 months of 2004. The data include approximately 100 different UPC's selected within a wide variety of grocery store food and beverage categories. In each product category, the data set includes the top 1-3 UPC's by national dollar sales volume.[5] In total, the data set consists of about 50 million observations.

Few papers have studied the comovement of prices across retailers, perhaps because most price data available to academic researchers covers only a narrow cross-section of retailers.[6] The most closely related work to the present analysis is Hosken and Reiffen (2004). They show that sales account for a large fraction of the variation in prices, and find support for the view that these transitory price fluctuations reflect temporary changes in retail margins rather than wholesale price changes.[7]

The huge cross-sectional dimension of my data allows me to carry out a more detailed analysis of price variation across products, stores and cities than has been possible using other data sources. In the case of the U.S. Bureau of Labor Statistics (BLS) CPI Research Database data studied by Hosken and Reiffen (2004), on average 7 price quotes are collected per month for each item category

[4]The data contain weekly dollar sales and sales volume. I construct average prices by dividing dollar sales by sales volume. This could lead to an underestimate of the extent of price rigidity if prices change more than once a week.

[5]The categories are Beer, Bread, Cereal, Cheddar cheese, Crackers, Cream cheese, Canned soup, Coffee, Flour, Frankfurters, Ice cream, Apple juice, Margarine, Marinara, Oil, Peanut butter, Ravioli, Lime diet soft drinks, Cola, Diet cola, Lime soft drinks, Other soft drinks, Other diet soft drinks, Spaghetti, Sugar, and Tuna. I selected these categories to roughly correspond to the food-at-home categories (excluding fresh food) in the U.S. Consumer Price Index (CPI). A list of all included UPC's and stores is available from the author upon request.

[6]A substantial amount of academic research has focused on the Dominick's Finer Foods database from the University of Chicago Graduate School of Business which covers a single retail chain.

[7]In a related exercise, Leibtag et al. (2007) study the synchronization of manufacturer price changes in the U.S. coffee industry. They find substantial comovement in the timing of price changes across major coffee manufacturers.

and area. In many cases, BLS price collectors collect different UPC's at different stores for the same product category, implying that often only a single observation is available for a unique UPC at a given point in time.[8] As a consequence, Hosken and Reiffen (2004) analyze the role of manufacturers by studying price comovement across products in narrowly defined product categories rather than particular UPC's. My data also have a much greater number of price quotes for identical UPC's at a given point in time than AC Nielsen "scanner panel" data based on household surveys.[9]

It is important to note that the sample of stores included in the present data set is not randomly selected. First, not all stores agree to provide AC Nielsen with data, and to share this data in disaggregated form. It is well-known that Walmart does not share its data with AC Nielsen. Second, the data included in the data set were selected to represent the largest U.S. supermarket chains. Supermarket chains accounting for a small fraction of retail sales, such as independent supermarkets, are not included.

2 Results

I begin by documenting some basic properties of the price dynamics in the data. Figure 1 depicts a typical weekly price series from the data set along with a "regular" price series that excludes sales. Since there is no variable in the raw data indicating whether a product is on sale, I identify sales here and elsewhere in the paper using a crude "sale filter". The sale filter labels as a sale any price change that returns either to the original regular price or to a new (repeating) regular price.[10]

Columns 1 and 2 of table 1 present summary statistics on the monthly frequency of price change for sale and non-sale price changes. The mean monthly frequency of price change across categories is 42.7% while the median is 43.9%. In most sectors over half of price changes are associated with the temporary sales identified by the sale filter. The mean frequency of price change for regular prices across product categories is 17.5% while the median is 19.0%.

Columns 3 and 4 of table 1 present statistics on price variability. The statistic presented here is

[8]See Broda and Weinstein (2007) for a discussion of the BLS sampling frame.

[9]See e.g. Broda and Weinstein (2007) for a discussion of the AC Nielsen scanner panel data. While these data contain a huge numbers of observations, they reflect a much smaller cross-section of prices for identical items, due to the modest size of the household panel, and the fact that consumers select from a huge array of different UPC's and often buy slightly different items.

[10]The sale filter requires that the price return to the original regular price, or to a new repeating regular price, within 6 weeks. The sale filter is described in greater detail in the Appendix to Nakamura and Steinsson (2007). The parameters used in the filter are L=3, K=3 and J=6 for weekly data.

the standard deviation of prices for the weekly price series. The underlying prices are first logged and demeaned by the average price for the store and UPC. The standard deviation therefore reflects time series variability of prices in percentage terms. The table presents the mean and median of these statistics across product categories.

The time series variation in prices over the course of a year is extremely large. The average standard deviation of log prices for the typical product (relative to its mean) is about 15.3%. A comparison between the two columns reveals that a large fraction of the variance in prices is accounted for by temporary sales. The mean standard deviation of regular prices is 9.2%, about two-thirds of the standard deviation including sales.[11]

Do these large fluctuations in prices reflect the pass-through of costs from some earlier stage of production? A simple way of studying this question is to consider how the time series variability of individual prices compares to the variability of UPC-level averages.[12] Column 1 of Table 2 presents the standard deviation of UPC-level average prices (including sales). The underlying data are monthly average prices, at the level of individual UPC's and stores. To reduce the sample to a more manageable size, these statistics are calculated using a restricted sub-sample of the data, including only the top 10 stores (if 10 exist) within a particular retail chain and city, and the top twenty cities by sales over all product categories in the data set. Column 2 presents the ratio of the standard deviations of the raw price data to the standard deviations of the UPC-level averages. The table reports the mean and median statistics across product categories.

Table 2 shows that the time series variation in raw prices is far greater than the variation in the UPC-level averages. This suggests that the large shocks driving retail prices do not arise at the manufacturer level. Indeed, the common UPC-level component is likely to be even less variable than is suggested by the analysis above since some of the idiosyncratic store-level price movements do not average out, even in this very large sample. In the next section, I consider a more sophisticated procedure for decomposing the sources of variation in prices.

[11] These statistics likely underestimate the role of "sales" in the data. The sale filter is conservative in identifying price patterns as "sales", particularly toward the end of the data set, where future prices are not observed.

[12] This exercise is similar to the exercise carried out in Hosken and Reiffen (2004). The main difference is that Hosken and Reiffen instead consider averages at the level of product categories, rather than UPC's.

2.1 Variance Decompositions

I next consider a simple variance decomposition of prices (including sales). I decompose the variation in prices into two broad classes: 1) price variation common to all items within a product category (e.g. Beer) and 2) price variation idiosyncratic to particular UPC's. Within each of these broad classes, I decompose the price variation into variation that is common across all stores, variation that is common only to stores within the same retail chain, and variation that is completely idiosyncratic to particular stores.[13]

I estimate the variance decomposition using panel data on prices, where each observation is the monthly average price for an individual UPC at an individual store (e.g. A 12 pack of 12 ounce Diet Pepsi at the Pathmark on 125th street in New York City).[14] These price observations are demeaned by the UPC and store-level mean so that all of the variability is time series variation. The sub-sample used in the estimation is the same one used to estimate the statistics in Table 2.

I estimate the variance decomposition separately for each product category and city in the data set for which a sufficient amount of data are available.[15] The categorization described above implies 6 distinct sources of price variation (3 sources of variation each within of the two categories described above). These components are estimated using a standard maximum likelihood estimator.[16]

Table 3 reports the results of the variance decomposition. Columns 1-3 report the fraction of price variation that is common within a product category. Column 1 reports the fraction that is common both across all UPC's within a product category and across all stores in the data set. Column 2 reports the fraction of the variation that is common within the product category and across stores in particular retail chain (but not across retail chains). Finally, column 3 reports the fraction of the variation that is common only to a product category and store (but not across

[13] I do not adjust the prices for inflation. CPI inflation was 2.9% between January 2004 and January 2005 and therefore has little effect on my results. Over longer time periods, it would be essential to consider a model allowing for trend inflation.

[14] I consider prices averaged over months because this allows the variance decomposition to capture correlations between price changes at retailers in slightly different weeks so long as the price changes occur in the same month. The results from the variance decomposition are very similar if I use prices for the first week of each month rather than monthly average prices.

[15] For the model to be identified, there must be at least two retail chains that sell products in the city/product category, and at least two UPC's in the product category.

[16] The variance decomposition is implemented using a random effects model with i.i.d. shocks for each of the 6 components. These estimates do not account for dynamic correlations, though I analyze monthly average prices to allow for correlations across weeks within a month. Alternative approaches to estimating variance components models include ANOVA and REML. See e.g. Baltagi (2005) for an excellent survey of these methods. See the table 3 for a listing of the variance components.

stores).

Columns 4-6 report a similar set of statistics for the components of price variation that are idiosyncratic to particular UPC's. Column 4 reports the fraction of UPC-level variation that is common across all stores within the same city. Column 5 reports the fraction of UPC-level variation that is common within a particular retail chain (but not across retail chains). Finally, column 6 reports the fraction of UPC-level variation that is idiosyncratic to a particular store and UPC. All statistics are calculated by first averaging the variance components across stores in the sample, and then calculating the mean fractions over all product categories.[17]

I now aggregate these components into somewhat more user-friendly categories. The fraction of price variation common across all retail stores is the sum of: 1) the fraction due to variation at the category level over all stores (7.1%) and the fraction at the UPC-level over all stores (9.4%). These estimates imply that total product-level variation is 16.4%. The component due to chain-level dynamics is the sum of: 1) chain-level variation for product categories (9.8%) and 2) chain-level variation for particular UPC's (55%). Together, these estimates imply that the chain-level variation is 64.8%. Finally, the store-level component of price variation (common to a product category in a store) is estimated to be 2.1%, and the component of price variation idiosyncratic to both a particular store and to a particular UPC is estimated to be 16.6%.

To summarize, the variance decomposition shows that retail-level shocks drive an important wedge between the retail prices we observe, and manufacturer costs. Only 16% of the price variation is common to all stores selling an identical product. The majority of price variation is coordinated at the level of the supermarket chain. Though I do not present these results here, I find similar results for the timing of price changes.[18]

Do variations in retail-level demand and supply factors explain the price variation? Let us first consider variations in retailer costs. BLS estimates the gross margins of "Food and Beverage" stores are only 28.3%.[19] Since the time series standard deviation of weekly prices is approximately 15%, this implies that retail costs such as labor and rent would need to be hugely variable to explain the retailer-specific variation in prices. Moreover, shocks to retail labor or rent are likely to affect

[17] A detailed table of the variance decomposition by product category is available in the online appendix to this paper.

[18] I estimated an analogous variance decomposition for the monthly frequency of price change and obtained similar results regarding the relative importance of the different variance components.

[19] See http://www.brookings.edu/es/research/projects/productivity/workshops/20031121_chapter4.pdf for a discussion of these estimates.

all the UPC's in a given category at the same time. Yet, Table 3 shows that the majority of price variation (71.6%) is common neither across all UPC's within a product category nor across retail chains.

An alternative explanation of the retail chain-level variation in prices is demand shocks. Demand shocks specific to particular UPC's and retail chains could explain the observed price variation. This is, however, difficult to reconcile with the fact that only a small fraction of price variation (19%) is common to all products in a category at a given retail store. For example, shocks to seasonal demand for particular product categories seem likely to affect the demand for all UPC's in the product category at the same time. It is important to note that while advertising and promotional activity may be highly correlated with the timing of price adjustments, these endogenous demand factors must themselves be explained by a successful retail pricing theory.

2.2 Sales and Price Volatility

Temporary sales play a dominant role in explaining price fluctuations in the retail price data. (See Figure 1 and Table 1). Some of the most common theories of sales in the industrial organization literature are dynamic pricing theories. These include models that present sales as a means of price discriminating between different types of consumers (e.g. Varian, 1980; Sobel, 1984), and those that emphasize the role of store inventories (e.g. Lazear, 1986; Aguirregabiria, 1999). These theories generate variations in prices independent from shocks to the marginal cost of production or exogenous shocks to demand.

A natural question is, therefore, whether the large amount of idiosyncratic price variation I observe in the data is related to the prevalence of temporary sales. Figure 2 presents a scatter plot of the relationship between the fraction of price variation explained by the residual component in the variance decomposition and the frequency of price changes due to sales.[20] Each point in the scatter plot corresponds to a particular product category. The figure shows that there is a robust positive relationship between these variables. Product categories with a large number of sales, such as soft drinks, also have a disproportionately large fraction of residual price variability.

[20]Sales are identified using the sale filter described above.

2.3 Dynamic Behavior of Retail and Wholesale Prices

Finally, I consider how the dynamics of individual retail prices differ from the dynamics of price series averaged across retailers. Table 4 presents the autoregressive coefficients in regressions of prices (including sales) on their one month lags. The first specification is based on prices in the first week of each month for a particular store and UPC. The second specification is based on average prices over a month for a particular store and UPC. The third specification is based on average monthly prices for all retailers selling a particular UPC. In all cases, the underlying prices are logged and demeaned by the log average price for the store and UPC. The sub-sample used in the estimation is the same one used to estimate the statistics in Table 2.

These estimates reveal important differences between the dynamics of the individual price series and the UPC-leel averages. The autoregressive coefficient for individual prices is -0.04. Thus, individual prices are close to serially uncorrelated at a monthly frequency. The serial correlation rises to 0.19 if one considers monthly averages rather than the price in the first week of each month. The third column presents the results for averages across all retailers selling a given UPC. These series are far more persistent: the autoregressive coefficient is 0.40.[21] The empirical properties of retail prices—both the remarkably low persistence, and the high fraction of idiosyncratic variation—make clear that individual retail prices are not closely linked to standard price determinants in macroeconomics and international economics such as wages, productivity and exchange rates. The substantially lower volatility and greater persistence of average prices across stores leaves greater scope for a close link between manufacturer-level prices and factors such as wages, productivity and exchange rates.

3 Who Adjusts Prices?

One can use the results of the variance decomposition to analyze the question of whether retailers or manufacturers play a dominant role in price-setting. This question has important implications for how we model price rigidity. For example, if manufacturers determine the timing of all temporary sales, then there cannot be much price rigidity at the manufacturer level for the products I consider.

[21] This estimate is likely to be biased downward because not all idiosyncratic shocks wash out in the UPC-level average. This idiosyncratic variation remains significant, despite the large number of stores, due to the huge variability in individual prices.

The evidence presented above has two potential interpretations in this regard.

On the one hand, if manufacturers have a limited ability to price discriminate to retailers within the same city, then the empirical evidence I have presented suggests that retailers play a dominant role in price-setting. This assumption may be justified for two reasons: 1) The Robinson-Patman Act formally restricts the ability of manufacturers to price discriminate across retailers in the same geographical area[22] and 2) There are arguably greater search frictions in sales to households than to large retailers. On the other hand, there may be huge amount of retailer-specific price discrimination on the part of the manufacturer despite the Robinson-Patman Act. In this case, manufacturer prices may be highly responsive to retail-level shocks.

One would like to distinguish between these explanations using direct evidence on manufacturer prices. A number of factors make it important to interpret manufacturer prices with care. Manufacturers often offer complex trade deals to retailers. A retailer may be required carry out advertising, or sell a particular number of units during a time period, to receive a trade discount. Manufacturers often offer multiple trade deals simultaneously, allowing retailers to select which deal to take (and when) to take them up.[23] Indeed, Maratou (2006) reports, based on a survey of 43 supermarket chains, that in 49.8% of cases the retailer "initiates the trade promotion" and in 58.9% of cases the retailer "selects the trade promotion type". These factors make wholesale prices substantially more difficult to interpret than retail prices. This remains an important topic for future research.

[22]The Robinson Patman Act states that a manufacturer cannot charge different prices for an identical item to retailers that are located less than 200 miles apart. Volume discounts are allowed, though this may be less relevant for the large stores in my sample.

[23]See Neslin (2002) and Dreze and Bell (2003) for recent discussions of trade dealing practices.

References

Aguirregabiria, V. (1999): "The Dynamics of Markups and Inventories in Retail Firms," *Review of Economic Studies*, 66, 275–308.

Baltagi, B. H. (2005): *Econometric Analysis of Panel Data*. John Wiley & Sons, Chichester, England.

Broda, C., and D. E. Weinstein (2007): "Product Creation and Destruction: Evidence and Implications," NBER Working Paper 13041.

Burstein, A. T., J. C. Neves, and S. Rebelo (2003): "Distribution costs and real exchange rate dynamics during exchange-rate-based stabilizations," *Journal of Monetary Economics*, (50), 1189–1214.

Dreze, X., and D. R. Bell (2003): "Creating Win-Win Trade Promotions: Theory and Empirical Analysis of Scan-Back Trade Deals," *Marketing Science*, 22(1), 16–39.

Engel, C. (1993): "Real exchange rates and relative prices: An empirical investigation," *Journal of Monetary Economics*, 32, 35–50.

Goldberg, P. K., and R. Hellerstein (2007): "A Framework for Identifying the Sources of Local-Currency Price Stability with an Empirical Application," Working Paper.

Hosken, D., and D. Reiffen (2004): "Patterns of Retail Price Variation," *Rand Journal of Economics*, 35(1), 128–146.

Kehoe, P., and V. Midrigan (2007): "Sales, Clustering of Price Changes, and the Real Effects of Monetary Policy," Working Paper, University of Minnesota.

Klenow, P. J., and O. Kryvtsov (2007): "State-dependent or time dependent pricing: Does it matter for recent u.s. inflation?," Working Paper.

Lazear, E. P. (1986): "Retail Pricing and Clearance Sales," *American Economic Review*, 76, 14–32.

Leibtag, E., A. Nakamura, E. Nakamura, and D. Zerom (2007): "Cost Pass-Through in the U.S. Coffee Industry," *Economic Research Report USDA*, 38.

Maratou, L. M. (2006): "Bargaining power impact on off-invoice trade promotions in U.S. grocery retailing," Cornell University Thesis.

Nakamura, E. (2007): "Accounting for Incomplete Pass-Through," Working Paper.

Nakamura, E., and J. Steinsson (2007): "Five Facts about Prices: A Reevaluation of Menu Cost Models," Working Paper.

Neslin, S. A. (2002): *Sales Promotion*. Marketing Science Institute, Cambridge, MA.

Sobel, J. (1984): "The Timing of Sales," *Review of Economic Studies*, 51, 353–368.

Varian, H. R. (1980): "A Model of Sales," *American Economic Review*, 70, 651–659.

TABLE 1
Basic Statistics on Prices

	Monthly Frequency		Price Variability	
	All	Regular	All	Regular
Median	42.7	17.5	15.3%	9.1%
Mean	43.9	19.0	15.3%	9.2%
Sample Size	43,006,064	43,006,064	43,006,064	43,006,064

The underlying data are weekly price observations. The "regular'' price series is calculated based on the filter described in the text. The monthly frequency of price change is the fraction of the time that the weekly price (for a particular store and UPC) differs from the price 4 weeks earlier. Price variability is calculated by first logging and demeaning at the store-UPC, level and then calculating the standard deviation of this series. The statistics above are means across stores and product categories.

TABLE 2
Volatility of Prices vs. UPC Av.

	Price Variability	
	UPC Av.	Ratio to Av.
Median	3.51%	2.7
Mean	4.28%	3.1
Sample Size	1008	346,930

For each store and UPC, the raw weekly prices (including sales) are averaged within months, then logged and demeaned at the store-UPC level. The "UPC Av." is constructed by averaging this series across all retail stores. "Price Variability" is the standard deviation of this series. "Ratio to Av." is the ratio of price variability for the UPC-store series to the price variability for the UPC Av. series. The statistics above are means across product categories.

TABLE 3
Variance Decomposition of Prices

Category-Level			UPC-Level		
All Stores	Chain	Indiv..	All Stores	Chain	Indiv.
7.1	9.8	2.1	9.4	55.0	16.6

The variance decomposition is estimated using monthly average prices including sales. For each store and UPC, the raw weekly prices (including sales) are averaged within months, then logged and demeaned at the store-UPC level. The variance decomposition is carried out using this monthly demeaned series. The statistics above are means across product categories. The variance decomposition is based on 279,718 observations.

TABLE 4
Dynamic Properties of Prices

	Raw	Av. Monthly	Av. UPC Monthly
AR(1) Coef.	-0.04	0.19	0.40
	(0.002)	(0.002)	(0.03)
Sample Size	317,500	319,286	969

The table gives the autoregressive coefficients AR(1) regressions. In the first specification, the dependent variable is the price in the first week of each month for a particular store and UPC. In the second specification, it is the average price over a month for a particular store and UPC. In the third specification, is the average monthly price across all retailers selling the UPC. In all cases, the prices are logged and demeaned at the store-UPC level.

Figure 1: Prices and Regular Prices

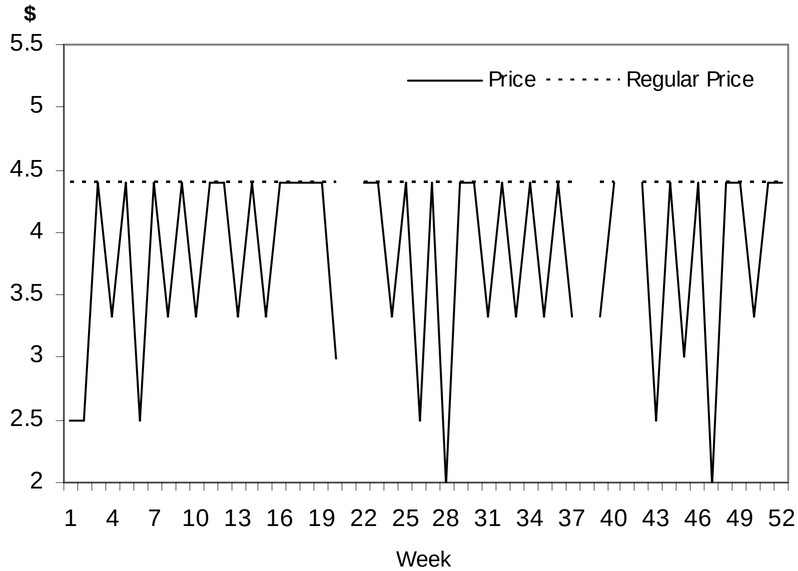

The figure plots a price series for a 12 Pack of 12 Ounce Diet Pepsi from a particular store in the data set. The regular price is constructed according to the ``sale filter'' algorithm described in the text. The missing observations correspond to weeks when no units were sold.

Figure 2: Sales vs. Residual Variance

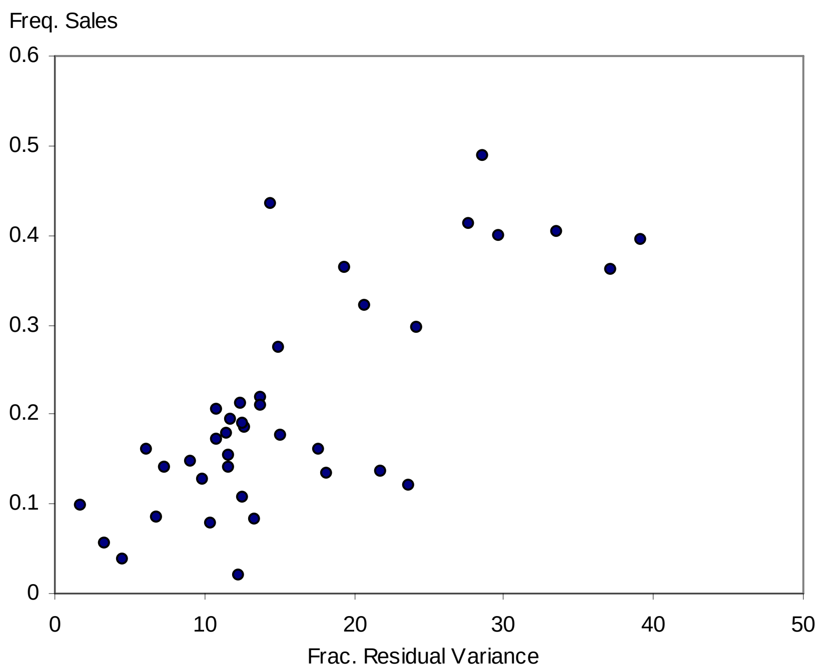

The figure presents a scatter plot of the frequency of price changes due to temporary sales (where sales are identified by the sale filter described in the text) versus the fraction of the ``residual'' variation in prices in the variance decomposition. Each point corresponds to a unique product category.

TABLE A1

Variance Decomposition of Prices: Category Detail

	Category-Level			UPC-Level		
	All Stores	Chain	Indiv..	All Stores	Chain	Indiv.
Beer	5.5	4.7	0.3	24.1	50.8	14.6
Bread	6.5	2.3	2.3	12.8	70.8	5.3
Cereal	3.6	12.1	1.1	7.5	64.8	11.0
Cheddar cheese	4.1	1.8	0.6	15.9	66.7	10.8
Crackers	15.2	2.2	0.8	11.5	58.4	11.9
Cream cheese	38.1	38.5	7.7	0.5	9.1	6.1
Canned soup	13.1	10.8	1.7	8.4	55.1	10.9
Coffee	1.8	5.7	0.7	9.4	59.4	23.0
Flour	32.7	19.2	0.9	6.5	34.2	6.5
Frankfurters	1.3	2.3	1.3	7.5	64.2	23.4
Ice cream	0.7	8.6	6.5	6.0	56.3	21.8
Apple juice	0.1	0.0	2.2	11.4	73.7	12.7
Margarine	1.6	8.4	1.2	12.5	66.7	9.6
Marinara	3.2	0.8	0.3	10.8	72.6	12.2
Oil	3.4	3.8	0.8	8.4	66.1	17.4
Peanut butter	2.3	1.1	0.9	8.9	70.1	16.7
Ravioli	18.8	68.6	6.3	0.1	0.7	5.5
Soft drinks lime diet	1.3	3.3	1.3	8.2	50.7	35.2
Soft drinks cola	3.3	0.8	1.3	8.1	54.2	32.3
Soft drinks cola diet	3.4	3.3	1.6	10.4	52.6	28.8
Soft drinks lime	1.6	3.2	0.9	7.6	58.2	28.5
Soft drinks other	5.8	20.3	5.6	5.0	41.2	22.1
Soft drinks other diet	5.4	12.2	5.8	3.5	39.8	33.4
Spaghetti	4.7	11.2	0.6	3.3	66.8	13.3
Sugar	5.1	8.7	2.1	28.2	47.3	8.7
Tuna	1.1	1.8	0.4	7.3	79.0	10.4
Weighted Mean	7.1	9.8	2.1	9.4	55.0	16.6

The variance decomposition is estimated using monthly average prices including sales. For each store and UPC, the raw weekly prices (including sales) are averaged within months, then logged and demeaned at the store-UPC level. The variance decomposition is carried out using this monthly demeaned series. The statistics above are means across stores for individual categories. The last row is a mean across product categories. The variance decomposition is based on 279,718 observations.