THE ECONOMICS OF SALES

By

Philip John Glandon, Jr

Dissertation

Submitted to the Faculty of the

Graduate School of Vanderbilt University

in partial fulfillment of the requirements

for the degree of

DOCTOR OF PHILOSOPHY

in

Economics

August, 2011

Nashville, TN

Approved:

Professor Ben Eden

Professor Jacob S. Sagi

Professor Kevin X.D. Huang

Professor Mario J. Crucini

To Grey and Sala,

who have given me an endless supply of inspiration, joy, and comfort.

ACKNOWLEDGMENTS

This work would not have been possible without the financial support of the Vanderbilt University Graduate School Fellowship, the Dornbush Research Assistanceship, and a grant from the Council of Economics Graduate Students. My research has benefited from many conversations with the members of my committee: Professors Mario Crucini, Jacob Sagi, and Kevin Huang. I would like to thank Matt Jaremski for his willingness to work with me on my third chapter. He is an excellent colleague and working with him was enjoyable and educational. I am especially thankful for the guidance and encouragement of my adviser, Dr. Ben Eden, who demonstrated an incredible amount of patience and genuine interest in my work as a researcher. I take with me many fond memories of conversations over lunch at Sitar.

I would like to thank my parents who have been, and continue to be, the best role models one could hope for. Through all phases of life, they have exemplified the delicate balance of parenting. In making choices, they have given me guidance without overbearance. In times of disappointment, they've given consolation while allowing me to learn from my mistakes. In my success, they have celebrated while encouraging me to look forward.

The years I have spent in graduate school have been immensely enjoyable and fulfilling because of the hard work of my wife. She has tirelessly supported our family both emotionally and financially through it all. Any success I achieve as an economist will be largely due to the support of my lovely wife.

TABLE OF CONTENTS

LIST OF TABLES

LIST OF FIGURES

CHAPTER 1

FACTS ABOUT SALES

1.1 Introduction

The long running debate about price stickiness has recently focused on temporary price reductions ("sales" hereafter). Many of the items that compose the CPI "basket" exhibit frequent price changes. This posed a challenge to New Keynesian models that rely heavily on explicitly rigid prices to generate real effects of money (Bils and Klenow, 2004; Nakamura and Steinsson, 2008). However, a large portion of price changes are temporary reductions followed by a return to the "regular" price and are commonly referred to as sales (Klenow and Kryvtsov, 2008).

Several recent articles analyze whether sales are necessarily inconsistent with sticky price models (Eichenbaum, Jaimovich and Rebelo, 2011; Kehoe and Midrigan, 2010; Guimaraes and Sheedy, 2011, which I will refer to as EJR, KM, and GS hereafter). These studies uniformly conclude that sales are not an important source of aggregate price flexibility. Eichenbaum, et al. conclude "...a lot of high-frequency volatility in prices and quantities has little to do with monetary policy and is perhaps best ignored by macro economists." This raises two empirical questions that, to the best of my knowledge, have not been addressed: 1) How do aggregate characteristics of sales vary over time? and 2) How important are sales in understanding aggregate price movements?

These questions are motivated by a simple thought experiment: if households face an unexpected reduction in nominal income, would they tend to buy more items at a sale price? Alternatively, if a store faced a reduction in demand, would managers change the frequency and/or depth of their price promotions? If the answer to either of these questions is "yes", then we ought to reconsider the conclusion that sales are unimportant for macro economists. This idea is closely related to the analysis in Chevalier and Kashyap (2011), who show that stores may respond to certain types of demand shocks by adjusting the frequency or size of sales.

In this chapter I use scanner data to document several facts about the aggregate behavior of sales. My primary source of data is the IRI research database, which contains six years of weekly scanner data on 31 categories from a sample of grocery and drug stores located in 50 US markets.[1] I supplement the analysis

[1] To learn more about this database, see Bronnenberg, Bart J., Michael W. Kruger, Carl F. Mela. 2008. Database paper: The IRI marketing data set. Marketing Science, 27(4) 745-748. I wish to thank SymphonyIRI for making this data available.

with the Dominick's Finer Foods (DFF) database, a publicly available scanner data set from the University of Chicago's Booth School of Business.[2] I document the following five facts about about the nature of sales:

1. Sales account for as much as 45% of revenue (depending on the category).
2. Sale frequency is positively correlated with expenditure share (across categories and across products).
3. Most monthly variation in average unit price is due to changes in the frequency and size of sales and is *not* reflected in the CPI.
4. When demand faced by a multiproduct retailer falls, the fraction of revenue from sales tends to increase.
5. Sales are generally staggered (a) across stores (b) across products produced by the same manufacturer, and (c) across products within a store.

Taken together, these facts suggest that sales play an important role in the evolution of average unit price and nominal expenditure. They imply that we should think carefully about how we interpret the CPI, particularly when it is used to convert nominal magnitudes into real magnitudes. The facts are not consistent with models of sales in which idiosyncratic and transient cost shocks are the primary reason that sales occur. Rather, they support alternative models that rely on some combination of inventory management, consumer heterogeneity, search costs, and strategic behavior to generate intertemporal price dispersion. In Chapter 2, I discuss how these facts pertain to various models of sales and then sketch a model in which sales tend to dampen the effect of demand shocks on output.

Road map for Chapter 1

Section 1.2 describes both data sources and includes a discussion of how I selected the sample for this analysis. In Section 1.3 I describe the algorithm used to identify regular prices and sales. I present highly aggregated summary statistics and establish facts (1) and (2) in Section 1.4. Section 1.5 looks at aggregate dynamics of sales by comparing indexes of average regular price and average unit price over time. Fact (3) comes from this section. Fact (4) is established in Section 1.6, which studies how grocery and drug chains use sales to respond to large changes in demand. Section 1.7 analyzes the extent to which sales are synchronized and section 1.8 concludes.

Information Resources, Inc. ("IRI") has changed its name to SymphonyIRI Group, Inc. All estimates and analyses in this paper based on SymphonyIRI Group, Inc. data are by the author and not by SymphonyIRI Group, Inc.

[2]To obtain data from the Dominick's Finer Foods Database, visit the University of Chicago's Booth School website: http://research.chicagobooth.edu/marketing/databases/dominicks/

1.2 Data

Scanner data have three important advantages over BLS survey data for studying pricing behavior. First, quantity sold is available which indicates (among other things) how "important" a particular price is.[3] Second, scanner data contain weekly observations and tend to pick up almost all movements in prices since intra week price movements are rare. BLS survey data are sampled monthly or bimonthly. Finally, with scanner data we observe the price and quantity sold of *all* products available within a category-store.

The scanner data I use come from two sources: Dominick's Finer Foods via the University of Chicago Booth School of Business and the IRI marketing data set provided by Information Resources Inc (Bronnenberg et al., 2008).

Dominick's Finer Foods

The Dominick's Finer Foods (DFF) database is available publicly via the University of Chicago Booth School of Business. The database contains weekly scanner data (price and quantity sold) for 29 categories sold in each of the chain's 90 different stores located near Chicago, IL from 1989 until 1996.

These data cover all stores operated by a single chain and provide the exact location of each store. Unfortunately, the stores are located in a small geographic area and tell us only about the pricing behavior of a single firm. However, they do span the recession of 1991 so we can analyze how DFF reacted to this particular recession.

There are also some gaps in the data. Several categories are missing for certain months. I am able to construct a balanced panel of 22 categories covering October 1989 to August 1994 (58 months). In most of the analysis that follows I will uses this subset of categories.

IRI Marketing Data

Information Resources, Inc. have recently made a much larger set of scanner data available for academic research (Bronnenberg et al., 2008). These data cover 31 categories in 50 different markets and contain both grocery stores and drug stores from several different chains. The data have detailed information about product attributes and also contain information about other marketing activity such as feature and display. The primary disadvantage to this data is that we cannot determine the exact location of the stores or the demographics of the shoppers who typically visit them. The data span the relatively mild recession of 2001, but contains very few observations prior to its onset.

[3] Stigler and Kindahl (1970) showed that transaction prices are far more informative and often behave differently than list prices.

Table 1.1: IRI Sample Markets

Market	Stores	Chains	Items	Observations (millions)	Revenue ($million/year)
Dallas	17	5	181,628	23.3	67
Hartford	11	5	129,921	17.1	66
Houston	17	4	168,303	21.4	64
Los Angeles	57	5	555,632	69.5	238
New England	24	4	255,051	33.9	157
New York	35	6	392,104	49.9	203
Philadelphia	19	6	212,600	26.9	90
Phoenix	16	4	175,750	22.1	77
Raleigh	19	4	178,943	23.2	63
Sacramento	14	5	138,184	17.5	52
San Francisco	27	5	250,950	31.0	115
Seattle/Tacoma	15	4	149,220	18.7	54
St. Louis	14	4	148,557	19.9	93
Syracuse	7	4	75,827	9.5	42
Total	292	65	3,012,670	383.9	1,380

Notes: For most of the analysis in this chapter, I use only the markets listed above. It includes all IRI markets with at least four different chains that provided a balanced panel for all six years.

Only some of the chains found in the IRI research database provided data for all six years (2001-2006). I do not want changes in the composition of panels to be a source of variation in aggregate measures, so I only include stores with data for all six years in this analysis. Also, for practical purposes, I only present statistics from markets with at least four chains in sections two and three. The resulting panel used in sections two and three contains 14 markets, 65 chains, and 292 stores. Table 1.1 summarizes the markets included in the sample. For the price index construction and for all of section four I include all 50 markets.

Missing Item-Weeks

If a UPC was not scanned at a particular store in a particular week, then that store-UPC-week is missing from the data. There are two reasons that an item was not scanned: 1) the item was not available at the time, or 2) the item was available but not purchased. If an observation is missing because of (2), then I want an observation for that week (with the actual price and zero units sold). Otherwise it should remain as missing.

I assume that if an item was not scanned for 9 weeks or more, then it was not available. If a store-UPC is missing for fewer than 9 weeks, I set the quantity sold to zero and the price equal to the most recently observed regular price. This is based on the assumption that an item is unlikely to sell zero units if it is on sale.

Loyalty Cards and Store Coupons

One difficulty with using scanner data arises from the use of loyalty cards and store coupons. These promotions result in different prices for different shoppers within the week. Scanner data only allow us to compute the average price paid during the week. This makes identifying regular price less straightforward because store specific discounts that don't apply to all shoppers show up as small week-to-week variation in price. The algorithm used above will typically not identify these average price fluctuations as sales unless they result in a temporarily large drop in price. The effect on my analysis will be to overestimate the flexibility of regular price and to underestimate the use of sales.

1.3 Defining the Regular Price and Identifying a Sale

To proceed, I need to construct two variables: 1) a binary indicator of whether the item was on sale and 2) the non-sale (or "regular") price (which is different from the observed price when an item is on sale). Ideally these variables would come straight from the store, but they are not provided in the data on a consistent basis so I construct them myself using an algorithm discussed in detail below.[4]

There are several operational definitions of sale and regular prices used in the literature. I find that they are not well suited for this study because they tend to miss sales that are more complicated than a one or two week drop in price followed by a return to the previous price (see Appendix A).

An Operational Definition of Regular Price and Sale

Part of the difficulty in settling on an operational definition of a sale is that there is no widely accepted theoretically based definition of a sale. My starting point for settling this issue is the fact that an item's price is usually one of two frequently observed prices (Eichenbaum et al., 2011). In chapter two, I present a simple model in which stores select randomly between two prices. In this model, the higher of the two prices is the "regular price" and the lower of the two is analogous to the sale price. Thus, I use the following algorythm to create a series of regular prices for each product.

I start by setting the regular price to the observed price whenever it is larger than or equal to the 13 week centered moving average price. If the observed price does not change for six weeks or more, then the observed price in those weeks is the regular price. Also, if two adjacent prices are equal and one is the regular price (as determined above), then they are both the regular price.[5] Following Kehoe and Midrigan

[4]The DFF data contain a "deal" flag that is widely understood to be incomplete. The IRI data contain a price reduction indicator that is computed using IRI's proprietary algorithm for identifying sale prices. For consistency and transparency I chose to create my own sale indicator. I use the IRI sale flag as a benchmark for evaluating the algorithm.

[5]The purpose of this is to ensure that regular price changes in the week that we observe the change. It is needed when regular price falls because the moving average is slower to fall.

Figure 1.3.1: Regular Price Filter Example from the Data

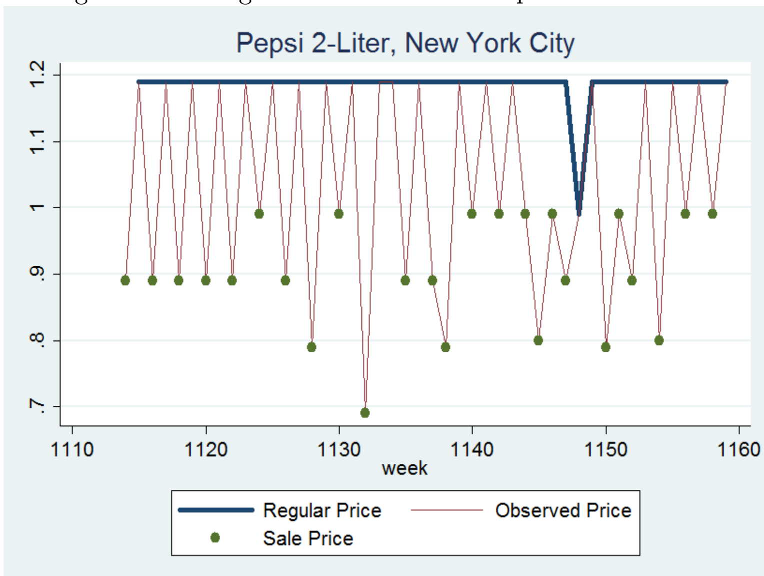

Notes: This is an example (Pepsi 2-Liter bottles at a store in New York City) from the IRI data set. The thin red line represents the scanned price and the thick blue line is the regular price, as determined by the regular price algorithm described in the text. All prices that are 5% or more below the regular price line are considered sale prices, which are represented by green dots.

(2010), I set the remaining unspecified regular prices equal to the previous period's regular price.

Now I have an observed price series and a regular price series. An item is considered to be on sale if the observed price is at least five percent below the regular price.[6]

Figure 1.3.2 shows an example of the regular price series that I constructed using the algorithm described above. The dots represent the observed price series. This particular example is Pepsi 2 Liter Bottles from a store in New York City. It happens to be the number one revenue generating store-UPC in the sample of carbonated beverages sold in New York City. Notice that there are several instances in which the price drops in one week and then drops further the next week before returning to the regular price. This illustrates that a sale is often more complicated than a simple price drop followed by a return to the previous price.

Comparison versus Other Methodologies

To give the reader some idea of how various definitions of regular price and sale prices compare, I provide several examples from the data as well as some summary statistics. Figure 1.3.2 compares the regular price filter I use to the filter used in Kehoe and Midrigan (2010). I select a top selling UPC from each of four

[6]Chevalier and Kashyap use a similar tolerance to allow for the fact that there are occasional small price measurement errors. These errors result from the fact that price is calculated from total dollar sales and total unit sales.

Table 1.2: Sale Definition Comparison for Top 10 Categories

Correlation Matrix	IRI	Glandon	K-M	C-E	Sales % of Observations
IRI	1.00				38.5%
Glandon	0.75	1.00			30.9%
Kehoe-Midrigan	0.64	0.71	1.00		28.6%
Campbell-Eden	0.40	0.50	0.45	1.00	11.4%

Notes: Columns two through four contain the correlation between pairs of sale definitions indicated by the column and row headers. Column 6 indicates the fraction of observations that are identified as a sale using the definition specified in column 1. The sample used for this comparison is the top ten revenue generating categories in the Los Angeles market. I also limit the comparison to the top 1,000 UPCs in each category for computational purposes. Sale % of observations is weighted by revenue.

categories to show how the two definitions differ from each other. The Kehoe-Midrigan filter will occasionally select the lowest frequently charged price as the "regular price".

Table 1.2 provides the total fraction of price quotes identified as a sale by each of the following definitions of sales: 1) IRI's proprietary definition, 2) The "Glandon" definition (used in this paper), 3) The Kehoe-Midrigan adaptation (5% below the Kehoe-Midrigan regular price), 4) the definition used in Campbell and Eden (2005).[7] I also include a correlation matrix of the binary sale indicators. The IRI definition finds the most sales (39 percent) and the Glandon definition has the highest correlation with the IRI definition (0.75) but finds eight percentage points fewer sales in the sample analyzed.[8]

1.4 Summary Statistics

There are three characteristics of sales that deserve our attention:

1. Frequency (measured by fraction of items on sale per week)

2. Size (percentage discount off of the regular price)

3. Quantity Ratio (quantity sold per sale relative to quantity sold per regular price).

Frequency and size are often referred to in the marketing literature as the breadth and depth of price promotions and are determined by the store's managers in order to maximize profit given expectations about shopper behavior (Blattberg et al., 1995). The quantity ratio is the quantity sold per sale relative to the quantity sold at the regular price. It measures the consumer's response to sales.

[7]The IRI research database includes a binary sale indicator variable called PR. Campbell and Eden (2005) defines a sale as an x% or larger drop in price that is completely reversed within two weeks.

[8]For computational purposes, I limit the sample to the top 10 categories and the top 1,000 UPCs in each category for the Los Angeles market.

Figure 1.3.2: Regular Price Filter Comparison: Four Examples

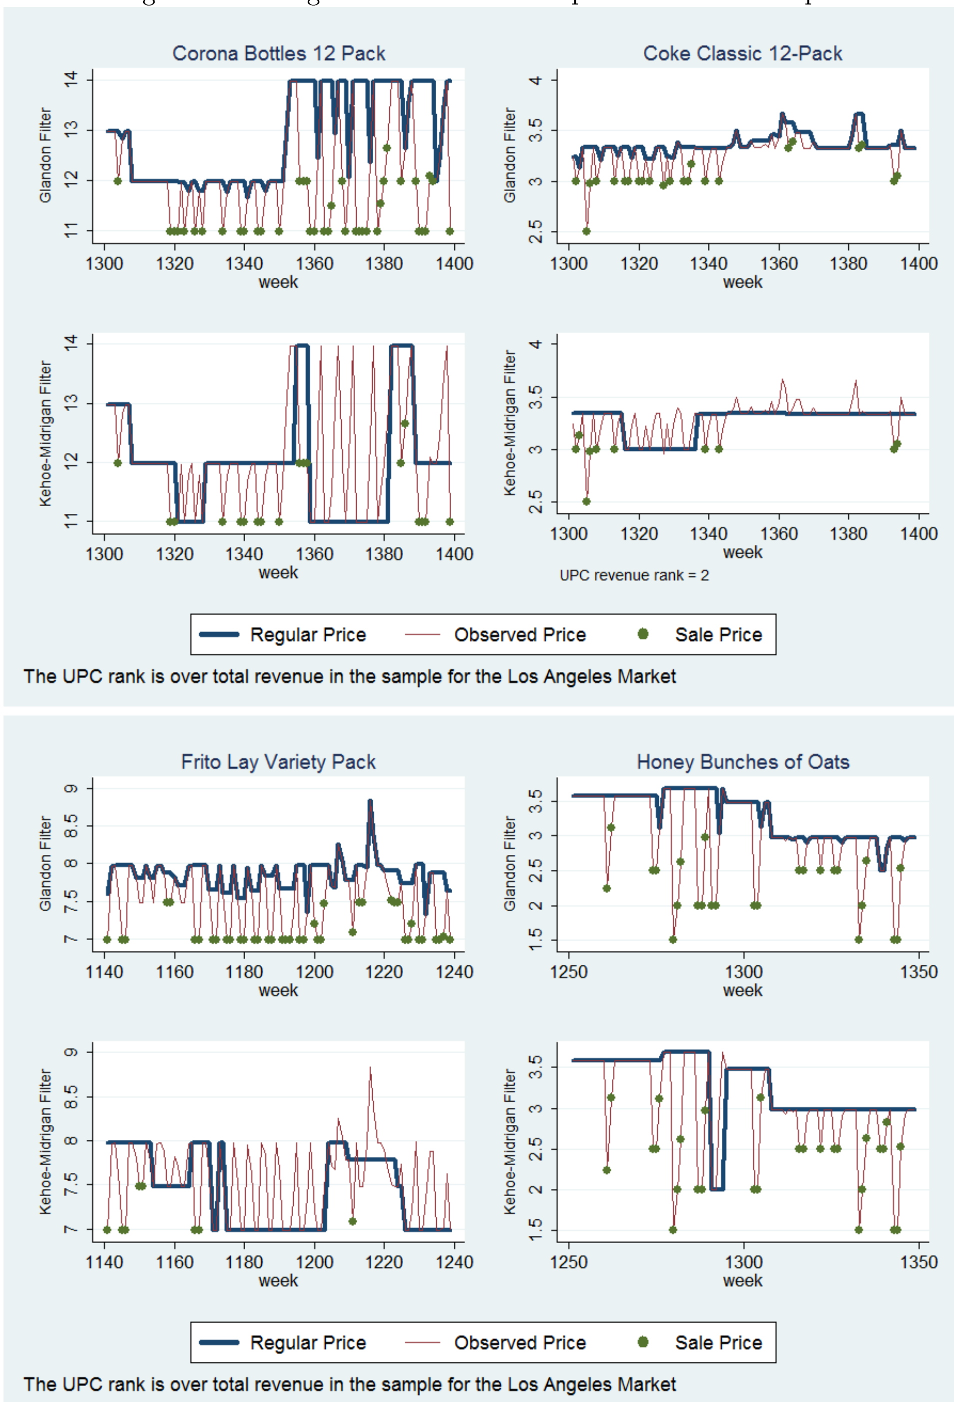

Notes: The figures above compare the regular price filter described in the text to the Kehoe-Midrigan regular price filter. I select one UPC from each of four categories: Beer, Carbonated Beverages, Cold Cereal, and Salty Snacks. Each UPC is one of the top four selling UPCs from the Los Angeles market. The top chart for each category uses the Glandon Filter and the bottom chart uses the Kehoe-Midrigan filter. These examples were selected to illustrate how the two different filters differ from each other.

Table 1.3: Summary Statistics by Category: Dominick's Finer Foods

Category	Products	Months	Annual Revenue $Millions	Sales fraction of total: Weeks	Revenue	Units	Median Discount	Median Quantity Ratio
Soft Drinks	162	90	51.3	34.2%	70%	67%	20%	5.5
Cereals	110	86	29.0	9.6%	23%	29%	12%	2.8
Cheeses	134	91	25.0	25.1%	40%	46%	9%	1.9
Refrigerated Juices	39	91	16.1	34.6%	60%	68%	14%	3.5
Laundry Detergent	139	91	15.6	10.7%	41%	43%	12%	4.7
Frozen Entrees	220	91	14.9	18.4%	43%	53%	20%	4.7
Cookies	188	90	13.1	17.6%	38%	42%	13%	3.0
Beer	88	71	10.8	28.2%	58%	57%	14%	3.7
Bottled Juice	117	91	10.1	20.3%	34%	39%	10%	2.4
Bathroom Tissue	31	91	9.6	23.9%	49%	57%	10%	4.1
Canned Soup	130	91	9.2	14.5%	26%	30%	10%	2.0
Frozen Juices	44	91	7.4	22.3%	46%	53%	14%	3.6
Snack Crackers	85	89	7.3	22.8%	38%	42%	12%	2.1
Paper Towels	27	91	6.7	22.6%	38%	43%	9%	2.8
Canned Tuna	50	91	5.3	21.0%	41%	50%	9%	2.5
Cigarettes	4	91	5.0	4.2%	5%	5%	2%	1.5
Shampoos	453	52	5.0	11.3%	24%	31%	20%	3.7
Fabric Softeners	73	91	4.7	12.7%	27%	26%	9%	2.5
Dish Detergent	70	91	4.6	12.2%	27%	33%	10%	3.0
Grooming Products	229	63	4.5	11.7%	21%	25%	16%	2.8
Frozen Dinners	84	59	4.3	22.9%	42%	51%	17%	4.0
Soaps	72	64	4.2	14.7%	25%	26%	9%	2.0
Analgesics	119	91	4.0	7.3%	15%	16%	13%	2.7
Front End Candies	102	91	4.0	11.3%	17%	24%	15%	2.0
Crackers	60	88	3.7	22.8%	36%	41%	11%	2.3
Toothpaste	115	91	3.2	13.0%	29%	33%	15%	3.5
Oatmeal	24	71	3.2	13.4%	25%	33%	12%	2.7
Toothbrushes	118	91	1.2	12.2%	24%	28%	20%	3.0
Bath Soap	75	63	0.3	6.0%	14%	18%	18%	3.9
All Categories	3,162		283.3	17.0%	43%	47%	13.7%	3.5

Notes: Column 2, is the number of unique products contained in the sample. A product corresponds to a UPC unless there are several UPCs that DFF identifies as the same item, in which case the product is identified by the variable "nitem" from the Dominick's UPC file. Column 5 is the fraction of product-store-weeks that contained a sale price. Columns 6 (7) are the units sold (revenue) from sales divided by the total units sold (revenue). Column 8 is the median discount for products that were on sale $(p_{reg} - p_{sale})/p_{reg}$. Column 9 is the median of the quantity ratio. The quantity ratio is q_{sale}/q_{reg} where q_{sale} and q_{reg} are the quantity of items sold per week on sale and quantity of items sold per week at the regular price.

Table 1.4: IRI Summary Statistics by Category

Category	Store x Products	Annual Revenue ($Millions)	Sales Fraction of Total: Weeks	Sales Fraction of Total: Revenue	Sales Fraction of Total: Units	Median Discount	Median Quantity Ratio
Carbonated Beverages	225,408	180.0	21.9%	43.0%	44.5%	22.1%	1.9
Beer	125,690	136.0	16.6%	35.0%	30.1%	11.7%	1.8
Cold Cereal	160,894	107.3	16.1%	31.8%	41.5%	30.3%	2.9
Salty Snacks	317,814	104.5	17.6%	32.1%	35.8%	21.4%	2.0
Frozen Dinners	237,906	68.7	25.7%	42.5%	48.1%	28.3%	2.6
Yogurt	147,471	57.7	20.2%	28.0%	39.3%	23.9%	2.0
Toilet Tissue	35,925	49.7	17.5%	39.6%	35.3%	23.7%	3.2
Laundry Detergent	87,963	48.5	16.5%	35.7%	44.8%	24.1%	3.1
Soup	138,809	44.8	14.1%	29.6%	36.9%	27.1%	3.0
Frozen Pizza	81,255	38.0	25.7%	42.7%	49.6%	24.1%	2.7
Coffee	103,386	37.0	13.8%	32.2%	39.4%	19.3%	2.7
Paper Towels	31,729	35.5	14.7%	36.4%	31.0%	22.6%	2.6
Hot dogs	28,152	26.7	22.2%	36.8%	48.3%	29.8%	2.4
Spaghetti Sauce	78,634	26.5	17.5%	33.5%	41.9%	22.7%	2.7
Diapers	80,693	23.2	14.5%	23.8%	27.3%	15.4%	2.5
Margarine and Butter	31,970	19.3	16.3%	22.6%	29.0%	22.7%	1.7
Mayonnaise	24,292	17.3	12.0%	26.7%	31.1%	21.6%	2.1
Facial Tissue	28,251	16.3	19.6%	31.7%	38.8%	23.6%	2.1
Toothpaste	122,630	16.0	13.0%	27.6%	34.3%	22.4%	2.7
Shampoos	223,410	14.1	11.4%	21.6%	29.0%	23.0%	3.0
Peanut Butter	21,693	13.6	14.0%	24.6%	34.0%	19.2%	2.2
Mustard and Ketchup	42,846	12.1	9.4%	23.2%	30.5%	19.1%	2.3
Deodorant	213,178	12.0	10.1%	21.3%	28.1%	25.1%	3.1
Blades	58,811	11.4	8.6%	13.3%	18.7%	20.2%	2.3
Household Cleaners	33,456	6.9	13.8%	22.4%	27.7%	19.5%	2.1
Toothbrushes	76,042	5.5	12.6%	25.8%	32.1%	25.1%	3.1
Sugar Substitutes	12,911	5.1	6.8%	9.9%	11.9%	13.1%	1.7
Photo Supplies	14,725	3.2	7.2%	20.0%	23.0%	21.3%	2.7
Razors	16,179	1.5	9.5%	21.2%	22.9%	16.7%	2.6
All Categories	**3,012,670**	**1,380.5**	**16.7%**	**34.2%**	**39.5%**		

Notes: Same as Table 1.3 notes. Column 2 contains the number of store-UPC combinations in the data (larger than the number of UPCs since most UPCs are sold at several stores). These summary statistics aggregate across all of the markets listed in Table 1.1.

Tables 1.3 and 1.4 contain summary statistics by category for the DFF and IRI samples respectively. The table includes the following measures of sales: fraction of weeks on sale, fraction of revenue from sales, fraction of units sold at a sale price, the median discount, and the median quantity ratio. I also provide the number of products (store-products for the IRI sample) and average annual revenue for the entire sample. There are several things worth mentioning about these statistics.

First, although sales account for a small fraction of price quotes, they account for a relatively large fraction of revenue. About 15 percent of price quotes are sales (14 percent for DFF and 17 percent for IRI) while more than a third of revenue is generated from sales (43 percent for the DFF sample and 34 percent for the IRI sample). Two to three times as many units are sold during a sale week compared to a regular price week (as measured by the quantity ratio).

Second, the larger categories tend to have sales more frequently. Figure 1.4.1 illustrates this by plotting annual revenue against fraction of items on sale. There is a positive relationship between a category's share of expenditure and its frequency of sales.

This relationship also holds at the UPC level as well. To show this, I calculate average weekly revenue for each UPC in the sample and separate UPCs into deciles by store. I plot the average fraction of weeks on sale by revenue decile in Figure 1.4.2. There is a strong positive relationship between an items *long run* revenue share and the frequency with which it is on sale. This relationship is consistent with the idea that sales are often used to attract trips to the store (Chevalier et al., 2003; Hosken and Reiffen, 2004; Lal and Matutes, 1994).

Third, the discount is often quite large. In the IRI data, the median discount versus regular price is over 20 percent for 21 of the 29 categories and over 25 percent for six of them. If shoppers are willing to search for deals, hold some inventory, and selectively substitute, substantial savings are available from buying on sale. One study has found that households in the UK save an average of 6.5 percent of annual expenditure by buying on sale(Griffith et al., 2008).

Finally, sales result in far more purchases than a regular price. In the IRI sample 2.6 times as much revenue is generated per sale price than per regular price. The spike in quantity sold during a sale is much higher than can be explained by a simple model of supply and demand (Hendel and Nevo, 2006; **?**; Feenstra and Shapiro, 2003).

To summarize, sale prices occur most frequently on popular items, discounts are often 20% or more, and a disproportionate share of purchases occur at sales prices. These facts indicate that sales are an important determinant of average price paid (unit price) and also nominal consumption expenditure. Chevalier and Kashyap (2011) suggest that average unit price (or more practically, an average of best price and regular price) provides a better picture of the price consumers are paying over time than any individual price series.

Figure 1.4.1: Annual Revenue versus Fraction of Items on Sale by Category

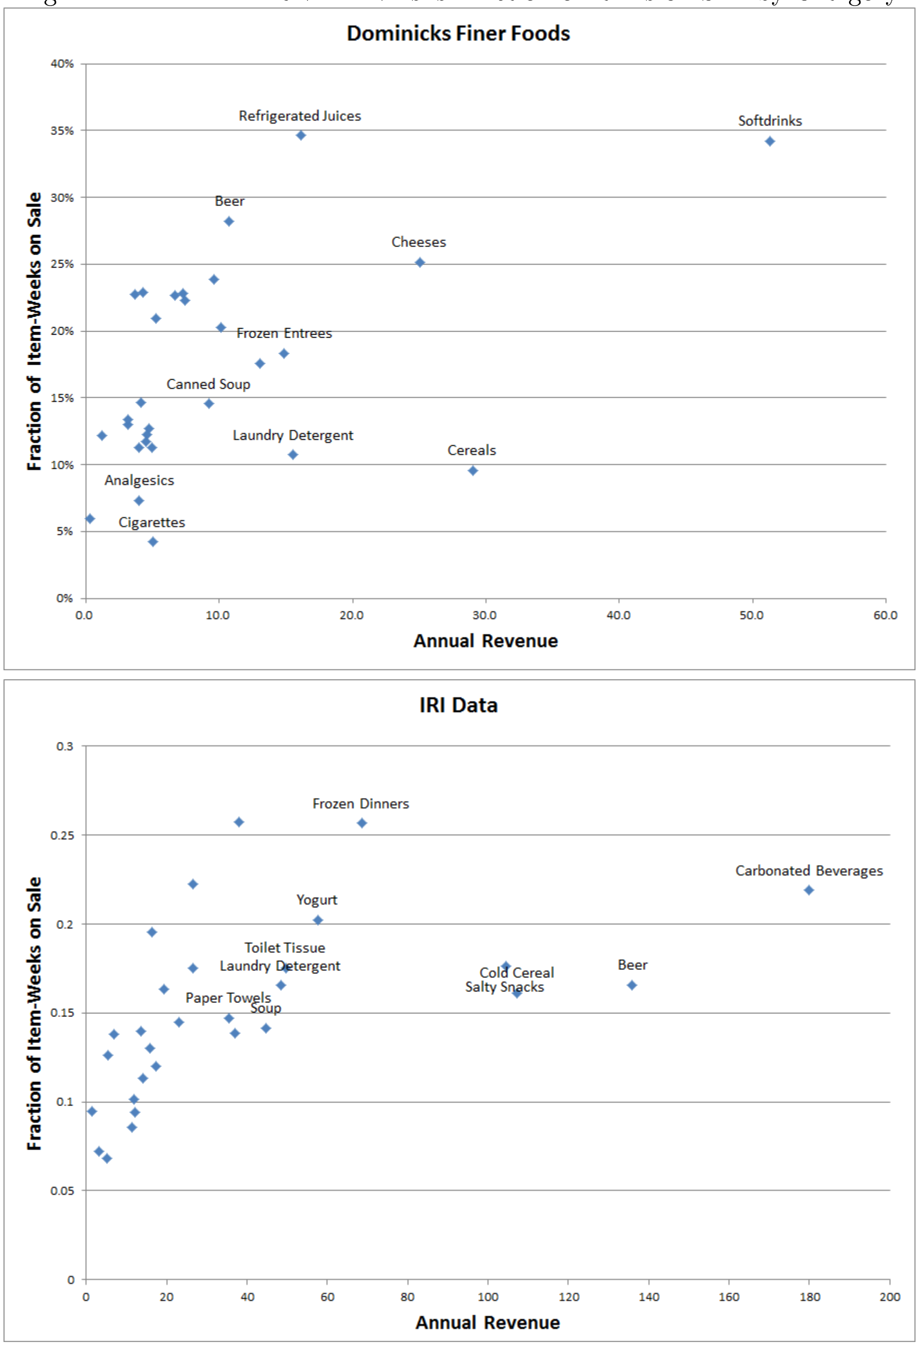

Notes: Each point represents a category. The vertical axis is fraction of item-weeks on sale and the horizontal axis is average annual revenue (as reported in Tables 1.3 and 1.4).

Figure 1.4.2: Fraction of Weeks on Sale by Average Weekly Revenue Decile

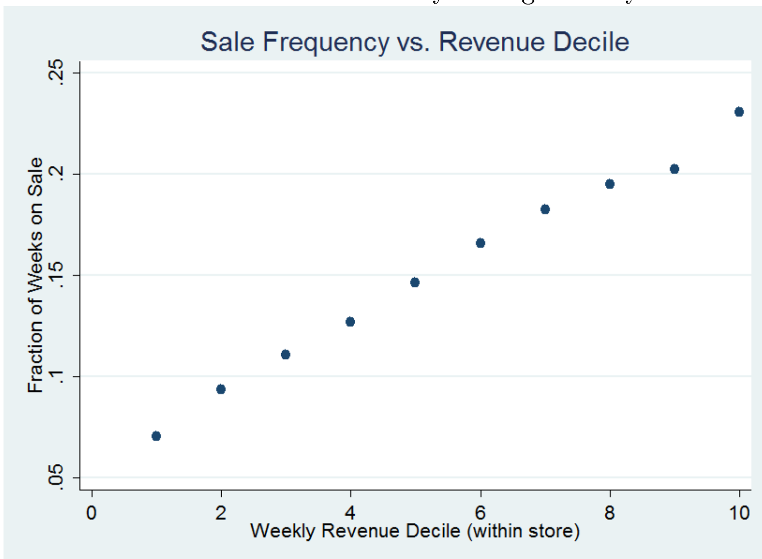

Notes: This figure plots average fraction of weeks on sale by (within store) average weekly revenue decile.

In the following section I will quantify the effect of sales on the dynamics of average unit price.

1.5 Sales and Aggregate Price Adjustment

Most of the macro literature cited above concludes that sales are not a significant source of aggregate price flexibility. A corollary to this claim is that the quantity response to sales and the frequency and depth of sales are static features of the economy. In this section, I test the hypothesis that the characteristics of sales are constant over time. I find that sales do change over time and as a result, average unit price is much more volatile than regular price.

Ultimately we would like to know whether and how the characteristics of sales are related to business cycles and monetary policy shocks. Unfortunately, the data I have cover only six years during which there was one mild recession (2001) so a conclusive study of the behavior of sales business cycles is not possible using these data. Nevertheless, I show that a large increase in the fraction of items on sale coincides with an increase in the unemployment rate.

Unit Price versus Regular Price

I begin by looking at the behavior of sales within each store-product over time. Consider the following expression of the average price per unit sold of store-product i during quarter t:

$$U_{it} = (1 - w_{it})\, R_{it} + w_{it} S_{it}$$

where U_{it}, R_{it}, and S_{it} are, respectively, the average unit price, average regular price, and average sale price during quarter t. The term w_{it} is the quantity sold at a sale price divided by the total quantity sold during the quarter. This equation simply says that average unit price is the weighted average of regular price and sale price. We can rearrange this equation the following way:

$$\frac{U_{it}}{R_{it}} = (1 - w_{it} d_{it})$$

Here, $d_{it} \equiv \frac{R_{it} - S_{it}}{R_{it}}$ is the average sale discount expressed as a percent of the regular price. Note also that we can write $w_{it} = z_{it} f_{it}$, where z_{it} is the quantity response to a sale (the ratio of units sold per sale week to units sold per week) and f_{it} is the fraction of weeks on sale in quarter t. Thus, we have the following expression:

$$1 - \frac{U_{it}}{R_{it}} = z_{it} f_{it} d_{it} \tag{1.5.1}$$

Let us call the LHS of equation 1.5.1 the realized discount. This is the amount saved from buying on sale expressed as a percentage of regular price. The strongest form of the hypothesis that sales are a static feature of the economy predicts that the ratio of unit price to regular price $\frac{U_{it}}{R_{it}}$ should vary little over time. This assumes that stores do not vary their sale plans and consumers, in aggregate, do not vary their sale purchases over time.

I compute the standard deviation of $\frac{U_{it}}{R_{it}}$ across time for each store-UPC in the sample and provide a histogram of these values in Figure 1.5.1. About nine percent of (revenue share weighted) store-UPCs have essentially no variation in the ratio of unit price to regular price. These products happen to be those that are almost *never* on sale ($U_{it} \approx R_{it}$). Most items exhibit substantial quarter to quarter variability. The average (across store-UPCs) of the standard deviation of $\frac{U_{it}}{R_{it}}$ is 7.1 percent. This evidence is inconsistent with the hypothesis that sales are a static feature of the economy.

Of course, it is possible that seasonal variation in sales is the primary source of the variation documented above. To see if this is the case, I repeat the analysis, but calculate the standard deviation conditional on the quarter of the year. The variation captured in the bottom panel of Figure 1.5.1 is the variation across time, holding the quarter of the year constant. The distribution shifts to the left a bit (versus not controlling for seasonal variation), but only a small fraction of the variation in $\frac{U_{it}}{R_{it}}$ is due to seasonal fluctuations in sales.

Figure 1.5.1: Time Variation in $\frac{U_{it}}{R_{it}}$

Empirical Distribution of $St.Dev_t\left(\frac{U_{it}}{R_{it}} | i\right)$

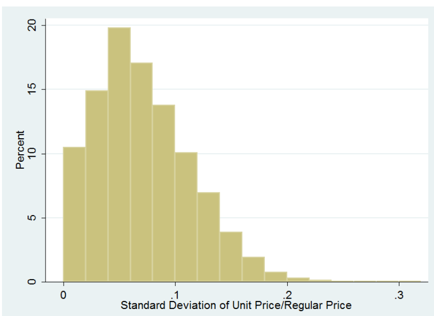

Sample Statistics of $St.Dev_t\left(\frac{U_{it}}{R_{it}} | i\right)$

	All Quarters
Average	7.1%
Median	6.5%
25th Percentile	4.0%
75th Percentile	9.8%
Observations	236,521

Empirical Distribution of $St.Dev_t\left(\frac{U_{it}}{R_{it}} | i, s\right)$

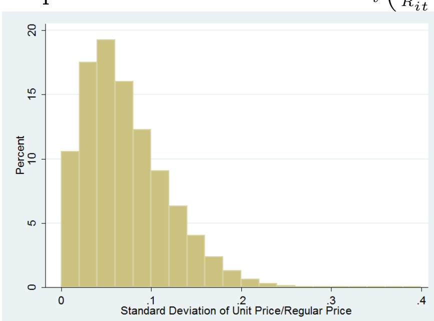

Sample Statistics of $St.Dev_t\left(\frac{U_{it}}{R_{it}} | i, s\right)$

	Q1	Q2	Q3	Q4
Average	7.3%	7.2%	7.1%	7.2%
Median	6.4%	6.3%	6.1%	6.4%
25th Percentile	3.7%	3.7%	3.7%	3.8%
75th Percentile	10.0%	9.9%	9.8%	9.9%
Observations	281,819	275,972	275,816	272,928

Notes: I calculate $\frac{U_{it}}{R_{it}}$, the average unit price divided by the average regular price, for each store-UPC i and quarter t. Next, I compute the standard deviation of this ratio for each store-UPC, $St.Dev_t\left(\frac{U_{it}}{R_{it}} | i\right)$. The graph above is a histogram of the standard deviations with observations weighted by total revenue. The statistics to the right give the average of the standard deviations (the sample analog of $E_i\left[St.Dev_t\left(\frac{U_{it}}{R_{it}} | i\right)\right]$ as well the 25th, 50th, and 75th percentiles. The sample consists of store-UPC cells with observations for all 24 quarters and each quarter must have at least 7 weeks of observations. I exclude products that had no sales. In the bottom panel, I compute the standard deviation conditional on the quarter of the year, s. Thus, I compute four standard deviations for each store-UPC, one for each quarter.

Table 1.5: Sale Variance Decomposition

Source of Variation	Mean	Median
$\left(\frac{1}{\sigma^2_{RD_i}}\right) Var_t\left(\ln z_{it}\right)$	18.7%	7.2%
$\left(\frac{1}{\sigma^2_{RD_i}}\right) Var\left(\ln f_{it}\right)$	44.3%	28.0%
$\left(\frac{1}{\sigma^2_{RD_i}}\right) Var\left(\ln d_{it}\right)$	23.9%	20.8%
$\left(\frac{1}{\sigma^2_{RD_i}}\right) 2\sum_{i\neq j}\sigma_{ij}$	13.1%	24.0%

Notes: I decompose$\sigma^2_{RD_i} \equiv Var\left[\ln\left(1-\frac{U_{it}}{R_{it}}\right)\right]$for each UPC-store with data in all 24 quarters and at least one sale. I present each term as a percent of $\sigma^2_{RD_i}$. The sources of variation are: z_{it}(average units sold per sale week relative to all weeks), d_{it} (average discount versus regular price), and f_{it} (the fraction of weeks on sale in quarter t). I present both the mean and median (across the UPC-stores). A few outliers cause large differences between the mean and median.

I now turn to a description of how the characteristics of sales vary over time. The log of the realized discount is additive in the logs of z_{it}, f_{it}, and d_{it}. Thus, we can decompose variation in the log of the realized discount into variation in the logs of z_{it}, f_{it} , and d_{it}, plus a covariance term:

$$\sigma^2_{\ln RD_i} = \sigma^2_{\ln z_i} + \sigma^2_{\ln f_i} + \sigma^2_{\ln d_i} + covariances \tag{1.5.2}$$

This allows us to see whether variation in $\frac{U_{it}}{R_{it}}$ is the result of firm decisions (variation in f_{it} and d_{it}) or shopper behavior (variation in z_{it}). In Table 1.5, I report mean and median standard deviations expressed as a percent of $\sigma^2_{\ln RD_i} \equiv Var_t\left[\ln\left(1-\frac{U_{it}}{R_{it}}\right)|i\right]$. About two thirds of the variation in the realized discount comes from changes in the frequency and depth of sales.

We have not ruled out the possibility that sales are negatively correlated across products within a quarter, which would result in relatively low aggregate variation in sales. To see if this is the case, I calculate the mean and median of $\frac{U_{it}}{R_{it}}$, z_{it}, f_{it}, and d_{it} (across store-UPCs) for each quarter. I plot the results in figure 1.5.2.

Unit price relative to regular price exhibits a substantial decline over the course of 2002, after which, it remains fairly constant with some seasonal fluctuation. The cause of this change is evidently a large increase in the frequency of sales that occurred over the same time period. In the third quarter of 2003, the median frequency of sales shifted up from twice per quarter to three times per quarter. The average frequency of sales continues to trend upward for the rest of the sample, but the effect on unit price is offset by a steady decline in the quantity response to sales and the average discount.

Figure 1.5.2: The Characteristics of Sales Over Time

Panel A: $\dfrac{U_{it}}{R_{it}}$

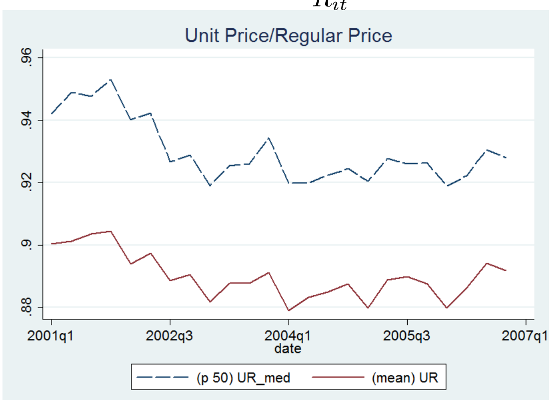

Panel B: f_{it}

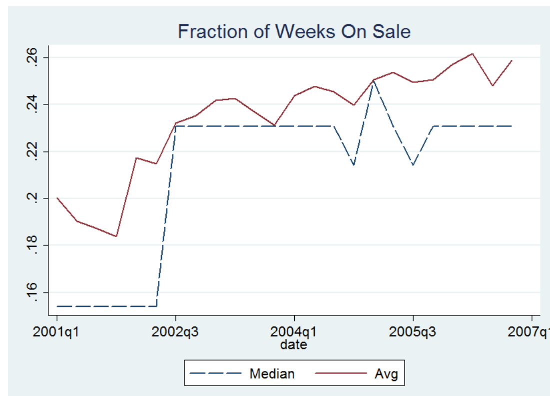

Panel C: d_{it}

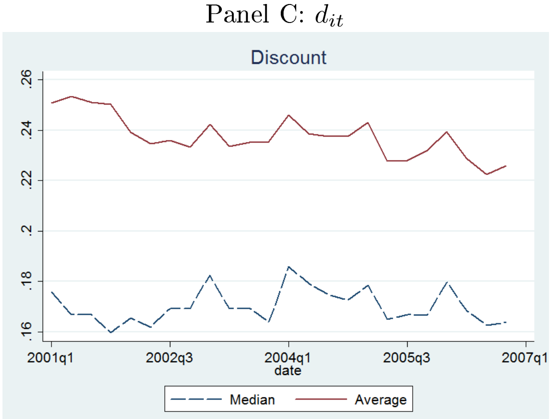

Panel D: z_{it}

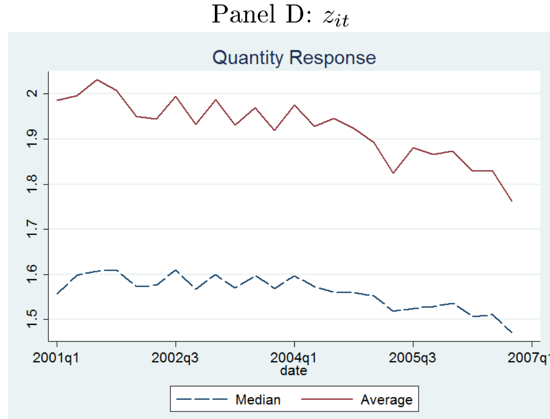

Panel A: Revenue weighted mean and median of $\frac{U_{it}}{R_{it}}$, average unit price divided by average regular price for item i in quarter t. Panel B: Revenue weighted mean and median of f_{it}, the fraction of weeks item i is on sale during quarter t. Panel C: Mean and median of $d_{it} = \frac{R_{it} - S_{it}}{R_{it}}$, the average percent discount as a percent of regular price, *conditional on at least one sale*. Panel D: Mean and Median of the quantity response, z_{it} which is the average units sold on sale relative to average units sold in all weeks.

Sales and Aggregate Price

In this subsection, we will study the impact of sales on the price level by comparing an index of regular price to an index of average unit price. By doing so, we will see how changes in sales and regular price affect average price paid over time.

Price Index Construction

In order to construct indexes of regular price and unit price, I employ the same approach used by the BLS to calculate the consumer price index with one important difference. The BLS samples a single price quote each month. I have weekly price and quantity data and wish to make use of this extra information. Therefore, I take an average of all prices within a month rather than randomly selecting a single price.

For the index of regular price, I take a simple average of regular price for the month (usually there are 4 and they are often all the same). For the index of unit price, I take total revenue divided by units sold to get the average price paid. I then use these monthly average prices the same way the BLS uses a sampled price. Specifically, I calculate a price relative for each month (ratio of current price to last month's price) and then aggregated using a fixed weight geometric average. The details are explained in the appendix.

Figure 1.5.3 plots the two indexes for each of the markets in the sample. Both of the indexes begin with a base value of one in February of 2001.[9] The two indexes usually track each other closely but there are several examples in which they diverge substantially for several years. Gaps between the regular price and the unit price can be attributed entirely to *changes* in one or more of the characteristics of sales.[10] The results in section 4.1, and in particular, the graphs in Figure 1.5.1, indicate that we should expect to see regular price and unit price deviate from each other beginning in 2002. San Francisco and Seattle are the clearest examples of this phenomenon.

Explaining Unit Price Volatility

As the graphs in Figure 1.5.3 indicate, unit price is substantially more volatile than regular price. The standard deviation of the log change of unit price is two to four times that of regular price (see Table 1.5.3). Changes in regular price only explain a fraction of the variation in unit price.

As a descriptive exercise, I project the log difference of the unit price index, r_t^u, onto the log change of the regular price index r_t^r and a vector of monthly dummies $\mathbf{m_t}$ to capture seasonal variation in unit price:

$$r_t^u = \mathbf{m_t}\lambda + \alpha_1 r_t^r + e_t \tag{1.5.3}$$

[9] The first month is dropped due to difficulties in identifying sales and regular price in the first four weeks of the sample.
[10] The Appendix contains a technical explanation for why this is the case.

Figure 1.5.3: Unit Price versus Regular Price

The figures above plot the indexes of regular price and unit price for each market in the sample. The calculation of these indexes is discussed in Section 1.5. The first month plotted is March 2001 and both indexes equal 1 in February 2001.

Table 1.6: OLS Projection: $r_t^u = \mathbf{m_t}\lambda + \alpha_1 r_t^r + e_t$

Weekly Frequency (t indexes weeks)	α_1	R-Squared	$\sigma_{Unit\ Price}$	$\sigma_{Reg\ Price}$
New York	0.42	0.19	1.1%	0.2%
Los Angeles	0.39	0.20	1.1%	0.2%
San Francisco	0.58	0.36	1.8%	0.3%
Dallas	0.36	0.16	1.0%	0.3%
Houston	0.35	0.15	0.9%	0.2%
All Markets	0.54	0.33	0.7%	0.1%

Monthly Frequency (t indexes months).	α_1	R-Squared	$\sigma_{Unit\ Price}$	$\sigma_{Reg\ Price}$
New York	0.16	0.31	1.1%	0.3%
Los Angeles	0.33	0.39	1.2%	0.4%
San Francisco	0.36	0.51	1.7%	0.5%
Dallas	0.40	0.39	1.1%	0.6%
Houston	0.41	0.44	1.0%	0.4%
All Markets	0.51	0.44	0.8%	0.2%

Notes: The top panel uses the log change in a weekly price index and the bottom panel uses log change in a monthly price index. Each row contains the coefficient estimates from an OLS projection of the log difference of unit price (r_t^u) on the log difference of regular price (r_t^r) and a vector of month dummies ($\mathbf{m_t}$). Each market is estimated separately (listed in Column 1). Both $r_t^u$and r_t^r are standardized (demeaned and divided by the standard deviation) prior to the estimation.

The purpose of estimating equation 1.5.3 is to see how much of the variation in r_t^u is explained by seasonal variation (captured by the term $\mathbf{m_t}\lambda$) and variation in r_t^r. In Table 1.6 I provide the OLS estimates of Equation 1.5.3 using weekly and monthly frequencies. The R-squared statistics indicate that between ten and thirty percent of the week-to-week change in aggregate price can be explained by changes in regular price and seasonal variation in sales. Aggregating up to a monthly index increases the explanatory power of regular price but leaves over half of the variation in unit price unaccounted for. Non-seasonal fluctuations in sales (frequency, size, and quantity ratio) explain half or more of the variation in unit price.

The reader may wonder how much of the volatility in unit price disappears if the aggregation period is increased from a month to a quarter. I construct the indexes described above taking quarterly averages of unit price and regular price instead of monthly averages. The results are reported in Table 1.7. Averaging across quarters smooths things out considerably, but non-seasonal variation in sales still accounts for 35 to 50 percent of the quarter-to-quarter variation in unit price.

Sales and Aggregate Price Adjustment

I now compare the two indexes to the (corresponding components of the) CPI-U. I construct the relevant CPI by gathering item indexes from the BLS that correspond to each of the IRI categories. I then average

Table 1.7: Quarterly OLS Projection: $r_t^u = \mathbf{q_t}\lambda + \alpha_1 r_t^r + e_t$

	α_1	R-Squared	$\sigma_{Unit\ Price}$	$\sigma_{Reg\ Price}$
New York	0.31	0.71	1.2%	0.4%
Los Angeles	0.55	0.49	1.5%	0.7%
San Francisco	0.34	0.50	1.5%	0.6%
Dallas	0.80	0.66	1.2%	0.8%
Houston	0.87	0.73	1.4%	0.8%
All Markets	0.80	0.66	0.6%	0.4%

Notes: Each row contains the coefficient estimates from an OLS projection of the log difference of unit price (r_t^u) on the log difference of regular price (r_t^r) and quarterly dummy variables ($\mathbf{q_t}$). Each market is estimated separately (listed in Column 1). All variables are all standardized (demeaned and divided by the standard deviation) prior to the estimation.

across these item indexes using the same weights that I use to aggregate across the IRI categories (revenue share). For the IRI indexes, I include all 50 markets and aggregate across markets using the same weights used by the BLS to construct the CPI-U. I plot all three indexes in figure 1.5.4 to see how they compare.[11]

The CPI, the regular price index, and the unit price index all track each other fairly closely for the first four quarters of the sample, and then diverge for several years.[12] Regular price behaves much more like the CPI than does unit price. The correlation between the CPI and the regular price index is 0.90 while the correlation between the CPI and the unit price index is 0.60. This is to be expected because the CPI uses fixed weights so it does not fully capture the effect of *changes* in sales on unit price.[13] In section 1.5 we found that there was a large increase in the frequency of sales that began in the third quarter of 2001.

The most striking feature of figure 1.5.4 is that unit price falls well below the regular price index (and the CPI) early in the sample period. The gap between regular price and unit price grows to about one percentage point over the first 12 months and then grows to more than three percentage points by the end of 2002.

In the middle panel of figure 1.5.4, I plot the "gap" between regular price and unit price along with the unemployment rate. Vertical gray bars highlight business cycle contractions (peak to trough) as defined by the NBER. Following the contraction of 2001, average unit price fell relative to regular price by about three percentage points. This change is fairly persistent. By the end of 2006, the gap between unit price and regular price remained two percentage points higher than it was at the beginning of 2001.

[11] Details of this calculation can be found in the Appendix

[12] The indexes I calculate will deviate from the CPI for two main reasons other than the fact that I use weekly data. First, I do not deal with new product substitutions as the CPI does because I only include products that were available in all sample periods (though the results are robust to relaxing this requirement). Second, the CPI includes far more outlets than I include in this sample (eg Mass Merchandisers).

[13] For a full discussion of why the CPI under reports the effect of changes in the features of sales, see Appendix B.

Figure 1.5.4: Index Comparison: CPI vs. Regular Price vs. Unit Price

IRI Sample	DFF Sample
CPI, Regular Price Index, Unit Price Index	Regular Price Index Unit Price Index

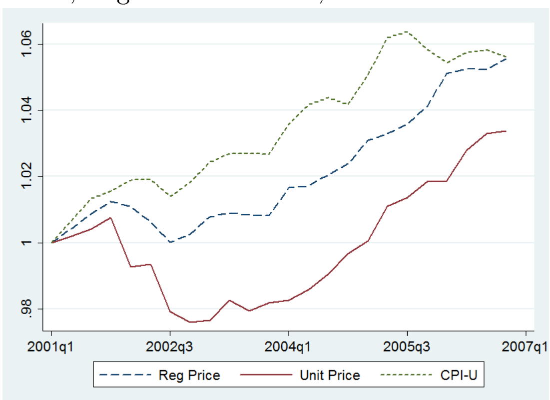

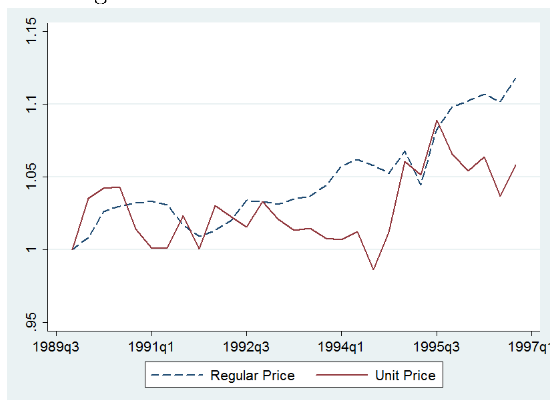

Index Gap: Regular Price - Unit Price

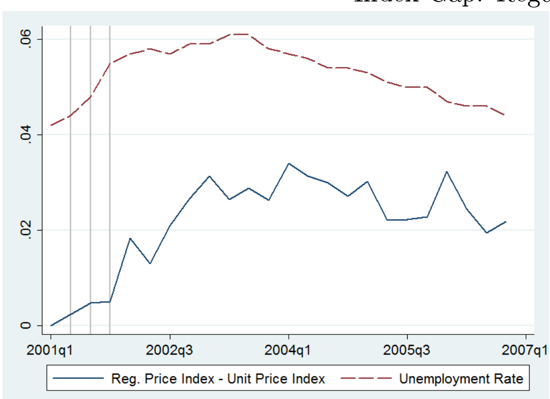

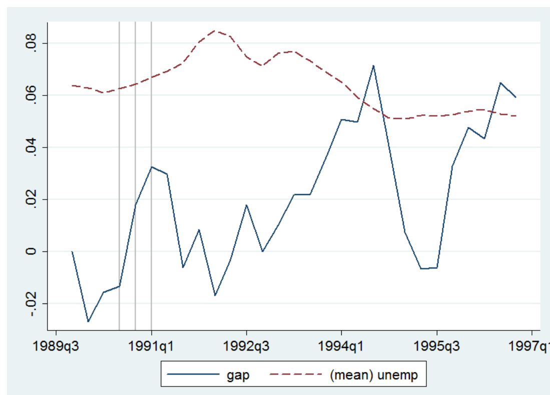

Fraction of Revenue From Sales

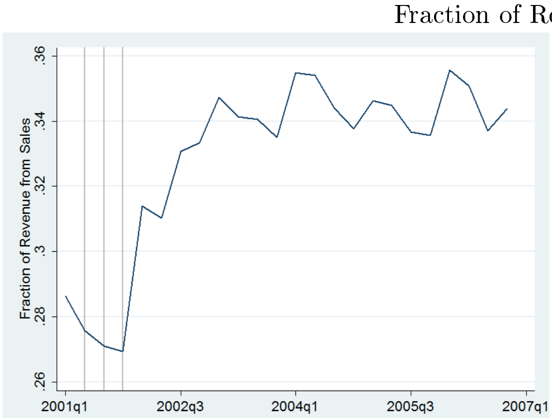

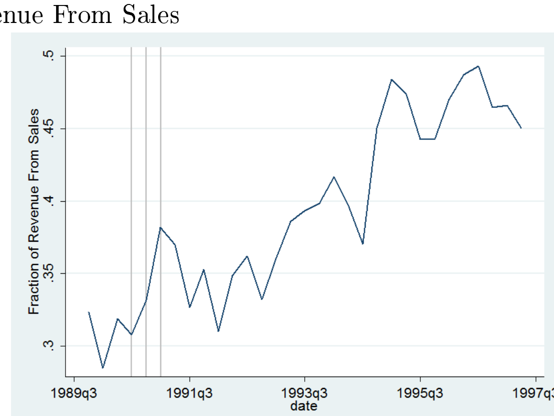

Panel A plots quarterly indexes of unit price and regular price over the full sample (left side is the IRI sample and the right side is the DFF sample). The CPI for comparable items is also included on the IRI chart for comparison. Panel B plots the difference between the two indexes over time. The initial month is equal to zero in both cases because of the way the indexes were calculated. Panel C plots the fraction of revenue from sales for each of the samples. Vertical lines indicate NBER contraction dates.

In the DFF data, the gap between unit price and regular price is much more volatile. This is probably due to changes in the environment that are specific to DFF. For instance, in Chapter 3 I show that DFF appears to have responded to the entry of Wal-Mart by having more frequent sales. Nevertheless, we do see a four percentage point increase in the difference between regular price and unit price during the recession of 1991. This change is almost entirely reversed within 5 quarters.

With such short sample periods, I am unable to make strong statements about the role of sales in aggregate price adjustment. However, this evidence suggests that average unit price exhibits larger and faster responses to downward pressure on prices. In the figures above, we observe two occasions during which there would have been downward pressure on prices (coming from the demand side) and in both cases, unit price fell relative to regular price. However, differences in the timing of these changes relative to contraction dates leaves us with an unclear picture of the role that sales play in aggregate price adjustment.

1.6 Store Response to Reductions in Revenue

In this section, I investigate whether stores use sales to respond to changes in residual demand. The concept of residual demand I have in mind for a grocery store is a vector of functions (one for each product offered) that maps the set of price plans into quantities sold. Events that would cause this mapping to shift include the entry of a competitor, changes in another store's price plan, and changes in the incomes of regular shoppers, just to name a few. I would like to determine whether large shifts in this mapping result in changes to the frequency and size of sales. Unfortunately, I have little hope of estimating such a system of equations (given the available data) so I'll have to come up with an alternative approach.

In the IRI sample, there are several instances in which grocery stores faced large reductions in revenue over a short period of time. I am unable to identify the causes of large revenue changes because very little is known about the stores.[14] I assume that sudden and large changes in revenue are due to shifts in demand and test whether firms respond with changes to the frequency and/or size of sales. My assumption seems reasonable because I find that stores tend to increase the frequency of sales following a large reduction in revenue.

Data

For this analysis, I use data from all chains with balanced panels for all six years in the data. Previous literature (mentioned by EJR) establishes that store level pricing is largely controlled at the chain level and I assume that a chain's price plan would be market specific, so I aggregate the data into a panel of 132

[14] IRI's vendors required confidentiality in order to make the data available for academic research

Figure 1.6.1: Demand Shock Identification Example from the Data

Negative Demand Shock Example	Positive Shock Followed by Negative Shock

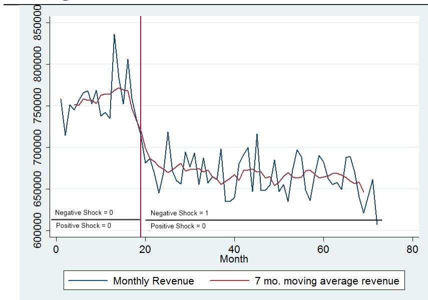

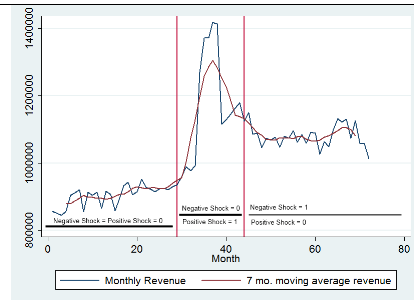

Notes: The graphs above indicate actual examples from the data of (a) a store that appears to have experienced a large negative demand shock, and (b) a store that experienced a large positive demand shock followed by a negative demand shock. The red lines indicate how I identify the timing of changes in the binary demand shock variables used in the regressions.

market-chains (e.g. Chain1-Chicago). The week-to-week and month-to-month variation in all of the relevant variables is quite large so I aggregate to the quarter level in order to reduce measurement error caused by holiday shopping weeks that overlap months (e.g. Fourth of July, Labor Day, and Thanksgiving).

Demand Shock Identification

My objective is to find large and persistent (non-seasonal) shifts in revenue. I compute a two quarter moving average of revenue to smooth out transitory fluctuations. I consider a demand shock to have occurred in the quarter prior to a 10% year-over-year drop in the moving average of revenue. I apply a symmetric approach to identify positive demand shocks.

I construct an indicator variable for each of two shock types. Neg_{it} (Pos_{it}) takes a value of zero prior to the occurrence of a negative (positive) demand shock and then takes a value of one in the quarter in which the shock occurs and every quarter after, unless a shock in the opposite direction occurs later. Figure 1.6.1 provides two actual examples from the data to illustrate how the binary shock variables are determined.

Using the procedure described above, I find that 70 of the 124 chains in the data experience negative demand shocks at some point during the sample period. These shocks occur over the course of 16 quarters. The quarter with the most shocks is the fourth quarter of 2001 with 17 chains experiencing negative demand shocks in this quarter. These 17 chains are not concentrated in a single market. The remaining 53 negative demand shocks are spread out more or less evenly over time. I conclude that these sudden drops in revenue are not due to large changes in marginal costs.

Estimation

My objective is to evaluate whether demand shocks result in changes to sales (temporary price reductions). To do this, I use the fixed effects panel estimator and test several different model specifications. The first two specifications use a pre/post shock indicator variable to estimate effect of a sudden change in revenue on sales. They are both essentially panel difference-in-differences estimators. The third specification uses a continuous measure of the state of demand facing the chain.

With each specification, I consider three different measures of sales as the dependent variable. The first is realized discount ($RealDisc_{it}$) which is the percentage difference between revenue at regular price and actual revenue (and described at length in section 1.5). Realized discount combines all three characteristics of sales (frequency, size, and quantity ratio). The second dependent variable I consider is the fraction of items on sale ($frac_{it}$). This is the revenue weighted average fraction of UPC-weeks on sale. The final dependent variable I consider is the fraction of revenue from sales ($SaleRevFrac_{it}$).

In the first specification, the demand shocks are captured by the binary variables Neg_{it} and Pos_{it}. I estimate the following equation:

$$y_{it} = \lambda_t + \delta_n Neg_{it} + \delta_p Pos_{it} + c_i + e_{it} \tag{1.6.1}$$

where y_t is one of the three dependent variables described above, λ_t controls for aggregate time effects (by including a dummy for each quarter), c_i is the unobserved (chain specific) effect, and e_{it} is idiosyncratic error. The fixed effects estimator is consistent under any correlation structure between c_i and the other covariates, assuming e_{it} is strictly exogenous. The purpose of defining the shock variable as binary is to minimize the possibility that it is correlated with unexplained changes in sales behavior (e_{it}). I report estimates of δ_p and δ_n (the effect of positive and negative demand shocks) along with other relevant estimation results in table 1.8.

Other Specifications

As a robustness check, I estimate a second specification (for each of the three dependent variables described above) which allows for individual specific time trends:

$$\Delta y_{it} = g_i + \eta_t + \delta_n \Delta Neg_{it} + \delta_p \Delta Pos_{it} + \Delta e_{it} \tag{1.6.2}$$

In this specification, g_i is the individual time trend and η_t captures aggregate time effects (via a vector of quarter dummies). The estimates from this specification are more precise and broadly consistent with the original specification.

For the final specification, I abandon the binary classification of shocks. Instead, I use a one quarter lag of the year over year log change in two quarter moving average revenue. To allow for asymmetric effects, I separate this variable into two: one containing positive values and one containing negative values (they are zero otherwise and sum to the total). The advantage of this specification is that it does not require an arbitrary cutoff for determining demand shocks. The equation is the same as 1.6.1 but the shock variables are defined as described above.

The results indicate that realized discount responds to negative demand shocks. Following a 10 percent or larger decline in revenue, grocery chains increased the fraction of items on sale by an average of 2.3 percentage points. The fraction of reveue from sale increased by an average of 4.9 percentage points and the realized discount increased by an average of 1.2 percentage points. It is not clear whether the effect is symmetric or not. Specification 1 finds no effect of a positive demand shock on measures of sales. On the other hand, specifications 2 and 3 indicate that there is. This evidence suggests that multiproduct retailers use sales to respond to chain-wide demand shocks. In Chapter 3, I show that DFF increased the fraction of items on sale following the entry of Wal-Mart.

Impulse Response Functions

Another way to describe the dynamics of revenue and sales is to estimate a system of equations using Vector Autoregression. For this exercise, I consider three endogenous variables: $\Delta \ln UnitPrice = r_t^u$, $\Delta \ln RegPrice = r_t^r$, and $\Delta \ln(Revenue)$. The model I have in mind is one in which the steady state is characterized by constant revenue growth and a constant ratio of unit price to regular price. One way to respond to a demand shock is to change the regular price (and hold the unit price/regular price ratio constant). Another way would be to adjust the frequency and/or size of sales while leaving the regular price unchanged. To see what grocery chains actually do, I estimate the following VAR model and present the impulse response functions (IRFs).

$$\mathbf{y}_t = \sum_{l=1}^{p} \mathbf{A}_l \mathbf{y}_{t-l} + \mathbf{e}_t$$

$$\mathbf{y}_t \equiv \begin{pmatrix} \Delta \ln Rev \\ r_t^r \\ r_t^u \end{pmatrix}$$

To estimate the model, I stack the panels (24 quarters of data for each of 100 grocery chains) into

Table 1.8: Panel Regressions

	Specification 1: $y_{it} = \lambda_t + \delta_n Neg_{it} + \delta_p Pos_{it} + c_i + e_{it}$			Specification 2: $\Delta y_{it} = \eta_t + \delta_n \Delta Neg_{it} + \delta_p \Delta Pos_{it} + g_i + \Delta e_{it}$			Specification 3: (see notes)		
Dependent Variable y_{it}:	$RealDisc_{it}$	$Frac_{it}$	$RevFrac_{it}$	$\Delta RealDisc_{it}$	$\Delta Frac_{it}$	$\Delta RevFrac_{it}$	$RealDisc_{it}$	$Frac_{it}$	$RevFrac_{it}$
δ_N	0.012	0.023	0.049	0.011	0.016	0.036	-0.048	-0.016	-0.083
p-value	*0.000*	*0.000*	*0.000*	*0.002*	*0.001*	*0.000*	*0.000*	*0.336*	*0.001*
δ_P	-0.005	-0.003	-0.003	-0.011	-0.012	-0.025	-0.070	-0.061	-0.147
p-value	*0.298*	*0.618*	*0.746*	*0.009*	*0.015*	*0.004*	*0.038*	*0.012*	*0.024*
Aggregate Time Effects?	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Unit Specific Trend?				Yes	Yes	Yes			
Observations	3,168	3,168	3,168	3,036	3,036	3,036	2,375	2,375	2,375
Groups	132	132	132	132	132	132	125	125	125
R-Squared within	0.24	0.25	0.32	0.24	0.16	0.18	0.11	0.10	0.13

Notes: Each of three model specifications is estimated using three different dependent variables (total of 9 regressions) which are listed as the column header. I estimate all three specifications using the fixed effects estimator. Specifications 1 and 2 use a binary indicator of a demand shock described in the text. Specification 3 measures the state of demand as the one quarter lag of the year-over-year log change in two quarter moving average revenue, separated into two variables to allow for asymmetric effects.

Figure 1.6.2: Impulse Response Functions: Revenue, and Unit Price

Cumulative Response of r^u to $\Delta \ln Rev$ | Cumulative Response of r^r to $\Delta \ln Rev$

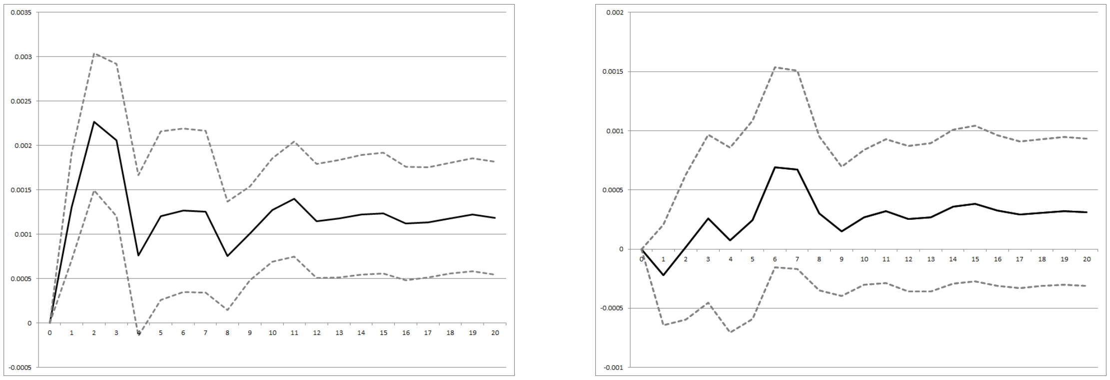

Notes: The figures above plot the response of $RealDisc$ and $\Delta \ln Rev$ to an orthogonalized impulse to $\Delta \ln Rev$. The Cholesky ordering is r^r, r^u $\Delta \ln Rev$.

a single time series. The idea is to treat each panel as a sample taken at intervals in time separated by several quarters (so that lags don't overlap different panels). I include time dummies as exogenous variables to control for nationwide cost shifts and seasonal fluctuations. I also include a vector of chain dummies to control for chain specific fixed effects. Lag length of seven quarters provides the best fit according to the SBIC.[15]

Figure 1.6.2 displays the cumulative response of the growth rates of unit price and regular price to an orthogonalized impulse to revenue growth. The Cholesky ordering of the endogenous variables is r_t^r, r_t^u, $\Delta \ln Rev_t$. In other words, I assume that neither regular price nor unit price respond to contemporaneous revenue growth shocks.

Looking at the IRFs, we can see that the contemporaneous effect of a revenue growth shock on unit price is to reduce it. This is because the contemporaneous correlation of the error terms is negative. In the following four quarters, unit price rises while regular price is mostly flat. That is, unit price rises after a positive shock to demand, but regular price does not. After about four quarters, regular price rises slightly (though not significantly) and unit price falls slightly. In the long run, the ratio of unit price to regular price returns to its original level. This pattern suggests that grocery chains adjust the frequency and size of sales as an early response to demand shocks.

[15] Both the HQIC and the SBIC agree that 6 lags is the best fit for the model. However, the AIC and BIC select a lag length of 11 quarters. Since each panel is only 23 quarters, I go with the shorter lag length. The results using either are qualitatively the same.

1.7 Synchronization of Sales

In this section I investigate the extent to which sales are synchronized. In the spirit of Nakamura (2008), we may be able to tell something about why sales typically occur by investigating the extent to which they are synchronized.[16] My results are consistent with Nakamura's and also Chevalier and Kashyap (2011) who find that sales are mainly retailer-product specific events that are staggered across stores and items.

This is important because many of the menu cost models referenced thus far model sales in a way that would imply synchronization.[17] In these models, sales are an inexpensive alternative for firms to react to large deviations between the frictionless optimal price and the inherited price. In Kehoe and Midrigan (2010), firms have a sale when a large enough gap opens between the inherited price and the frictionless price. Similarly, firms in Eichenbaum et al. (2011) can switch between two inherited prices for free or pay the menu cost to chose a different price plan. In either case, most of the week-to-week volatility in an individual price series results from large, temporary, idiosyncratic cost shocks.

Several papers have noted that sales on specific products tend to occur when demand for them is predictably *high*, not low (Chevalier et al., 2003; Warner and Barsky, 1995; DeGraba, 2006). The explanations for this phenomenon typically involve strategic behavior between competing sellers or between buyers and sellers. Second, if a sale is due to a temporary drop in marginal cost, then we should observe synchronization in the timing of sales. Below I explain why product specific cost/demand shocks would result in three types of synchronization: 1) across stores for a particular product and 2) across products for a particular manufacturer, and 3) across close substitutes within a store.

Across Store Synchronization

Most products (defined as a UPC) are sold at several different stores. A large portion of the marginal cost of retail goods is the acquisition price, which is more or less common to all retailers. Thus, one would expect that if a product is on sale in one store because of a temporary reduction in the wholesale price, then it ought to be on sale at several other stores for the same reason.

To measure this type of synchronization, I calculate the fraction of stores having a sale on product j in each week t ($frac_{j,t}$). I only consider UPC's that are sold in at least one third of the stores in the sample. Table 1.9 contains the 25th percentile, median, and 75th percentile of $frac_{j,t}$ conditional on it being

[16] Nakamura (2008) decomposes retail price variation in order to see how much can plausibly be explained by wholesale price changes. She finds that only a small portion of retail price variation appears to be due to changes in acquisition price.

[17] My interpretation of these models is that each commodity produced in the economy corresponds to a UPC (rather than a store-UPC). Thus, cost shocks that get passed through to price should occur simultaneously across many locations. If the correct interpretation is that a commodity corresponds to a store-UPC, then it is hard to imagine what would cause large, transitory, and frequent cost shocks that are specific to an item within a store.

larger than zero, for each category. If sales were perfectly synchronized, then the fraction of stores having a sale should be either 0 or 1. If sales are perfectly staggered then the fraction of stores holding a sale would always be equal to the probability of a sale. To get a clearer picture of synchronization across stores, I also provide histograms of $frac_{j,t}$ in Figure 1.7.1.

Some amount of synchronization is evident, particularly for carbonated beverages where it is not uncommon for a single product to be on sale in at least 40% of stores. However, it is extremely rare to find that a product is on sale in more than half of stores, regardless of the category. It is almost never the case that a single item is on sale in 75% or more of the stores. There is evidence that stores tend to put products on sale in the week following manufacturer coupon drops (Nevo and Wolfram, 2002) but this type of synchronization has little to do with idiosyncratic demand and supply shocks.

Across Product within Manufacturer Synchronization

Marginal cost driven sales should also result in synchronization across products from a particular manufacturer. The marginal cost of producing different sizes, flavors, scents, etc. of the same brand should be highly correlated. I measure synchronization within manufacturer by calculating the fraction of items on sale within each vendor-store-week combination. From the set of vendor-store-week combinations, I keep only those which contain at least 5 different items, one of which was on sale. I report the median of this measure for each category in 1.9.

For some of the categories, a certain amount of synchronization within a manufacturer is evident. This may be due to the fact that manufacturers often arrange for a coordinated promotional events of several items in their lineup. Without cost data from a manufacturer, we cannot distinguish between the two explanations. These results do not support the idea that sales are the result of temporary reductions in the marginal cost of production.[18]

Across Products within Store-Category Synchronization

The final type of synchronization I check for is within a category-store. If sales are responses to demand and/or supply shocks, then we should see some synchronization of sales across close substitutes. Since these models do not distinguish between purchase and consumption, a low price should induce lower prices for close substitutes.

I find that sales tend to be staggered across products within a store-category. The fraction of items on sale in a given store-category-week is usually below 25% and almost never above 40%. This result is

[18] It may be that sales are the result of temporary reductions in the wholesale price of items sold by grocery stores. However, US anti-trust law discourages manufacturers from charging their customers different prices.

consistent with Chevalier and Kashyap (2011) and Lal (1990). The literature has noted exceptions to this phenomenon during weeks of predictably high demand such as eggs in the week before Easter, and tuna during Lent (Hosken and Reiffen, 2004; Chevalier et al., 2003). But these exceptions work in the opposite direction that standard models predict.

1.8 Concluding Discussion

In the analysis above, I provide evidence that sales have a large influence on the average price paid and the rate at which it changes over time. Existing macro studies of sales emphasize the stickiness of regular price. I have shown that indexes of unit price and regular price can diverge substantially (up to 4 index points) and persistently.

Sales are generally staggered (not synchronized) across stores and products. This fact along with evidence from other work indicates that modeling sales as the result of large, transient, and idiosyncratic cost shocks probably misses the true motivation for most sales. Alternative motivations for sales such as strategic behavior (Guimaraes and Sheedy, 2011) or inventory management (Aguirregabiria, 1999; Campbell and Eden, 2005; Eden and Jaremski, 2010) seem to be a more promising route for modeling sales in the macro economy. This distinction may be important because there appears to be cyclical variation in the amount shoppers save by buying on sale. At this stage it seems premature for macro economists to ignore sales.

I have also shown that the CPI is less volatile than an index of unit price. One task often fulfilled by the CPI is to deflate nominal magnitudes into real magnitudes. Recommending changes to the CPI methodology falls well beyond the scope of this paper, but my results raise two important questions.[19] Does the CPI do a reasonable job of deflating nominal magnitudes into real magnitudes? and 2) How accurately are we measuring inflation over the course of the business cycle? The characteristics of sales change over time and tend to affect average unit price much more so than the CPI reflects. At a minimum, these results suggest that we consider carefully the ways in which we use the CPI.

[19]For a lengthy discussion of scanner data and the CPI, see the book that contains the article by Feenstra and Shapiro (2003).

Table 1.9: Synchronization of Sales

All measures are conditional on at least one item being on sale

Category	Fraction of Stores having a Sale (Observation = UPC-Week)			Fraction Manufacturer's Items on Sale (Observation = Mfg-Store-Week)			Fraction of category on sale (Observation = store-week)			Average Fraction of Items on Sale
	25th Pctle	Median	75th Pctle	25th Pctle	Median	75th Pctle	25th Pctle	Median	75th Pctle	
Beer	8%	15%	26%	9%	17%	31%	9%	14%	20%	16%
Blades	2%	6%	12%	3%	7%	15%	4%	7%	11%	8%
Carbonated Beverages	11%	24%	37%	10%	21%	38%	15%	21%	27%	22%
Cigarettes	1%	2%	4%	1%	2%	5%	1%	2%	4%	4%
Coffee	7%	13%	21%	8%	17%	33%	7%	12%	17%	14%
Cold Cereal	6%	13%	23%	7%	15%	30%	9%	14%	20%	16%
Deodorant	4%	8%	13%	4%	8%	14%	4%	9%	14%	10%
Diapers	6%	11%	18%	6%	11%	20%	7%	13%	19%	13%
Facial Tissue	8%	16%	28%	7%	16%	30%	10%	19%	30%	19%
Frozen Dinners	15%	25%	36%	13%	23%	35%	15%	24%	33%	26%
Frozen Pizza	16%	25%	34%	14%	26%	40%	15%	24%	33%	26%
Household Cleaners	6%	12%	21%	7%	15%	32%	7%	13%	21%	13%
Hot dogs	13%	22%	32%	14%	26%	46%	13%	21%	30%	23%
Laundry Detergent	8%	15%	23%	8%	15%	27%	9%	14%	21%	16%
Margarine and Butter	8%	14%	23%	7%	15%	27%	8%	15%	25%	16%
Mayonnaise	3%	10%	23%	5%	14%	30%	7%	12%	18%	12%
Milk	5%	11%	19%	8%	20%	40%	6%	10%	16%	12%
Mustard and Ketchup	2%	7%	15%	4%	10%	22%	5%	8%	12%	9%
Paper Towels	5%	11%	19%	7%	16%	34%	7%	13%	20%	14%
Peanut Butter	6%	13%	23%	6%	14%	26%	8%	14%	22%	13%
Photo Supplies	1%	4%	10%	2%	5%	11%	5%	7%	11%	7%
Razors	2%	6%	12%	3%	7%	14%	7%	11%	17%	9%
Salty Snacks	7%	16%	30%	9%	19%	36%	11%	16%	21%	17%
Shampoos	4%	8%	15%	4%	9%	16%	5%	8%	13%	11%
Soup	4%	11%	24%	5%	12%	26%	5%	12%	20%	14%
Spaghetti Sauce	10%	18%	28%	9%	18%	33%	10%	17%	25%	18%
Sugar Substitutes	1%	4%	10%	2%	6%	15%	4%	8%	13%	6%
Toilette Tissue	8%	15%	25%	7%	15%	28%	9%	16%	24%	17%
Toothbrushes	4%	9%	17%	4%	9%	18%	5%	10%	17%	13%
Toothpaste	4%	10%	18%	4%	10%	18%	6%	11%	17%	12%
Yogurt	10%	19%	28%	11%	22%	38%	11%	19%	28%	20%

Notes: This table presents the 25th, 50th and 75th percentiles for fraction of items on sale in a given week conditional on the fraction being larger than 0 (at least one item on sale). In columns 2 through 4, the unit of observation is a UPC-Week and the measure is the fraction of stores with the UPC on sale in a given week. In columns 5 through 7, the unit of observation is a Manufacturer-Store-Week and the measure is fraction of the manufacturer's items on sale. In Columns 8 through 10 the unit of observation is a store-week and the measure is fraction of items on sale. The last column provides the category average fraction of items on sale.

Figure 1.7.1: Fraction of Stores With Item j on Sale in Week t

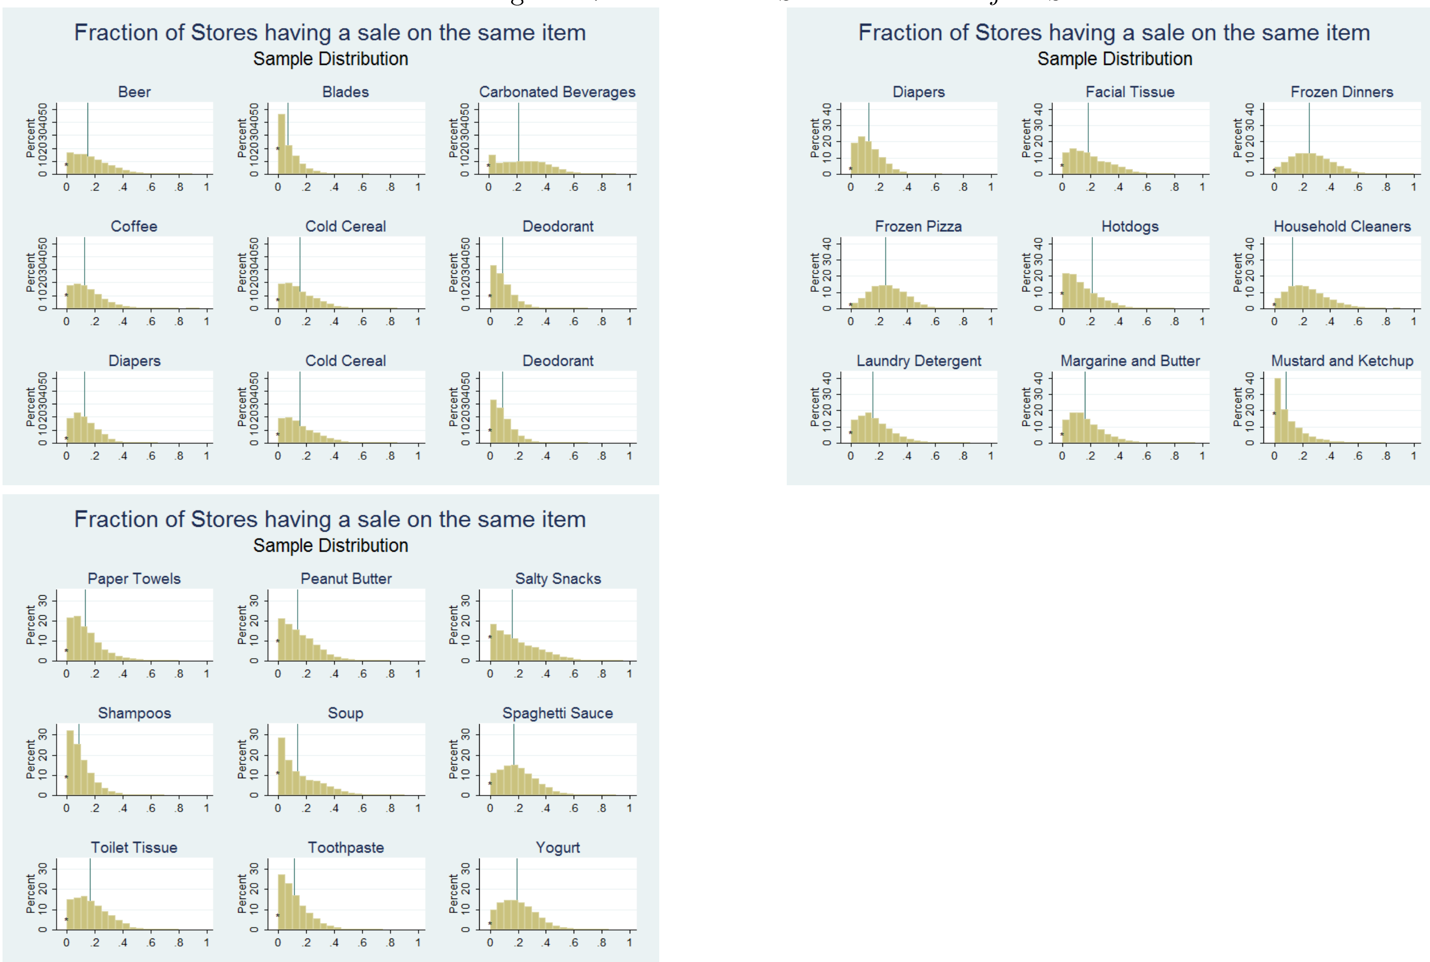

Notes: The figures above show the distribution of the fraction of stores with UPC j on sale during week t. Each observation is a UPC-Week. I only include UPCs that are sold in at least 40 percent of the stores in the sample.

REFERENCES

Aguirregabiria, Victor, "The Dynamics of Markups and Inventories in Retailing Firms," *Review of Economic Studies*, April 1999, *66* (2), 275–308.

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

Blattberg, Robert C., Richard Briesch, and Edward J. Fox, "How Promotions Work," *Marketing Science*, 1995, *14* (3), G122–32.

Bronnenberg, Bart J., Michael W. Kruger, and Carl F. Mela, "The IRI Marketing Data Set," *Marketing Science*, 2008, *27* (4), 745–748.

Campbell, Jeffrey R. and Benjamin Eden, "Rigid Prices: Evidence from U.S. Scanner Data," 2005.

Chevalier, Judith A. and Anil K. Kashyap, "Best Prices," 2011.

_ , _ , and Peter E. Rossi, "Why Don't Prices Rise during Periods of Peak Demand? Evidence from Scanner Data," *American Economic Review*, March 2003, *93* (1), 15–37.

DeGraba, Patrick, "The Loss Leader Is a Turkey: Targeted Discounts from Multi-product Competitors," *International Journal of Industrial Organization*, 2006, *24* (3), 613–628.

Eden, B. and M. Jaremski, "Rigidity, Dispersion and Discreteness in Chain Prices," *Working Papers*, 2010.

Eichenbaum, Martin, Nir Jaimovich, and Sergio Rebelo, "Reference Prices, Costs, and Nominal Rigidities," *American Economic Review*, 2011, *101* (1), 234–262.

Feenstra, R. C and M. D Shapiro, "High-Frequency Substitution and the Measurement of Price Indexes," *Scanner Data and Price Indexes*, 2003, *64*, 123–146.

Griffith, Rachel, Ephraim Leibtag, Andrew Leicester, and Aviv Nevo, "Timing and Quantity of Consumer Purchases and the Consumer Price Index," *SSRN eLibrary*, October 2008.

Guimaraes, Bernardo and Kevin D. Sheedy, "Sales and Monetary Policy," *American Economic Review*, 2011, *101* (2), 844–876.

Hendel, Igal and Aviv Nevo, "Measuring the Implications of Sales and Consumer Inventory Behavior," *Econometrica*, November 2006, *74* (6), 1637–1673.

Hosken, Daniel and David Reiffen, "How Retailers Determine Which Products Should Go on Sale: Evidence from Store-Level Data," *Journal of Consumer Policy*, 2004, *27* (2), 141–177.

Kehoe, Patrick and Virgiliu Midrigan, "Prices Are Sticky After All," *NBER Working Paper No. w16364*, 2010.

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

Lal, Rajiv, "Price Promotions: Limiting Competitive Encroachment," *Marketing Science*, 1990, *9* (3), 247–262.

_ **and Carmen Matutes**, "Retail Pricing and Advertising Strategies," *Journal of Business*, 1994, *67* (3), 345–370.

Nakamura, Emi, "Pass-Through in Retail and Wholesale," *American Economic Review*, 2008, *98* (2), 430–437.

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

Nevo, Aviv and Catherine Wolfram, "Why Do Manufacturers Issue Coupons? An Empirical Analysis of Breakfast Cereals," *RAND Journal of Economics*, 2002, *33* (2), 319–339.

Stigler, G. J and J. K Kindahl, "The Behavior of Industrial Prices," *NBER Books*, 1970.

Warner, Elizabeth J. and Robert B. Barsky, "The Timing and Magnitude of Retail Store Markdowns: Evidence from Weekends and Holidays," *Quarterly Journal of Economics*, May 1995, *110* (2), 321–352.

APPENDIX

1.A Regular Price Filter and Sale Definition Details

I use my own regular price filter because most of the other filters used in the literature miss certain types of sale episodes. The simple algorithms used in the literature that analyzes monthly BLS data do not work well when applied to weekly data because sales often last for several weeks during which the price changes.

EJR and KM work with concepts similar to regular price. However, both of the algorithms used in these papers often classify (what retailers and shoppers would consider) a sale price as the "reference price" or the "list price" when a sale price is the most frequently observed price over certain intervals. EJR simply define the reference price as the modal price for the quarter. KM use a more complicated algorithm for determining the "list" price that does not restrict the frequency with which such a price can change. However, I find that this algorithm will also select what appears to be the sale price as the "list" price. I wish to emphasize that neither EJR nor KM claim to be identifying the "regular price" as I have defined it above.

To get an idea of the difference between the two filters, I redo table 1.4 using the Kehoe-Midrigan filter to identify the regular price. The results are presented in Table 1.10. In general, the Kehoe-Midrigan price filter tends to find fewer sales than the regular price filter that I use.

Table 1.10: Sales Using Kehoe-Midrigan Regular Price Filter

Category	Sales Fraction of Total			Median Discount	Median Quantity Ratio
	Weeks	Revenue	Units		
Carbonated Beverages	20.2%	38.9%	38.6%	20.6%	1.8
Milk	15.2%	15.2%	16.6%	14.9%	1.3
Beer	13.9%	29.7%	25.0%	11.3%	1.7
Cold Cereal	17.1%	32.7%	41.6%	29.1%	2.6
Salty Snacks	19.2%	29.8%	32.1%	20.2%	1.7
Cigarettes	8.2%	9.0%	8.0%	12.5%	2.1
Frozen Dinners	22.7%	37.7%	41.2%	25.4%	2.2
Yogurt	21.3%	27.9%	37.7%	22.6%	1.8
Toilet Tissue	19.3%	39.8%	34.6%	23.1%	2.3
Laundry Detergent	15.9%	34.3%	42.7%	23.1%	2.7
Soup	14.1%	28.4%	33.9%	26.7%	2.6
Frozen Pizza	22.5%	36.8%	41.8%	22.2%	2.2
Coffee	13.4%	30.0%	36.6%	19.3%	2.5
Paper Towels	16.6%	37.2%	31.2%	21.9%	2.0
Hot dogs	22.7%	35.8%	45.0%	28.8%	2.2
Spaghetti Sauce	16.2%	30.4%	37.3%	22.2%	2.5
Diapers	14.9%	24.3%	26.7%	15.0%	2.5
Margarine and Butter	16.4%	22.1%	27.6%	21.5%	1.6
Mayonnaise	12.7%	27.3%	30.8%	20.8%	1.8
Facial Tissue	20.3%	30.7%	35.5%	23.1%	1.7
Toothpaste	14.2%	29.5%	35.3%	22.2%	2.7
Shampoos	13.3%	26.2%	31.6%	22.3%	3.3
Peanut Butter	13.9%	23.6%	32.5%	19.1%	1.9
Mustard and Ketchup	9.9%	21.8%	27.6%	18.3%	2.2
Deodorant	11.3%	23.5%	28.8%	24.9%	3.2
Blades	10.8%	16.1%	20.9%	20.2%	2.4
Household Cleaners	14.3%	22.9%	27.3%	18.5%	1.9
Toothbrushes	13.3%	27.2%	31.5%	24.6%	3.0
Sugar Substitutes	7.6%	10.4%	11.8%	11.9%	1.7
Photo Supplies	9.8%	21.3%	23.0%	21.0%	2.8
Razors	13.1%	29.2%	30.2%	16.2%	3.2
All Categories	**16.2%**	**28.9%**	**33.3%**		

Notes: This table is the same as Table 1.4 but it uses the Kehoe-Midrigan regular price filter to determine regular price. A sale is any price that is at least 5% below the regular price. These summary statistics aggregate across all of the markets listed in Table 1.1.

1.B Price Index Discussion

The point of this appendix is to show how the CPI does not fully reflect changes in unit price due to changes in sales. I begin with a brief description of how the BLS computes the CPI. In the simplest terms, the BLS collects a sample of prices each month and then aggregates them in two steps. The first step takes sampled prices and aggregates them into basic indexes which are specific to an item and a geographic area. An example of a basic index is salad dressing in Chicago-Gary-Kenosha. The second step of aggregation is to take averages of subsets of the basic indexes to form the various price indexes published by the BLS (for example, the CPI-U).

We focus on how sampled prices are aggregated into a basic index (the stage at which quantity sold could be incorporated). First, the BLS chooses several store-items whose prices will be recorded on a monthly basis. A 20 oz. box of Cheerios from the Dominick's Finer Foods on Lincoln Ave. is an example of a store-item. These prices are then aggregated into a price relative for each area-item-month. An area-item price relative is (in most cases) the expenditure share weighted geometric average of the ratio of adjacent period prices. The formula for a price relative of item i in area a between months t and $t-1$ is presented in (1.B.1) below. A basic index for period T is then formed by chaining the price relatives between the base period and period T.

$$ {}_{a,i}R_{[t,t-1]} = \prod_{j \in \{a \times i\}} \left(\frac{P_{j,t}}{P_{j,t-1}} \right)^{w_j} \tag{1.B.1} $$

The point I wish to emphasize is that the weights w_i are not adjusted to reflect the fraction of transactions that occurred at price $P_{j,t}$. This means that adjustments in average unit price due to high frequency substitution are not reflected in the CPI. Since a price index is primarily used to measure the *change* in the price level, fixing weights is not an issue if the characteristics of sales are static. However, if the characteristics of sales change over time, then it is unclear that a fixed weight geometric average will accurately reflect changes in the price level due to sales.

The following numerical example will help to illustrate this point. Let us suppose that instead of sampling a single price per month for a particular store-item, we are able to collect price and quantity data for each of 4 weeks in the month. This additional data will require another level of aggregation (assuming we wish to apply the same basic approach that the BLS currently takes). There are many possible approaches to aggregating scanner data into price indexes and this topic is addressed thoroughly in Feenstra and Shapiro (2003) and the citations therein.

For this example, I simply wish to highlight the effect of changes in the characteristics of sales on average price paid. I compare the growth rates of average menu price and average unit price under three

Table 1.11: Numerical Example of Change in Simple Average vs. Weighted Average

	Panel A Change in Frequency			Panel B Change in Size			Panel C Change in Importance		
Price	Mo. 1	Mo. 2	Mo. 3	Mo. 1	Mo. 2	Mo. 3	Mo. 1	Mo. 2	Mo. 3
Week 1	.90	.90	1.00	.90	.80	.90	.90	.90	.90
Week 2	1.00	1.00	1.00	1.00	1.00	1.00	1.00	1.00	1.00
Week 3	1.00	.90	1.00	1.00	1.00	1.00	1.00	1.00	1.00
Week 4	1.00	1.00	1.00	1.00	1.00	1.00	1.00	1.00	1.00
Quantity Sold									
Week 1	240	133	240	240	240	240	240	342	240
Week 2	120	67	120	120	120	120	120	86	120
Week 3	120	133	120	120	120	120	120	86	120
Week 4	120	67	120	120	120	120	120	86	120
Frequency of Sales	25%	50%	25%	25%	25%	25%	25%	25%	25%
Size of Sales	10%	10%	10%	10%	20%	10%	10%	10%	10%
Quantity Ratio	2.0	2.0	2.0	2.0	2.0	2.0	2.0	4.0	2.0
Average menu price	0.98	0.95	0.98	0.98	0.95	0.98	0.98	0.98	0.98
Inflation (monthly)		**-2.6%**	**2.6%**		**-2.6%**	**2.6%**		**0.0%**	**0.0%**
Average unit price	0.96	0.93	0.96	0.96	0.92	0.96	0.96	0.94	0.96
Inflation (monthly)		**-2.8%**	**2.9%**		**-4.2%**	**4.3%**		**-1.8%**	**1.6%**
Difference		**-0.2%**	**0.2%**		**-1.6%**	**1.7%**		**-1.8%**	**1.8%**

different scenarios in which one of the characteristics of sales changes while the other two are fixed. The results presented in Table 1.11 show that an index of average menu price and average unit price differ when any of the following three characteristics of sales *changes*:

1. Frequency (measured by fraction of items on sale per week)

2. Size (percentage discount off of the regular price)

3. Quantity Ratio (quantity sold per sale relative to quantity sold per regular price).

The index of average menu price understates the change in average unit price. In practice, an increase in the average discount would likely correspond to an increase in the quantity ratio. The effects of these two changes together would drive an even larger gap between the unit price and the menu price. On the other hand, an increase in the frequency of sales may reduce the quantity ratio and the effects would offset each other.

1.C Index Construction: Average Unit Price and Average Regular Price

To analyze the importance of sales on the dynamics of prices, I calculate two different price indexes using the scanner data. The only difference between these two indexes occurs in the aggregation across weeks within a store-product-month cell. The equations that follow show how I construct the average price for each store-product-month. Denote the average price of UPC j in month t as $P_{j,t}$ with a superscript to denote the two different averages: r and u indicating *regular*, and *unit* respectively:

Average regular price:

$$P^r_{jt} = \frac{1}{n_{jt}} \sum_{w \in t} p^r_{jw} \tag{1.C.1}$$

Average unit price:

$$P^u_{jt} = \frac{1}{q_{jt}} \sum_{w \in t} q_{jw*} p_{jw} \tag{1.C.2}$$

where w indexes the week, n_{jt} is the total number of price observations for UPC j in month t (usually the number of weeks in the month), and q_{jw} is the quantity sold during week w. In simple terms, the average regular price is the simple average of regular price and the average unit price is revenue divided by quantity sold.

Since we have several thousand UPCs in our sample, we need to aggregate across products into a single price index. To do so, I use the same technique as the BLS, which is to take the geometric average of monthly price relatives (using the UPC's average share of revenue as the weight). Importantly, I use the same aggregation procedure for both price indexes. The month t price index with base period 0 is:

$$I_{0,t-1} = \exp\left(\sum_j w_j \ln \frac{P_{j,t}}{P_{j,0}} \right) \tag{1.C.3}$$

1.D CPI Construction Details

I calculate a comparable index from the components of the CPI. Specifically, I aggregate the CPI item indexes that correspond to the product categories contained in the scanner data sample. Table 1.12 shows the CPI Item series used and the corresponding weights (which are the same as those used to aggregate the scanner data). Table 1.13 shows the weights I apply to each of the markets contained in the IRI scanner data sample.

Table 1.12: CPI Item Indexes and Corresponding IRI Scanner Data Categories

Series ID	Description	Weight	IRI Category(s)
CUUR0000SEFA02	Breakfast Cereal	0.11	coldcer
CUUR0000SS05011	Frankfurters	0.03	hotdogs
CUUR0000SEFN01	Carbonated Beverages	0.18	carbbev
CUUR0000SEFP01	Coffee	0.04	coffee
CUUR0000SEFR01	Sugar & Sweetners	0.01	sugarsub
CUUR0000SEFS01	Butter & Margarine	0.02	margbutr
CUUR0000SS16014	Peanut Butter	0.01	peanbutr
CUUR0000SEFT01	Soups	0.05	soup
CUUR0000SEFT02	Frozen and Prepared Foods	0.11	fzdinent, fzpizza
CUUR0000SEFT03	Snacks	0.11	saltsnck
CUUR0000SEFT04	Spices Seasoning Condiments	0.06	spagsauc, mayo, mustket
CUUR0000SEFW01	Beer	0.14	beer
CUUR0000SS61021	Photo Supplies	0.00	photo
CUUR0000SEGB01	Hair Dental Shaving	0.06	toothpaste, shamp, deod, blades, toothbr
CUUR0000SSGE013	Infants Equipment	0.02	diapers
CUUR0000SEFJ04	Other Dairy	0.06	yougurt

Notes: This table presents the BLS item indexes used to construct the CPI-U comparable to the indexes I created using the IRI data.

Table 1.13: Market Weights Applied to IRI Scanner Data Market Indexes

IRI Market	BLS Weights	Notes	IRI Market	BLS Weights	Notes
atlanta	1.37		omaha	0.88	(used Lincoln)
birmingham	0.89		philadelphia	2.73	
boston	2.52		phoenix	1.04	
buffalo	0.68		portland	0.83	
charlotte	1	No Match	raleigh	0.82	
cleveland	1.32		richmond	0.89	
dallas	1.87		roanoke	0.2	(no match 25% size of Birmingham)
desmoines	0.8	No Match	sacramento	0.83	(no match - population ~ Portland)
detroit	2.4		saltlake	1.01	(no match used Provo)
grandrapids	0.2	No Match	sandiego	1.16	
greenbay	0.2	No Match	sanfran	2.89	
harrisburg	1	No Match	seattle	1.37	
houston	1.73		southcarolina	0.92	
indianapolis	1	(no match 1.7mm people)	stlouis	1.15	
kansascity	0.73		washington	2.01	
knoxville	0.81	(no match used Chattanooga)	westtex	1	
la	4.1		syracuse	0.87	
milwaukee	0.74		hartford	0.83	
minneapolis	1.18		chicago	3.81	
newengland	0.74	(no match - used Burlington)			
nyc	3.39				

Notes: These are the weights applied to each market in constructing the regular price index and unit price indexes discussed in Section 4.

1.E Within Market Synchronization

Below I repeat the synchronization analysis within a single city. The 75th percentile measures are generally higher than those for the full sample. This is likely due to the fact that there is synchronization in sales within stores of a particular chain and the top 3 chains account for over half of the stores in this sample. I conclude that there is little evidence of synchronization other than within stores belonging to the same chain.

Table 1.14: Synchronization of Sales within Los Angeles

Category	Fraction of Stores having a Sale (Observation = UPC-Week) 25th Pctle	Median	75th Pctle	Average Fraction of Items on Sale	Store Count	Chain Count	Top 3 Chains % of Stores
Beer	8%	15%	26%	16%	127	10	57%
Blades	2%	6%	12%	8%	131	13	54%
Carbonated Beverages	11%	24%	37%	22%	132	13	55%
Cigarettes	1%	2%	4%	4%	131	12	55%
Coffee	7%	13%	21%	14%	130	11	55%
Cold Cereal	6%	13%	23%	16%	130	11	55%
Deodorant	4%	8%	13%	10%	132	13	55%
Diapers	6%	11%	18%	13%	131	12	55%
Facial Tissue	8%	16%	28%	19%	132	13	55%
Frozen Dinners	15%	25%	36%	26%	123	11	58%
Frozen Pizza	16%	25%	34%	26%	113	10	63%
Household Cleaners	6%	12%	21%	13%	132	13	55%
Hot dogs	13%	22%	32%	23%	112	10	63%
Laundry Detergent	8%	15%	23%	16%	131	12	55%
Margarine and Butter	8%	14%	23%	16%	124	11	57%
Mayonnaise	3%	10%	23%	12%	129	11	55%
Milk	5%	11%	19%	12%	130	11	55%
Mustard and Ketchup	2%	7%	15%	9%	129	11	55%
Paper Towels	5%	11%	19%	14%	131	13	54%
Peanut Butter	6%	13%	23%	13%	128	11	55%
Photo Supplies	1%	4%	10%	7%	132	13	55%
Razors	2%	6%	12%	9%	131	13	54%
Salty Snacks	7%	16%	30%	17%	132	13	55%
Shampoos	4%	8%	15%	11%	132	13	55%
Soup	4%	11%	24%	14%	130	11	55%
Spaghetti Sauce	10%	18%	28%	18%	130	11	55%
Sugar Substitutes	1%	4%	10%	6%	130	12	55%
Toilette Tissue	8%	15%	25%	17%	132	13	55%
Toothbrushes	4%	9%	17%	13%	132	13	55%
Toothpaste	4%	10%	18%	12%	132	13	55%
Yogurt	10%	19%	28%	20%	119	11	60%

Notes: This table presents synchronization data for Los Angeles in 2005. Again, the percentiles are calculated from those UPC-week cells in which at least one item was on sale (zeros are excluded). In columns 2 through 4, the unit of observation is a UPC-Week and the measure is the fraction of stores with the UPC on sale in a given week. I also present the number of stores and chains, as well as the percent of all stores that belong to the top 3 chains. This is to give the reader an idea of how much synchronization to expect given that chains typically have similar pricing plans across stores.

CHAPTER 2

A MODEL OF SALES

2.1 Introduction

Three important papers on the macroeconomics of sales have concluded that sales are, more or less, unimportant for macro economists (Eichenbaum et al., 2011; Guimaraes and Sheedy, 2011; Kehoe and Midrigan, 2010, EJR, GS, and KM, hereafter.). EJR and KM generate sales in modified menu cost models by adding frequent, large, and transient idiosyncratic cost shocks. In GS, on the other hand, sales arise in equilibrium as a result of strategic interaction between firms who face some consumers with very low price elasticities and others with high price elasticities. Policy experiments performed with calibrated versions of each model indicate that high frequency price changes due to sales do not result in neutrality of money.

However, each of these models is inconsistent with some important facts about sales discussed in the previous chapter. First, sales are not synchronized across stores or similar items. This suggests that a sale on an individual item is not caused by a temporary reduction in the manufacturer's price. Second, average unit price varies more than regular price does because the fraction of revenue from sales varies over time. I also provided evidence that unit price appears to be more responsive to demand shocks than regular price. The evidence suggests that sales have a role in aggregate price adjustment.

The objective of this chapter is to show that even if the primary purpose of sales is to price discriminate, sales may be important for mitigating the effect of a demand shock on quantity sold. I sketch a simple model in which sales arise as a result of differences in information about where to find the lowest price. Demand uncertainty is resolved after firms set price and in equilibrium, firms are indifferent between the sale price and the regular price, *ex ante*.

There are two channels through which sales, can facilitate price adjustment. The first channel is analogous to the additional units sold when a monopolist is able to price discriminate. Think in terms of a market in which a price sensitive group can obtain a good for a low price and the less sensitive group obtains it for a higher price. If there is a shift in the relative size of these two groups, the monopolist who can price discriminate will accommodate this shift better than one who cannot. The second channel results from the assumption that price and capacity are fixed before uncertainty about the state of demand is realized. The lowest priced items sell first. Thus when demand is high, there is available capacity at higher prices so the market clears. On the other hand, when demand is low, the low priced items sell first so the market performs better than if firms all selected a single (high) price.

The model is able to replicate several facts that we observe about price setting behavior. First, firms select between one of two prices. I interpret the high price as the regular price and the low price as the sale price. Second, sales are not synchronized across locations. Third, posted prices appear to be sticky, but average price paid is correlated with the state of demand. Finally, the fraction of items sold at a sale price tends to be high when aggregate demand is low. This is consistent with the finding that sales rise when unemployment is relatively high. An important result from this model is that sales help to clear the market when prices are sticky and demand turns out to be relatively low.

Below I provide a brief overview of macro and micro models of sales. The idea is to contrast several different approaches. Next I present the model and highlight some analytical results obtained from a parametrized version. I conclude with some suggestions for extending the model.

2.2 Macro Models with Sales

Three important articles have recently been published about the macroeconomics of sales. All three models generate individual price series that fluctuate a lot from week to week but commonly return to a modal price. They each conclude sales are more or less, unimportant for macroeconomics. None of the models result in cyclical changes in the aggregate fraction of revenue from sales.

Both EJR and KM augment a standard menu cost model by giving firms a cheaper, but imperfect alternative to incurring the menu cost of adjusting price. In EJR, firms can select a different price from the "price plan". KM allow firms to incur a small cost for a one period adjustment (after which price automatically reverts to the inherited price). Both models generate frequent week to week price changes by incorporating volatile *idiosyncratic* cost and/or demand shocks. In both cases, sales actually result in stickier "reference" or "list" prices because firms tend to take advantage of the alternative to incurring the menu cost.

This work convincingly demonstrates that sales, in the context of a modified menu cost model, do not imply money neutrality. This result depends crucially on the use of highly volatile idiosyncratic costs to generate frequent week to week price changes. EJR find that indeed, accounting data indicate substantial acquisition cost volatility. This simply pushes the question back to the manufacturer. Why are prices from the manufacturer so volatile? This is especially difficult to reconcile with the fact that most of the items analyzed are storable.

Large idiosyncratic cost shocks at the manufacturer should cause substantial co-movement of prices for individual items across retailers and stores.[1] This behavior is absent from the data. I find that sales tend to be staggered across locations and Nakamura (2008) reports that only 16 percent of price variation is common across stores. To summarize, it seems unlikely that sales result from cost shocks, and even if they do, the story is more complicated than what has been modeled thus far.

Guimaraes and Sheedy (2011) build a DSGE macro model in which sales arise out of strategic behavior between firms. In this model, sales are a technique for extracting more value out of a market composed of consumers with varying demand elasticities. Nevertheless, the effects of monetary policy are nearly identical to those of a benchmark model without sales. This is because sales are strategic substitutes. When there is shock common to everyone in the market, the effect on sales is small because all firms want to make the same change to the frequency and/or depth of sales, which in turn provides substantial incentive not to do so.

In this model, there is very little volatility in the fraction of revenue from sales because consumers do not actively take advantage of sale prices. Equilibrium pricing means that all consumers are indifferent between any given instant in which to purchase goods. As we have seen in the data from chapter 1, there are substantial fluctuations in the fraction of revenue from sales that do not appear to be purely random.

2.3 Micro Models of Sales

In the industrial organization (IO) literature, the topic of sales is often treated as a particular type of price dispersion. Upon reviewing this literature, it appears that over the course of the 1970's and 1980's, much effort was directed towards finding the bare essentials that result in equilibrium price dispersion. In general, IO models of sales (price dispersion) rely on imperfect information, heterogeneous consumers, and/or heterogeneous cost.

Reinganum (1979) generates price dispersion in a fairly simple model in which consumers engage in sequential search. Equilibrium price dispersion results from variation in marginal cost and downward sloping demand (as opposed to unit demand). On the other hand, Varian (1980) obtains equilibrium price dispersion in a model with identical firms and free entry by including two different types of shoppers (informed and uninformed). In equilibrium, firms select price from a continuous distribution and each sell to an equal fraction of uninformed shoppers regardless of price. One lucky firm will select the lowest price and sell to all of the informed shoppers in addition to its share of the uninformed shoppers.

These models are built upon the intuitively appealing idea that knowing when and where to buy

[1] This is because the Robinson-Patman Act of 1936 restricts the ability of manufacturers to price discriminate among retailers unless justifiable by differences in cost.

something at a low price is not trivial. However, basing empirical analysis on them is troublesome because they rely heavily on continuity in the equilibrium distribution of prices.

Sobel (1984) and Conlisk, Gerstner and Sobel (1984) model sales as a means of intertemporal price discrimination. Consumers with varying preferences enter the market in each period. Some are impatient with high reservation values while others are patient with low reservation values. The firm(s) will typically charge a high price that only induces the high types to buy. Over time, low types accumulate until market elasticity is high enough to induce a one period price cut. This model captures many of the features of sales that we observe (e.g. returning to a previous price after a temporary price cut), but it is not clear how to implement the idea when there are many close substitutes available right next to each other.

Consumer heterogeneity and search costs are at the heart of micro models of sales and should be taken into account by macro studies of sales.[2] These ideas do not rule out the possibility that sales are important for macroeconomics. For instance, shifts in the composition of consumers could result in fluctuations in the frequency or depth of sales. This idea is discussed in Chevalier and Kashyap (2011).

2.4 A Model of Price Dispersion with Two Prices

I try to replicate two important facts about sales with this model: 1) sales are not the result of week to week variation in marginal cost, and 2) price typically varies between two discrete prices. I follow the IO literature and generate price dispersion using differences in information about where to find the best price.

Firms sell identical goods and randomize between two prices. Much like Varian's (1980) model, there are two types of consumers: informed and uninformed. The informed consumers are able to locate stores charging a low price and always buy from low priced stores. As a result, low priced stores always sell all of their capacity. On the other hand, high priced stores sell some fraction of their capacity when demand is low. This model is closely related to the uncertain and sequential trade models in Eden (1994) and Eden (2005). Prices are not allowed to adjust in response to the realization of demand and capacity is fixed, but dispersion of posted prices facilitates market clearing.

The model reproduces two important facts about sales discussed in Chapter 1. First, the probability of a sale at one location does not depend on whether another store is having a sale. Sales are staggered rather than synchronized. Second, posted price does not vary with demand, but average unit price does. When demand is high, average unit price is also high. When demand is low, the fraction of items sold at a sale price is relatively high so unit price is relatively low.

[2]Guimaraes and Sheedy (2011) is an outstanding example of incorporating a strategic motivation for sales into a DSGE model.

Consumers

There are two types of consumers: shoppers and buyers. Shoppers learn each firm's posted price at no cost and chose a store charging the lowest price at random. In Varian (1980), the probability that there is a tie between stores charging the lowest price is zero. In this model, the distribution of prices will be discrete and a positive fraction of firms charge the lowest price. The fraction of consumers who are shoppers is $S \in (0,1)$.

Buyers do not learn the posted price of any firm. They only learn the maximum possible price that may be charged by their preferred store and the maximum possible price that they would pay if they visited another store at random. For simplicity, I make an ad hoc assumption about the way buyers select a store. Buyers will visit their favorite store as long as the highest possible price at this store is less than or equal to the highest possible price at other stores. Otherwise, they randomly select another store to visit. Each store has an equal number of "loyal" buyers.

The assumption about how buyers chose where to shop could be modeled explicitly through a cost of searching. One could also think of a behavioral motivation. Perhaps buyers use a simple rule that they believe ensures they pay a "fair" price on average. The assumption is closely related to the one made in Wilde and Schwartz (1979) in which a fixed proportion of consumers have a certain taste for shopping and sample a fixed number of prices.

Consumers are otherwise identical and demand the good according to the demand function $D(p)$, which is continuously differentiable with $D'(p) < 0$. Finally, the number of consumers is a random variable which takes two possible values: N and $(1+\delta)N$ with equal probability. I assume, without loss of generality, that $\delta > 0$.

Firms

There is a large number of equally sized firms which we will normalize to a measure of one. Firms sell a fixed quantity L, of a homogeneous good. Capacity is distributed evenly across firms.[3] The cost of selling a unit of capacity is constant and assumed to be zero for simplicity. Firms select a distribution of prices prior to the realization of demand in order to maximize expected profits given expectations about what other firms will do.

The Sequence of Events

To understand the model, it is best to carefully describe the sequence of events:

[3]The model can easily be extended to allow for endogenous capacity choice. See the Appendix for a brief discussion. Since I am primarily interested in the effect of sales on capacity utilization, we do not gain anything by adding this complexity to the model.

1. Nature determines whether there will be N or $N(1+\delta)$ consumers in the market. Neither firms nor consumers observe the outcome of this event.

2. Firm j chooses a price distribution $F_j(p) \sim [p_l, p_h]$ and informs everyone of p_h

3. Firm j selects a price randomly from $F_j(p)$

4. Shoppers observe each firm's price and select randomly from the set of stores posting the lowest price

5. Buyers who prefer store j decide to shop there if $\max\{p_h^i\}_{i \neq j} \geq p_h^j$

6. Buyers and shoppers form a line at the store they selected and are treated symmetrically. By this I mean that any segment of the line has the same ratio of buyers to shoppers.

7. Consumers who visit store j each buy $D(p_j)$ units.

8. If firm j does not have enough capacity to satisfy demand, the fraction of consumers who did not get to buy from firm j at price p_j go find another firm that still has capacity.

Equilibrium

Equilibrium is a price distribution for each firm that maximizes expected profits given the price distributions selected by other firms. I focus on a symmetric equilibrium in which the distribution of prices is discrete with two possible outcomes. This is motivated by the fact that stores typically charge one of two prices in any given week (EJR, 2011). The high and low prices (p_h and p_l) are analogous to the "regular" and "sale" price. Let μ be the probability assigned to p_l.

I provide the intuition behind a two price equilibrium and then state the equilibrium conditions formally. Firms must be unable to increase expected profits by changing their price distribution. This implies three things about the equilibrium price distribution.

First, firms must be indifferent between either price. Low priced firms sell all of their capacity regardless of the state of demand. High priced firms sell a fraction of their capacity (to their loyal buyers) in the low state of demand and all of their capacity in the high state of demand (since all "additional" consumers must find a high priced store in order to buy the good). Either price is expected to generate the same amount of revenue in equilibrium. Let ϕ be the average fraction of capacity sold by a firm posting the high price, p_h. Firms are indifferent between two prices $\{p_l, p_h\}$ as long as $p_l = \phi p_h$. Note that $1 > \phi$.

Second, the low price must equate demand to capacity available at low priced stores *when demand is low.*[4] Suppose not. If there is excess demand at stores charging p_l when demand is low, then some shoppers

[4]Note that this implies that low priced stores sell out regardless of the state of demand.

would be unable to find capacity at the lowest priced stores and would go to the next lowest priced store with capacity. In this situation, a store could increase expected profits by selecting a slightly higher "low" price, say $p_l+\epsilon$ because it would sell all of its capacity regardless of the state of demand. On the other hand, if there is excess supply (at low priced stores when demand is low) then a firm would sell more units (at zero marginal cost) by cutting price a small amount. In a two price equilibrium, the number of consumers that visit low priced stores in the low demand state is $\mu(1-S)N$ buyers plus SN shoppers. The total amount of capacity available at a low price is μL. Thus, one equilibrium condition is that $(\mu(1-S)N+SN)\,D(p_l)=\mu L$.

Our third equilibrium condition is that the high price, p_h, should equate capacity available at high priced stores to the demand of consumers who eventually visit a high priced store *when demand is high*. The number of consumers who visit a high priced store in the high demand state is $(1-\mu)\,(1-S)N+\delta N$ and the capacity of high priced stores is $(1-\mu)\,L$.[5] The third equilibrium condition is $(1-\mu)\,(1-S)N+\delta N=(1-\mu)\,L$. Posting a higher price is unprofitable because no consumers will visit such a store in the low demand state. On the other hand, selecting $p_h-\epsilon$ is a bad idea because it results in no additional consumers in either state of demand. I show formally that such deviations are not profitable in the appendix.

We can now define equilibrium as a vector $(p_l,p_h,\mu,\phi)\gg 0$ that satisfies the following conditions for a given level of capacity L:

1) $p_l=\phi p_h$

2) $\mu L=N\,(S+\mu(1-S))\,D(p_l)$

3) $(1-\mu)L=N\,((1-\mu)(1-S)+\delta)\,D(p_h)$

4) $\phi=\frac{N\left((1-\mu)(1-S)+\frac{1}{2}\delta\right)D(p_h)}{(1-\mu)L}$

To summarize, equation 1 is an arbitrage condition that requires expected revenue to be the same regardless of the posted price. Equation 2 is the market clearing condition for the low demand state. Capacity equals demand for the stores selling at the low price *when demand is low*. The market clearing condition for the high demand state is equation 3. It says that capacity equals demand for the stores selling at a high price *when demand is high*. Finally, equation 4 is the definition of ϕ, which is the expected fraction of capacity sold at the high price.

In the Appendix, I show that an equilibrium of the type described above exists as long as demand is not too elastic or inelastic at the high price. The equilibrium can only be solved in closed form if we assume certain types of demand functions. In the next section, I solve for the case of unit elastic demand and do some comparative statics.

[5] Some of these consumers will have initially visited a low priced store only to find that it was out of stock.

Parametrized Example

To analyze the model, I solve a simple parametrized version analytically. Let us suppose that $D(p) = \frac{a}{p}$, $S \in (0,1)$ and $\delta > 0$. We will assume that capacity is exogenously determined for now. The equilibrium vector (p_l, p_h, μ, ϕ) in the case of unit elastic demand is as follows:

$$p_l = a\frac{N}{L}\left(1+\frac{\delta}{2}\right)$$

$$p_h = a\frac{N}{L}\left(1+S+\delta\right)$$

$$\mu = \frac{2S}{2S+\delta}$$

$$\phi = \frac{2+\delta}{2(1+S+\delta)}$$

Notice that both μ and ϕ are less than zero. I would like to emphasize that this equilibrium only applies for $S \in (0,1)$. If S is zero or one, then the distribution of prices degenerates.

Fraction of prices quotes that are sales

The fraction of price quotes that are sale prices, p_l, is equal to the probability of a sale, μ. Notice that μ depends on the relative size of shoppers S and the volatility of demand, δ. Of course, by construction, it cannot depend on the realization of demand since it is specified before hand.

Not surprisingly, μ is increasing in the fraction of consumers that are shoppers, S. This is because *ceterus paribus,* more shoppers tend to drive up the price that equates demand to supply in the low demand state (p_l). Thus, firms have an incentive to increase the probability of a sale, which in turn, would push the equilibrium low price back to the point at which firms are indifferent between prices.

The probability of a sale decreases with the volatility of demand, δ. The upside of posting a high price increases with the size of the high demand state (relative to the low demand state) so firms have an incentive to increase the probability of posting a high price as δ gets larger.

Discount size

The percentage difference between p_l and p_h $(1-\frac{p_l}{p_h})$ (referred to as the discount in Chapter 1), is equal to $\frac{2S+\delta}{2(1+S+\delta)}$. The discount is increasing in both S and δ. Notice that the low price does not depend on the number of shoppers. When the relative size of shoppers increases, the high price needs to rise in order to compensate for relatively fewer consumers who will buy at the high price in the low demand state. An increase in the parameter δ increases both expected demand and the standard deviation of demand. As a result, both the low price and the high price are increasing in δ, but the high price is more sensitive. This result is also obtained in uncertain and sequential trade models discussed in Eden (2005).

Average posted price versus average unit price

This model also illustrates how aggregate price may appear to be unresponsive to changes in demand while at the same time, average unit price is responsive to demand. The consumer price index aggregates a survey of posted prices. In this model, the price level measured in this way is not affected by the state of demand. In either case, average posted price is $\mu p_l + (1-\mu)p_h$. On the other hand, average unit price does depend on the state of demand. More units are purchased at the high price when demand is high so average unit price is higher in the high demand state.

Output elasticity

In this model, sales tend to reduce the effect of demand fluctuations on output. To see this, consider two different economies: one in which the fraction of shoppers, S is zero, and the other in which shoppers make up some fraction of consumers ($S \in (0,1)$). When there are no shoppers, stores never have sales and the price is always p_h. Output, measured by the amount of capacity sold, is $ND(p_h)$ in the low demand state and $(1+\delta)ND(p_h)$ in the high demand state. When there are no sales, output elasticity is one (percentage change in demand = percentage change in output). When $S \in (0,1)$, output in the high demand state relative to output in the low demand state is $1 + \frac{\delta}{1+S} < 1 + \delta$ which means that output elasticity is $\frac{1}{1+S}$.

The effect of sales on the elasticity of output depends on the fraction of consumers that are shoppers, S. If the number of shoppers is small, then sales have little effect on the elasticity of output because they occur infrequently and the discount is small. The analysis above applies only to the case of constant unit-elastic demand. For future research, I propose using numerical analysis to study the model under more general demand functions.

2.5 Conclusion

The purpose of this chapter was to illustrate how sales can facilitate price adjustment even when prices are perfectly sticky. In the model outlined above, firms set price before the realization of demand so posted prices cannot adjust to changes in demand. However, because there is a distribution of prices in equilibrium, average unit price responds to demand because the lower priced items get purchased first and the higher priced items tend to be purchased in higher quantities when demand is relatively high. Therefore, sales, (or dispersed prices) are an important mechanism through which the quantity sold adjusts to fluctuations in demand.

Several features of this model are consistent with what we observe in the data. First, posted prices are uncorrelated across locations. In Chapter 1 we found that sales on specific items were not synchronized across stores. I interpret this as evidence that sales are not primarily the result of idiosyncratic cost shocks. In the model above, sales occur because stores randomize between two prices when consumers are heterogeneous in their information and/or willingness to shop. The probability of a sale does not depend on whether another store is having a sale. Instead, the probability of a sale depends on characteristics of consumers and the volatility of demand.

Second, average posted price does not respond to changes in demand but average *unit* price does. In the model, stores select price before demand is realized so posted prices cannot adjust. However, consumers tend to buy the lower priced items first so when demand is low, the fraction of revenue from sales is high and the average *unit* price is low. In the data, the CPI is slow to respond to monetary policy shocks. Evidence presented in chapter one indicates that unit price is more responsive to aggregate demand conditions. In Chapter 1 we also saw that a rise in the unemployment rate preceeded an increase in the fraction of revenue from sales.

The key difference between this model and others used for studying sales in the macroeconomy is that the fraction of revenue from sales is responsive to aggregate demand shocks. To the best of my knowledge, EJR, KM, and GS, all result in more or less static features of sales. The model sketched above is a prototype that could be developed further to conduct policy analysis. It suggests that there is a simple way in which sales might play an important role in aggregate price adjustment. For future research, I propose a cash-in-advance version of this model that could be calibrated and used to conduct monetary policy experiments.

REFERENCES

Chevalier, Judith A. and Anil K. Kashyap, "Best Prices," 2011.

Conlisk, John, Eitan Gerstner, and Joel Sobel, "Cyclic Pricing by a Durable Goods Monopolist," *Quarterly Journal of Economics*, 1984, *99* (3), 489–505.

Eden, Benjamin, "The Adjustment of Prices to Monetary Shocks When Trade Is Uncertain and Sequential," *Journal of Political Economy*, June 1994, *102* (3), 493–509.

_ , *A course in monetary economics: Sequential trade, money, and uncertainty*, Blackwell, 2005.

Eichenbaum, Martin, Nir Jaimovich, and Sergio Rebelo, "Reference Prices, Costs, and Nominal Rigidities," *American Economic Review*, 2011, *101* (1), 234–262.

Guimaraes, Bernardo and Kevin D. Sheedy, "Sales and Monetary Policy," *American Economic Review*, 2011, *101* (2), 844–876.

Kehoe, Patrick and Virgiliu Midrigan, "Prices Are Sticky After All," *NBER Working Paper No. w16364*, 2010.

Nakamura, Emi, "Pass-Through in Retail and Wholesale," *American Economic Review*, 2008, *98* (2), 430–437.

Reinganum, Jennifer F., "A Simple Model of Equilibrium Price Dispersion," *Journal of Political Economy*, 1979, *87* (4), 851–858.

Sobel, Joel, "The Timing of Sales," *Review of Economic Studies*, July 1984, *51* (3), 353–368.

Varian, Hal R., "A Model of Sales," *American Economic Review*, 1980, *70* (4), 651–659.

Wilde, Louis L. and Alan Schwartz, "Equilibrium Comparison Shopping," *Review of Economic Studies*, July 1979, *46* (3), 543–553.

APPENDIX

2.A Existence of Equilibrium

Below I will show that an equilibrium of the type described in the text above exists. I do so by solving for equilibrium under a parametrized version of the model that can be solved analytically.

Solving for the equilibrium vector.

I take L as given and show how to solve for (p_l, p_h, μ, ϕ). Using conditions (1) and (2) and assuming that demand is homogeneous of degree $\alpha < 0$ we can see that:

$$\phi = \left(\frac{\mu L}{D(p_h)(S + \mu(1-S))N_1} \right)^{\frac{1}{\alpha}}$$

Next we can use equation (3) to substitute for $D(p_h)$ above:

$$\phi = \left(\frac{\mu\left((1-\mu)(1-S) + \delta\right)}{(1-\mu)(S + \mu(1-S))} \right)^{\frac{1}{\alpha}} \tag{2.A.1}$$

Combining conditions (3) and (4) we can write a separate equation for ϕ in terms of μ

$$\phi = \frac{(1-\mu)(1-S) + \frac{1}{2}\delta}{(1-\mu)(1-S) + \delta} \tag{2.A.2}$$

Unless $\alpha = -1$, there is no analytical solution to the equations above. It can be shown that equations 2.A.1 and 2.A.2 are both decreasing in μ. So, I perform a numerical analysis to get an idea of what restrictions must be placed on S and δ in order for the solution to result in $\mu \in (0,1)$ and $\phi \in (0,1)$.

In general, the equilibrium exists except for very small values of δ and S (in the neighborhood of .01). Under these circumstances, the distribution of price degenerates.

Ensuring there are no profitable deviations

To ensure we indeed have a Nash equilibrium, I check to make sure that deviations from this strategy are not profitable given that other firms charge p_l with probability μ and p_h with probability $1 - \mu$.

Small reduction in p_h

Does a small reduction in p_h increase expected revenue? In the high demand state, marginal revenue is $(1-\mu)L$ for price cuts because the store will not sell additional units. In the low demand state, marginal revenue is $N(1-S)(1-\mu)\left(D'(p_h)p_h + D(p_h)\right)$. Thus, the change in expected revenue from this deviation is:

$$-N(1-S)(1-\mu)D(p_h)(1+e) - (1-\mu)L$$

which is weakly negative if:

$$1 + \frac{(1-\mu)L}{N(1-S)(1-\mu)D(p_h)} \geq |e|$$

Since $(1-\mu)L > N(1-S)(1-\mu)D(p_h)$, the LHS $>$ 2. Thus, as long as demand isn't too elastic at p_h (elasticity less than two is sufficient but not necessary) then increasing p_h is not profitable.

Small increase in p_h

If an individual firm chooses to post a high price that is higher than p_h, then it will sell to no one in the low demand state because its regular customers will search for a store with a lower high price. It will sell to $N\left((1-S)(1-\mu)+\delta\right)$ customers in the high demand state because all other stores will be sold out. Thus, the expected change in revenue from a small increase in p_h is:

$$N\left((1-S)(1-\mu)+\delta\right)D\left(p_h\right)(1+e) - N(1-S)(1-\mu)D\left(p_h\right)$$

where e is the elasticity of demand at p_h. This quantity is weakly negative when:

$$-e \geq \frac{\delta}{(1-S)(1-\mu)+\delta}$$

A sufficient condition for this deviation to be unprofitable would that demand is elastic at p_h (because the RHS is less than one).

Changes in p_l

A reduction in p_l reduces expected revenue because all capacity is sold at p_l regardless of the state of demand. A store increasing p_l by a small amount would no longer sell all of its capacity in the low demand state at the low price. Choosing a low price above p_l is like selecting two different high prices in terms of the expected number of shoppers that will show up. We have already shown above that no "high price" other than p_h can be optimal unless demand is highly elastic or inelastic. Of course, an individual store would be indifferent between charging p_h with probability 1 and charging p_l with probability 1.

2.B Endogenous Capacity

Allowing for capacity to be chosen endogenously is a fairly simple extension. I assume that firms select capacity at the moment they set the distribution of prices. The cost of L units of capacity is $C(L)$, which is increasing and convex. The cost of selling a unit of capacity is still assumed to be zero.

The optimal capacity choice will depend on the distribution of prices and the average fraction of capacity sold at the high price. Thus, the optimal capacity can be determined as:

$$L(p_l, p_h, \mu, \phi) = \max_L \mu L p_l + (1-\mu)\phi L p_h - C(L)$$

The necessary and sufficient condition for an optimal choice of L is :

$$\mu p_l + (1-\mu)\phi p_h = C'(L)$$

Equilibrium is now a vector $(p_l, p_h, \mu, L, \phi) \gg 0$ conditions 1 through 4 in section 2.4 as well as the following additional condition for determining capacity:

5) $\mu p_l + (1-\mu)\phi p_h = C'(L)$ (capacity choice is optimal given the price distribution and ϕ).

Note that conditions 1) and 5) imply that the optimal capacity is L^* such that $C'(L^*) = p_l$.

CHAPTER 3

SALES AND FIRM ENTRY: THE CASE OF WAL-MART[1]

3.1 Introduction

In the ongoing quest to understand pricing behavior, Nakamura and Steinsson (2008) found that the fraction of price quotes that are "sales" (i.e. temporary price reductions) has increased substantially over the last two decades. Certain product categories, such as breakfast cereal or potato chips, are now "on sale" twice as often as they were in the late 1980's. Determining the cause of this trend is important not only for understanding why firms have sales, but also for the ongoing debate about the role that sales play in aggregate price adjustment.[2] We examine one possible explanation for the rise in the frequency of sales: the diffusion of Wal-Mart stores. We show that frequent but temporary price reductions can be a rational response to firm entry and then show that a representative grocery chain appears to have responded this way to Wal-Mart's entry.[3]

The expansion of Wal-Mart dramatically altered the retail landscape. Since 1980, Wal-Mart have grown from 300 stores located in 11 states to over 3,700 stores with locations in every state. The chain's revenue is now about 8 percent of U.S. consumption expenditure on goods, and 80% of grocery stores cited Wal-Mart-type stores as their biggest concern.[4] Unlike traditional retailers who have periodic price reductions (i.e. sales), Wal-Mart attracts customers through "everyday low prices". **?** estimated that competition with this strategy was responsible for a 21% reduction in purchases at incumbent stores.

Many empirical studies have examined Wal-Mart's effect on the prices and revenue of incumbent retailers. Basker (2005) and Basker and Noel (2009) find that incumbents lower their average quarterly price over time, whereas Volpe and Lavoie (2008) find that the prices of national brands are lowered further than those of private-label brands. Singh et al. (2006) find that the majority of revenue lost to Wal-Mart is due to decreased customers rather than decreased baskets. More importantly, they argue that incumbents can significantly mitigate revenue losses by keeping just a few of their best customers.

To the best of our knowledge, only Ailawadi et al. (2010) has addressed Wal-Mart's effect on sales

[1]This chapter was co-authored with Matthew Jaremski, an economics graduate student at Vanderbilt University.

[2]See for example,Chevalier and Kashyap (2011); Eichenbaum et al. (2011); Guimaraes and Sheedy (2011); Kehoe and Midrigan (2010)

[3]Throughout the rest of the paper, the term "sales" will only refer to temporary price reductions. We use the term revenue when we address the price times quantity sold.

[4]National Grocers Association (2003)

behavior using high frequency data. They find that the number of sales decreases for supermarkets and increases for drug stores and mass format stores in response to Wal-Mart. The paper represents an impressive first step, but there are two drawbacks. First, it examines the entry of Wal-Mart supercenters even though many of the locations were already served by a Wal-Mart discount store. Despite a potential "cooling off" period, stores might have already adjusted to Wal-Mart, the small reactions to additional entry might be expected. Second, it focused on category-level data, while we find that changes to pricing strategy following Wal-Mart's entry depend on product specific characteristics.

We begin by showing that an increase in sales could be an optimal response to Wal-Mart by recasting the repeated price competition model in Lal (1990). In the model, two incumbent firms sell to loyal customers and customers who only buy from the lowest priced firm. Both firms charge a high price in duopoly and split the market. When a third firm with a lower marginal cost and no loyal customers enters, the incumbent's high prices are no longer optimal and they will do better by taking turns setting a low price. Similar to Wal-Mart, the entrant chooses a constant but low price strategy.

We use scanner data from the Dominick's Finer Foods database to test whether the stores in the grocery chain responded to Wal-Mart entry with more frequent sales.[5] The data span six years and consist of 3,828 products allowing us to control for unobserved heterogeneity at very fine levels (e.g. the UPC-store). The data contains each store's location, allowing us to isolate Wal-Mart's effect on individual stores, and the sample period corresponds to the initial entry of Wal-Mart. We find that stores significantly increased their sales frequency as their distance to Wal-Mart declined. Consistent with a "loss-leader" strategy, the increases in sales frequency were concentrated on the most popular products. The adjustment of sales thus seems to be a competitive response to Wal-Mart and not a secular trend.

3.2 A Repeated Game of Retail Price Competition with Firm Entry

The Industrial Organization literature presents several reasons for the existence of sales, but many of these models are unsuited for studying the frequency of sales. In Varian (1980), firms keep consumers (rationally) uninformed over time by choosing price randomly from a continuous distribution. However, the only unambiguous definition of a sale price in this model is that the lowest observed price is the sale price. This definition leaves no room for variation in the frequency of sales. In Conlisk et al. (1984), a monopolist generally charges a high regular price but is occasionally induced into charging a temporarily low price when enough low reservation price consumers accumulate in the market. The model provides clear predictions

[5] Among others, Hoch et al. (1994), Hoch et al. (1995), Peltzman (2000), Chevalier et al. (2003), and Kehoe and Midrigan (2010) have all used the DFF data to study price setting.

about the frequency of sales, but the assumptions do not approximate the market for consumer packaged goods which we wish to study.

The most compelling model for studying the frequency of sales in the context of firm entry is Lal (1990). He seeks to explain the peculiar fact that on any given week, either brand A or brand B could be found on sale in a single store, but never both. We recast this model to represent retailers who face the entry of a low cost competitor.[6] We use the model to show that firm entry may cause two incumbent firms to switch from charging the same price every period to a strategy of alternating between a high and a low price.

Model Setup

The model consists of a retail market in which there are initially two firms, A and B, (called "incumbents") engaged in repeated Bertrand price competition. A third firm, C, (called the "entrant") unexpectedly enters the market. Each firm maximizes discounted profits using a common discount rate of $\delta \in (0,1)$. Firms A and B have a marginal cost of $c > 0$ and firm C's marginal cost is normalized to zero.

There are two types of customers who purchase a homogeneous basket of goods from one of the firms in each period as long as the price is less than or equal to r. The first type of customer is loyal to one of the incumbent firms and will only purchase the basket from that firm. The second type of customer is a "switcher" who considers A and B to be perfect substitutes, but prefers them to C with varying intensity. The number of switchers is normalized to 1 and the number of loyal customers per incumbent is $\alpha > 0$.

Because switchers prefer the incumbents, firm C must charge a price lower than the minimum of the incumbents' prices to attract any customers. The lower C's price is relative to $\min\{p_A, p_B\}$, the more units C will sell. Assuming without loss of generality that $p_A \geq p_B$ the fraction of switchers that will buy from firm C is characterized the following way:

$$\text{Entrant's Share} = \begin{cases} 0 & \text{if } p_c \geq p_B \\ \frac{p_B - p_C}{d} & \text{if } p_B \geq p_c \geq p_B - d \\ 1 & \text{if } p_B - d > p_c \end{cases} \tag{3.2.1}$$

Here d is a demand parameter that reflects the opportunity cost of visiting firm C instead of A or B (e.g. the cost per unit of distance to get to C). Firms A and B will sell α baskets to their loyal consumers and the incumbent with the lower price of the two will sell to the switchers who do not buy from firm C.[7]

[6] We also take a different approach to proving the existence of the type of equilibrium we are interested in.

[7] For example, if $p_A \geq p_B \geq p_C$ then the revenue of A, B, and C will be $p_A\alpha$, $p_B\left(\alpha + 1 - \frac{(p_A - p_C)}{d}\right)$, and $p_C\left(\frac{p_B - p_C}{d}\right)$ respectively. This assumes that $p_C + d \geq p_B \geq p_C$, otherwise, B gets either none or all of the switchers.

Equilibrium

To understand how firm C's entry changes the pricing strategies of A and B, we first analyze how they behave *before* C's arrival. The maximum total profit in this duopoly occurs when both firms charge r every period and threaten to punish deviations with a finite period Nash reversion strategy. Proposition 1 describes this equilibrium and states the conditions under which the price of r can be supported.

Proposition 1. *If $\delta \geq \frac{\alpha-1}{\alpha+1}$, then the following symmetric strategy profile is a pareto-dominant sub-game perfect Nash equilibrium: Both firms charge a price of r in every period as long as both firms charged r in the previous period. If a firm deviates, both firms charge a price of c for the next t-1 periods where t is the largest positive integer such that $\delta^t \geq \frac{2\alpha}{2\alpha+1}$. In the t^{th} period following the deviation, firms resume charging a price of r. If either firm deviates, then the punishment restarts.*

Proof. See Appendix 3.B ☐

Once firm C enters the market, the game has several equilibria. We focus our attention on the pure strategy equilibria in which firm C plays a best response in each stage game. We argue that pure strategy equalibria are more plausible because they do not require firms to have a randomizing device or a mechanism for detecting deviations. Requiring the entrant to play a best response provides a simple equilibrium in which only the incumbents need punishments to support the equilibrium path.[8] From this set of pure strategy equilibria, we focus on the one that maximizes the discounted profits of the incumbents.

Once the entrant arrives, the equilibrium strategy that maximizes the incumbents' joint profits involves the incumbents staggering and alternating prices between the monopoly price, r, and a lower price, $\bar{r}$. That is to say that in any given period, one incumbent charges r and the other $\bar{r}$ and then in the following period they switch. This strategy can be supported without explicit collusion using a credible and effective punishment. The equations in (3.2.2) specify the punishment prices that A and B charge in the t^{th} period following a deviation from the equilibrium path:

$$\begin{aligned} p_A = p_B = c \quad & \text{for} \quad t \leq t^* - 1 \\ p_A = r,\ p_B = p \quad & \text{for} \quad t = t^* \\ p_A = \bar{r},\ p_B = r \quad & \text{for} \quad t = t^* + 1 \\ \\ p_A(t) = p_B(t-1) \quad & \text{for} \quad t \geq t^* + 2 \\ p_B(t) = p_A(t-1) \quad & \text{for} \quad t \geq t^* + 2 \end{aligned} \tag{3.2.2}$$

[8] This type of strategy is also consistent with Wal-Mart's slogan at the time: "Always low prices. Always."

On the equilibrium path, incumbent firms take turns charging r and $r \bar{<} r$. If either deviates, they punish each other by charging a price of c for the next $t^* - 1$ periods. In period t^*, firm A only sells to its loyal customers at a price of r whereas firm B sells to some switchers by charging $p < r$.[9] After period t^*, the incumbents return to alternating prices of r and $\bar{r}$. If either incumbent deviates, the punishment phase begins again. Proposition 2 states that the strategy profile described above is a sub-game perfect Nash Equilibrium as long as the discount rate is large enough.

Proposition 2. *As long as δ is large enough , there exists a sub-game perfect Nash Equilibrium strategy profile in which:*

1. Firm C charges a price of $\frac{\bar{r}}{2}$ in every period.

2. On the equilibrium path, firms A and B alternate between a price of r and $\bar{r} = \frac{2}{3}(\alpha(1 + d) + c) < r$.

3. Firms A and B punish each other as described in 3.2.2 for deviations from the equilibrium path.

Proof. See appendix 3.B. □

Most models of imperfect competition predict a reduction in average price when a competitor enters a market.[10] This model's contribution is to suggest that firms have periodic sales instead of permanently lowering price. In this way, they are still able to keep some price sensitive shoppers, while continuing to extract monopoly profits from their loyal customers some of the time.

3.3 Empirical Analysis

We wish to test the hypothesis that firms respond to entry by having sales. In practice, however, there are many reasons that firms have sales and they probably all operate simultaneously. We therefore examine the frequency of sales at 85 Dominick's Finer Foods (DFF) grocery stores before and after Wal-Mart's entry. Table 3.1 summarizes the DFF weekly scanner data, which contains 3,828 products sold from 1989 to 1996.

The DFF data are well-suited for testing our hypotheses. First, the sample begins when Wal-Mart's presence in the Chicago-area was limited to a single store and continues through the opening of 26 additional stores. The near absence of Wal-Mart prior to the sample period also allows us to view each store's first reaction to Wal-Mart rather than the later introduction of a larger supercenter. Second, the panel aspect of the data allows us to control for unobserved heterogeneity that is constant over time. For example, customer demographics vary from place to place but are unlikely to change enough during the sample period to affect

[9] For convenience, the price p is set to equalize firm A and B's discounted future profits following a deviation.

[10] We focus on Wal-Mart, but this result applies to any "big box" retailer that enters a market where incumbents have loyal customers as well as switchers.

Table 3.1: Summary of Dominick's Finer Foods Sample

Category	Sold At Wal-Mart?	# of UPCs	# of Stores	% of Sale Obs.	Profit Margin (%)	Avg. Quantity Sold	Avg. Price ($)
Analgesic	Yes	320	85	3.9	31.3	1.4	5.4
Bathroom Tissue	Yes	57	81	15.4	16.9	13.9	1.9
Bottled Juices	Yes	217	85	6.0	29.2	72.7	1.8
Cereals	Yes	227	85	10.7	16.6	9.5	2.2
Cookies	Yes	428	85	9.3	27.7	16.7	3.1
Dish Detergent	Yes	125	85	7.6	22.1	6.1	2.1
Fabric Softeners	Yes	156	85	13.5	24.9	9.0	2.3
Front End Candies	Yes	222	85	15.9	39.5	7.0	2.3
Frozen Entrees	No	378	85	9.5	33.7	5.2	2.6
Frozen Fruit Juices	No	95	85	8.6	34.9	13.1	0.6
Laundry Detergent	Yes	236	85	14.2	18.0	23.0	1.4
Paper Towels	Yes	76	80	21.2	22.1	5.6	5.4
Refrigerated Juices	No	102	85	22.7	30.4	43.5	1.5
Snack Crackers	Yes	180	85	17.0	27.1	7.6	2.2
Soft Drinks	Yes	564	85	13.3	16.5	33.1	2.3
Toothpaste	Yes	288	85	7.9	23.7	2.8	2.3
Tuna	No	157	85	17.7	28.2	23.3	2.6
All Products		3,828	85	12.5	25.8	17.8	2.6

Notes: The selection of the sample is described in Appendix 3.C. We consulted modern stores to determine whether a product category was sold in Wal-Mart.

the pricing strategy of a particular store. Third, the data set contains the specific location of each store, allowing us to identify the effect of Wal-Mart on individual stores through variation across time and stores.

Because the sample period is early in the chain's expansion, the Wal-Mart stores which entered Chicago before 1996 were discount stores rather than supercenters. This distinction is important because Wal-Mart's discount stores do not sell fresh grocery products. Nevertheless, there is substantial overlap between the products in the DFF sample and those sold by Wal-Mart discount stores. Based on observation of current Wal-Mart discount stores, we estimate that the ones covered in our sample sold at least 13 of the 17 DFF categories that we include in our analysis. Dominick's stores therefore would have directly competed with Wal-Mart discount stores on several of its products.

The chain's pricing structure is a potential problem that must be addressed. DFF sets pricing policy as a chain, but this does not necessarily mean that there will be no variation in the frequency of sales across stores. While Dominick's has pricing zones, Eden and Jaremski (2010) show that there is heterogeneity across stores. Each week, over a quarter of a store's prices differ from the chain's most common price. In addition, the median fraction of stores having a sale on a specific UPC is 24 percent. Stores thus seemed capable of responding to idiosyncratic shocks.

Wal-Mart and the Frequency of Sales at Dominick's

The DFF data contain flags indicating whether UPC i was on sale at a store j during week t. While a deal flag correctly indicates when there was a deal, the documentation suggests that some deals may have gone unflagged. To capture these missing sales, we separately flag any price that declined and returned back to its original price or higher within two weeks. Our final sale dummy variable, $Sale_{i,j,t}$, is the union of the two measures.

To proceed, we need an operational definition of firm entry. The extent to which firms compete depends on variables such as driving distance, traffic patterns, and other factors. Rather than selecting a binary or discrete measure of competition, we use the driving distance to the nearest Wal-Mart as a proxy for the intensity of competition with Wal-Mart. Using Thomas J. Holmes' Wal-Mart location data (Holmes, 2011), we compute $Dist_{jt}$, the shortest driving distance to a Wal-Mart for store j during week t.

Figure 3.3.1 illustrates Wal-Mart's growth by mapping the location and approximate entry date of every store in the Chicago-area prior to 1996. Expanding towards the city-center, new Wal-Mart stores opened near existing stores. Holmes argues that this dense network of stores allows the chain to sustain distributional efficiency during expansion. Wal-Mart's entry location and timing thus seems to be determined by logistical efficiency and may be considered exogenous to the time-varying unobserved factors that affect Dominick's frequency of sales.

Figure 3.3.1: Location and Approximate Entry Date of Wal-Mart Stores Near Chicago

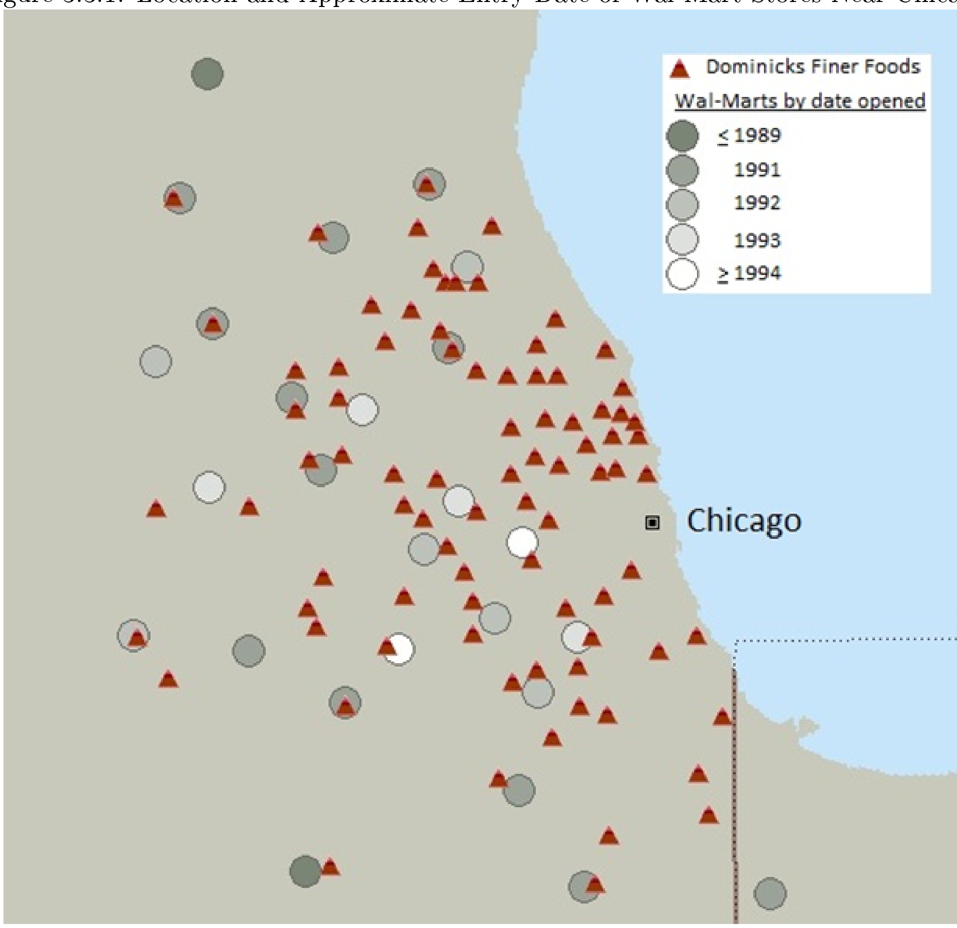

Note: Wal-Mart locations and entry dates were obtained from Holmes (2011). Dominick's locations come from the online documentation of the DFF database.

Figure 3.3.2: Average Distance to Wal-Mart vs. Fraction of Products on Sale

Notes: Average distance is the simple average across stores of the driving distance to the closest Wal-Mart. The other two series use two different smoothing techniques (moving average and HP-filter) to plot the revenue weighted fraction of products on sale. Only the subset of categories that have a significantly negative distance coefficient are used in the graph. The categories included are: bottled juices, breakfast cereal, frozen juices, soft drinks, canned tuna, and toilette tissue. A similar but less pronounced trend is observed for the entire set of categories.

Figure 3.3.2 plots the average driving distance to the nearest Wal-Mart and the average frequency of sales at DFF for selected categories. This figure illustrates that the frequency of sales rises as driving Wal-Mart enters the market. For example, during October of 1991, a 45 percent drop in the average distance to Wal-Mart (from 20 to 11 miles) corresponds to a 50 percent increase in the trend component of the sales (from 12 to 18 percent). This graph indicates that the chain-wide frequency of sales increased rapidly following Wal-Mart entry.

Two additional conclusions are visible in Figure 3.3.2. First, as shown by Singh et al. (2006), a response in sales does not begin until Wal-Mart moves into a reasonable competitive distance (i.e. within 30 miles). Second, the response begins to slightly dissipate after three years. Franklin (2001) finds that Wal-Mart's market share grows over time, suggesting that some customers might have become loyal to Wal-Mart or more sensitive to prices. Nevertheless, the frequency of sales remains at least 7 percentage points higher than its initial value.

Store Level Effect of Wal-Mart Entry

Building on the aggregate picture, we proceed with store-week panel regressions that control for unobserved store-level heterogeneity. The dependent variable ($\%Sale_{jt}$) is the percentage of products on sale in store j during week t, and the independent variable ($Dist_{jt}$) is the driving distance in miles from store j to the nearest Wal-Mart in week t. The regression is:

$$Sale_{jt} = \alpha_1 Dist_{jt} + \alpha_{\mathbf{2}} \mathbf{Q}_t + c_j + e_{jt} \tag{3.3.1}$$

where Q_t is a vector of quarter dummies to control for seasonal variation, c_j is the unobserved store heterogeneity that is fixed over time, and e_{jt} is the error term. Observations are weighted by the store's average share of chain revenue and standard errors are clustered by store to account for within-group serial correlation.

The coefficients in Table 3.2 show that individual stores increase their sales as Wal-Mart moved closer. The estimates imply that the average drop in driving distance to the nearest Wal-Mart (35 miles) increased the fraction of products on sale in a store by 1.05 percentage points. As the average frequency of sales was small (around 10%), this effect is both statistically and economically significant. The remaining columns provide the coefficients of models that include a time trend or control for unobserved effects at the category-store level. Including a linear time trend increases the magnitude of the estimates of α_1.

Table 3.2: Linear Panel Regressions of $\%Sale_{jt}$ on $Dist_{jt}$

	% of UPCs on Sale in week t			
	Store		Store-Category	
Distance	-0.030***	-0.061***	-0.037***	-0.122***
	[0.008]	[0.009]	[0.008]	[0.009]
Linear Trend		-0.005***		-0.013***
		[0.001]		[0.001]
Observations	31,856	31,856	542,758	542,758
Groups	85	85	1,519	1,519
R-squared	0.048	0.059	0.010	0.027

Notes: The first two columns report results of a fixed effects panel estimate of two different models that use the store as the unit of analysis. The second column controls for a linear trend while the first column does not. The second two columns report analogous results from a random effects estimate of two models in which a category-store is the unit of analysis. The store level model includes a vector of quarter dummies and the store-category model includes a vector of category x quarter dummies to control for seasonal effects for the chain and category respectively. The Distance coefficients are reported in percentage points per mile. T-Statistics are in brackets. Standard errors are clustered by store, and regressions are weighted by Revenue Share of Cell. * denotes significance at 10%; ** at 5% level and *** at 1% level.

Product Level Effect of Wal-Mart Entry

The results presented above indicate that DFF increased sale frequency in response to Wal-Mart's entry. If the purpose of these additional sales was to induce certain customer groups to make a trip to the store, then we would expect the additional sales to be focused on certain products. The data allow us to evaluate which categories and products experienced an increase in sales. To do this, we estimate linear probability regressions for each category.[11] Each observation is a UPC-store-week, and the dependent variable $Sale_{ijt}$ is a binary indicator of whether product i was on sale in store j during week t. We measure an individual product's popularity as its share of category revenue over its life and across all stores.[12]

We begin with a simple model that averages Wal-Mart's effect across all products:

$$Sale_{ijt} = \beta_1 Dist_{jt} + \beta_{\mathbf{3}} \mathbf{Q_t} + c_{ij} + e_{ijt} \tag{3.3.2}$$

where c_{ij} is unobserved UPC-store heterogeneity that is fixed over time. A negative β_1 coefficient implies that the average frequency of sales across the entry category would rise in response to a decrease in the distance to the nearest Wal-Mart. Next, we add the interaction of $Share_i$ and $Dist_{jt}$ to evaluate whether

[11] Results from probit or logit models are qualitatively similar to those found in our linear probability model.

[12] Our contention is that $Share_i$ is determined by consumer preference rather than by store-level weekly promotion fluctuations. Hosken and Reiffen (2004) use the same procedure to deal with endogeneity.

stores selected popular products to discount in response to Wal-Mart. The model becomes:

$$Sale_{ijt} = \beta_1 Dist_{jt} + \beta_2 Dist_{jt} \times Share_i + \beta_{\mathbf{3}} \mathbf{Q_t} + c_{ij} + e_{ijt} \tag{3.3.3}$$

Here, the effect of competition with Wal-Mart depends on the category (through β_1) as well as the product's popularity (through β_2).

Table 3.3 shows the results of the two models described above. In the model without the share interaction, seven out of 17 categories have a significantly negative coefficient on $Dist_{jt}$. However, when the interaction is included, the category-level effects all but disappear. β_1 remains significantly negative for only 3 categories (Bathroom Tissue, Bottled Juice, and Frozen Juice), while β_2 is significantly negative for 15 of the 17 categories. We conclude that popular products not only experienced the largest increase in sales, but the frequency of sales on less popular products seems to have decreased.

The coefficient estimates from the model summarized by equation 3.3.3 are summarized in Figure 3.3.3, which displays the average effect of a 35 mile reduction in the distance to Wal-Mart for the 5th, 50th, and 95th percentiles of $Share_i$. The median response is generally close to zero, and the only significant responses are found for the most popular products. The UPC-level approach provides additional evidence that the rise in sales across DFF stores was the result of competitive behavior rather than a general increase in sales.

The results discussed above are also consistent with "loss-leader" models that suggest firms will advertise low prices on only a few products (often below marginal cost) to attract shoppers who purchase other profitable products. For example, DeGraba (2006) illustrates how a low price on turkeys during Thanksgiving will attract a Thanksgiving dinner host who will also purchase a long list of other products needed for the dinner. Lal and Matutes (1994) argue that "loss-leaders" should be purchased frequently and costly to store. The pattern of coefficient estimates across categories match quite well with these characteristics, in that, the majority of categories that were sold often or were costly to store had negative and significant signs whereas the rest had positive and significant signs. Taking the "loss-leader" hypothesis a step further, Hosken and Reiffen (2004a) find that stores tend to put popular products on sale. Once again this matches our results, suggesting that Dominick's employed a "loss-leader" type strategy in selecting which items to put on sale.

The Effect of Wal-Mart on Other Aspects of Sales

While we have focused on the frequency of sales, Dominick's could also have adjusted their pricing strategy on other margins, such as sales depth, in response to Wal-Mart entry. The depth of a sale (the percentage

Table 3.3: Linear Panel Regressions of *Sale* by Category

	Analgesic		Bottled Juices		Cereals		Cookies		Dish Detergent		Front End Candies	
Distance	0.071***	0.100***	-0.083***	-0.043***	-0.074***	0.012***	0.017***	0.014***	-0.005	0.063***	0.074***	0.090***
	[0.002]	[0.003]	[0.004]	[0.005]	[0.002]	[0.003]	[0.002]	[0.002]	[0.004]	[0.006]	[0.003]	[0.004]
Dist.*Share		-4.231***		-4.709***		-13.99***		0.754***		-5.166***		-2.018***
		[0.307]		[0.446]		[0.316]		[0.281]		[0.408]		[0.480]
	Frozen Entrees		**Frozen Fruit Juices**		**Fabric Softeners**		**Laundry Detergent**		**Paper Towels**		**Refrigerated Juices**	
Distance	0.016***	0.089***	-0.082***	-0.029***	0.009**	0.092***	0.050***	0.119***	-0.014	0.071***	-0.031***	0.084***
	[0.003]	[0.004]	[0.005]	[0.006]	[0.004]	[0.007]	[0.003]	[0.005]	[0.011]	[0.018]	[0.007]	[0.010]
Dist.*Share		-16.9***		-3.825***		-7.629***		-10.11***		-2.810***		-6.531***
		[0.694]		[0.231]		[0.610]		[0.599]		[0.368]		[0.251]
	Soft Drinks		**Snack Crackers**		**Tuna**		**Toothpaste**		**Bathroom Tissue**			
Distance	-0.037***	0.061***	0.167***	0.104***	-0.057***	0.007	0.045***	0.118***	-0.238***	-0.195***		
	[0.004]	[0.004]	[0.004]	[0.005]	[0.004]	[0.005]	[0.003]	[0.004]	[0.010]	[0.018]		
Dist.*Share		-37.2***		7.397***		-5.912***		-9.568***		-1.360***		
		[0.769]		[0.385]		[0.343]		[0.366]		[0.385]		

Summary of Results

Distance Level	With No Share Interaction	With Share Interaction	Share Interactions	
Significantly Negative	7	3	Significantly Negative	15
No Significance	2	1	No Significance	0
Significantly Positive	8	13	Significantly Positive	2

Notes: The columns report results of a fixed effects panel estimate that use the UPC-Store level as the unit of analysis. The regression is run seperately for each catagory. Each regression contains a vector of quarter dummies to control for seasonal effects. The coefficients are reported in percentage points per mile. T-Statistics are in brackets. Standard errors are clustered by store, and regressions are weighted by Revenue Share of Cell. * denotes significance at 10%; ** at 5% level and *** at 1% level.

Figure 3.3.3: Change in Frequency of Sales For 35 Mile Drop in Distance to Wal-Mart

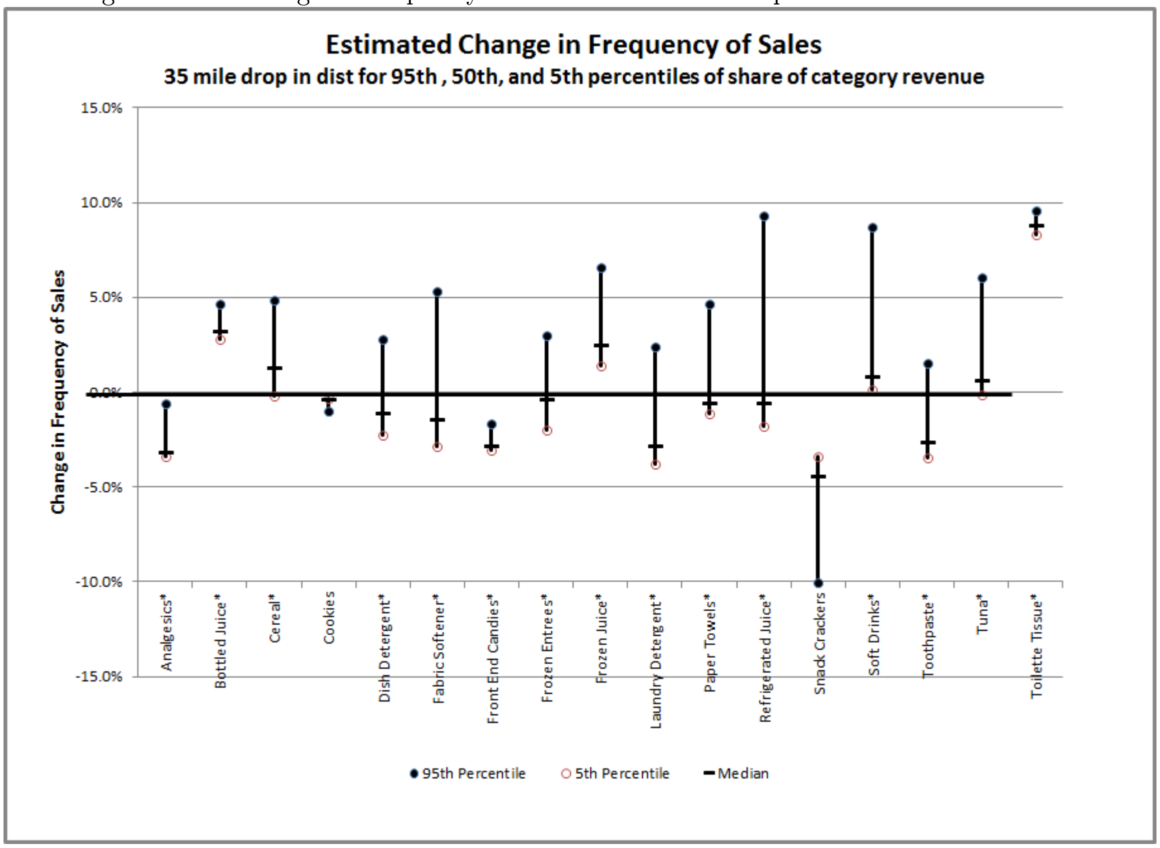

Notes: This figure plots the estimated effect of a 35 mile decline in distance to Wal-Mart (approximately the sample average) on the frequency of sales for each of three share percentiles, by category. The values are calculated by evaluating Equation (6) at different revenue share percentiles (5%, median, and 95%) for each category. Stars denote categories that have a negative and significant coefficient on share. The underlying coefficients and standard errors are available upon request.

Table 3.4: Regression Results Using Alternative Measures

	Store				Store-Category			
	Sales Discount		Markup		Sales Discount		Markup	
Distance	-0.051***	0.100***	0.025***	0.028***	-0.094***	0.059***	0.006	0.013**
	(0.007)	(0.006)	(0.003)	(0.003)	(0.008)	(0.006)	(0.004)	(0.004)
Linear Trend		0.024***		0.001		0.025***		0.001*
		(0.001)		(0.000)		(0.001)		(0.000)
Observations	30,977	30,977	31,856	31,856	519,673	519,673	542,758	542,758
	0.036	0.289	0.089	0.090	0.025	0.092	0.004	0.005

Notes: The first two columns report results of a fixed effects panel estimate of two different models that use the store as the unit of analysis. Sales discount is the percentage difference between sales price and regular price. Markup is the percentage difference of all prices from their cost. The store level model includes a vector of quarter dummies and the store-category model includes a vector of category x quarter dummies to control for seasonal effects for the chain and category respectively. The distance coefficients are reported in percentage points per mile. T-Statistics are in brackets. Standard errors are clustered by store, and regressions are weighted by Revenue Share of Cell. * denotes significance at 10%; ** at 5% level and *** at 1% level.

discount versus regular price) is not mechanically related to the frequency. DFF may have reduced sales depth in order to minimize the effect on average price of increasing frequency. On the other hand, they may have also increased the depth of sales as well as the frequency with the hope of attracting more shoppers.

We regress the sale discount for the store and store-category levels in Table 3.4. Contrary to the previous results, the effect of distance on sale depth depends on whether a trend is included. Wal-Mart's distance has a negative relationship with discounts when a trend is not included, but a positive relationship when it is. The sales discounts, therefore, increased across the entire Dominick's chain. Unless this increase was a chain-wide response to Wal-Mart, then the depth of sales seems to have declined as the frequency increased.

We also examine each store's average price markup $\frac{p-c}{p}$ to provide a view of the entire price response to Wal-Mart. Holding regular price constant, an increase in the frequency or depth of sales reduces average markup. However, DFF may have chosen to increase the *regular price* markup to mitigate the effect on of more frequent sales on average markup. We estimate that the effect of Wal-Mart on markup is unambiguously negative. This result is consistent with other studies of Wal-Mart entry that find incumbent stores reduced average prices in response to Wal-Mart. Moreover, the ambiguous results on the depth of sales lead us to conclude that much of the decline in markup was the result of an increase in the frequency of sales.

Conclusion

Drawing from related strands of research in the marketing and economics literature, we find that an increase in the frequency of sales can be a rational response to competition with a low cost retailer. The data from a representative chain of grocery stores support this strategy: stores which came into competition with Wal-Mart significantly increased their average frequency of sales. Moreover, the increased price promotion activity was focused on "loss-leader" products, providing additional evidence that the behavior was a strategic response to Wal-Mart entry rather than a coincident change in some other factor (e.g. promotion activity initiated by manufacturers).

This study has implications for two other areas of research. First, there have been several macroeconomic studies that evaluate the role of sales in price adjustment. The topic was initiated with the observation that prices change frequently, but that many of these changes are the result of sales (Bils and Klenow, 2004). Several recent studies attempt to reconcile the frequent adjustment of prices, that is largely due to sales, with the cornerstone assumption of price stickiness embedded in New Keynesian macroeconomic models (Eichenbaum et al., 2011; Guimaraes and Sheedy, 2011; Kehoe and Midrigan, 2010). These studies find that nominal rigidities are still important in spite of the frequent price adjustments associated with sales. Our results, however, caution against concluding that sales are unimportant for aggregate price adjustment because we show that temporary price reductions may be used in response to a persistent shock.

Second, Nakamura and Steinsson (2008) find that the fraction of price quotes that are sales has increased substantially over the last two decades. Certain product categories, such as breakfast cereal or potato chips, are now "on sale" twice as often as they were in the late 1980's. As the expansion of Wal-Mart took place over the same period, our results suggest that Wal-Mart could be at least partially responsible for the rise in sales.

REFERENCES

Ailawadi, K. L, J. Zhang, A. Krishna, and M. W Kruger, "When Wal-Mart Enters: how incumbent retailers react and how this affects their sales outcomes," *Journal of Marketing Research*, 2010, *47* (4), 577–593.

Basker, Emek, "Selling a Cheaper Mousetrap: Wal-Mart's Effect on Retail Prices," *Journal of Urban Economics*, 2005, *58* (2), 203–229.

_ and Michael Noel, "The Evolving Food Chain: Competitive Effects of Wal-Mart's Entry into the Supermarket Industry," *Journal of Economics and Management Strategy*, 2009, *18* (4), 977–1009.

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

Chevalier, Judith A. and Anil K. Kashyap, "Best Prices," 2011.

_ , _ , and Peter E. Rossi, "Why Don't Prices Rise during Periods of Peak Demand? Evidence from Scanner Data," *American Economic Review*, March 2003, *93* (1), 15–37.

Conlisk, John, Eitan Gerstner, and Joel Sobel, "Cyclic Pricing by a Durable Goods Monopolist," *Quarterly Journal of Economics*, 1984, *99* (3), 489–505.

DeGraba, Patrick, "The Loss Leader Is a Turkey: Targeted Discounts from Multi-product Competitors," *International Journal of Industrial Organization*, 2006, *24* (3), 613–628.

Eden, B. and M. Jaremski, "Rigidity, Dispersion and Discreteness in Chain Prices," *Working Papers*, 2010.

Eichenbaum, Martin, Nir Jaimovich, and Sergio Rebelo, "Reference Prices, Costs, and Nominal Rigidities," *American Economic Review*, 2011, *101* (1), 234–262.

Franklin, A. W, "The impact of Wal-Mart supercenters on supermarket concentration in US metropolitan areas," *Agribusiness*, 2001, *17* (1), 105–114.

Guimaraes, Bernardo and Kevin D. Sheedy, "Sales and Monetary Policy," *American Economic Review*, 2011, *101* (2), 844–876.

Hoch, Stephen J., Byung-Do Kim, Alan L. Montgomery, and Peter E. Rossi, "Determinants of Store-Level Price Elasticity," *Journal of Marketing Research*, 1995, *32* (1), 17–29.

—, **Xavier Drèze, and Mary E. Purk**, "EDLP, Hi-Lo, and Margin Arithmetic," *The Journal of Marketing*, 1994, *58* (4), 16–27.

Holmes, T. J, "The Diffusion of Wal-Mart and Economies of Density," *Econometrica*, 2011, *79* (1), 253–302.

Hosken, Daniel and David Reiffen, "How Retailers Determine Which Products Should Go on Sale: Evidence from Store-Level Data," *Journal of Consumer Policy*, 2004, *27* (2), 141–177.

Kehoe, Patrick and Virgiliu Midrigan, "Prices Are Sticky After All," *NBER Working Paper No. w16364*, 2010.

Lal, Rajiv, "Price Promotions: Limiting Competitive Encroachment," *Marketing Science*, 1990, *9* (3), 247–262.

— **and Carmen Matutes**, "Retail Pricing and Advertising Strategies," *Journal of Business*, 1994, *67* (3), 345–370.

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

Peltzman, Sam, "Prices Rise Faster Than They Fall," *Journal of Political Economy*, 2000, *108* (3), 466–502.

Singh, V. P, K. T Hansen, and R. C Blattberg, "Market entry and consumer behavior: An investigation of a Wal-Mart supercenter," *Marketing Science*, 2006, *25* (5), 457.

Varian, Hal R., "A Model of Sales," *American Economic Review*, 1980, *70* (4), 651–659.

Volpe, R. J and N. Lavoie, "The effect of Wal-Mart supercenters on grocery prices in New England," *Applied Economic Perspectives and Policy*, 2008, *30* (1), 4.

APPENDIX

3.A Demand in a Hotelling Model

Suppose there is a measure 1 of switchers who are distributed uniformly across the unit interval and differ only in their cost of visiting the entrant. Denote a switcher's type as $i \in [0,1]$. Switchers of type i face a cost of visiting the entrant of $d(i) = i\bar{d}$ where $\bar{d}$ is the highest cost any switcher incurs to visit the new store. The marginal type who would be indifferent between visiting the new store or not is $\tilde{i} = \frac{p_B - p_c}{\bar{d}}$. All switchers of type $i < \frac{p_B - p_c}{\bar{d}}$ purchase from the entrant, and the rest purchase from the lowest priced incumbent.

3.B Proofs

Proof of Proposition 1

To prove that choosing a price of r in every period is part of an SPNE in pure strategies, we propose a punishment for deviating and then check to make sure it is credible and effective. Because there is no pure strategy equilibrium in the stage game (except in very special cases), the punishment cannot involve reverting to a "bad" equilibrium forever.

Suppose the punishment for deviating is to charge a price equal to marginal cost, c, for t periods. In period $t+1$, both firms resume charging a price of r unless another deviation occurs. If either firm deviates during the punishment, the punishment starts over from the beginning. The duration of the punishment, t, is chosen to be as large as possible such that: $\frac{\delta^t\left(\alpha+\frac{1}{2}\right)(r-c)}{1-\delta} \geq \frac{\alpha(r-c)}{1-\delta}$ This inequality ensures that the punishment is credible. The RHS is the continuation value of charging r forever assuming that the opponent charges something less than r. The LHS is the present value of profits assuming that after the punishment, both players go back to charging r every period. Rearranging terms, we can see that the most severe punishment that is credible would be to choose the largest t such that:

$$\delta^t \geq \frac{2\alpha}{2\alpha+1} \tag{3.B.1}$$

For this threat to deter deviations, we must ensure that a one shot deviation is unprofitable. Therefore, the punishment will prevent deviations if:

$$\frac{\left(\alpha+\frac{1}{2}\right)(r-c)}{1-\delta} \geq (\alpha+1)(r-c) + \frac{\delta^t\left(\alpha+\frac{1}{2}\right)(r-c)}{1-\delta}$$

Which will be satisfied as long as:

$$\frac{1-\delta^t}{1-\delta} \geq \frac{\alpha+1}{\alpha+\frac{1}{2}} \tag{3.B.2}$$

Which implies that the lower bound for δ is $\frac{1}{2}$ for the SPNE to exist. Combining 3.B.1 and 3.B.2, we have:

$$\delta \geq \frac{\alpha-1}{\alpha+1} \tag{3.B.3}$$

The conditions under which the monopoly price can be supported are summarized by 3.B.3. The RHS of 3.B.3 is bounded above by 1 and increasing for $\alpha > 0$. We assume that the discount factor $\delta \in (0, 1]$. Thus, a largerα, requires that firms must be more patient support the monopoly price.

Maximum Profits Attainable by Incumbents

We claim that having a single incumbent charge a low price while the other charges the monopoly price results in the highest possible joint profit level for the incumbents. We prove this claim by contradiction. Suppose that the incumbents explicitly collude and set prices jointly. Suppose that the cartel found it optimal to set both prices to $p < r$. Let $\epsilon = r - p$. The cartel's profits in this case would be $2p\alpha + p\left(1 - \frac{p-p_c}{d}\right)$. If instead the cartel had a single firm charge $p < r$ and the other charge r, its profits would be $r\alpha + p\alpha + \frac{p(1-(p-p_C)}{d} = \epsilon\alpha + 2p\alpha + \frac{p(1-(p-p_C)}{d}$, which is $\epsilon\alpha$ more than if the cartel set both prices to p. Therefore, setting both prices below r cannot be optimal.

Proposition 2

We prove Proposition 2 in two sections. First, we establish the prices charged by all firms in the equilibrium proposed in Section 3.2. Next we show that the punishment strategy is both credible and effective.

Equilibrium Prices

Here we establish the price that C will charge in every period as well as the "sale" price that A and B will alternate with the monopoly price r. Recall that we assume C plays a best response to the lowest priced incumbent and the "sale" price is assumed to maximize the single period profits of the firm having a sale, given the price that C is charging.

Without loss of generality, we begin by assuming that $p_A \geq p_B$. Because firm C will always choose a price $p_C \in [p_B - d, p_B]$, its profits are:

$$\pi_C = p_C\left(\frac{p_B - p_C}{d}\right) \; if \; p_C \in [p_B - d, p_B]$$

π_C is maximized as long as one of the following conditions hold:

$$\frac{d\pi_C}{dp_C} = \begin{cases} \frac{p_B - 2p_C}{d} \leq 0 & \iff p_C = p_B - d \\ \frac{p_B - 2p_C}{d} = 0 & \iff p_C > p_B - d \end{cases}$$

which implies that firm C's optimal response function is:

$$p_C = p_B - d \text{ iff } p_B \geq 2d \tag{3.B.4}$$

$$p_C = \frac{p_B}{2} \text{ iff } p_B < 2d \tag{3.B.5}$$

W assume that only one of the incumbents will select a low price in any given period. It has been shown that such a strategy results in the maximum possible profits between the two incumbents. Now I wish to determine the "sale" price. Suppose, without loss of generality, $r = p_C \geq p_B$. The profits earned by firm B are:

$$\pi_B = \begin{cases} (p_B - c)\left(\alpha + 1 - \frac{p_B - p_C}{d}\right) & \text{iff } p_C + d \geq p_B \geq p_C \\ (p_B - c)(\alpha + 1) & \text{iff } p_C \geq p_B \geq 0 \\ (p_B - c)\,\alpha & \text{iff } r \geq p_B \geq p_C + d \end{cases}$$

We can disregard the second case because we have already argued that C would never find such a scenario optimal. Thus the relevant best response function for firm B is characterized by:

$$p_B = \frac{d(1+\alpha) + p_C + c}{2} \text{ iff } d(\alpha - 1) \leq p_C \leq \min\left\{d(\alpha - 1) + c,\, 2r - c - d(\alpha + 1)\right\} \tag{3.B.6}$$

$$p_B = r \text{ iff } p_C < r - d \tag{3.B.7}$$

Now we proceed to identify the prices p_B^* and p_C^* that are mutually best responses. These equilibrium prices will depend on the parameters r, d, c, and α. We are interested in the case in which $p_B < r$ which is only possible when condition 3.B.6 holds. There are two possible scenarios to consider. The first is when 3.B.4 also holds:

$$p_B = \frac{d(1+\alpha) + p_B - d + c}{2}$$

Which implies:

$$p_B = d\alpha + c\ ,\ p_C = d(\alpha - 1) + c$$

If this were an equilibrium, then C's share would be $\frac{p_B^* - p_C^*}{d} = 1$, which implies that the cartel would sell to none of the switchers. This is only optimal if $p_B^* = r$ because the cartel would only charge a price strictly less than r if they could sell to some of the switchers by doing so. Thus, the combination of 3.B.6 and 3.B.4 cannot represent an equilibrium where $p_B^* < r$.

The second scenario involving $p_B^* < r$ occurs when 3.B.5 and 3.B.6 hold:

$$p_B = \frac{d(1+\alpha) + \frac{p_B}{2} + c}{2} \text{ } and \text{ } p_B < 2d$$

This implies that:

$$p_B^* = \frac{2}{3}\left(d(1+\alpha) + c\right)\ ,\ p_C^* = \frac{1}{3}\left(d(1+\alpha) + c\right)$$

Next we ensure that the inequalities are satisfied. Condition 3.B.5 requires that:

$$\frac{2}{3}\left(d(1+\alpha) + c\right) < 2d \Rightarrow 2 - \frac{c}{d} > \alpha$$

Condition 3.B.6 requires:

$$\frac{1}{3}\left(d(1+\alpha) + c\right) \leq 2r - c - d(\alpha + 1) \Rightarrow \frac{3r}{2d} - \frac{c}{d} - 1 \geq \alpha$$

and

$$\frac{1}{3}\left(d(1+\alpha) + c\right) \geq d(\alpha - 1) + c \Rightarrow 2 - \frac{c}{d} \geq \alpha$$

Note that the first and last condition are the same. Thus, equilibrium price for firm C and the "sale price" for the incumbent are:

$$p_B^* = \frac{2}{3}\left(d(1+\alpha) + c\right)\ ,\ p_C^* = \frac{1}{3}\left(d(1+\alpha) + c\right)$$

$$iff\ \alpha \leq \min\left\{\left(2 - \frac{c}{d}\right)\ ,\ \left(\frac{3r}{2d} - \frac{c}{d} - 1\right)\right\}$$

Proof that punishment is credible and effective

To complete the proof of Proposition 2, we must show that the punishment outlined in 3.2.2 is credible and harsh enough to prevent the incumbents from deviating. In the analysis below, we take a different approach than Lal (1990) in order to add in marginal costs and address some technical issues. To compress the notation we define the following additional variables:

Δ_p , Δ_r =profit to an incumbent charging $\bar{r}$ when the other incumbent charges p, r respectively

δ_p , $\delta_{\bar{r}}$ = profits to a defecting firm when the other firm charges p, r respectively

$\Pi = \sum_{t=0}^{\infty} \delta^t r\alpha = \frac{r\alpha}{1-\delta}$ =discounted profits of selling only to your loyal customers forever

The punishment strategies given in 3.2.2 are credible ((1) and (2)) and effective ((3) - (5)) under the following conditions:

1. $\pi_A = r\alpha\delta^{t^*-1} + \Delta_r\delta^{t^*} + r\alpha\delta^{t^*+1} + ... = \frac{\delta^{t^*-1}}{1-\delta^2}(r\alpha + \delta\Delta_r) \geq \Pi$
2. $\pi_B = \Delta_p\delta^{t^*-1} + r\alpha\delta^{t^*} + \Delta_r\delta^{t^*+1} + ... = \pi_A \geq \Pi$
3. $\delta_p\delta^{t^*-1} + \delta^{t^*}\pi_A \leq \pi_B$
4. $r\alpha\delta^{t^*-1} + \Delta_r\delta^{t*} + \delta_{\bar{r}}\delta^{t^*+1} + \pi_A\delta^{t^*+2} \leq \pi_A$
5. $\Delta_r\delta^{t^*-1} + \delta^{t^*}\pi_B \leq \pi_B$

Conditions (1) and (2) state that the continuation value of the punishment sequence for the two incumbents (π_A and π_B) must be at least as large as the discounted profits from serving only loyal customers, Π. Conditions (1) and (2) also indicate how t^* and p are selected. t^* is chosen so that it is as large as possible without violating inequality (1), ensuring that the threat is as severe as it can be and still be credible. The price p is chosen to satisfy (2), that $\pi_A = \pi_B$.

Condition (3) is required so that A will not deviate in period t^*. Condition (4) ensures that A will not deviate in period $t^* + 2$. Finally, (5) ensures that B will not deviate in period t^*.

We now analyze when it is possible for conditions (1) – (5) hold. First notice that (3) always holds when (5) holds. This is simply because $\delta_p < \Delta_r$. Δ_r is the profit one incumbent makes when the other charges a price of r. The quantity δ_p is the profit that an incumbent could make if it were to deviate when the other is charging a price $p < r$. Since Δ_r is the best a firm can do when the other charges r, we know that $\delta_p < \Delta_r$. Therefore if (v) holds, so does (iii).

Table 3.5: Numerical Examples of Existence Conditions

$\frac{\Delta_r}{r\alpha}$	1.1	1.5	2	5	10
$\min\delta$	.96	.82	.83	.74	.69
$t^*(\min\delta)$	2	2	3	3	4

Because (3) is redundant, we analyze conditions (1) (2) (4) and (5) to find the parameter values for which they can be satisfied. Rearranging terms in (1) we can see that the duration of the punishment, t^* depends on $\frac{\Delta_r}{r\alpha}$. Specifically, t^* will be the largest integer that satisfies:

$$\frac{\Delta_r}{r\alpha} \geq \left(\frac{1+\delta-\delta^{t^*-1}}{\delta^{t^*}}\right)$$

To interpret this condition, notice that the RHS is increasing in t^* (because $\delta \in (0,1)$). Secondly, the ratio $\frac{\Delta_r}{r\alpha} > 1$ can be interpreted as a measure of the temptation to cheat when the firm is supposed to be charging r, the high price. The larger the temptation, the harsher is the punishment must be.

The next step is to use inequalities (4) and (5) to determine what values of δ make the threat severe enough. It turns out that the lower bound on δ also depends on the ratio $\frac{\Delta_r}{r\alpha}$. First, notice that $\delta_{\bar{r}}$ will be less than but arbitrarily close to Δ_r. If the incumbent deviates when her opponent charges $\bar{r}$, then the best she can do is to slightly undercut her opponent's price and obtain a profit slightly less than Δ_r. Because Δ_r is the upper bound on the single period profits earned by A if she deviates, then we can substitute Δ_r for $\delta_{\bar{r}}$ in (4) and still be certain that A will be deterred from deviating in period t^*+2. After the substitution, (4) and (5) can be written as:

$$\frac{\Delta_r}{r\alpha} \leq \frac{1-\delta^{t^*}}{1-\delta-\delta^2+\delta^{t^*+1}} \; if \; 1 > \delta+\delta^2-\delta^{t^*+1}$$

This is a bit tricky to interpret. When the RHS of the inequality is positive, then it must be larger than $\frac{\Delta_r}{r\alpha}$. When the RHS is negative, conditions (4) and (5) are always satisfied. The existence of an SPNE of the form described above depends on the magnitude of $\frac{\Delta_r}{r\alpha}$ which we know is larger than 1. In the table below, we provide different levels of $\frac{\Delta_r}{r\alpha}$ with the corresponding t^* for the minimum level of δ that make the strategy a credible and effective threat.

These results differ from those reported in Lal (1990). He claims that the strategy profile given in 3.2.2 is an SPNE as long as $\delta > .62$, regardless of the level of $\frac{\Delta_r}{r\alpha}$. We show here that the minimum possible discount factor depends on the size of the gains from alternating sales relative to the "outside option". Nevertheless, even small levels of profits gained by selling to the switchers will result in alternating sales if the interest rate is low enough.

The final step is to show that charging a price of r and $\bar{r}$ in alternating periods using the proposed punishment strategies is a Nash equilibrium. Said differently, a one shot deviation cannot be profitable. This will be true if:

$$\frac{1}{(1-\delta^2)}\left(r\alpha + \Delta_r\delta\right) \geq \delta_{\bar{r}} + \delta\pi_A \tag{3.B.8}$$

It is easy to show that 3.B.8 is satisfied if (5) is satisfied. Thus, the strategy profile in 3.2.2 constitutes a sub-game perfect Nash equilibrium as long as δ is large enough given $\frac{\Delta_r}{r\alpha}$.

3.C Dominick's Finer Foods Sample Selection

The DFF sample offers a large number of products and stores, but there are several stores, UPCs, and UPC-store cells with very few observations. The main concern is that we cannot be sure of the reason for sparsely populated data. For instance, a UPC-store cell with only one year (out of a possible seven) may represent a product deletion or incomplete data records. We want to make sure that the variation in the fraction of products on sale at a particular store is not affected by changes in the mix of available data. Attempting to balance this objective with the desire to use as much data as possibly, we implemented the following selection procedure:

1. Drop the final 18 weeks of the sample because it represents only a partial-year of data.

2. Drop any category that does not span the entire sample period.

3. Next, we break up categories based on the relative number of products.

   (a) For smaller categories (Bathroom Tissue, Bottled Juices, Cereals, Dish Detergent, Fabric Softeners, Front End Candies, Frozen Fruit Juices, Laundry Detergent, Paper Towels, Refrigerated Juices, Snack Crackers, Toothpaste, and Tuna)

      i. Drop all UPC-store cells with less than 165 observations

      ii. Drop any store with less than 40 products in a category

   (b) For large categories (Analgesics, Cookies, Frozen Entrees, and Soft Drinks)

      i. Drop all UPC-store cells with less than 180 observations

      ii. Drop any store with less than 50 products in a category

To illustrate how much of the data is excluded, Table 3.6 presents summary statistics before and after the sample selection is taken. Although we delete nearly 2/3 of the UPC-store cells, we still analyze 80% of the raw sample's revenue.

Table 3.6: Sample Selection Summary

Category	**Stores**		**UPCs**		**Store-UPCs**		**Revenue $Millions**		
	Selected	Raw	Selected	Raw	Selected	Raw	Selected	Raw	% of Total
Analgesics	85	93	320	641	15,163	40,775	29	39	75%
Bottled Juice	85	93	217	511	14,598	36,656	80	100	80%
Cereals	85	93	227	490	16,221	36,620	227	268	85%
Cookies	85	93	428	1,126	27,669	78,731	98	128	76%
Dish Detergent	85	93	125	287	8,428	22,005	33	46	73%
Front End Candies	85	93	222	503	14,956	34,450	35	43	81%
Frozen Entrees	85	93	378	898	25,133	67,459	130	162	80%
Frozen Juices	85	93	95	175	7,796	13,803	69	75	91%
Fabric Softeners	85	93	156	318	10,244	25,184	35	47	76%
Laundry Detergent	85	93	236	581	14,143	45,047	95	153	62%
Paper Towels	80	93	76	164	4,183	11,741	50	67	75%
Refrigerated Juices	85	93	102	227	7,023	17,212	134	157	85%
Soft Drinks	85	93	564	1,720	36,614	112,017	438	526	83%
Snack Crackers	85	93	180	425	12,425	30,603	58	71	82%
Tuna	85	93	157	278	9,757	19,969	45	53	85%
Toothpaste	85	93	288	608	14,281	39,263	23	31	73%
Toilette Tissue	81	93	57	128	3,627	9,867	81	100	81%
Total			3,828	9,080	242,261	641,402	1,660	2,066	
Percent			42%		38%				80%

Notes: Comparison of selected sample to raw sample. See description above for selection criteria