Informational Rigidities and the Stickiness of Temporary Sales

Eric Anderson
Northwestern University

Emi Nakamura
Columbia University

Duncan Simester
Massachusetts Institute of Technology

Jón Steinsson
Columbia University

December 31, 2013

Using a unique dataset from a large retailer, we study two types of retailer response to a demand or cost shock: (1) a change in the regular price, or (2) a change in the depth or frequency of discounts. We identify cost shocks using changes in the wholesale price together with changes in commodity prices, and identify demand shocks using cross-sectional variation in unemployment rates. We find that if the retailer responds it does so through the regular price rather than through changes in the depth or frequency of discounts. To explain these findings we present several institutional facts that highlight differences in the ways that regular prices and discounts are planned, funded and promoted. A key conclusion is that the frequent price changes associated with discounts are based on old information, and so they do not reflect a rapid response of prices to economic shocks. We argue that information "stickiness" for temporary sales means that temporary sales are "sticky plans" that are updated infrequently. Even though temporary sales lead to frequent price changes, the underlying mechanism that controls the frequency and depth of these price changes is sticky.

Evidence that a retailer only responds to macro shocks through its regular price and not through its sale prices suggests that we should exclude price changes due to sales when evaluating macro price flexibility. This sharply lowers the observed frequency of price changes. We also show that the decision to include or exclude price changes due to sales affects other salient features of pricing behavior. We conclude that including price variation from sales when measuring macro price flexibility may be misleading.

Keywords: Regular Retail Prices, Retail Sales, Trade Deals.

JEL Classification: E30, L11, M30.

This research has been supported by National Science Foundation grants SES 0922011 and SES 105610. The paper was previously circulated under the title: "“Temporary Sales: On Autopilot and “Funded” by Frequent Flyer Accounts”. We thank Ariel Burstein, Gee Hee Hong, Nicholas Li, and seminar participants at the 2012 NBER's Summer Institute, Price Dynamics Workshop.

1. Introduction

The speed of price adjustment to aggregate shocks is a central determinant of the effect of demand shocks on output and in particular the effects of monetary and fiscal policy. In an influential study, Bils and Klenow (2004) showed that consumer prices adjust quite frequently. Subsequent empirical work has shown, however, that much of this price flexibility is due to temporary sales, which have vastly different empirical characteristics than “regular” price changes (Nakamura and Steinsson, 2008). An important question is to what extent price changes associated with temporary sales contribute to the adjustment of the aggregate price level to aggregate shocks.

A growing literature on sticky information points out that even if prices do change, they may fail to respond to recent economic shocks if the information set the price changes are contingent on is old (e.g., Mankiw and Reis, 2002; Burstein, 2005). In these cases, the prices may be flexible but they follow “sticky plans” whereby pricing decisions are made only periodically. We show, both institutionally and empirically, that temporary sales are indeed based on such sticky plans.

We analyze this using an exceptionally detailed dataset on retail prices and wholesale prices from a large US retailer over the period 2006-2009. The data reflects the internal pricing database of the firm and explicitly identifies regular prices and sale (discounted) prices. This unique feature of our dataset avoids having to identify temporary sales using a “sale filter,” as in other papers in the literature. We augment this quantitative evidence with narrative evidence on the institutions used by retailers and manufacturers to set manufacturer trade deals and retailer temporary sales.

We begin by studying how this retailer responds to cost and demand shocks. We identify cost shocks using changes in the wholesale price together with changes in commodity prices, while demand shocks are identified using cross-sectional variation in unemployment rates. We use these shocks to compare two types of possible retailer response: (1) a change in the regular price, or (2) a change in the depth or frequency of discounts.

The findings indicate that if the retailer responds it does so through the regular price rather than through changes in its discounts. For example, we find that, in response to an increase in the base wholesale price, regular prices increase rapidly and completely. If temporary sales are also flexible, one might expect retailers would respond to an increase in marginal cost by *both* raising regular retail prices and reducing the frequency and size of temporary sales. However, temporary sales are completely unresponsive to wholesale price changes.

It is crucial in carrying out this exercise that our dataset allows us to distinguish between movements in the base wholesale price and wholesale price changes associated with trade deals. The “wholesale price” variable recorded in scanner price datasets often includes price reductions due to trade deals. However, such price reductions may not reflect a true change in the retailer’s marginal cost. As we will discuss, trade deal funds are often deducted from a fixed budget—implying that a promotion today comes at the cost of future promotions. If we were instead to focus on the relationship between retail price changes and wholesale prices including trade deals, we would come to quite different conclusions, since retailers are, in general, required to offer discounts and promotions to collect trade deal funds. This explains the difference in our results versus Eichenbaum et al. (2011), who use a measure of wholesale prices including trade deals, and find that temporary sales almost always coincide with wholesale price declines. Our analysis shows that retailers do not adjust the frequency or size of sales in response to changes in the base wholesale price, which is a true shock to the retailer’s marginal cost.

We next study how prices respond to manufacturer cost shocks. Our data include both a period of relative stability in manufacturer costs – 2006 and the first half of 2007 – and a period of highly volatile manufacturer costs due to the large run-up and subsequent fall in commodity costs in late 2007 through 2009. We find that the frequency of regular price increases roughly quadrupled in response to the sharp commodity price increases in 2008. In contrast, we find no response of the frequency and depth of temporary sales to these cost shocks. The pricing policy we observe is, therefore, consistent with a state-dependent rule for regular prices, and an unresponsive plan for temporary sales.

We use the cross-sectional aspect of our data to study the relationship between pricing behavior and the unemployment rate, allowing for time fixed effects. This identification strategy relies only on relative movements in prices and unemployment rates across different geographic locations, and is analogous to the one used by Coibon et al. (2013). Like Coibon et al. (2013), we find that an increase in unemployment does not lead to an increase in temporary sales.[1] However, we do find that this retailer increases Regular Retail Prices more frequently as unemployment rates increase. One reason for this result may be that in times of low demand managers raise prices to meet short term revenue and profit targets. In the short-run, this strategy can be effective because demand is relatively non-responsive to changes in the regular retail price (Hoch, Drèze and Purk 1994).

The distinction between temporary sales and regular price changes would be less important if both sources of price changes led to the same conclusions about macro price flexibility. This is not the case; including price variation due to temporary sales as a measure of price flexibility may lead to misleading conclusions about the response of the aggregate price level to macro shocks. In particular, the frequency of price changes in retail data depends critically upon whether we include or exclude price changes due to sales; if we exclude price changes due to sales then we observe a sharp reduction in retail price flexibility. Our evidence that a retailer only responds to macro shocks through its regular price and not through its sales prices suggests that we should exclude price changes due to sales.

We also investigate the extent to which temporary sales contribute to two other features of pricing behavior: cross-sectional heterogeneity in price flexibility, and the slope of the hazard function of price change. Excluding temporary sales dramatically reduces the extent of cross-sectional heterogeneity in the frequency of price change for retail products. Most products adjust their regular and base wholesale prices roughly once every eighteen months during tranquil times and more frequently when costs move rapidly. In contrast, the frequency of temporary sales varies dramatically across products even in tranquil times (but not at all over time). We also show that including temporary sales dramatically changes estimates of the hazard function of price adjustment. While the hazard function of Retail Prices (including sales) is downward sloping, the hazard function of Regular Retail Prices is upward sloping.[2]

---

[1] Coibon et al. (2013) find that the frequency and size of temporary sales actually drops when the unemployment rate rises. We find that the frequency drops and size rises. However, both effects are statistically insignificant.

[2] Our accurate measure of the (true) Regular Retail Price provides an opportunity to investigate different "sale-filters" designed to purge a Retail Price series of temporary sales. These findings reveal that the accuracy of

To help interpret and explain these findings we conclude the paper by presenting several institutional facts about pricing and trade promotions. We focus first on the scheduling process. Due to the logistical complexity of executing successful temporary sales and associated promotional activity, manufacturers and retailers jointly set a schedule for temporary sales – a promotion calendar – through an annual planning process. This means that temporary sales follow sticky plans.[3]

The most common way that manufacturer trade deal budgets are determined is via accrual accounts, which are analogous to frequent flyer accounts. Just as consumers accumulate "miles" when they fly on, say, United Airlines, retailers accrue funds in a manufacturer trade deal budget for the total volume purchased from a manufacturer. Temporary sales are "funded" out of these accrual accounts. In other cases, the manufacturer's trade deal contributions are established as part of the annual negotiation process. As a result, trade deals typically do not represent commensurate drops in the retailer's marginal cost, because the retailer is "spending down" a finite trade deal budget.

Payment from the accrual account is typically contingent on execution of a trade deal. Retailers typically receive money after there is verification of a price discount, in-store signage or advertising of the manufacturer's project. This means that the trade deal money funding may be more appropriately interpreted as a contracted payment for services rendered (including the change in the price), rather than a reduction in the marginal cost that the retailer decides whether to respond to through a temporary sale. In some scanner datasets used in the macroeconomics and industrial organization literatures the wholesale price variable is confounded by trade deal funding. If trade deals are included in the wholesale price series, the data must be interpreted

---

the filtered price measure depends critically on the design of the sale-filter. We design our own filter and also investigate the quarterly mode filter recently proposed by Eichenbaum et al. (2011). Both filters are effective at recovering an upward-sloping hazard function from the Retail Price series. In contrast, the simple V-shaped filter used elsewhere in the literature is not as effective at filtering out temporary sales and yields a downward-sloping or flat hazard.

[3] In a separate study, two of the authors were involved in manipulating the depth of temporary discounts on a sample of items. Even though they were merely varying the prices (and not which items were to be discounted), the lead time on making these decisions was almost four months. We discuss these issues in greater detail in Section 8.

with great caution.

Our paper is related to several recent papers that study the behavior of regular prices and sales. Eichenbaum et al. (2011), show that "price plans" consisting of a small set of prices are quite inertial and argue that this inertial behavior can be represented by the behavior of a "reference price." Coibon et al. (2013) argue that the frequency of temporary sales decreases when unemployment rates increase, yet the inflation rate for effective price paid by consumers decreases. They document consumer switching between retailers that reconciles these two seemingly contradictory facts. Chevalier and Kashyap (2011) consider the implications of temporary sales for price indexes. Klenow and Willis (2007) argue that the size of sale-related price changes is more responsive to inflation than the size of regular price changes. Hendel and Nevo (2011) estimate a micro-founded model in which temporary sales arise out of a motive to price discriminate between more and less price sensitive consumers. Chu and Nevo (2013) show that households had a higher propensity to buy on sale during the Great Recession.

Beyond these studies, a recent literature investigates other arguments, complementary to ours, for why it may be important to distinguish between regular prices and sales in measuring the flexibility of prices. Kehoe and Midrigan (2012) point out that even if temporary sales were completely responsive to movements in underlying costs, the temporary nature of sales implies that they contribute much less to the adjustment of the aggregate price level than regular price changes. Guimares and Sheedy (2011) develop a model in which temporary sales are strategic substitutes, implying that they tend to average out in the cross-section, and again limiting their impact on aggregate statistics.

The paper proceeds as follows. Section 2 describes the data. Section 3 presents summary statistics on price change. Section 4 presents our analysis of retail price movements in the days surrounding a base wholesale price change. Section 5 presents time-series analysis of the frequency of regular price changes, wholesale price changes and sales during the tranquil pre-2008 period in comparison to the subsequent commodity price run-up. Section 6 contains analysis describing how prices change as the unemployment rate changes across geographic regions. Section 7 describes the implications of temporary sales for cross-sectional heterogeneity and the hazard function of price adjustment. Section 8 discusses the institutions of manufacturer trade deals and retailer temporary sales, and the implications of the institutions for how to interpret wholesale cost series in standard datasets. Section 9 concludes.

2. Data

Our data come from a large retailer that sells products in the grocery, health and beauty, and general merchandise categories. It contains 195 weeks (15 quarters) of store transactions at a sample of 102 stores. The 195 weeks extend from the first quarter of 2006 through the end of the third quarter of 2009.[4]

For this sample, we have data on the number of units sold each week for each product at the Stock Keeping Unit (SKU) level at each store.[5] The dataset reports three price measures: (1) the Regular Retail Price, (2) the Retail Price that was actually paid (including any temporary sales), and (3) the Base Wholesale Price of the item.

A unique advantage of our data is the availability of the Regular Retail Price variable. Related data sets, such as the Dominick's data, track revenue but do not typically record the regular (or "shelf") price. As a consequence, syndicated data providers such as SymphonyIRI and Nielsen have created algorithms to impute the regular price from the observed average prices. Analogous sale filters have been adopted by academics (e.g., Nakamura and Steinsson, 2008; Chahrour, 2011, Kehoe and Midrigan, 2012). These imputation algorithms will, however, naturally introduce some noise into the regular price variable.

Differences between the Retail Price and the Regular Retail Price are attributable to temporary sales. The overwhelming majority of temporary sales are "advertised promotions" (88%). These are advertised in the store's weekly flyer. Seasonal markdowns account for 4% of temporary sales. Only 1% of temporary sales are due to clearance sales.

---

[4] The stores were selected as a control group for a pricing test conducted by the retailer and are considered representative of the retailer's stores. The stores are located in 14 Mid-West and East Coast states. Because they are in different "price zones," the Regular Retail Price and the Retail Price (including temporary sales) for an SKU in a given week is not always the same at all stores.

[5] We exclude "Direct Store Delivery" (DSD) categories (primarily alcohol, beverages and dairy), which represent approximately 2% of the observations and have very different institutional features. We discuss the implications of excluding the direct store delivery (DSD) categories in Appendix A.

An important departure from previous studies is that we measure unanticipated cost shocks using a Wholesale Price measure that does not include discounts associated with trade promotions. In the late 1960's and early 1970's trade promotions could be interpreted as unanticipated cost shocks to retailer. As we will explain in detail in Section 8, the way that trade promotions are now planned and funded means that this interpretation is generally no longer correct. Instead, trade promotions are now planned events that are predictable well in advance as part of a joint manufacturer-retailer annual promotion planning cycle. The retail response to the trade promotion is also generally negotiated in advance. Interpreting retail price changes that coincide with trade promotions as a retail response to an unanticipated cost shock is clearly inappropriate. In contrast, changes in the Base Wholesale Price are negotiated events that are immediately reflected in the Base Wholesale Price measure (unlike datasets in which the wholesale price is an "average acquisition cost"). Moreover, these Base Wholesale Price changes are not conditional on a retail price response. As a result, Base Wholesale Price changes can be thought of as unanticipated cost shocks that the retailer decides whether (and how) to respond to.[6] We believe that the availability of this Base Wholesale Price measure is a key advantage of the data that we use in this paper. It will serve an important role in explaining why our results differ in places from previous findings.

As in other scanner price datasets, each of the price measures in our dataset represents a weighted average over all of that week's retail transactions for the item in that store. This implies that if the Regular Retail Price changed in the middle of the week then the price we observe is an average of the price before and after the change, weighted by the number of items purchased at the different price levels. Similarly, if an item was temporarily discounted but customers only received the discount when they presented the retailer's frequent shopping card then the Retail Price would represent a weighted average of the price paid by customers who received the discount and those who did not. To avoid double-counting price changes that occur mid-week when calculating the frequency of price changes, we exclude price changes less than 1-cent in magnitude, and price changes that are in the same direction as a price change in the immediately

[6] The retailer that we study uses the Base Wholesale Price (without trade promotions) as its measure of marginal product cost. It tracks promotion "funding" separately through a stand-alone system that is used entirely for trade promotion planning and evaluation.

preceding week.[7] If there were no transactions for a SKU at a given store in a given week then we do not observe any price information. To avoid biases that arise because we are more likely to observe prices if products are on sale (there is more likely to be a transaction), we restrict attention to observations in which there are at least one quarter (13 weeks) of consecutive observations.[8] This results in a sample of 1,255,832 weekly observations, where the unit of observation is the price of a SKU at a store in a given week.

Many items have multiple color or flavor variants (e.g. orange versus mint flavored 1oz tic tac candy). The individual flavors of 1oz tic tac candy are identified at the SKU-level, while all of the flavors of 1oz tic tac candy share a common item-number. All SKUs under a single item-number have the same Regular Retail Price at a store in any week. They also share any temporary discounts and generally have the same Wholesale Price (although in some cases there is small variation in the Wholesale Price across different SKU numbers that share the same item-number). In our analysis we will cluster standard errors at the item level to account for the interdependence in price movements across stores and/or across SKUs that share the same item number.

Finally, in our analysis of the reaction of retail and wholesale prices to underlying costs and regional unemployment, we use additional data on the spot price of a gallon of diesel and CBSA-level unemployment rates. The diesel price data was downloaded from the US Energy Information Administration website and is for the Los Angeles price of a gallon of ultra-low-sulfur number 2 diesel fuel. The CBSA-level unemployment rates were obtained from the Bureau of Labor Statistics Local Area Unemployment Statistics program. In cases where the stores in our main dataset were located in rural areas that were not part of a CBSA for which an unemployment rate was available, we manually matched the store with the closest CBSA, and used the unemployment rate for that CBSA.

---

[7] If a price reduction occurs in the middle of week t and continues in week t+1 then the average price paid will be *both* lower in week t than week t-1, and lower in week t+1 than week t. This introduces a risk of double-counting price changes. Excluding price changes in the same direction as a price change in the immediately preceding week addresses this.

[8] Figure A1 shows the impact of varying the required number of contiguous observations. Including shorter price sequences causes the frequency of retail price decreases to go up, consistent with the idea that we are oversampling sales.

3. Summary Statistics

Using a representative item in our dataset, we illustrate in Figure 1 the three price measures in our dataset: Retail Prices, Regular Retail Prices and Base Wholesale Prices. The figure reveals that Wholesale and Regular Retail Prices exhibit similar dynamics, adjusting both infrequently and persistently, while Retail Prices including sales adjust much more frequently due to the presence of temporary sales.

Table 1 quantifies these (well-known) facts for a broader sample of products. Table 1 reports the frequency and size of price changes for all three price measures, weighted by total unit sales for each item. These statistics are all calculated as weighted averages of the 1,255,832 SKU x Store level observations, using total unit sales (across the 195 week data period) as weights. While not shown in Table 1, it is important to note that most temporary sales are very short in duration. Over 52% of temporary sales last just 1 week and the average duration of temporary sales is 1.9 weeks.

The weekly frequencies of price change for Base Wholesale and Regular Retail Prices are 1.13% and 1.22% respectively (this implies durations of 20 and 19 months respectively). As has been observed in many other datasets, the frequency of retail price change including sales is an order of magnitude higher—26.35% per week (implied duration of 0.9 months). Retail price changes including sales are also far larger than Base Wholesale or Regular Retail price changes (roughly 27% vs. 8%).[9] These findings are robust to using unweighted averages and to relaxing the restriction on price changes in the immediately preceding week.

Although these findings are well-established, they provide important motivation for the analysis that follows. When evaluating retail price flexibility in response to cost or demand shocks, macro-economists are faced with an important decision: whether to include or exclude price changes resulting from sales. If a retailer only responds to these shocks through its Regular Retail Prices and not through its sale prices, then we should exclude price changes due to sales.

[9] The sizes of the price changes are calculated as a percentage of the midpoint of the current and previous weeks' prices, ensuring that the calculation does not introduce an asymmetry in the magnitude of price increases versus price decreases.

As we illustrated in Table 1, this will sharply reduce measure of observed price flexibility. In the next three sections of the paper (Sections 4 through 6) we investigate whether this retailer responds to cost and demand shocks by changing its Regular Retail Prices and/or its sale prices.

4. Do Sales Respond to Wholesale Costs?

If sales represent an additional dimension of flexibility for retailers to respond to underlying movements in costs, then when wholesale costs increase we should expect to see changes in both the Regular Retail Prices and the size or frequency of sales.[10] In this section, we investigate the evolution of retail prices surrounding changes in the Base Wholesale Price to investigate whether this, in fact, occurs.

Figure 2 presents the average change in Regular Retail Prices and Base Wholesale Prices in the 13weeks before and after a Base Wholesale Price increase (Panel A) and decrease (Panel B).[11] Panel A shows that Regular Retail Prices respond to Wholesale Price increases immediately, and the immediate "pass-through" is roughly "cent-for-cent."[12] Panel B shows that Wholesale Price decreases also yield an immediate reaction in Regular Retail Prices, but the extent of pass-through is much lower (Anderson et al., 2012 and McShane et al. 2013).

Figure 3 plots the corresponding movements in the frequency and depth of temporary sales: Panels A and B present the results for increases and decreases, respectively. The Discount Frequency is the proportion of times that the Retail Price is less than the Regular Retail Price, and the Discount Depth is the percentage difference in these prices. In contrast to the rapid and

[10] See, for example, Kehoe and Midrigan's (2012) model of sales for an example of this logic.

[11] The figure is constructed as follows. For each product, the Wholesale and Regular Retail Price series are first divided by the average Regular Retail Price in the 13 weeks prior to the Wholesale Price change. We then take an average of each of these rescaled prices across products. Finally, for each of the resulting average price series, we subtract from the series the value in the week prior to the Wholesale Price change, so that all three series have a value of zero in the week prior to the Wholesale Price change. For this analysis, we restrict attention to items without missing values in the 13 weeks before and after the Base Wholesale Price change. We weight the observations by the total units sold and randomly select a single Wholesale Price change to represent each SKU x Store (to avoid over-weighting a single SKU x Store that has many Wholesale Price changes). We exclude Base Wholesale Price changes of less than 1-cent.

[12] This result is consistent with the results of Nakamura and Zerom (2010) and Goldberg and Hellerstein (2013), who find that retail prices respond quickly to wholesale price changes.

complete response of Regular Retail Prices to Wholesale Price increases, sales appear unresponsive, though quite noisy.

Perhaps what appear to be small movements in the frequency of discounts nevertheless may have economically important implications for quantities because of an extreme sensitivity of purchases to discounts. To investigate this issue, we constructed a measure of the Percent Saved through Discounts, defined as the percentage by which revenues over a given period would have been higher if the retailer had sold all the units at the Regular Retail Price instead of the Retail Price. Figure 3 also plots this measure. Even measured in this way, there is no evidence that temporary sales respond to movements in the wholesale cost.

Next we present analogous results in regression form. Table 2 presents results from estimating the following Weighted OLS regression using the same sample of observations used to construct Figures 2 and 3:

$$Y_{ist} = \sum \mu_i + \beta_1 \textit{Post Wholesale Price Change}_{ist} + \beta_2 \textit{Trend}_t + \varepsilon_{ist}, \qquad (1)$$

where Y_{ist} is the weekly Regular, Retail or Base Wholesale price change measure, the Discount Frequency, Discount Depth, or Percent Saved through Discounts (for item *i* in store *s* in week *t*). The μ_i terms are item fixed effects, *Post Wholesale Price Change* is a dummy variable set to 1 in the weeks after the Base Wholesale price change, and $\textit{Trend}_t$ is a linear trend. The effects are presented as a percent of the average Regular Retail Price in the 13 weeks prior to the Wholesale Price change. The unit of analysis is a SKU x store x week, and the observations are weighted by the total unit sales of the SKU in the store. We cluster the standard errors by item to account for any correlation in the errors across stores and/or SKUs that share the same item number. The top half of Table 2 presents results for wholesale price increases, while the bottom half presents results for wholesale price decreases. Each column in the table reports results for a different dependent variable.

The estimated change in the Regular Retail, Retail and Base Wholesale prices are close to identical following an increase in the Base Wholesale Price.[13] In contrast, there is no statistically significant increase in Discount Frequency, Discount Depth or the Percent Saved through

[13] Recall that changes in all three price series are measured as a percentage of the baseline (the average Regular Retail Price before the Wholesale Price change).

discounts. In the case of Wholesale Price decreases, none of the retail price variables change significantly.[14]

We conclude that for both Wholesale Price increases and decreases there is no change in the depth or frequency of temporary discounts. In response to Base Wholesale Price increases, almost all of the temporal variation in the Retail Price (including sales) can be attributed to changes in the Regular Retail Price and not to changes in sale activity. Base Wholesale Price decreases lead to increases in retail markups.

One notable feature of Figure 2 is that Regular Retail prices start to increase in anticipation of a Base Wholesale Price increase, and continue to increase slightly after the Base Wholesale Price increase. Upon discussing this issue with managers at the Retailer, we learned that the Retailer requires 60 to 90 days' notice of Wholesale Price increases and sometimes uses this advance warning to increase the Regular Retail Price in anticipation of the Wholesale Price increase. In almost of the cases in which the Regular Retail Price increased in the weeks before the Wholesale Price increase, the Regular Retail Price was not changed again in week 0 (the week of the Wholesale Price increase). Similarly, the increases in the Regular Retail Price after the week 0 almost all appear to be delayed price increases—for almost all of these items the Regular Retail Price did not change in week 0.

5. Do Sales Respond to Production Costs?

Could it be that sales respond to underlying production costs even though they do not respond to Base Wholesale Prices? To investigate this, we leverage the fact that our sample period incorporates a period of rapid rise and fall in the price of oil and other commodities in 2007-2009. To the extent that temporary sales are used to respond to underlying movements in production costs, we should expect to see the frequency and depth of discounts change in response to the commodity cost fluctuations.

Figure 4 (Panel A) plots the average weekly frequency of Regular Retail and Base

[14] This result lines up well with other evidence that prices are more responsive to underlying cost increases than decreases (e.g., Peltzman, 2000).

Wholesale Price increases (left axis), along with changes in diesel prices (right axis) on a biannual basis. Panel B presents analogous statistics for temporary sales. The frequencies of price change are weighted by total unit sales, and are adjusted for the seasonal pattern observed in 2006.[15] The diesel price variable is the 12-month change in diesel prices, lagged by one quarter.[16]

Figure 4 (Panel A) shows that the increase in diesel prices in 2008 was matched by a sharp rise in the frequency of Wholesale Price increases (the correlation between the two series is 0.72). In conversations with managers at the retailer, they attributed the spike in the frequency of Wholesale Price increases in 2008 to the commodity price changes. The frequency of Regular Retail Price increases also spikes sharply at this time. In stark contrast, Panel B shows that the frequency and depth of temporary sales as well as the Percent Saved through Discounts were seemingly unaffected by the huge run up and subsequent fall in diesel and other commodity prices.

Table 3 quantifies these findings using Weighted OLS regression,

$$Y_{ist} = \sum\mu_i + \beta_1 \mathit{Change\ in\ Diesel\ Price}_t + \beta_2 \mathit{Qtr\ 2}_t + \beta_3 \mathit{Qtr\ 3}_t + \beta_4 \mathit{Qtr\ 4}_t + \beta_5 \mathit{Trend}_t + \varepsilon_{ist}, \qquad (2)$$

where Y_{ist} is the weekly frequency of Regular Retail or Base Wholesale price changes, the Discount Frequency, Discount Depth, or Percent Saved through Discounts (for item i in store s in week t). The μ_i terms are item fixed effects. The *Change in Diesel Price* is measured in dollars and is constructed as above (without adjusting for seasonality). In particular, we calculate the difference in the average price in the current week compared to the previous year, and then lag this difference by a quarter (to account for any lead-time in the timing of pricing decisions). *Trend* is a weekly time-trend and the *Qtr* variables are quarter dummies. The unit of observation is a SKU x store x week, and the 1,255,832 observations are weighted by the total unit sales of the SKU in the store. The standard errors are again clustered by item.

[15] The frequency of price change is considerably higher in the first quarter of the year than in other quarters, consistent with the seasonal pattern found in Nakamura and Steinsson (2008). To adjust for this pattern, we subtract from each frequency statistic the average frequency of price change for that time period in 2006, relative to the overall frequency for that year.

[16] Lagging the change by a quarter recognizes that there is a lead-time between the timing of the diesel price change, the timing of the pricing decisions, and the implementation of those decisions.

Table 3 shows that a $1 increase in the price of a gallon of diesel (compared to 12-months earlier) is associated with a 0.844 percentage point increase in the weekly frequency of Wholesale Price increases and a 0.695 percentage point increase in the weekly frequency of Regular Retail Price increases. These correspond to roughly a 92% lift in the frequency of Wholesale Price increases, and a 63% lift in the frequency of Regular Retail Price increases relative to their average values (the observed weekly frequencies of price change in Table 1 are 0.92% and 1.11%).

In contrast, changes in the price of oil are not associated with any statistically significant changes in the frequency or depth of discounts. A $1 increase in the price of diesel is associated with an 0.852 percentage point reduction in the weekly frequency of discounts and a 0.232 percentage point increase in their average depth (the average frequency of sales over the sample period is 21.74% and the average depth is 24.01%). Neither of these effects approaches statistical significance. The Retail Price response to the oil shock can, therefore, be attributed entirely to changes in the Regular Retail Price, as opposed to the level of sale activity.

We might again worry that small movements in the frequency or depth of sales could have economically important implications. However, when using the Percent Saved through Discounts as a dependent variable in our regression we see no evidence that this measure is associated with the commodity cost fluctuations.

We note that our conclusion that sales are unresponsive to diesel prices relies on a null result, rather than a statistically significant finding. While it is generally difficult to draw strong conclusions from null results, we can be somewhat confident in doing so as our analysis uses such a large dataset. With over 1.2 million observations, even small effects can be estimated precisely.

6. Cross-Sectional Evidence on the Responsiveness of Sales

Next we study the responsiveness of the different forms of price change to variation in unemployment across different Core Based Statistical Areas (CBSAs).[17] Table 4 presents results from the following Weighted OLS regression,

$$Y_{ist} = \sum\mu_{it} + \beta_1 \textit{Change in Unemployment}_{st} + \beta_2 \textit{Qtr } 2_t + \beta_3 \textit{Qtr } 3_t + \beta_4 \textit{Qtr } 4_t + \beta_5 \textit{Trend}_t + \varepsilon_{ist}, \qquad (3)$$

where μ_i is an item x time (week) fixed effect. *Change in Unemployment*$_{st}$, is defined in a similar way to the change in diesel prices in the previous section. In particular, it is calculated as the 12-month change in the monthly unemployment rate in that CBSA, lagged by one quarter. We use the same sample of 1,255,832 observations weighted by the total unit sales of the SKU in the store. This identification approach is analogous to the one used by Coibon et al. (2013). It identifies the impact of changes in unemployment on retail prices solely using variation in regional unemployment (where the regions are represented by the store locations). The Base Wholesale Price does not vary across stores at a given point in time, and so we only estimate equation (4) for the retail price measures.

We find a statistically significant effect of the CBSA-level unemployment rate on the frequency of Regular Retail Price increases. Notice that the effect goes in the "wrong" direction from what a macroeconomist might have predicted: a 1 percentage point increase in the CBSA-level unemployment rate is associated with a 0.048% *increase* in the frequency of Regular Retail Price increases. The retailer may be using these Regular Retail Price increases to help meet short-term revenue or profit goals.[18] In the short-run, demand for consumer package goods is relatively inelastic to increases in the Regular Retail Price. This inelasticity is well-documented and is explained either by customers not immediately noticing the price increases, or by the cost of switching retailers when they do notice the price changes. For example, Hoch, Drèze and Purk (1994) report the findings from a randomized field experiment in which Regular Retail Prices were increased (or decreased) by 10% in 26 product categories. The price increases led to a demand decrease of just 3% over the next 16 weeks, resulting in

---

[17] Core Based Statistical Areas are geographic areas consisting of a county or set of counties that include a core urban area with a population of at least 10,000 people and the surrounding areas that are linked to the core by a high degree of social and economic integration as measured by commuting patterns. Core Based Statistical Areas are defined by the Office of Management and Budget.

[18] Another possibility is that the firm increases prices more frequently, but the price increases are smaller in magnitude. However, further investigation reveals that this is not the case - there is no change in the average size of the Regular Retail Price increases (or the frequency of Retail Price decreases).

short-run increases in both revenue and profit. In our data, increased unemployment was associated with a reduction in purchase quantities. Increasing the Regular Retail Price is a simple solution for managers who want to ensure that they meet short-run revenue and profit goals.

We estimate that an increase in unemployment is associated with a decrease in the Frequency of Discounts and an increase in the Depth of Discounts. However, neither effect is statistically significant. Our estimates, furthermore, indicate that changes in the unemployment rate lead to virtually no change in the Percent Saved through discounts. Our estimates thus support the finding of Coibon et al. (2013) that and increase in unemployment does not lead to an increase in temporary sales. Coibon et al. (2013) find the opposite effect of changes in unemployment on temporary sales, while our results indicate no effect.

Institutional features provide a straightforward explanation for why Regular Retail Prices are more responsive to regional economic variation than temporary discounts. Almost all discounts are advertised, while Regular Retail Price changes are not, and so the discounted price is the same for all stores within a single advertising zone. These advertising zones are large, while Regular Retail Prices are fixed within pricing zones that are considerably smaller.[19] This provides more degrees of freedom with which to tailor Regular Retail Prices to local conditions.[20]

In Section 8 we provide describe several additional institutional features that help to explain why retailers do not respond to cost or demand shocks by changing the depth or frequency of sales. However, before presenting these features, we first complete our empirical analysis by asking whether the source of price changes (changes in the Regular Retail Price versus temporary sales) matters when using retail data to evaluate macro price flexibility.

---

[19] The interdependence of prices due to these zones provides an additional reason for clustering the standard errors at the item level. This accounts for any correlation in the error term across stores.

[20] In Table 4 the Discount Frequency result approaches significance ($p < 0.15$) and so we further investigated this model by isolating whether there may have changes in the depth of discounts on some items. We did find some support for this. However, the effects were small and were limited to unadvertised discounts (recall that these represent less than 10% of the discounts in the sample). Notably we did not find any evidence that changes in unemployment were associated with any changes in the Percent Saved through Discounts.

7. Sales, Heterogeneity and the Hazard Function of Price Adjustment

Perhaps the most important consequence of the results in the previous sections are their implications for our understanding of retail price flexibility. If a retailer only responds to macro shocks through its regular price and not through its sales prices, then we should not consider price changes that result from sales when measuring the extent of price flexibility. As we illustrated in Figure 1, excluding sale price changes sharply reduces the frequency of observed price changes.

Excluding price changes due to sales also changes other features of pricing behavior that macroeconomists have emphasized. In particular, we will re-examine the level of cross-sectional heterogeneity in the frequency of price change, and whether the slope of the hazard function of price changes is downward-sloping. These features of the data both imply that price changes are more "bunched" than if they were distributed randomly.

Such bunching matters for the macroeconomic effects of price rigidity because it amplifies effects of price rigidity. Intuitively, bunching of price changes implies that price changes are "wasted" in terms of being opportunities for prices to adjust because many of them occur soon after another price change and thus at times when prices have already adjusted to most shocks. Carvalho (2006) and Nakamura and Steinsson (2010) show that cross-sectional heterogeneity amplifies the effects of monetary shocks. Similarly, Kehoe and Midrigan (2012) show that the "bunching" in the timing of price changes associated with temporary sales implies that sales contribute much less to aggregate price flexibility than "regular" price changes.

7.1 Cross-Sectional Heterogeneity in the Frequency of Price Change

Temporary sales dramatically amplify the extent of cross-sectional heterogeneity in the frequency of price change. Figure 5 presents a histogram of the weekly price change frequency in the 2007 calendar year for Wholesale and Regular Retail Prices (Panel A) and for Retail Prices (Panel B). The unit of analysis in these histograms is a SKU in a store and we weight the observations by total unit sales.[21] In the case of Regular Prices, 99% of products have a weekly

[21] The 1,255,832 observations are first aggregated into 18,344 SKU x Store measures of price change frequency. The histogram is then constructed using these 18,344 observations (weighting by total unit sales).

frequency of price change between 0 and 10%. In contrast, when sales are included, 65% of products have a weekly frequency of price change greater than 10%.

Another way of quantifying the effect of sales on cross-sectional heterogeneity is to compare the standard deviation of the frequency of price change with and without sales. Across SKUs, the cross-sectional standard deviation in the frequency of Regular Retail Price changes is 2.38%, whereas this statistic jumps to 16.80% when sales are included. Across categories, the cross-sectional standard deviation of the frequency of Regular Retail Price changes is 2.97%, whereas this statistic rises to 12.84% when sales are included.

7.2 The Hazard Function of Price Change

The hazard function plots the probability that a price spell will end after *x* weeks conditional on that price spell having survived at least *x-1* weeks. An upward sloping hazard function therefore indicates that prices are more and more likely to change the longer they have remained unchanged; whereas a downward sloping hazard indicates that prices are most likely to change again soon after they change.

We estimate the following model for the hazard function,

$$\lambda_i(t) = v_i \lambda_0(t), \qquad (3)$$

where v_i is a product specific random variable that reflects unobserved heterogeneity in the level of the hazard, and $\lambda_0(t)$ is a nonparametric hazard function common to all items.[22] We assume that $\lambda_0(t)$ is a step function with dummies for each quarter. We assume that v_i is distributed Gamma($1,\sigma_v^2$). We estimate the model using maximum likelihood and the price sequences from a common set of 27,788 SKUs and stores.[23]

[22] It is crucial to control for cross-sectional heterogeneity in estimating the hazard function, since hazard functions estimated using a panel of data from multiple heterogeneous products may yield a downward slope even if the true hazard is flat or upward sloping (see for example Kiefer, 1988). The presence of multiple price spells for each product in our dataset is important since it helps in identifying the extent of cross-sectional heterogeneity. See Lancaster (1979) for a discussion of this model.

[23] For each price spell, we truncate the duration if it exceeds 52 weeks and drop the first week to avoid capturing price changes that occurred in the middle of the previous week (see discussion in Section 2). In Figure 6 we use all of the 27,788 SKU x stores for which we have valid sequences of both Regular Retail Price changes and

Figure 6 illustrates the hazard function for Regular Retail Prices and Retail Prices (including sales), along with 95% confidence intervals. The Retail Price change hazard is downward sloping. In contrast, the hazard function for Regular Retail Price changes is distinctly upward sloping. This later result contrasts with the results in Nakamura and Steinsson (2008) and Klenow and Kryvtsov (2008). Those authors found flat or even slightly downward sloping hazard functions for regular retail prices. One possible explanation for the difference is that these other papers use a possibly imperfect sales flag to identify the Regular Retail Price, while we observe it directly in our dataset. We show below that the slope of the hazard function depends on the effectiveness of these filters.

7.3 Filtering Out Temporary Sales

A key advantage of our dataset is the accuracy of our measure of Regular Retail Prices. All other papers in the literature that we are aware of study temporary sales using imputed measures of regular prices based on a "sale filter" or noisy measures of whether a discount was in effect at a given point in time.[24] In contrast, our dataset contains the true Regular Retail Price for the retailer we study – this is the only measure of prices that senior managers at the retailer use to track markups.

Table 5 compares the Regular Retail price we observe in our data to "sale filter" based measures similar to those that have previously been used in the literature. We present results for three different sale filters. First, we construct a V-shaped filter used by Klenow and Kryvtsov (2008) and others. This filter eliminates dips in the price level relative to the prices 2 weeks in the past and 2 weeks in the future.[25] Second, we constructed a longer version of the V-shaped filter to eliminate dips relative to the prices 12 weeks before and 12 weeks after the price

---

Retail Price changes. Because we drop the week after the price change a valid sequence requires an observation in the second week after the initial price change.

[24] For example, in the BLS data, temporary sales are identified either using sale filters or using a "sale flag" that identifies whether there was a "sale sign" next to the product when the BLS price collector visited the store (e.g. Nakamura and Steinsson, 2008; Klenow and Kryvtsov, 2008). In the Dominick's data, temporary sales are identified primarily using sale filters, since the sale flag in that dataset is known to be quite noisy (e.g., Midrigan, 2011).

[25] Specifically, the V-shaped filter removes any price increase between period t-1 and t for which the Retail Price in week t is not higher than the Retail Price in week t-2, and removes any price decreases between period t-1 and t for which the Retail Price in week t-1 is not higher than the Retail Price in week t+1.

change.[26] Third, we used the sale filter proposed by Eichenbaum et al. (2011 henceforth EJR), which constructs a "reference price" as the modal price over a calendar quarter (13 weeks). The comparison with the V-shaped and 12-week filters is conducted at the weekly level and uses a common set of 67,974 weekly observations for which we have valid observations for both of these filters. The comparison with the EJR filter is (necessarily) conducted at the quarterly level, and uses a common set of 19,070 quarterly observations for which we have valid observations for the EJR filter. The observations are all weighted by total unit sales. As a baseline we also report results for both the (true) Regular Retail Price and unfiltered Retail Prices (including temporary sales).

Table 5 shows that while the V-shaped filter performs relatively poorly, both the 12-week sale filter and the EJR filter do a reasonable job of matching the frequency and size of price changes. The 12-week sale filter yields a weekly frequency of price change of 2.77%, compared to 2.86% for the true Regular Retail Price series. The EJR filter yields a quarterly frequency of price changes of 17.83%, compared to 14.77% for the true Regular Retail Price series. The correlation of binary indicator variables for price changes for the true Regular Retail Price versus the sales filters is 0.78 for the 12-week filter and 0.62 for the EJR filter.

A related statistic is the fraction of the time that the filters are able to identify the precise timing of a price change. The 12-week filter correctly identifies a true price change 80.70% of the time, and correctly identifies no price change 99.22% of the time. The EJR filter is slightly less successful in identifying the precise timing of price changes. Most of the discrepancies in the EJR filter arise for precisely the reasons that EJR anticipate in their discussion of the "reference price": (1) when there is a Regular Retail Price late in the prior quarter, so that the modal price in the prior quarter is different than in the current quarter, and (2) when there are frequent or multi-week discounts, so that the modal price is different from the Regular Retail Price. The differences between the sale-filter based measures and our true measure of Regular Prices underscores the value of having an accurate measure of Regular Retail Prices in our dataset.

---

[26] The 12-week filter removes any price increase between period t-1 and t for which the Retail Price in week t is not higher than max(Retail Price_{t-1}, … Retail Price_{t-n}), and removes any price decreases between period t-1 and t for which the Retail Price in week t-1 is not higher than max(Retail Price_{t}, … Retail Price_{t+n-1}), where n=12. We considered various different values of n, and n=12 provides the best fit to the true Regular Retail Price series.

Figure 7 investigates to what extent the filtered price series recovered from the Retail Prices using the three sale filters yield hazard functions that resemble those for the true Regular Retail Price. Panel A presents the results for the V-shaped filter, Panel B presents the results for the 12-week sale filter, while Panel C presents the results for the EJR filter. In each panel both curves are estimated using a common sample of SKUs x stores. In the case of the EJR filter, we estimate the hazard function on quarterly data, and rescale the hazard function to be comparable to the weekly hazard function for our 12-week sale filter. The figure reveals that both the 12-week and EJR filters are successful at recovering an upward-sloping hazard function. However, the more “conservative” V-shaped sale filter, which filters out fewer sale-related price changes, leads to a downward-sloping or flat hazard function. While previous studies have typically found flat or somewhat downward-sloping hazard functions of price changes (e.g., Nakamura and Steinsson, 2008; Klenow and Kryvtsov, 2008), our findings suggest that this may reflect reliance on inaccurate or incomplete sale filters.

8. Institutions of Manufacturer Trade Deals and Retailer Temporary Sales

To help interpret and explain why this retailer does not respond to demand or supply shocks by changing the depth or frequency of its discounts, we focus on the institutions of manufacturer trade deals and retailer temporary sales. The information in this section is based on interviews with both the firm that provided data for this study and a convenience sample of manufacturers and retailers. We should note that the details of these promotion funding mechanisms differ across manufacturers and retailers.[27] However, we know from surveys of manufacturers that a large fraction of these mechanisms share the key features that we emphasize (see, e.g., Acosta, 2012). We have organized our findings into six facts.

---

[27] Precisely documenting how the promotion funding mechanisms work for every manufacturer and every retailer is extremely difficult. For example, in 2002 two of us (Anderson and Simester) sent an MBA student to intern for 10 weeks at a retailer and document promotion funding. We learned that there was no uniform promotion funding practice among manufacturers selling to the retailer and the retail category managers could not easily document the flow of promotion funds. The most senior retail managers admitted that the promotion funding process had become extremely complicated and difficult to trace. At that time, determining the true marginal cost of a promoted item proved almost impossible.

A. Temporary Sales Follow "Sticky Plans"

From a logistical point of view, temporary sales are complicated events that require a substantial amount of planning and coordination between retailers and manufacturers. For example, when a promotion is run at a retail chain it may be accompanied by coupons, radio-television advertising, digital marketing, in-store displays, feature advertising, or product sampling. Retailers and manufacturers both understand that these demand generating activities are highly complementary with temporary sales and thus need to be coordinated carefully. In addition, there is often coordination with the retailer to ensure that sufficient inventory is available.[28] As a result, most retailers and manufacturers jointly set a schedule for temporary sales and associated promotional activity – a promotion calendar – through an annual planning process. In other words, temporary sales follow sticky plans that are generally updated at an annual frequency. A detailed example of a manufacturer's annual promotion plan is provided by Blattberg and Neslin (1990, p. 392). The plan provides details on the specific price promotions and the associated cost to the manufacturer.

B. Temporary Sales Are "Paid For" Out of Trade Deal Budgets

A central feature of promotion funding mechanisms for understanding the incentives faced by retailers is that manufacturer "funding" of promotions at a given retailer is determined by a trade deal budget. For example, suppose a manufacturer's product normally has a regular retail price of $2.49, but the manufacturer wants to encourage the retailer to lower the price to $1.99 for one week eight times during the year. To "fund" the $0.50 discount the retailer may be "paid" $0.35 per unit sold at $1.99 during the eight promotion weeks.[29] This amount will be deducted from a manufacturer trade deal budget that is specific to each retail account. In addition to "funding" temporary sales, the trade deal budget can be used to fund advertisements, in store displays and other demand generating activity associated with the temporary sale.

Importantly, the amount of funds in a trade deal budget limits the overall amount of support

[28] See for example Anderson and Simester (2001) and Anderson, Fitzsimmons and Simester (2006).

[29] In many cases, the amount that is paid out of the accrual account (e.g., $0.35 per unit) is designed to keep the retailer satisfied with the total dollar margin during the promoted weeks. So, if the retailer earns a 33% gross margin per unit at the regular price then the retailer may be happy to run a promotion that yields 25% gross margin but substantially greater unit volume.

that the manufacturer provides for temporary sales and associated promotional activity. Thus, if the retailer wants support for frequent or deep discounts early in a particular planning period, he must recognize that this will have the consequence that there will be less funds in the trade deal budget to support discounts later in the planning period. Reductions in the wholesale price associated with trade deals therefore may not reflect reductions in the retailer's marginal cost – since the retailer is spending down a finite resource (the trade deal budget). An implication of this is that the wholesale price variables in scanner price datasets used in the macroeconomics and industrial organization literatures must be interpreted with great caution, since these variables often include trade deal funding, and therefore cannot be viewed as measures of the retailer's marginal cost.

C. Manufacturers and Retailers Jointly Determine the Timing and Depth of Temporary Sales

The joint planning of the annual promotional calendar implies that manufacturers and retailers collaborate to determine the timing and depth of temporary sales. While the overall level of discount activity is constrained by the size of the manufacturer trade deal budget for each retailer, the retailer can influence the exact timing and depth of each discount. For many promotions, the manufacturer plan may allow for a "trade deal window" of several weeks where retailers can execute a promotion (see Blattberg and Neslin 1990, p. 319). This flexibility allows retailers to adjust their promotion plan to local market conditions. For example, if a competing retailer is expected to offer a deep discount on Coke then a retailer may decide to promote a different carbonated soft drink, such as Pepsi. In a subsequent week, the retailer may take advantage of the trade deal window to promote Coke. In contrast, if this is a week with high store traffic, such as 4th of July, the competing retailers may promote multiple soft drink brands (e.g., both Coke and Pepsi may be promoted that week). These facts contrast with stylized theoretical models that assume a manufacturer makes a take-it-or leave-it offer to all retailers in a market in a given week. Instead, it is important to recognize that the observed timing and depth of promotions is the result of a flexible, joint planning process.

D. Manufacturer Trade Deals are Contingent Contracts

An important challenge faced by manufacturers is how to induce retailers to promote their products effectively. In the late 1960s and early 1970s, manufacturers would simply offer retailers a temporary discount to the wholesale price. So, instead of a base wholesale price of $10 per case, a manufacturer may have temporarily lowered the price to $8 per case. The intent of the manufacturer was to induce the retailer to hold a temporary sale.

This funding strategy by manufacturers for sales was not incentive compatible. Anticipating the return of the higher normal base wholesale price, some retailers would “forward buy,” i.e., purchase large quantities of the product at the discounted price. They would also not reduce the retail price (i.e., not hold a temporary sale), presumably because they understood that a unit of inventory sold at a sale price during the week when the manufacturer prices was temporarily low was just as likely to have to be replenished in the following weeks at the normal base wholesale price as a unit sold a few days earlier or later. In addition, intermediaries would arbitrage geographic and temporal price differences; a new market emerged where retailers could purchase “diverted” goods from intermediaries. Finally, retailers would fail to execute in-store programs, such as in-store displays, that had been agreed upon in the promotion planning process.

In response to these incentive problems, manufacturers have begun requiring retailers to verify “performance” in order to receive trade deals funds. In the case of a temporary sale, the retailer must prove that it has indeed put the product in question on sale in order to release the funds from the trade deal budget. Several different mechanisms are used in the industry to verify performance. For example, the retailer might have to submit scanner price data showing that it lowered the price to $1.99 in the example above. Alternatively, the manufacturer’s sales-force may be responsible for verifying retailer compliance with a promotion.

The contingent nature of funding out of trade deal budgets provides an explanation for the otherwise puzzling feature of some scanner datasets that sharp drops in observed acquisition costs of inventory are associated with sharp contemporaneous drops in retail prices (Eichenbaum et al., 2011). As we discuss above, interpreting the drops in observed acquisition costs of inventory as straight discounts by manufacturers might lead one to conclude that retailers should instead respond to such discounts by stocking up on inventory for the product in question.

E. *Trade Deal Budgets Accumulate Funds Like Frequent Flyer Accounts*

When you fly on, say, United Airlines, you accrue “miles” in your frequent flyer account for every mile that you fly. Sometimes you accrue double miles or triple miles; sometimes you accrue miles for adopting a credit card or buying groceries. But, in the end you have an account of miles that you can redeem for free travel, upgrades and other offers provided by United Airlines. The accumulation of funds in trade deal budgets often works in a similar manner.

A manufacturer typically establishes an accrual rate, perhaps 6%, and then this determines how many dollars (points) accrue for the retailer. If a retailer buys $10,000 worth of product that accrues at 6% then the retailer has “earned” $600 in accruals. Accruals can also be based on a dollars-per-case metric, but this is less common. Just like frequent flyer programs, the accrual rates may vary over time. So, during some periods a manufacturer may offer 9% accruals rather than 6% accruals.

The goal of an accrual mechanism is to enhance the incentives of the retailer to support a manufacturer’s brand in both promoted and non-promoted periods. Volume purchased during the non-promoted periods accrues trade funds that can be used to support promotional activities such as price discounts, in-store displays and feature advertising. Incentives are aligned in the sense that more volume leads to greater trade funds that can be used to fund future promotions, which in turn will drive additional volume.

Not all trade promotion budgets are accrual accounts. In some cases, trade promotion budgets are simply fixed budgets—i.e., the manufacturer allocates a fixed dollar amount to a retail account each year. However, accrual accounts are very common for large retailers such as the one we study. Acosta (2012) reports that 83% of manufacturers with sales over $1 billion indicated that trade funding is based on accruals.

F. *Managers Pay Attention to Regular Prices*

A key concern that macroeconomists have had about analyzing data on regular prices is: How do we know that regular prices are relevant in determining the ultimate price faced by the consumer? What if actual prices are set entirely independently from regular prices—making regular prices a vacuous concept?

One way of investigating this question is to ask which price measures firms’ senior managers believe are most important. At the firm we study, Regular Retail Prices and Base Wholesale Prices are clearly viewed as the primary measures of the firm’s pricing policy and costs. These two metrics are summarized in monthly pricing reports that are shared among senior leaders in the company.[30] The monthly pricing report lists every change in the Base Wholesale Price and the Regular Retail Price (in the "main" pricing zone) that occurred in the calendar month. It then summarizes the impact on profit margins by category and at the aggregate firm level. In this report, the Base Wholesale Prices and the Regular Retail Prices are interpreted as the true variable cost of a unit and the true price of a unit. Notably there is no reference to temporary sale prices or the funding of temporary sales by manufacturers. In several years of conducting research with this firm we (Anderson and Simester) have never observed a regular management report describing temporary sale prices or the amount of manufacturer trade deal funding.

These findings are relevant to interview studies on pricing such as the seminal work by Blinder et al. (1998) and the many follow-up studies using similar methodologies. These studies interview managers about their pricing practices and yield a frequency of price change of close to one year. Bils and Klenow (2004) note that these studies yield frequencies of price change much lower than those observed in consumer price data including sales, and suggest that this difference may arise from a difference in industry composition. An alternative explanation, consistent with our institutional findings, is that when managers are asked about changes in the firm's "price" they interpret this question as referring to changes in the firm's regular price—which changes much less frequently than prices including sales.

9. Conclusion

Sale prices can result in extremely high frequency price changes in retail price data. A key question is whether these frequent price changes facilitate rapid responses to changing economic conditions, or whether they are merely part of a "sticky plan" that is determined substantially in advance and therefore not responsive to changing conditions. We use an exceptionally detailed dataset on retail and wholesale prices to investigate this question. We find that temporary sales are unresponsive to both demand shocks and supply shocks. Instead, the retail response is accurately measured through changes in the Regular Retail Price.

[30] Doubtless, the use of regular prices varies across firms. We believe, however, that the importance of the regular price is likely to hold true for other retailers of consumer packaged goods.

These empirical findings are consistent with the institutional features of price-setting, whereby temporary sales are typically "funded" via trade promotion budgets, and are orchestrated according to trade promotion calendars set substantially in advance. The institutions of trade deals imply that the wholesale price variables that appear in macroeconomic and industrial organization studies of price-setting must be interpreted with great caution as measures of the marginal cost. Furthermore, since trade promotion budgets are often "accrual accounts" that accumulate funds in proportion to sales, these institutions can help explain otherwise puzzling features of the cyclicality of temporary sales.

Our analysis suggests that Regular Retail Prices are sticky prices that change infrequently but are responsive to macroeconomic shocks, such as the rapid run-up and decline of oil prices. In contrast, temporary sales follow sticky plans. These plans include price discounts of varying depth and frequency across products. But, the plans themselves are relatively unresponsive in the near term to macroeconomic shocks. We believe that this characterization of regular and sale prices as sticky prices versus sticky plans substantially advances our understanding of how retailers adjust prices in response to macroeconomic shocks.

Appendix A: Direct Store Delivery

Direct Store Delivery (DSD) refers to the practice of manufacturer of bypassing the retailer's distribution system and delivering certain goods directly to individual retail stores. In the analysis we omit DSD categories, primarily alcohol, beverages, and dairy.[31] Collectively these categories contribute approximately 2% of the observations. There are important institutional differences in how pricing decisions are made in DSD categories, which imply that both the Wholesale Price and the Regular Retail Price measures – key features of our dataset – cannot be interpreted in the same way in these categories as in other categories.

Most important for our purposes, accrual accounts are not used to "fund" temporary discounts in DSD categories. In the case of alcohol, this is the result of legal restrictions. In the other categories, it may be because the incentive problems accrual accounts are designed to solve are not as severe in the case of DSD items since retailers hold no inventory apart from what is on the shelf at each time (so "forward buying" in response to discounts is not possible) and since the manufacturer can better monitor performance.

Manufacturers play a more direct role in setting retail prices in DSD categories.[32] Temporary sales on DSD items are often funded by temporary reductions in the Wholesale Price. Moreover, temporary price fluctuations are often coded as movements in the Regular Retail Price in the DSD category and discounts are more persistent. Approximately 20% of the Retail Price changes on DSD items arise because of long sales of 13 weeks or more, compared to just 1% to 2% for non-DSD items.

For non-DSD categories, we have argued that institutional features of how prices are set – arising because of incentive problems associated with the implementation of temporary sales – imply that Regular Prices respond to cost and demand shocks, while temporary sales do not and

---

[31] Firms cannot legally store alcohol unless they are a bonded wholesaler, which in practice requires that wholesalers deliver directly to stores. As a consequence, alcohol items are always DSD items.

[32] Perhaps for this reason, in DSD categories, both Wholesale and Retail prices vary much more across stores in response to regional competition. Regular Retail Prices and Retail Prices may vary at the region level in DSD categories, rather than at the "pricing zone" level for other products. In addition, whereas for almost all other products Wholesale Prices are constant across the national chain, Wholesale Prices for DSD items may also vary at the regional level.

are instead purely associated with intertemporal price discrimination. It may be that managers use a similar two-part approach to setting prices in DSD categories. Unfortunately, for the institutional reasons discussed above, this decomposition does not coincide with the Regular Retail Price vs. Retail Price distinction as it does in non-DSD categories, so our data are not able to speak to this issue.

References

Acosta (2012): “The Trend behind the Spend,” AMG Strategic Advisors, http://www.amgstrategicadvisors.com/thoughtLeadership.aspx?id=120, accessed on July 11, 2012.

Anderson, E., G. Fitzsimmons, and D. Simester (2006): “Measuring and Mitigating the Cost of Stockouts,” *Management Science*, 52(11), 1751-1763.

Anderson, E., N. Jaimovich and D. Simester (2012): “Price Stickiness: Empirical Evidence of the Menu Cost Channel,” Working Paper, Northwestern University.

Anderson, E. and D. Simester (2001): “Are Sale Signs Less Effective When More Products Have Them?” *Marketing Science*, 20(2), 121-142.

Bils, M. J. and P. Klenow (2004): “Some Evidence on the Importance of Sticky Prices,” *Journal of Political Economy*, 112(5), 947-985.

Blattberg, R.C. and S.A. Neslin (1990): *Sales Promotion: Concepts, Methods, and Strategies*, Prentice Hall, New York, NY.

Blinder, A. S., E. R. D. Canetti, D. E. Lebow, and J. B. Rudd (1998): *Asking about Prices*, Russell Sage Foundation, New York, NY.

Burstein, A.T. (2006): “Inflation and Output Dynamics with State-Dependent Pricing Decisions,” *Journal of Monetary Economics*, 53(7), 1235-1257.

Carvalho, C. (2006): “Heterogeneity in Price Stickiness and the Real Effects of Monetary Shocks,” *Frontiers of Macroeconomics*, 2.

Chahrour, R. (2011): “Sales and price spikes in retail scanner data,” *Economics Letters*, 110(2), 143-146.

Chevalier, J. A. and A. K. Kashyap (2011): “Best Prices,” NBER Working Paper No. 16680.

Chu, A. and A. Nevo (2013): “The Elasticity of Substitution Between Time and Market Goods: Evidence from the Great Recession,” Working Paper, Northwestern University.

Coibon, O., Y. Gorodnichenko and G.H. Hong (2013): “The Cyclicality of Sales, Regular, and Effective Prices: Business Cycle and Policy Implications,” Working Paper, University of California at Berkeley.

Eichenbaum, M, N. Jaimovich and S. Rebelo (2011): “Reference Prices and Nominal Rigidities,” *American Economic Review*, 101(1), 234-262.

Goldberg, P. and R. Hellerstein (2013): “A Structural Approach to Identifying the Sources of Local-Currency Price Stability,” *Review of Economic Studies*, 80(1), 175-210.

Guimaraes, B. and K. Sheedy (2011): “Sales and Monetary Policy,” *American Economic Review*, 101(2), 844-876.

Hendel, I. and A Nevo (2011): “Intertemporal Price Discrimination in Storable Goods Markets,” Working Paper, Northwestern University.

Hoch, S.J., X. Drèze, and M.E. Purk (1994): “EDLP, Hi-Lo, and Margin Arithmetic,” *Journal of Marketing*, 58, October, 16-27.

Kehoe, P., and V. Midrigan (2012): “Prices Are Sticky After All,” Working Paper, New York University.

Kiefer, N.M. (1988): “Economic Duration Data and Hazard Functions,” *Journal of Economic Literature*, 26, 646-679.

Klenow, P. J., and O. Kryvtsov (2008): “State-Dependent or Time-Dependent Pricing: Does It Matter for Recent U.S. Inflation,” *Quarterly Journal of Economics*, 123(3), 863-904.

Klenow, P.J. and J.L. Willis (2007): “Sticky Information and Sticky Prices,” *Journal of Monetary Economics*, 54, 79-99.

Lancaster, T. (1979): “Econometric Methods for the Duration of Unemployment,” *Econometrica*, 47, 939-956.

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

McShane, Blakely, Chaoqun Chen, Eric Anderson, Duncan Simester (2013), “Decision Stages and Asymmetries in Regular Price Pass-through”, Working Paper, Kellogg School of Management, Northwestern University.

Midrigan, V. (2011), “Menu Costs, Multi-Product Firms and Aggregate Fluctuations,” *Econometrica*, 79, 1139-1180.

Nakamura, E., and J. Steinsson (2008): “Five Facts about Prices: A Reevaluation of Menu Cost Models,” *Quarterly Journal of Economics*, 123(4), 1415-1464.

Nakamura, E., and J. Steinsson (2010): “Monetary Non-Neutrality in a Multi-Sector Menu Cost Model,” *Quarterly Journal of Economics*, 125(3), 961-1013.

Nakamura, E. and D. Zerom (2010): “Accounting for Incomplete Pass-Through,” *Review of Economic Studies*, 77(3), 1192-1230.

Peltzman, S. (2000): “Prices Rise Faster than They Fall,” *Journal of Political Economy*, 108(3), 466-502.

Table 1
Frequency and Size of Wholesale and Retail Price Changes

	Average Weekly Frequency	Absolute Size
Base Wholesale Price Changes		
Any Change	1.13%	7.58%
Increases	0.92%	7.28%
Decreases	0.21%	8.84%
Regular Retail Price Changes		
Any Change	1.22%	8.78%
Increases	1.11%	8.75%
Decreases	0.11%	9.10%
Retail Prices (including temporary sales)		
Any Change	26.35%	27.45%
Increases	13.41%	28.65%
Decreases	12.94%	26.21%

The table reports the average weekly frequency of price changes and the average absolute percentage size of the price changes. The unit of observation is a SKU at a store in a week and for the frequency measures the sample size is 1,255,832. Not all items have price changes in every week and so the sample sizes for the absolute size measures range from 1,428 (Regular Retail Price decreases) to 280,958 (Retail Price changes). The observations are weighted by total unit sales for the SKU in that Store (across all 195 weeks). We restrict attention to observations in sequences with at least 13 weeks of consecutive transactions, and exclude "Direct Store Delivery" categories.

Table 2
Price Response Before and After a Change in the Base Wholesale Price

	Base Wholesale Price	Regular Retail Price	Retail Price	Discount Frequency	Discount Depth	% Saved Through Discounts
Wholesale Price Increases						
Post Wholesale Price Change	4.78%**	5.03%**	4.81%**	1.02%	0.33%	0.14%
	(0.30%)	(1.03%)	(1.26%)	(2.41%)	(1.13%)	(0.42%)
Trend	0.06%	2.11%**	0.83%	0.74%	2.67%**	0.96%**
	(0.15%)	(0.38%)	(0.50%)	(1.60%)	(0.61%)	(0.35%)
Adjusted R^2	0.90	0.54	0.17	0.14	0.21	0.13
N	35,262	35,262	35,262	35,262	7,583	35,262
Wholesale Price Decreases						
Post Wholesale Price Change	-4.07%**	-0.44%	-0.50%	1.25%	0.55%	0.04%
	(1.29%)	(0.44%)	(1.08%)	(5.32%)	(1.23%)	(0.98%)
Trend	0.70%	0.43%*	0.68%	-3.45%	0.71%	-0.22%
	(0.54%)	(0.23%)	(0.49%)	(2.61%)	(1.35%)	(0.46%)
Adjusted R^2	0.87	0.34	0.24	0.14	0.51	0.18
N	10,989	10,989	10,989	10,989	2,013	10,989

The table presents coefficient from estimating Equation 1 on each dependent variable. All the effects are presented as a percent of the average Regular Retail Price in the 13 weeks prior to the Wholesale Price change.The unit of analysis is an item x week and we include the 13 weeks before and after a Wholesale Price change. The items are weighted by total unit sales and the standard errors are clustered by Item. Item fixed effects (and a constant) are included but omitted from the table. *Significantly different from zero, $p < 0.05$, ** significantly different from zero, $p < 0.01$.

Table 3
Frequency of Price Increases and Diesel Price Movements

	Base Wholesale Price Increase Frequency	**Regular Retail Price Increase Frequency**	**Discount Frequency**	**Discount Depth**	**% Saved Through Discounts**
Lagged Change in Diesel Price	0.844** (0.114)	0.695** (0.126)	-0.852 (0.558)	0.232 (0.202)	-0.167 (0.147)
Trend	0.008** (0.001)	0.005** (0.002)	0.049** (0.013)	-0.003 (0.009)	0.011* (0.005)
Quarter 2	-0.206 (0.135)	-0.207 (0.118)	2.141 (1.552)	1.493 (1.371)	0.853 (0.510)
Quarter 3	-0.662** (0.186)	-0.975** (0.183)	2.803 (2.139)	1.922 (0.994)	1.163* (0.526)
Quarter 4	-1.055** (0.191)	-0.744** (0.117)	1.081 (1.161)	1.713* (0.796)	0.662 (0.345)
R^2	0.01	0.01	0.16	0.27	0.16

The table reports coefficients from estimating Equation 2 on each dependent variable. The coefficients reflect the percentage point increase in the dependent variable. Standard errors are clustered by Item and reported in parentheses and the observations are weighted by total unit sales. Item fixed effects (and a constant) are included but omitted from the table. The unit of observation is a SKU x store x week and the sample size is 1,255,832, except for Model 4 where the sample size is 218,229 (not all SKUs have discounts each week). *Significantly different from zero, p < 0.05, ** significantly different from zero, p < 0.01.

Table 4
Regional Variation in Unemployment and the Frequency of Price Changes

	Regular Retail Price Increase Frequency	**Discount Frequency**	**Discount Depth**	**% Saved Through Discounts**
Lagged Change in Unemployment	0.048** (0.009)	-0.166 (0.112)	0.076 (0.044)	-0.005 (0.035)
R^2	0.58	0.73	0.82	0.79

The table reports coefficients from estimating Equation 3 on each dependent variable. The coefficients reflect the percentage point increase in the dependent variable. Standard errors are clustered by Item and reported in parentheses and the observations are weighted by total unit sales. Item x week fixed effects (and a constant) are included but omitted from the table. The unit of observation is an item x week and the sample sizes are 1,255,832, except for Model 4 where the sample size is 218,229 (not all SKUs have discounts each week). *Significantly different from zero, $p < 0.05$, ** significantly different from zero, $p < 0.01$.

Table 5
Comparing Regular Retail Price Changes with the Filtered Retail Price Change Series

	(True) Regular Retail Price	Unfiltered Retail Price (Including Sales)	V-Shaped (2-Week) Sale Filter	12-Week Sale Filter	EJR Sale Filter (Quarterly Mode)
Pair-wise Correlation with Regular Retail Price changes					
Weekly		0.22	0.30	0.78	
Quarterly		0.24			0.62
Quarters with Price Changes					
Weekly	2.86%	18.48%		2.77%	
Quarterly	14.77%	74.89%			17.83%
Mean Absolute Size of Price Changes					
Weekly	7.03%	26.34%	22.56%	16.82%	
Quarterly	6.49%	21.62%			18.81%
Probability Regular Retail Price Change is Correctly Identified					
Weekly		99.19%	84.10%	80.70%	
Quarterly		99.97%			74.87%
Probability No Change in Regular Retail Price is Correctly Identified					
Weekly		62.63%	83.18%	99.22%	
Quarterly		29.45%			92.05%

The unit of observation is either a week or quarter in a SKU x Store. The weekly analysis uses a common sample of 67,974 observations and the quarterly analysis uses a common sample of 19,070 observations. Not all SKUs have price changes in every week or quarter and so the sample sizes for the “average size of price changes” are smaller. The observations are all weighted by total unit sales for that SKU x store.

Figure 1
Price Series: Example

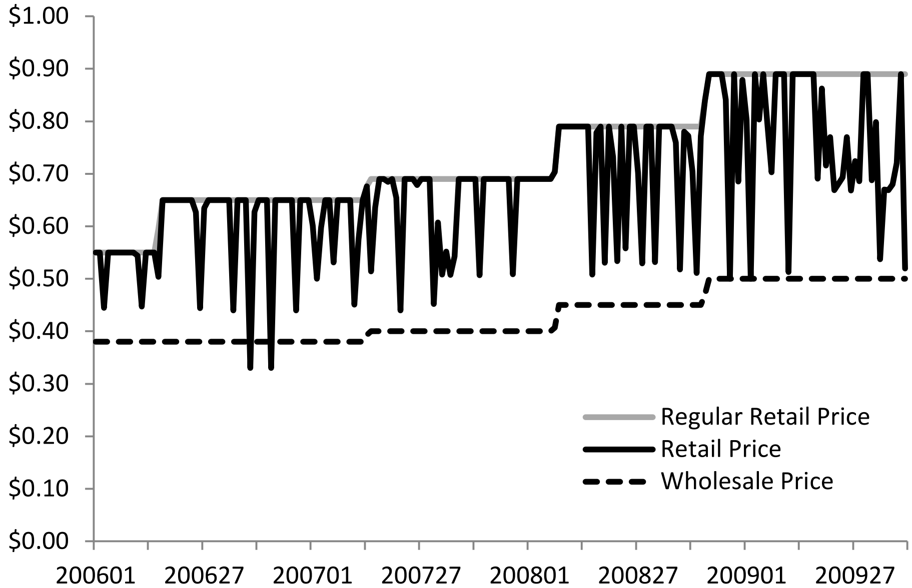

The figure reports the price trends for the three price variables for an arbitrarily chosen SKU at a single store that had sales of SKU in all 195 weeks of the data period.

Figure 2
Retail Price Response to Base Wholesale Price Change

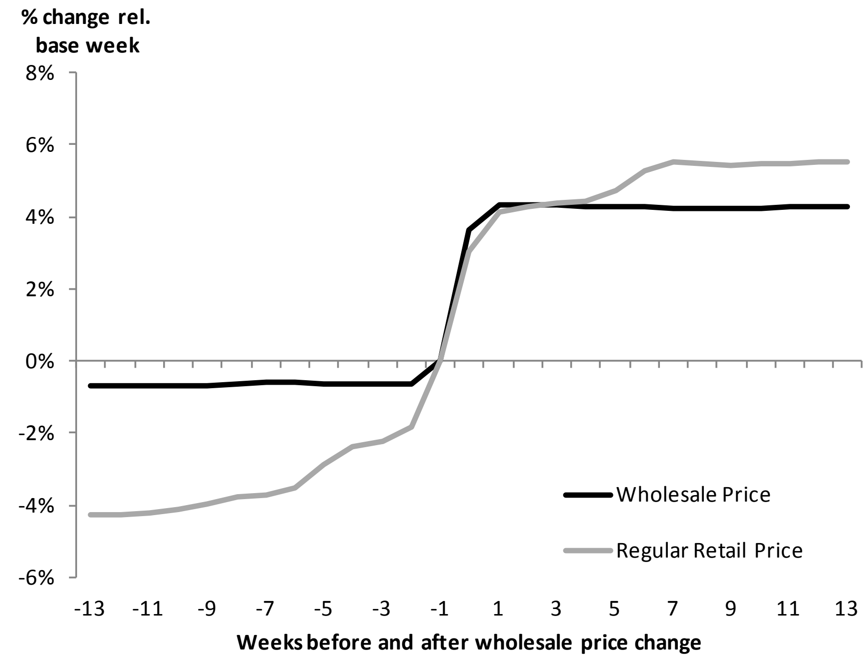

Panel A: Wholesale Price Increases

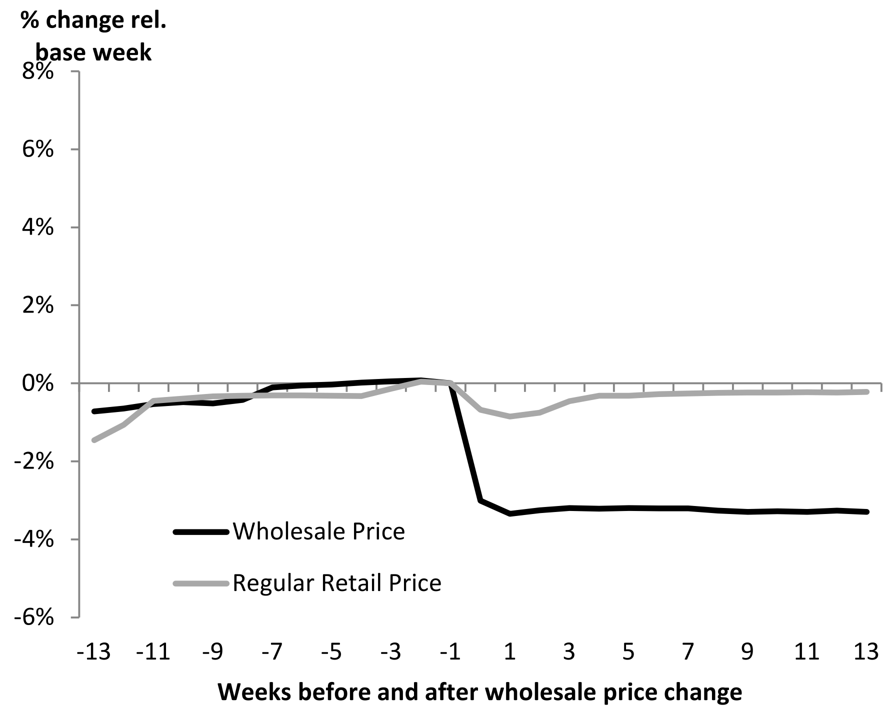

Panel B: Wholesale Price Decreases

The figure shows the average change in prices relative to the week before the Base Wholesale Price change. All series are presented as percentages of the average Regular Retail price in the 13 weeks prior to the cost change, and are weighted by total unit sales. For the Wholesale Price increases the sample includes 1,306 SKUs and for the Wholesale Price decreases the sample includes 407 SKUs.

Figure 3
Response of Temporary Sales to a Wholesale Price Change

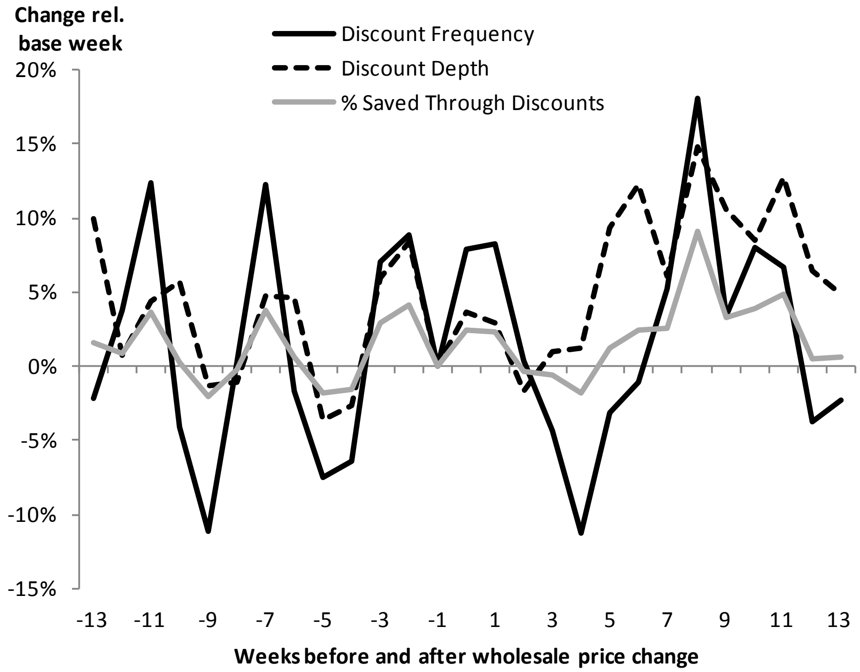

Panel A: Wholesale Price Increases

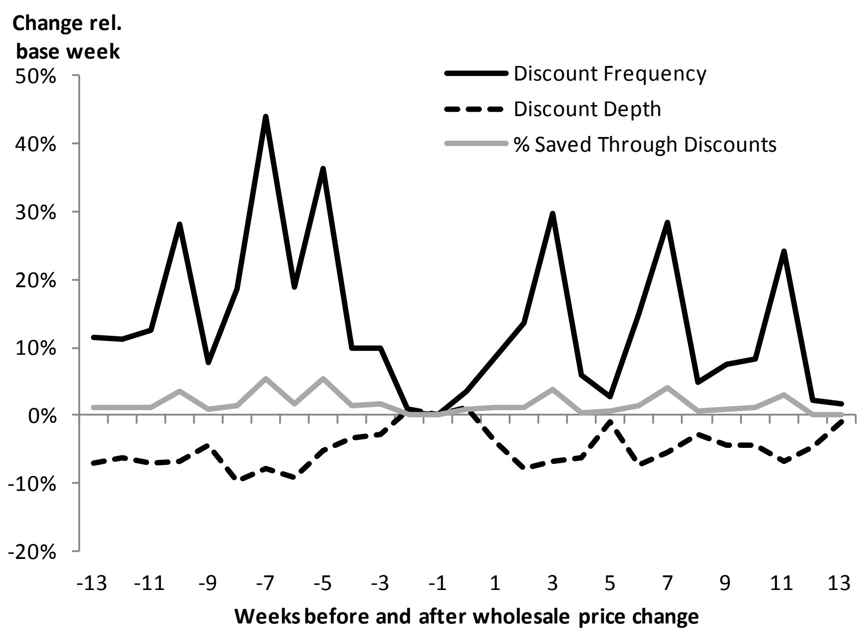

Panel B: Wholesale Price Decreases

Figure 4
Price Adjustment and Diesel Prices

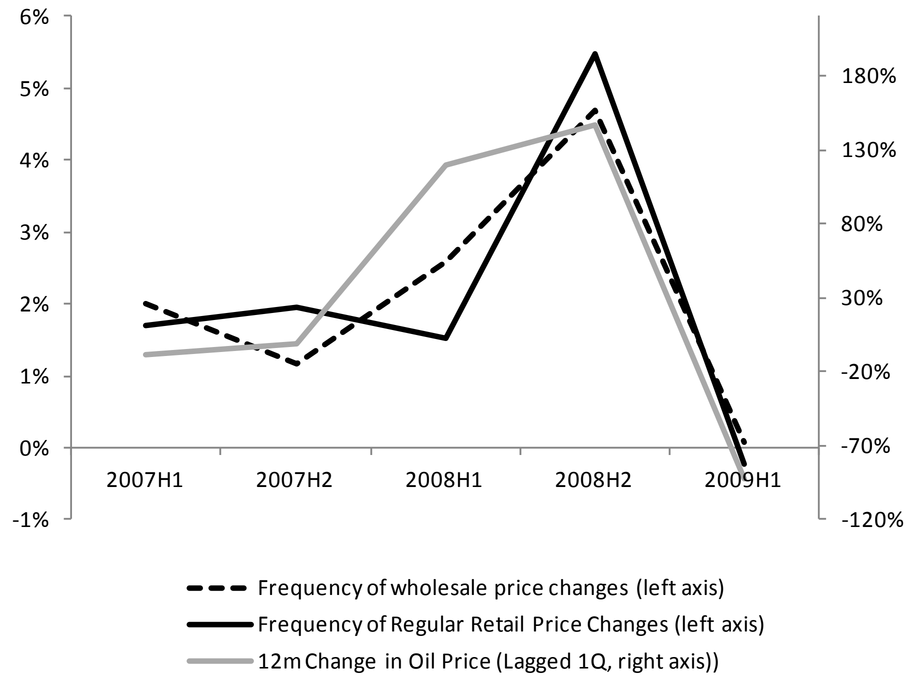

Panel A: Regular Prices and Base Wholesale Prices

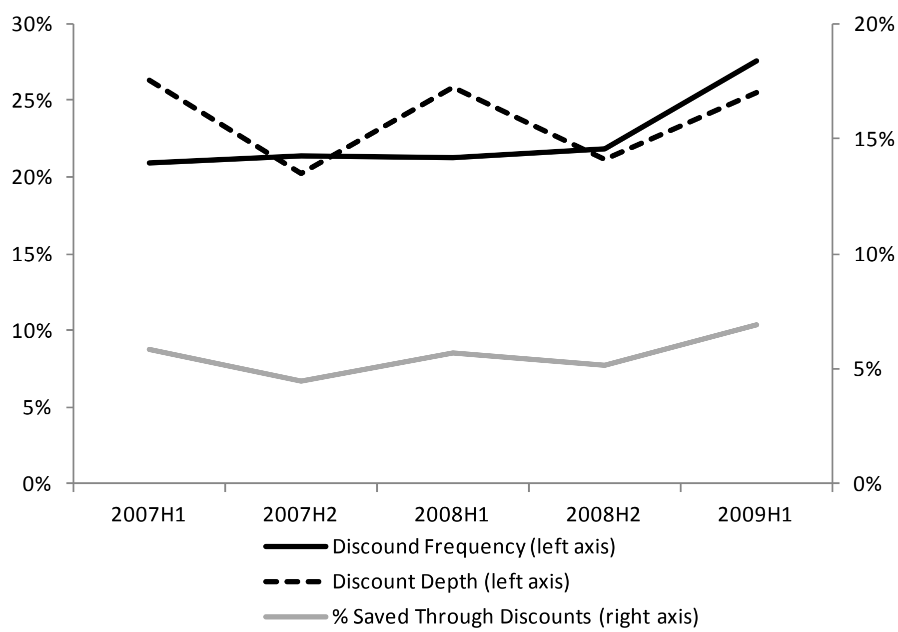

Panel B: Temporary Sales

Figure 5
Cross-Sectional Dispersion in Frequency of Price Change (in 2007)

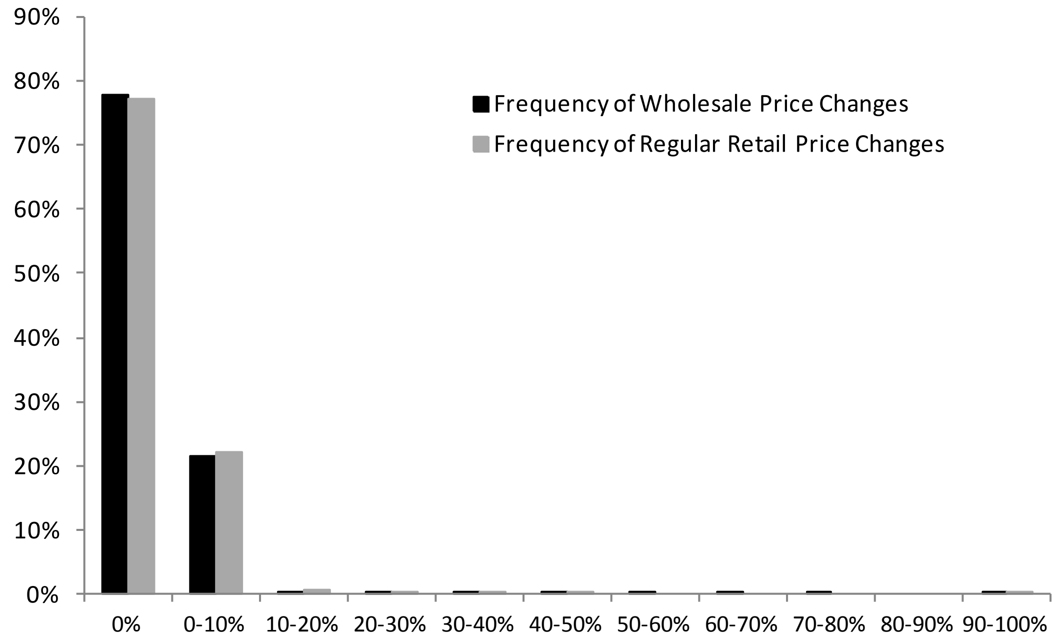

Panel A: Regular Retail Price and Base Wholesale Price

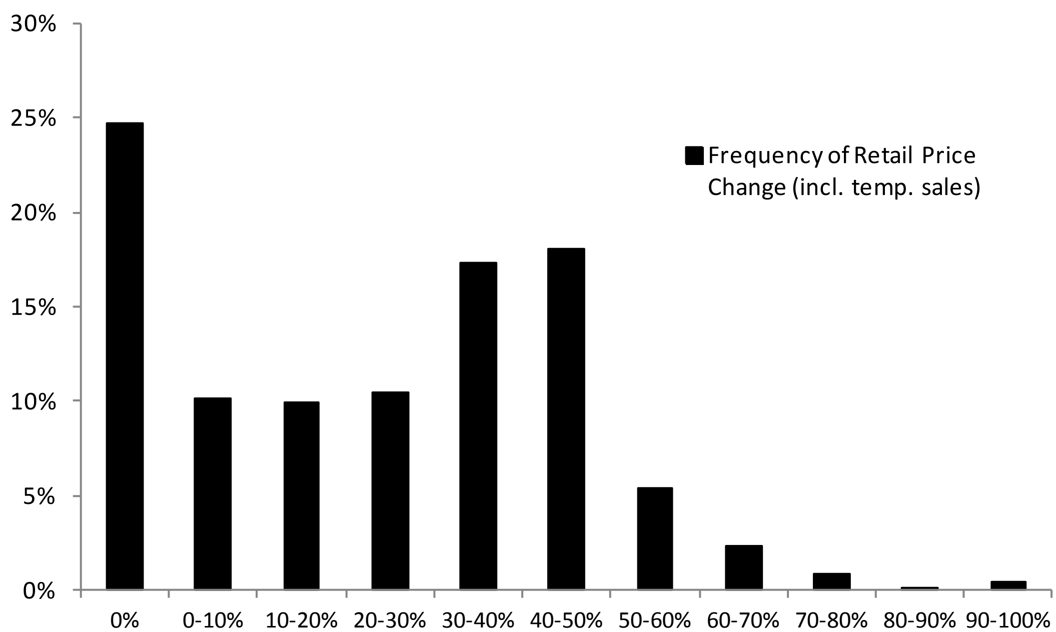

Panel B: Retail Price (Including Temporary Sales)

The figures report the distribution of the average weekly price change measures during the 2007 calendar year (similar findings are obtained when using data for different years). The unit of observation is a SKU x Store and the sample size for all three price series is 18,344. The observations are weighted by total unit sales.

Figure 6
Hazard Functions

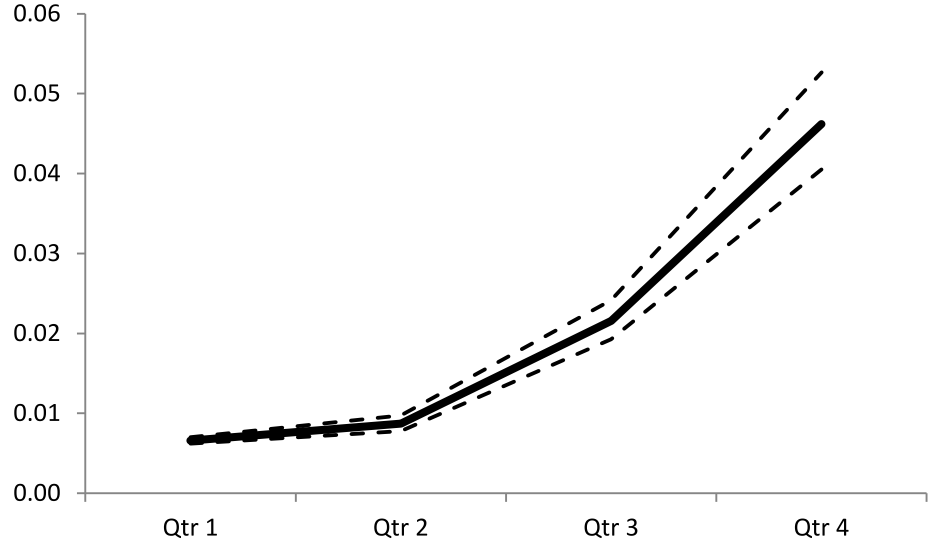

Panel A: Regular Retail Price Changes

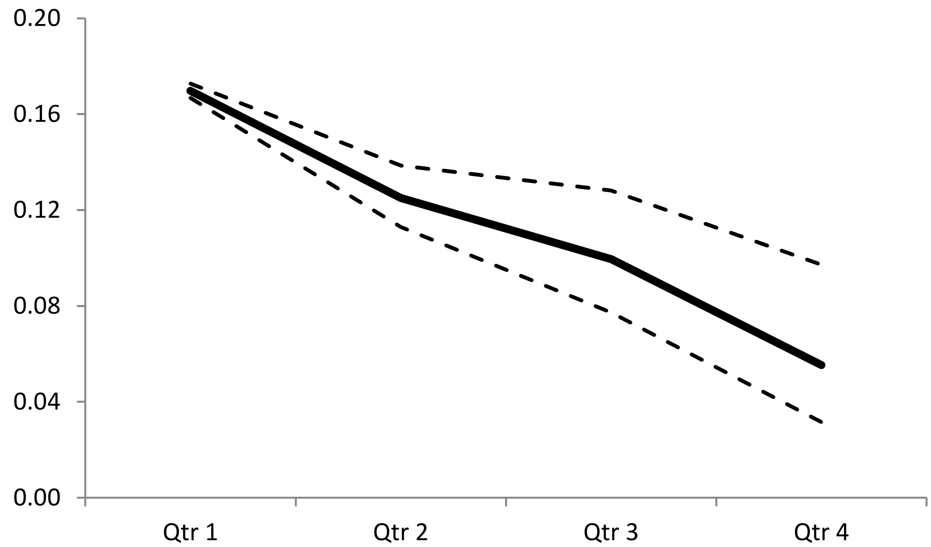

Panel B: Retail Price Changes

The figures report the weekly hazard of a price change according to the number of quarters since the last price change (estimated using Equation 2). The sequence starts in the second week after the initial price change. The sequence ends when there is a failure (a price change), a censored observation (a week without a transaction), or the sequence reaches 52 weeks after the initial price change. The sample for both curves includes price sequences from a common sample of 27,788 SKUs x stores.

Figure 7

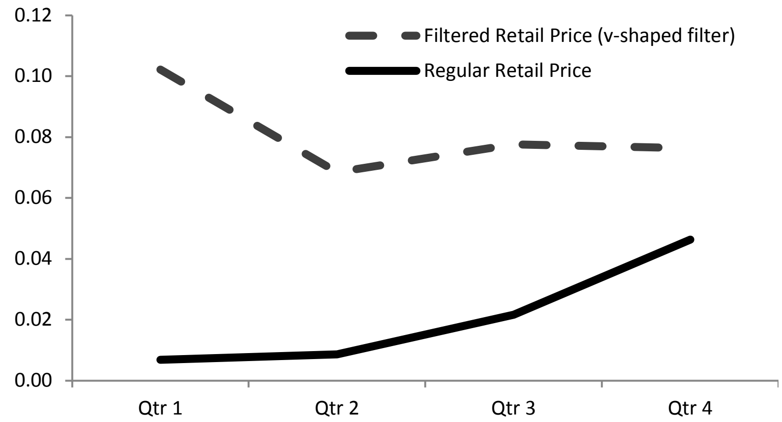

Panel A. Regular Retail Price and V-Shaped Filter

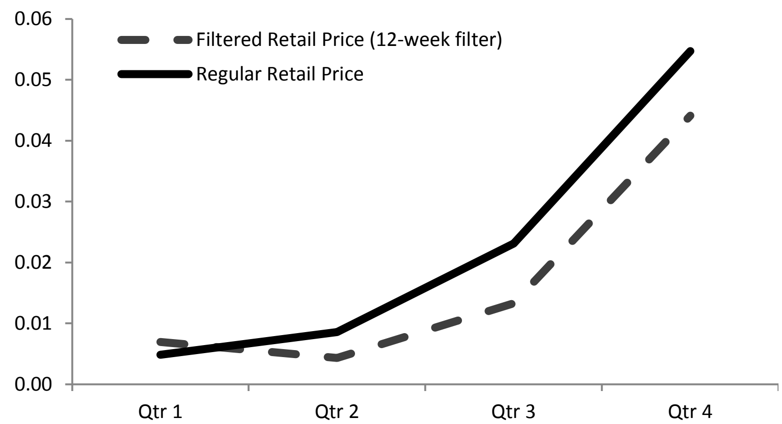

Panel B. Regular Retail Price and the 12-Week Filter

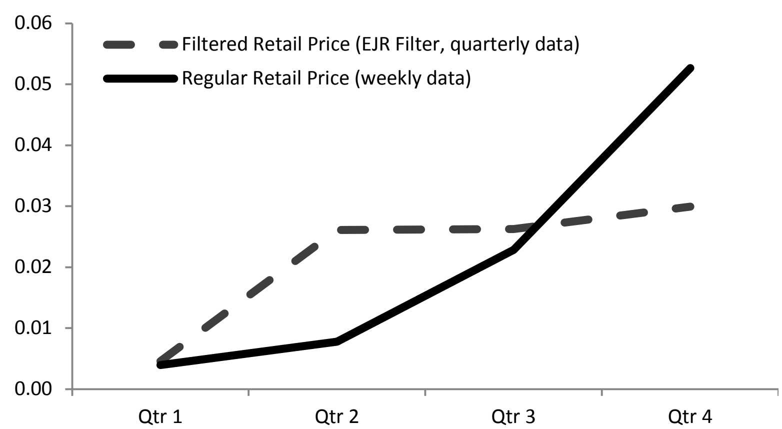

Panel C. Regular Retail Price and the EJR Filter

The figures report the weekly hazard of a price change according to the number of quarters since the last price change (estimated using Equation 2). The sequence starts in the second week after the initial price change. The sequence ends when there is a failure (a price change), a censored observation (a week without a transaction), or the sequence reaches 52 weeks after the initial price change. In each panel the sample for both curves includes all of the valid price sequences in a common sample of SKUs x stores. In Panel A the sample includes 22,604 SKUs x stores in Panel B the sample includes 1,805 SKUs x stores, in Panel C the sample includes 21,764 SKUs x stores. The findings are robust to using the same sample of 1,805 SKUs x stores in all three panels.

Figure A1
Frequency of Price Change and Length of Transaction Sequences

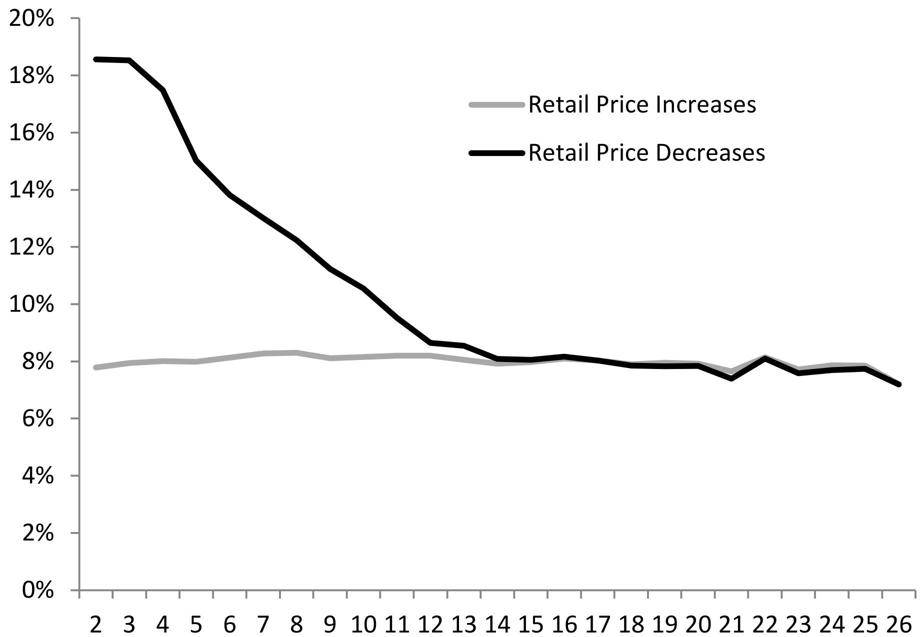

The figure reports the average weekly frequency of Retail Price increases and decreases using different samples of observations. The unit of observation is an item x week and the observations are grouped according to the number of consecutive weeks of transactions observed for the SKU at that store. The sample size is 196,498.