Vanderbilt University Department of Economics Working Paper Series
Nashville, TN 37235-1819

Vanderbilt University Department of Economics Working Papers 20-00001

TEMPORARY SALES IN RESPONSE TO AGGREGATE SHOCKS

Benjamin Eden
Vanderbilt University

Maya Eden
Brandeis University

Oscar Oflaherty
Vanderbilt University

Jonah Yuen
Vanderbilt University

Abstract

Using scanner data from supermarkets, we establish some stylized facts about temporary sales and argue that temporary sales play an important role in the reaction of prices to small demand shocks. We use a model in which temporary sales are reactions to aggregate shocks and the accumulation of unwanted inventories to account for our empirical findings.

The second author is from Brandeis University. All other authors are from Vanderbilt University. We would like to thank IRI for making the data available. All estimates and analysis in this paper based on data provided by IRI are by the authors and not by IRI.
Citation: Benjamin Eden and Maya Eden and Oscar Oflaherty and Jonah Yuen, (2020) "TEMPORARY SALES IN RESPONSE TO AGGREGATE SHOCKS", *Vanderbilt University Department of Economics Working Papers*, VUECON-20-00001.
Contact: Benjamin Eden - ben.eden@vanderbilt.edu, Maya Eden - meden@brandeis.edu, Oscar Oflaherty - oscar.oflaherty@vanderbilt.edu, Jonah Yuen - jonah.j.yuen@Vanderbilt.Edu.
Submitted: April 09, 2020. **Published:** April 15, 2020.
URL:http://www.accessecon.com/Pubs/VUECON/VUECON-20-00001.pdf

TEMPORARY SALES IN RESPONSE TO AGGREGATE SHOCKS*

Benjamin Eden Maya Eden Oscar O'flaherty Jonah Yuen

March, 2020

Abstract

Using scanner data from supermarkets, we establish some stylized facts about temporary sales and argue that temporary sales play an important role in the reaction of prices to small demand shocks. We use a model in which temporary sales are reactions to aggregate shocks and the accumulation of unwanted inventories to account for our empirical findings.

Key Words: Temporary Sales, Unwanted Inventories, Sequential Trade
JEL Code: D40

* The second author is from Brandeis University. All other authors are from Vanderbilt University. We would like to thank IRI for making the data available. All estimates and analysis in this paper based on data provided by IRI are by the authors and not by IRI.

1. INTRODUCTION

Temporary sales — defined as large price drops that quickly rebound — are a puzzling phenomenon from a macroeconomic perspective. It is unlikely that a 20% drop in prices that rebounds within a month is indicative of a sharp temporary contraction in demand, or of a sharp temporary increase in total factor productivity. In light of this, it has become standard practice in macroeconomics to focus on the behavior of "regular" prices, which are defined as prices that are not temporary sales (see Nakamura and Steinsson (2008), Eichenbaum, Jaimovich, and Rebelo (2011), and Kehoe and Midrigan (2015)). Sales are viewed as a tool for price discrimination that cancels out at the aggregate and contains no information on macroeconomic conditions (see Chevalier and Kashyap (2019), Guimaraes and Sheedy (2011), Salop and Stiglitz (1977), Shilony (1977) and Varian (1980)).

Here we argue that small demand shocks can lead to large temporary price reductions. We start by establishing some stylized facts emerging from a panel dataset of prices in supermarkets. Our main findings are as follows.

1. The fraction of weeks with no sales in any of the stores is much larger than the fraction predicted by the hypothesis that stores use a mixed strategy to choose temporary sales.
2. Goods with more fluctuations in regular prices have also more temporary sales.
3. Temporary sales contribute substantially to the weekly variation of the average cross-sectional price of the typical good.
4. Temporary sales do not contribute to the weekly variation of the average quantity sold.
5. Stores with relatively high average regular price tend to have more sales.
6. A temporary sale price may not necessarily be cheap relative to the price in other stores.

7. The upper part of the cross sectional price distribution appears to be more rigid than the lower part.
8. Sales play a critical role in the reaction of prices to a demand shock.

We propose a model in which temporary sales are reactions to unwanted inventories, which accrue in response to aggregate demand shocks, where an aggregate shock is good specific, but not store specific. The main insight of the model is that, if storage is associated with depreciation (as is the case with perishable goods that have expiration dates), then sharp, temporary reductions in prices can occur even in response to moderate shocks.

Our model is a flexible price version of Prescott (1975) hotels model: The Uncertain and Sequential Trade (UST) model in Eden (1990).[1] Most closely related is Bental and Eden (BE, 1993) that allows for storage and assumes exponential decay. In their model, there are demand and supply shocks, and the equilibrium price distribution depends on the current cost shock and the beginning of period level inventories. Inventories are accumulated when demand in the previous period was low. The accumulation of inventories leads to a reduction in prices (the entire price distribution shifts to the left) and as a result the quantity sold increases on average. Roughly speaking, the reduction in prices lasts until inventories are back to their "normal" level.

We adopt here the feature emphasized by Eden (2018), who assumes that units close to their expiration date are offered at a low price to minimize the probability that they will reach the expiration date before being sold. A store may therefore start at a relatively high "regular price" and then if it fails to make a sale switch to a low price until the level of inventories get back to "normal".

Our model is a flexible price model in which sales emerge in response to demand shocks. To gain some insight into the role of sales we consider the case in which temporary sales are not possible, say because of high menu type costs. We show that eliminating sales may affect average regular prices, average consumption and average

[1] For rigid price versions of the model see, Dana (1998, 1999) and Deneckere and Peck (2012).

production. It may also affect the variance of production over weeks. Thus in our model, temporary sales cannot be treated as "noise".

Looking at the data through the lenses of UST models is very different from looking at the data through the lenses of sticky price models. In some sticky price models, regular prices are costly to change but temporary change in prices cost much less. A shock to demand does not lead to an immediate change in regular prices and therefore if there is no cost for changing prices temporarily we may expect an increase in the frequency and size of sales until regular prices fully adjust. Coibion et.al. (2015) and Anderson et. al. (2017) find that temporary sale fail to react to changes in unemployment rate and conclude that temporary sales do not play an important role in the response to macroeconomic shocks.

In UST models, a shock to demand leads to the accumulation of unwanted inventories and to a reduction in prices. In most cases the shock and the amount of inventories accumulated are small and therefore the change in price lasts for a short time. The type of large aggregate shocks that produce changes in unemployment may lead to the accumulation of relatively large amount of inventories and to a relatively long spell of low prices that will be registered as changes in regular prices.

We find that goods with more variability (over weeks) in the average (across stores) regular price tend to have more sales. This is consistent with the view that goods with more demand uncertainty face more large shocks that lead to changes in regular prices and more small shocks that lead to temporary sales.

In the BE model, changes in wholesale prices produce an immediate change in retail prices as was found by Anderson et.al. (2017). Sales that occur in response to temporary changes in wholesale prices will occur in all stores and may explain why sometimes sale prices are not cheap relative to the price in other stores.

Aguirregabiria (1999) used a unique data set from a chain of supermarket stores in Spain and found a significant and robust effect of inventories at the beginning of the month on the current price. He also provides a description of the negotiation between the chain's headquarter and its suppliers. The toughest part of the negotiation with suppliers is about the number of weeks during the year that the brand will be under promotion, and about the percentage of the cost of sales promotions that will be paid by the wholesaler (e.g. cost of posters, mailing, price labels). A similar description is in Anderson et.al (2017) who present institutional evidence that sales (accompanied by advertising and other demand generating activities) are complex contingent contracts that are determined substantially in advance. There is also some flexibility. For many promotions, manufacturers allow for a "trade deal window" of several weeks where the seller can execute the promotion. These descriptions are consistent with the hypothesis that temporary sales are used to respond to high inventories. Sometimes the delivery schedule allows the firm to predict the level of inventories and as a result, temporary sales are set in advance. The flexibility in the timing of sales may reflect the need to respond to inventories that were accumulated as a result of demand shocks. We find large variations in the frequency of sale across stores that sell the same item. This suggests that the store manager has a say on the frequency of sales. We also find that sales play an important role in the price response to a demand shock and this suggests some flexibility in the timing of the sales.

We focus on the average posted price, rather than the average price paid by the consumer. This is different from Coibion et.al. (2015) and Glandon (2018). We argue that sale prices are used by the stores to react to negative demand shocks and are not merely a discrimination device. Our focus is thus on the behavior of the store rather than the welfare of the consumer.

Recently, Sheremirov (In press) studied the relationship between the cross sectional price dispersion and inflation and argued that sales should be included in aggregate models used by central banks for quantitative predictions. Here we focus on the behavior of the average price over weeks rather than on the cross sectional price dispersion.

Chevalier and Kashyap (2019) raise the issue of close substitutes. For example, should we look at Miller beer or at the category "beer"? If buyers do not care much about the kind of beer they buy (MIller, Miller light or Heineken) we should of course look at the category beer rather than at Miller beer. We think that our filter does not allow very close substitutes. We look at goods (UPCs) that were sold in all of the sample's week. Many UPCs do not survive this filter. Those who survive the filter are fairly popular items. These popular items are not likely to have close substitutes. We expect that a store manager will not care if a good with very close substitutes is on the shelf or not and therefore goods with very close substitutes will not be on the shelf in all of the sample weeks and will not survive our filter.

Will our empirical results survive aggregation of goods into categories and of time into monthly or quarterly frequencies? We think that the answer is in the negative. We are not likely to find that in about 40% of the weeks there are no sales in any of the stores if the definition of a good is wide enough. Similarly we are not likely to find that in about 40% of the quarters there are no sales in any of the stores. Our results may survive aggregation only if the shocks to demand are correlated across goods and time.

Are sales relevant for macro? If we define macro as the study of the effects of large correlated shocks then the answer is probably in the negative. We do not think that temporary sales are reactions to such shocks. Nevertheless the effect of large correlated shocks will depend on whether prices are rigid or flexible. The study of small shocks may be relevant for this question.

In the BE model the reaction to small and uncorrelated shocks is qualitatively the same as the reaction to large and correlated shocks. Both lead to the accumulation of unwanted inventories and to a decline in prices. Therefore it makes sense in the context of this model to argue that if prices react to small shocks they are also likely to react to large shocks.

From the point of view of UST models, there is no tension between flexible prices and money non-neutrality. See Eden (1994) and Bental and Eden (1996). These monetary models are more complicated than real UST models but the results are very similar: It does not matter much if the uncertainty is about the number of buyers that will arrive at the marketplace or the number of dollars that will arrive. We therefore think that a monetary shock that affect the demand for all goods will work in the same way as a good specific real demand shock.

2. DATA

We use a rich set of scanner data from Information Resources Inc. (IRI).[2] The complete data set covers 48 markets across the United States, where a market is sometimes a city (Chicago, Los Angeles, New York) and sometimes states (Mississippi). There are 31 diverse categories of products found in grocery and drug stores, such as carbonated beverages, paper towels, and hot dogs. We define goods by the Universal Product Code (UPC). The data provide information about the total number of units and total revenue for each UPC-store-week cell. We obtain the price for each cell by dividing revenue by the number of units sold. We use data from grocery stores in Chicago during the years 2004 and 2005. We use 3 samples. The 52 weeks in the year 2004, the 52 weeks in the year 2005 and the 104 weeks in the combined sample of 2004-2005.

We apply the following filtering (in a sequential manner):

(a) We drop all UPC-Store cells that do not have positive revenues in all of the sample's weeks.[3]

(b) We drop all UPCs that were sold by fewer than 11 stores.

(c) We drop all categories with less than 10 UPCs.

(d) We drop UPC-Week observations with no price dispersion.

The first exclusion is applied because we cannot distinguish between zero-revenue observations that occur when the item is not on the shelf and zero-revenue observations that occur when the item is on the shelf but was not sold. It is also required for identifying "temporary sale" prices. The second exclusion is aimed at reliable measures of the average cross-sectional price distribution. The third economizes on the number of category dummies. After applying (a)-(c) we obtain "semi balanced" samples in which the number of stores varies across UPCs but for each UPC the number of stores does not vary over weeks. [4]

---

[2] A complete description of the entire data set can be found in Bronnenberg, Bart J., Michael W. Kruger, Carl F. Mela. 2008. Database paper: The IRI marketing data set. Marketing Science, 27(4) 745-748.

[3] We also dropped observations in which the quantity sold was zero but revenues were positive.

[4] To get a sense of the effect of the sample exclusion on the result, Eden (2013) studies one week in detail. See the working paper version of Eden (2018): Vanderbilt University Department of Economics

The requirement that the product be sold continuously by more than 11 stores leads to a sample of fairly popular brands and as explained in the introduction the goods that survive these requirements are not likely to have very close substitutes. [5]

Summary statistics are in Table 1. The rows are the number of UPCs and the number of observations for individual categories. In the 2004 sample there were 32 UPCs in the beer category. The number of observations (UPC-Week cells) is (32)(52)=1664. In 2005 there were 56 UPCs in the beer category. The number of observations is not equal to (56)(52) because in 3 cells there was no price dispersion. The total number of observations for each sample is in the bottom of the Table. The combined 04-05 sample has fewer UPCs because criterion (a) in our filtering procedure is harder to satisfy when there are 104 weeks. As a result, the combined sample includes relatively more popular brands.

Working Papers 13-00015. Indeed, there is a difference between the sample of 8602 UPCs that were sold by more than 1 store during that week and the sample of 4537 UPCs that were sold by more than 10 stores. Relative to the larger sample, price dispersion in the smaller sample is lower. The highest price dispersion was found in an item that was sold by 2 stores and for this item the ratio of the highest to lowest price was 15.

[5] This is not unique to this paper. Sorenson (2000) has collected data on 152 top selling drugs. Lach (2002) excluded products that were sold by a small number of stores. Kaplan and Menzio (2015) exclude UPCs with less than 25 reported transactions during a quarter in a given market.

Table 1*: Summary statistics

	2004		2005		2004-2005	
Category	# UPCs	Obs.	# UPCs	Obs.	#UPCs	Obs.
Beer	32	1,664	56	2,909	20	2,080
Carbonated Beverages	86	4,472	144	7,471	58	6,032
Cold Cereal	93	4,836	133	6,900	53	5,512
Facial Tissue	12	624	18	893	-	-
Frozen Dinner Entrees	36	1,871	75	3,765	-	-
Frozen Pizza	25	1,300	53	2,744	12	1,248
Hot Dogs	14	728	21	1,091	-	-
Margarine & Butter	25	1,300	40	2,060	18	1,872
Mayonnaise	17	884	19	988	-	-
Milk	32	1,664	64	3,294	23	2,392
Mustard & Ketchup	14	728	21	1,092	-	-
Paper Towels	-	-	19	901	-	-
Peanut Butter	18	936	24	1,245	11	1,144
Salty Snacks	94	4,887	120	6,226	42	4,368
Soup	49	2,548	74	3,826	22	2,288
Spaghetti Sauce	13	676	32	1,660	-	-
Toilet Tissue	13	676	19	958	-	-
Yogurt	92	4,783	152	7,870	65	6,760
Totals	665	34,580	1084	56,368	324	33,696

* An observation is a UPC - Week cell. The first column is the category name. The two columns that follow are about the 2004 sample. The first is the number of UPCs in each category and the second is the number of UPC-Weeks in that category. The next two columns are for the 2005 sample and the last two columns are for the combined 2004-05 sample. Totals are in the last row.

Temporary Sales.

We assume that a temporary sale occurs when a price drop of at least 10% is followed by a price equal to or above the pre-sale price within four weeks. In the data there is an indication of whether the price is regular or sale. The frequency of sale according to the IRI definition is 35%; that is almost twice the frequency here. It may include smaller price reduction and may require promotion activities rather than just a price change. Our definition is similar to Coibion et. al. (2015) who use the IRI data set.[6]

Table 2 provides summary statistics about sales. The first row after the sample name, repeats the number of UPCs in each sample. The second is the average number of

[6] It may be the case that the wider definition of sales adopted by the IRI reflect the need to satisfy the requirement imposed by the supplier as described in Aguirregabiria (1999) and Anderson et.al. (2017). For example, if there are no unwanted inventories and the manager wants to have a sale just to satisfy the supplier's requirement he may cut the price by only 5% and declare a sale. We are interested in temporary sales that are made in response to unwanted inventories and we therefore adopted the definition based on the behavior of prices.

stores per UPC-Week cell with the minimum and the maximum number of stores in parentheses. In the 2005 sample, there are 21 stores on average. The minimum number of stores is 11 and the maximum number of stores is 35. The third row is the average number of stores after eliminating sales observations. The average number of stores for the 2005 sample is now 17. The minimum number of stores is 1 and the maximum is 35. The fourth row (Freq Sales) is the frequency of sales calculated as the fraction of UPC-Week-Store cells with a sale price (in parentheses are the frequency of sale for the UPC with the lowest frequency and the UPC with the highest frequency). For the 2005 sample the average frequency of sale is 0.2, the minimum is zero (there are UPCs with no sale prices) and the maximum is 0.45. The fifth row (No Sales) is the fraction of weeks in which there is no sale in any store, averaged over UPCs. For the 2005 sample, the average fraction of weeks with no sale is 0.45. The sixth (Sales in all stores) is the fraction of weeks in which there are sales in all stores (averaged over UPCs). For the 2005 the average is 0.5 percent.

The seventh and the eighth rows reports standard deviations over weeks. We first calculate the average log price for each UPC-week cell. We then calculate the standard deviation of these averages for each UPC across weeks. The seventh row reports the average (over UPCs) standard deviation when using the entire sample. This is 9% in the 2005 sample. The eighth reports the standard deviation when using the sample of regular prices which we obtain after removing sale observations. This is 7.5% in the 2005 sample. The ninth row reports the ratio of the standard deviation of the average price to the standard deviation of the average regular price. This is 9/7.5=1.2 in the 2005 sample. Thus sales increase the standard deviation by 20%. The tenth row is the contribution of sales to the standard deviation relative to the average frequency of sale. For the 2005 sample this is 0.2/0.2 = 1. Thus, on average sales in 1% of the weeks (or 5.2 weeks) increases the standard deviation by 1%. The following row reports the correlation between the frequency of sales and the standard deviation of the average log regular price. This correlation is 0.65 in the 2005 sample. Thus UPCs with more fluctuations in regular prices tend to have more sales. The last three rows report the standard deviations of the quantity sold over week and the ratio of the standard deviation when using all observations to the standard deviation when using only observations with regular prices.

Somewhat surprisingly, the standard deviation of quantity does not change much when we remove sale observations.

Table 2*: Statistics about sales

Sample	2005	2004	2004-2005
# UPC	1084	665	324
# stores, All	20.98 (min=11,max=35)	15.42 (11,21)	14.56 (11,19)
# stores, Regular	16.95 (min=1,max=35)	12.50 (1, 21)	11.51 (1, 19)
Freq Sales (Av. Freq)	0.2 (0,0.45)	0.19 (0, 0.47)	0.21 (0, 0.41)
No Sales	0.45	0.42	0.39
Sales in all stores	0.0051	0.002	0.002
SD Prices All	0.0908	0.0795	0.0837
SD Prices Regular	0.0755	0.0623	0.0708
Ratio	1.20	1.28	1.18
Relative	1	1.47	0.86
Corr	0.65	0.69	0.58
SD Quantity All	0.304	0.313	0.305
SD Quantity Regular	0.308	0.3	0.304
Ratio	0.99	1.04	1.00

* The first row after the sample name, repeats the number of UPCs in each sample. The second is the average number of stores (per UPC-Week cell – the minimum and the maximum number of stores are in parentheses). The third is the average number of stores after eliminating sales observations. The fourth (Freq Sales) is the frequency of sales calculated as the fraction of UPC-Week-Store cells that their price is a sale price (in parentheses are the frequency of sale for the UPC with the lowest frequency and the UPC with the highest frequency). The fifth (No Sales) is the fraction of weeks in which there is no sale in any store, averaged over UPCs. The sixth (Sales in all stores) is the fraction of weeks in which there are sales in all stores (averaged over UPCs). The seventh and eighth rows reports standard deviations over weeks. We first calculate the average log price for each UPC-week cell. We then calculate the standard deviation of these averages for each UPC across weeks. The seventh row reports the average (over UPCs) standard deviation when using the entire sample. The eighth reports the standard deviation when using the sample of regular prices which we get after removing sale observations. The ninth row reports the ratio of the standard deviation of all prices to the standard deviation of regular prices. The tenth row is the contribution of sales to the standard deviation relative to the average frequency of sale. For the 2005 sample this is 0.2/0.2 = 1. The following row reports the correlation between the frequency of sales and the standard deviation of the average regular price. The last three rows report the standard deviations of the quantity sold over week and the ratio of the standard deviation when using all observations to the standard deviation when using only observations with regular prices.

Table 2 suggests three main observations. Temporary sales contribute to the variation of the average price over weeks, UPCs with more fluctuations in the average regular price tends to have more sales, and the fraction of weeks in which there is no sale in any of the stores is high.

The fact that temporary sales increase the standard deviation of the average price (by about 20%) suggests that temporary sale prices play an important role in price flexibility and is consistent with the view that sales reduce the average price in some weeks more than in other weeks.

The correlation between the frequency of sale and the standard deviation of the average regular price says that UPCs with more fluctuations in regular prices also have more sales. This suggests that some underlying feature of the product determines both regular price variability and sale frequency. For example, it is possible that UPCs with more demand uncertainty have both relatively more small shocks and relatively more large shocks. The store use temporary sales to react to small shocks that lead to small amount of inventory accumulation and regular price changes to react to large shocks. Therefore UPCs with more sales tend to have larger variation in regular price.

Table 2 suggests that the fraction of weeks in which there was no sale in any of the stores is high. If all stores use a mix strategy to choose sales and if they all use the same probability that is equal to the observed frequency of sale, the probability of no sale in any of the store is: (1-0.2)^21=0.009 or 0.9 percent for the 2005 sample. This is very different from the observed frequency of 45%. Heterogeneity may lead to this result. It is possible that the phenomenon of no sale in any of the stores occurs primarily in UPCs with low sale probabilities. To examine this possibility, we divide the UPCs into five bins according to the frequency of sale.

This is done in Table 3. The first row defines the bin. The first bin contains all UPCs with frequency of sales between zero and 10 percent. The second bin contains all UPCs with frequency of sales between 11 and 20 percent and so on. The second row is the fraction of UPCs that are in the bin. Twenty-six percent are in the third bin which contains all UPCs with frequency of sale between 21 and 30 percent. The average frequency of sale is in the third row. For the third bin the average frequency of sale is 25%. The average number of stores is in the fourth row. It is close to 22 in the third bin. The fifth row is the fraction of weeks in which there were no sales in any of the stores. This fraction is 0.29 for the third bin. The sixth row is the probability of no sale in any of the stores calculated under the assumption that stores follow a mixed strategy and choose sales with probability equal to the average frequency in the bin. For example, in the 21-30 bin the average frequency is 0.25 and the average number of stores is close to 22. The probability of no sale in any of the 22 stores is 0.75^22 which is 0.2 percent. The average fraction of weeks in which there was no sales in any of the stores is 29 percent and is about 15 times higher than the predicted probability.

The last rows of Table 3 are about the standard deviation of the cross-sectional average price over weeks. The differences between the bins are large. For example, the average standard deviation of the average regular price for UPCs that are in the highest frequency of sale bin is 0.124, which is more than three times the average standard deviation for UPCs that are in the lowest frequency of sale bin. The ratio reported in the last row suggests that the contribution of sales to the standard deviation does not increase with the frequency of sale. For the highest frequency bin the standard deviation of the average price is 7% higher than the standard deviation of the average regular price (0.133/0.124=1.07). For the lowest frequency bin it is 17% higher. In the lowest frequency bin the average frequency of sale is 4%. Per percentage point the contribution of sales to the standard deviation is 17.2/4 = 4.3%. As can be seen from the last row (labelled "relative") the contribution per percentage point declines with the average frequency. It thus seems that the contribution of a sale to price flexibility is relatively large for UPCs with less sales.

Table 3*: By frequency of sale bins

Freq of sale	0-10	11-20	21-30	31-40	41-50	All bins
Frac UPCs	0.27	0.25	0.26	0.20	0.01	1
Av. Freq	0.04	0.15	0.25	0.34	0.42	0.2
# of stores	20.25	19.44	21.82	23.17	14.71	21
No Sales	0.79	0.46	0.29	0.21	0.19	0.45
Prob.	0.47	0.04	0.00	0.00	0.00	0.01
SD All	0.046	0.093	0.113	0.116	0.133	0.091
SD Regular	0.040	0.070	0.093	0.105	0.124	0.076
ratio	1.172	1.321	1.221	1.110	1.069	1.203
relative	4.300	2.140	0.884	0.324	0.164	1.015

* This Table uses the 2005 sample. UPCs are divided into 5 bins. The first bin contains all UPCs with frequency of sales between zero and 10 percent. The second bin contains all UPCs with frequency of sales between 11 and 20 percent and so on. The first row is the fraction of UPCs that are in the bin. The second row is the average frequency of sale for the bin. The third is the average number of stores in the bin. The fourth is the fraction of weeks in which there were no sales in any of the stores. The fifth is the probability of no sale in any of the stores calculated under the assumption that stores follow a mixed strategy and choose sales with probability equal to the average frequency. The last four rows are about the standard deviation of the average price over weeks. The row before the last is the ratio of the standard deviation of the average price to the standard deviation of the average regular price. The last row is the per percentage point contribution to the standard deviation. For example, the average frequency of sale in the first bin is 4%. Sales increase the standard deviation by 17.2%. Per percentage point the contribution is 17.2/4 = 4.3%.

Variation across stores

The frequency of temporary sales varies over UPCs and over stores. We now turn to the variations over stores. In Figure 1 we used the 2005 samples with 1084 UPCs. For each UPC we calculated the range of the frequency of sale across stores. The UPCs are ordered by the average frequency of sale (average across store-week cells) so that the first UPC in the graph has the lowest frequency of sale (zero) and the last UPC has the highest frequency of sale (44%). The average range after excluding UPCs with no sale is 0.27. This is huge. For example, a UPC with the average range may have some stores with a frequency of sale of 10% and some stores with a frequency of sale of 37%. There is also large variation across UPCs. A small fraction (3.7%) of the UPCs (like milk, some beer and mayo) have no sales. There are UPCs with frequency of sales that are over 40% (frozen dinners and frozen pizza). The correlation between the mean and the range is 0.6. Thus UPCs with higher frequency of sales tend to have more variations across stores.

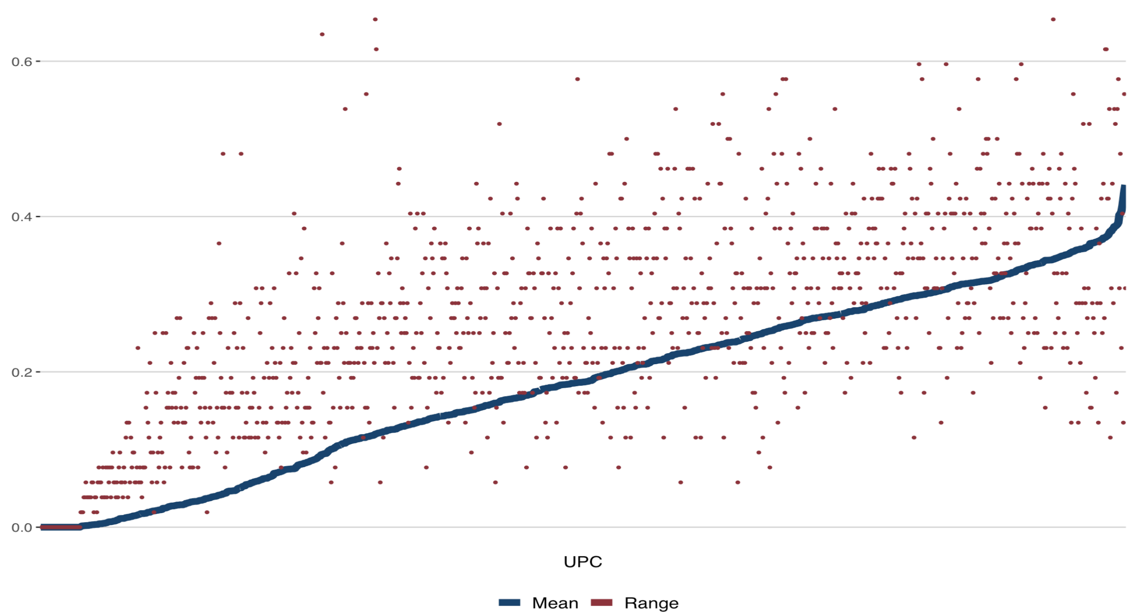

Figure 1: The mean and the range of the frequency of sale

Table 3a aggregate UPCs into bins. The first row is the range of the frequency of sale across the stores in the bin. The average (across UPCs) range for UPCs that have an average (across stores) frequency of sale between 31% and 40% is 0.356. This says that on average the frequency of sale's range across stores that sell a UPC in that bin is 36%. We may have for example a store that have a sale for a UPC in that bin in 10% of the weeks and a store that have a sale for the same UPC in 46% of the weeks. As was said above, this is a large variation over stores. The second row in the table is the standard deviation of the frequency of sale across stores that are in the bin. The average standard deviation for UPCs that are in the 31-40 bin is 10%. The last row is the correlation between the average regular price and the frequency of sale. Stores that sell a UPC in the 30-40 bin will have a correlation of 0.5 between the average regular price and the frequency of sale. This suggests that stores with a higher average regular price have more sales.

Table 3a*: Variations across stores by frequency of sale bins

Freq of sale	0-10	11-20	21-30	31-40	41-50	All bins
Range	0.131	0.266	0.322	0.356	0.346	0.262
SD	0.041	0.083	0.095	0.100	0.096	0.079
Correl	0.136	0.346	0.403	0.500	0.723	0.344

* The range is the maximum frequency of sale minus the minimum across stores that are in the bin. SD is the standard deviation of the frequency of sale across stores that are in the bin. The last row is the correlation between the average regular price and the frequency of sale across stores that are in the bin.

The large range and standard deviation measures suggest that the store manager has a significant role in choosing the sale frequency. The correlation between the average regular price and the frequency of sale suggests that managers who choose high regular prices tend to choose more sales. The correlation is higher for UPCs with high sale frequencies.

3. THE BEHAVIOR OF THE CROSS SECTIONAL PRICE DISTRIBUTION

We now turn to describe the behavior of the cross sectional price distribution over weeks. To allow for the description of the behavior of the cross sectional distribution of regular prices we use smaller samples in which there are at least 11 stores that post the regular price. That is, UPCs were dropped if after removing temporary sales observations there were fewer than 11 stores in any week. This is an additional filter that drastically reduced the number of UPCs in the sample. In the 2005 sample it reduced the number of UPCs from 1084 to 215. It also reduced the frequency of sales in that sample from 0.2 to 0.05.

Table 4 provides some statistics about the smaller samples and should be compared with Table 2. The standard deviation of the average price in the smaller samples is much lower than the standard deviation in the larger samples. In the smaller 2005 sample (with 215 UPCs) the standard deviation of all prices is 0.043. In the larger sample of 2005 the standard deviation of all prices is 0.09, which is more than double that magnitude. The standard deviation of the average regular price in the smaller sample of 2005 is 0.039. In the larger sample it is 0.076 that is almost double that magnitude. Since the smaller sample has a much lower frequency of sales, this is consistent with the positive correlation between the standard deviation of the average regular price and the frequency of sales. The contribution of sales to the standard deviation is large relative to the frequency of sale. In the smaller 2005 sample sales increase the standard deviation by 2% per percentage point of frequency. In the larger 2005 sample it is only 1% per percentage point. Thus, although there are fewer sales in the smaller sample the effect of each sale on the standard deviation is relatively large.

Table 4*: The smaller samples

Sample	2005	2004	2004-2005
# UPC	215	80	18
# stores, All	23.54 (min=11,max=35)	16.375 (11,21)	16.06 (13,19)
# stores, Regular	22.34 (min=11,max=35)	15.89 (11,21)	14.70 (11,19)
Av. Freq	0.05	0.03	0.02
No Sales	0.74	0.84	0.84
SD All	0.0434	0.0263	0.0347
SD Regular	0.0393	0.0220	0.0318
Ratio	1.1	1.2	1.09
Relative	2.00	6.67	4.50
Corr	0.51	0.52	0.11

* Similar to Table 2 but here we delete UPCs with less than 11 stores that post regular prices in any week.

We use the smaller samples described in Table 4 and split the stores in each UPC-Week cell into bins of approximately equal size. For example, in the two bins case, we have high and low-price stores, where the price of the stores in the high price bin is greater than or equal to the median. Our main findings are:

(a) The average price charged by the stores in any given bin fluctuates over weeks, but the variations in the average price are larger for low price bins. This holds also after removing sales observations.
(b) There are "sale prices" in all bins, but the fraction of sale prices is larger in low price bins.

Table 5 is about bin size. As was said before, the bins are only approximately of the same size because of the discrete nature of the data. In the 2 bins division, 60% of the stores are in bin 1 (the highest price bin) and 40% in bin 2 (the low-price bin). Later, when we control for store effects, the sizes of the bins are more similar.

Table 6a is about the frequency of temporary sales by bins. This is calculated by dividing the number of "sale prices" in the bin (aggregating over all UPCs and weeks) by the number of prices in the bin. When using the 2005 sample and the 2 bins division, 10% of the prices in the higher price bin 1 are "sale prices". The number for bin 2 is 34%. Using the 2005 sample and the 5 bins division, 42% of the prices in the lowest price bin (bin 5) are sale prices. The number for the highest price bin (bin 1) is 5%. This says that being on sale does not guarantee low relative price. The fraction of prices on sale is increasing with the index of the bin suggesting that the probability that an item is cheap relative to other stores given that it is on "sale" is higher than the unconditional probability.

Table 6b estimates the conditional probabilities: The probability that a price is in bin i given that it is a "sale price". For example, when using the 2005 sample and a 2 bins division, the probability that a "sale price" is in bin 1 is 0.3. This conditional probability is calculated as follows. The unconditional probability that a price is in bin 1 is: $\Pr ob(bin1) = 0.6$. The unconditional probability that a price is a "sale price" is: $\Pr ob(sale) = 0.2$. The probability that a price in bin 1 is a "sale price" is: $\Pr ob(Sale \mid bin = 1) = 0.1$. The probability that a price is in bin 1 and it is a "sale price" is: $\Pr ob(bin1 \cap Sale) = \Pr ob(bin1) \Pr ob(Sale \mid bin = 1) = (0.6)(0.1) = 0.06$. The probability that a price is in bin 1 given that it is a sale price is:

$\Pr ob(bin1 \mid price = "sale") = \frac{\Pr ob(bin1 \cap Sale)}{\Pr ob(Sale)} = \frac{0.06}{0.2} = 0.3$. There is a remarkable agreement about the estimates of the conditional probabilities across samples.

Table 5*: Bin size

	bin 1	bin 2	bin 3	bin 4	bin 5
1 bin					
All samples	1				
2 bins					
2004	0.60	0.40			
2005	0.60	0.40			
2004-2005	0.60	0.40			
3 bins					
2004	0.47	0.25	0.28		
2005	0.47	0.24	0.29		
2004-2005	0.46	0.25	0.29		
5 bins					
2004	0.34	0.16	0.16	0.15	0.19
2005	0.34	0.16	0.15	0.15	0.20
2004-2005	0.33	0.16	0.16	0.15	0.20

* The average fraction of stores in each bin. Averages are over weeks and UPCs.

Table 6a: Frequency of temporary sales by bins

	bin 1	bin 2	bin 3	bin 4	bin 5
2 bins					
2004	0.09	0.35			
2005	0.10	0.34			
04-05	0.10	0.37			
3 bins					
2004	0.06	0.22	0.38		
2005	0.07	0.24	0.37		
04-05	0.07	0.25	0.41		
5 bins					
2004	0.04	0.16	0.24	0.33	0.43
2005	0.05	0.18	0.27	0.32	0.42
04-05	0.05	0.18	0.27	0.35	0.45

* The first 5 columns are the frequency of "temporary sales" by bins. These frequencies are obtained by dividing the number of "temporary sale prices" in the bin (aggregating over UPCs and weeks) by the total number of prices in the bin.

Table 6b: The probability that the price is in bin i given that it is a "sale price"

	bin 1	bin 2	bin 3	bin 4	bin 5
2 bins					
2004	0.28	0.72			
2005	0.30	0.70			
2004-2005	0.28	0.72			
3 bins					
2004	0.15	0.29	0.56		
2005	0.16	0.30	0.54		
2004-2005	0.15	0.30	0.56		
5 bins					
2004	0.07	0.12	0.18	0.24	0.39
2005	0.07	0.13	0.18	0.23	0.39
2004-2005	0.07	0.12	0.18	0.23	0.39

Table 7 provides the averages of the main variables using the 2 bins division. Here and in the rest of the paper we use the larger samples described in Tables 1 -3 to estimate magnitudes labelled "All prices". To estimate magnitudes labelled "Regular prices" we delete "sale observations" from the smaller samples described in Table 4. These smaller samples have at least 11 stores per UPC-week cell and therefore we can divide the cell into bins.

The difference in average log price between the high price stores and the low-price stores (P1-P2) is about 20%. (It is 21% for the 2004 sample, 18% for the 2005 sample and 21% for the combined 04-05 sample). For regular prices, the average price is about 15% higher in the high price bin. Thus, it seems that temporary sales contribute to cross sectional price dispersion.

Table 7*: Means

	All Prices		Regular Prices	
	P1	P2	P1	P2
2004	0.81	0.59	0.9	0.76
2005	0.86	0.68	1.08	0.93
2004-05	0.76	0.55	1.17	1.03

* The Table uses the 2 bins division. P1 is the average log price for high price stores and P2 is the average log price for low price stores (average across UPCs). The first two columns use the larger samples of all prices described in Tables 1-3. The last two columns use the sample of regular prices obtained by deleting observations that are labeled as "sale prices" from the smaller samples described in Table 4.

Table 8 computes the standard deviation of the average price over weeks. We first calculate the average (over stores) price for each UPC-week-bin cell. We then calculate the standard deviation of these averages for each UPC-bin across weeks. Table 9 reports the average of these standard deviations across UPCs. In the two bins case, the standard deviation of P2 (the average weekly log price in the low-price bin) is more than 30% larger than the standard deviation of P1. It is larger by 54% for the 2004 sample, by 30% for the 2005 sample and by 40% for the 04-05 sample.

The following 3 rows in Table 8 describe the standard deviations when dividing each UPC-Week cell into three bins: High, medium and low. Also here, the standard deviation of the price in the low price bin is higher than the standard deviation of the price in the high price bin. The last rows in Table 8 are the standard deviations when dividing each UPC-Week cell into 5 bins. The standard deviations in bin 5 (the lowest price bin) are higher than the standard deviations in bin 1 (the highest price bin). The ratio of the standard deviations of the average price in bin 5 to the standard deviation in bin 1 is 1.8 on average (2 for 2004, 1.6 for 2005 and 1.76 for 2004-05).

Table 8*: Standard deviations over weeks

	2004	2005	2004-2005
One bin			
P	0.08	0.09	0.08
Two bins			
P1	0.07	0.09	0.08
P2	0.11	0.11	0.11
Three bins			
P1	0.06	0.08	0.07
P2	0.09	0.11	0.10
P3	0.12	0.11	0.12
Five bins			
P1	0.06	0.07	0.07
P2	0.08	0.10	0.09
P3	0.10	0.11	0.10
P4	0.11	0.11	0.11
P5	0.12	0.11	0.13
Regular prices			
	2004	2005	2004-2005
One bin			
P	0.02	0.04	0.03
Two Bins			
P1	0.02	0.04	0.03
P2	0.04	0.05	0.04
Three Bins			
P1	0.03	0.04	0.03
P2	0.04	0.05	0.05
P3	0.04	0.06	0.05
Five Bin			
P1	0.03	0.04	0.03
P2	0.04	0.05	0.05
P3	0.04	0.06	0.05
P4	0.05	0.06	0.05
P5	0.04	0.06	0.05

* This Table reports standard deviations over weeks. We first calculate the average price for each UPC-week-bin cell. We then calculate the standard deviation of these averages for each UPC-bin across weeks. The Table reports the average of these standard deviations over UPCs. The first rows report the standard deviation for the 2 bins case. The next rows report the standard deviation for the 3 bins case and the rows in the bottom report the standard deviation for the 5 bins case. The second half of the table repeats the calculations after eliminating all "temporary sale" observations from the smaller samples described in Table 4.

3.1. Store effect

Stores that are similar in price may be similar in other ways. For example, stores in rich neighborhoods may charge on average a price that is higher than the price charged by stores in poor neighborhoods. In an attempt to address this problem, we remove the store effect by running each UPC on a store dummy and using the residuals from these regressions instead of the original prices.

When using the residuals, the bins are more equal in size because the residuals are different across stores and the problem of lack of price dispersion is less common. The conditional probabilities in Table 6b' are not very different from the conditional probabilities in Table 6b.

Table 6a': Frequency of temporary sales by bins

	bin 1	bin 2	bin 3	bin 4	bin 5
2 bins					
2004	0.07	0.32			
2005	0.09	0.31			
2004-2005	0.08	0.35			
3 bins					
2004	0.05	0.15	0.38		
2005	0.06	0.18	0.35		
2004-2005	0.05	0.18	0.41		
5 bins					
2004	0.03	0.08	0.15	0.27	0.42
2005	0.05	0.10	0.17	0.27	0.39
2004-2005	0.04	0.09	0.17	0.31	0.44

Table 6b': The probability that a price is in bin i given that it is a "sale price"

	bin 1	bin 2	bin 3	bin 4	bin 5
2 bins					
2004	0.18	0.82			
2005	0.23	0.77			
2004-2005	0.20	0.80			
3 bins					
2004	0.08	0.25	0.67		
2005	0.11	0.29	0.60		
2004-2005	0.09	0.26	0.64		
5 bins					
2004	0.04	0.07	0.15	0.27	0.47
2005	0.05	0.10	0.17	0.27	0.42
2004-2005	0.05	0.08	0.15	0.27	0.45

Table 8' is comparable to Tables 8. For the sample of all prices, it shows the same pattern: The standard deviation across weeks is increasing with the index of the bin. The relationship between the index of the bin and the standard deviation is weaker for the samples of regular prices.

Table 8': Standard deviations over weeks

	2004	2005	2004-2005
One bin			
P	0.08	0.09	0.08
Two bins			
P1	0.06	0.09	0.07
P2	0.11	0.11	0.11
Three bins			
P1	0.06	0.08	0.07
P2	0.08	0.10	0.09
P3	0.12	0.12	0.13
Five bins			
P1	0.06	0.07	0.07
P2	0.07	0.09	0.08
P3	0.09	0.11	0.09
P4	0.10	0.12	0.11
P5	0.13	0.12	0.13
Regular prices			
	2004	2005	2004-2005
One bin			
P	0.02	0.04	0.03
Two Bins			
P1	0.02	0.04	0.03
P2	0.03	0.04	0.04
Three Bins			
P1	0.02	0.04	0.03
P2	0.02	0.04	0.03
P3	0.03	0.05	0.04
Five Bin			
P1	0.02	0.04	0.04
P2	0.02	0.04	0.03
P3	0.02	0.04	0.03
P4	0.02	0.04	0.03
P5	0.04	0.05	0.05

Table 9 computes the ratio of the average standard deviation in the lowest price bin to the average standard deviation in the highest price bin (averaged across samples). The first column (All) reports the ratio of the standard deviations of prices when using the sample of all prices. When using the 2 bins division, the standard deviation in the low-price bin is 42% larger than the standard deviation in the high price bin. This difference is 52% when controlling for a store effect and 53% when controlling for a UPC specific store effect. When using the 3 and 5 bins divisions the differences are larger. The percentage differences in the standard deviations are lower when using the sample of regular prices (Regular in the second row).

Table 9*: Ratio of the average standard deviation in the lowest price bin to the average standard deviation in the highest price bin.

	Original prices			UPC specific Store effect		
	2 bins	3 bins	5 bins	2 bins	3 bins	5 bins
All	1.421	1.65	1.8	1.53	1.81	2.01
Regular	1.420	1.57	1.62	1.22	1.3	1.36

* The Table reports the ratio of the average standard deviation in the lowest price bin to the average standard deviation in the highest price bin. Averages are over samples. The first row (All) is the ratio of the standard deviations of prices in the samples of all prices. The second row (Regular) is this ratio in the samples of regular prices.

Table 9 shows that the average price in the low-price bin fluctuates more than the average price in the high price bin. It supports the hypothesis that high prices fluctuate less than low prices even after removing “sale” observations.

4. IMPULSE RESPONSE FUNCTIONS

Our main hypothesis is that stores drop prices in reaction to the accumulation of unwanted inventories. When the shock to demand is small the drop in prices will last for a short time and will qualify as a "temporary sale". When the shock to demand is large and the amount of accumulated inventories is large the drop in prices will last for a longer time and will qualify as a change in regular price. We now turn to use Vector Auto Regression (VAR) analysis to test this hypothesis directly.

We examine the effect of two kinds of shocks to the demand for a given UPC: Store specific shocks and aggregate shocks. We start with store specific shocks.

4.1 Store specific demand shock

We ran a VAR with three variables: log prices, log quantity and a dummy variable with the value of 1 if the price is a "sale price" and zero otherwise. Prices were ordered first, then the sale dummy and then quantity. Our definition of a sale uses some information from the future and we therefore adopted a slightly different definition of a sale. In our new definition a sale starts with a drop in the price of more than 10% and continues until either there is a price increase or the number of weeks since the initial drop is 4. This definition of sale is based only on the past and current price observations and therefore it fits the VAR framework. We ran the VAR with both definitions. The results are very similar. We report here the results that use the new definition. The standard errors of the estimated impulse response functions are small and therefore we report here only the mean estimates.[7]

We ran a four lag VAR with UPC by store and UPC by week fixed effects. In the 2005 sample we have 1084 UPCs and we can think of the fixed effects as using the residuals from (3)(1084) regressions. For each UPC we run three regressions: Price on a store dummy and a week dummy, Quantity on a store dummy and a week dummy and Sale on a store dummy and a week dummy. We use the residuals from these regressions to estimate the impulse response functions.

The impulse response functions (IRFs) tell us the reactions of prices and sale to a store and week specific quantity shock. A shock is thus a decrease in demand relative to what can be predicted on the basis of the average demand across stores for the specific week and the quantity sold by the store in the past weeks.

The impulse response functions estimated in Figure 2 show that indeed prices go down in response to a negative demand shock. This is also the case if we remove sale observations (and run the VAR only on the two remaining vectors). The price response to a demand shock reaches a peak after 1 or 2 weeks (1 week in the 2004 and 04-05 samples and 2 weeks in the 2005 sample). After eliminating sale observations we get a similar pattern but the size of the response is smaller. The peak response after eliminating sale observation is about half of what we get when sale observations are included.

The results about the effect of a quantity shock on the sale dummy are in Figure 3. The predicted level of the sale dummy may be interpreted as the probability of having a sale. The probability of having a sale increases after a negative shock to demand. The peak response of about 0.5%, occurs after 1 or 2 weeks and it then declines.

---

[7] In the 2005 sample, the frequency of sale under the old definition is 0.195 and under the new definition it is 0.218. The correlation between the two is 0.682.

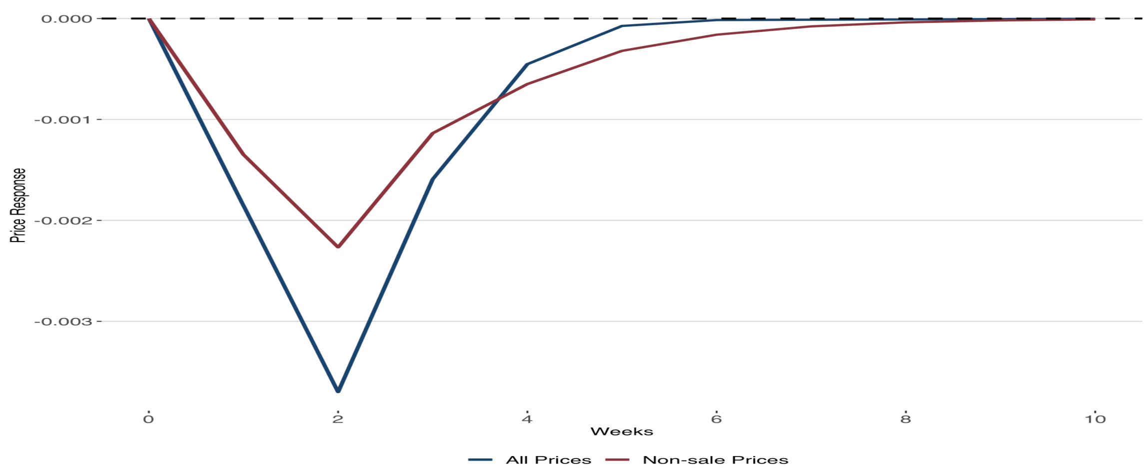

A. 2005

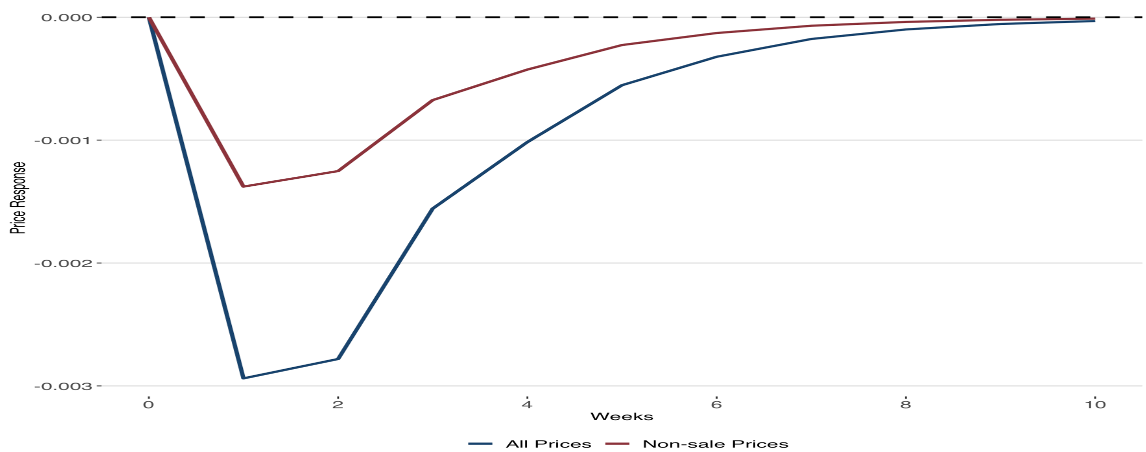

B. 2004

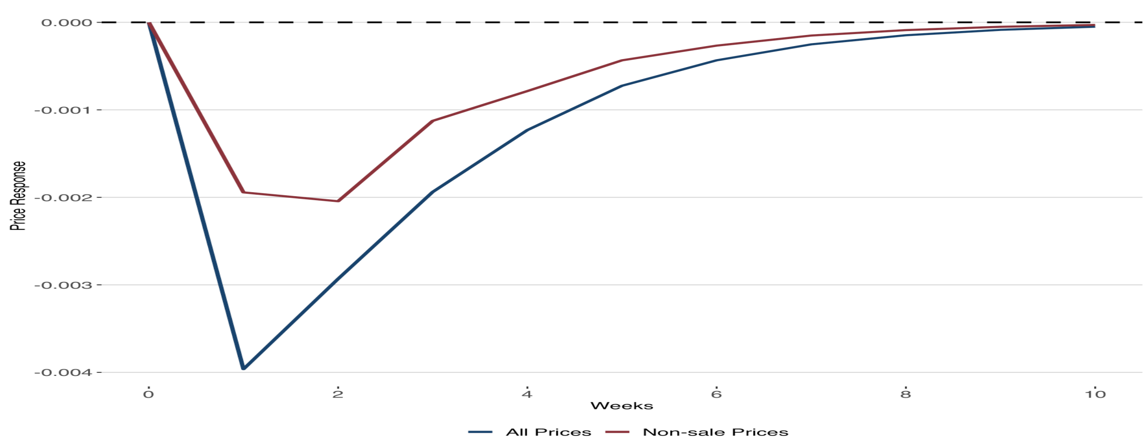

C. 2004-2005

Figure 2: The effect of a store/UPC specific demand shock on prices

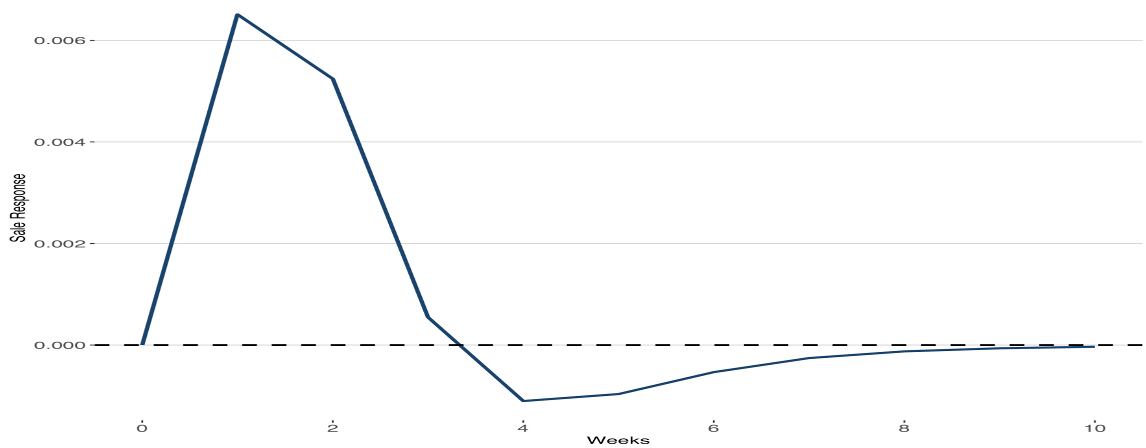

A. 2005

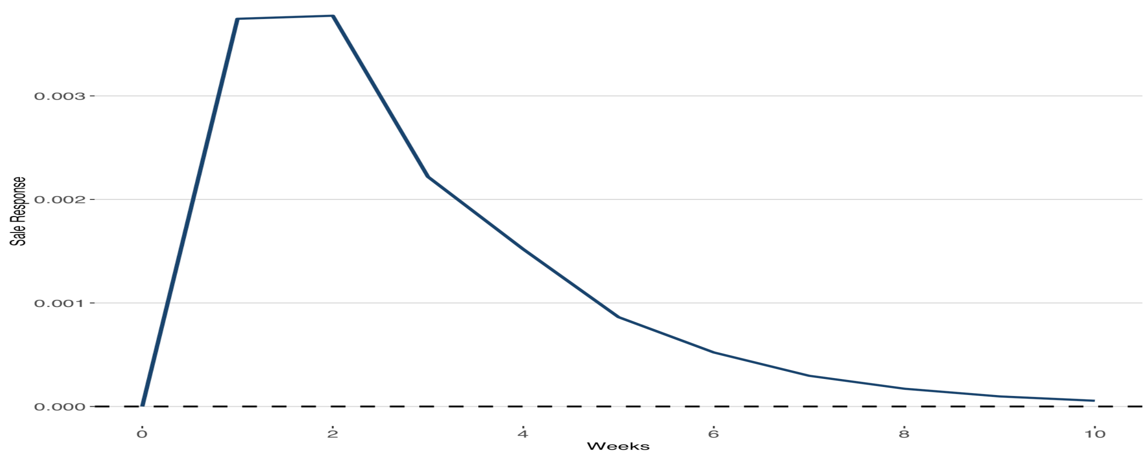

B. 2004

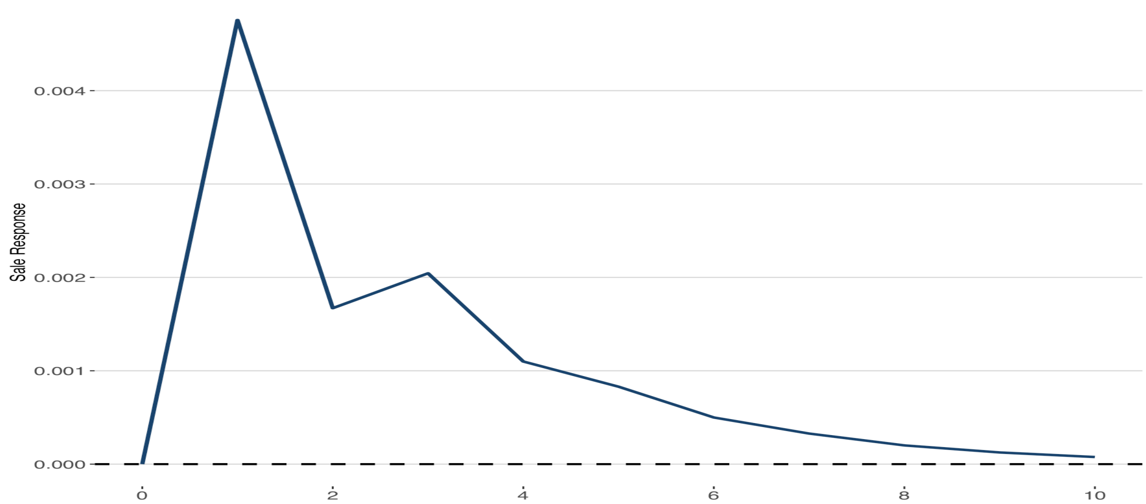

C. 2004-2005

Figure 3: The effect of a store/UPC specific demand shock on the "probability of a sale"

4.2 Aggregate demand shock

Do sales occur in response to shocks that affect the aggregate amount sold by all stores? To examine this question we use the 04-05 sample that allows for a week fixed effect. The 04-05 sample has 104 weeks. We use 52 weeks dummies where dummy 1 gets the value of 1 if it is the first week of the year and zero otherwise. Dummy 2 gets 1 if it is the second week of the year and zero otherwise and so on.

The variables in the VAR are the average log price (average over stores), the fraction of stores that had a sale in the particular week, and the average log quantity. We ran the VAR with UPC by week fixed effects. This is equivalent to using the residuals from a regression of the variables in each UPC on the week dummy.

Figure 4A describes the effect of a negative quantity (demand) shock on prices. We see that a negative demand shock leads to a reduction in prices even after eliminating sale observations. But sales seem to play an important role in the price response.

Figure 4B describes the response of sales to a negative demand shock. A negative demand shock leads to an increase in the fraction of stores that have a sale immediately after the shock and after that the effect declines gradually over time.

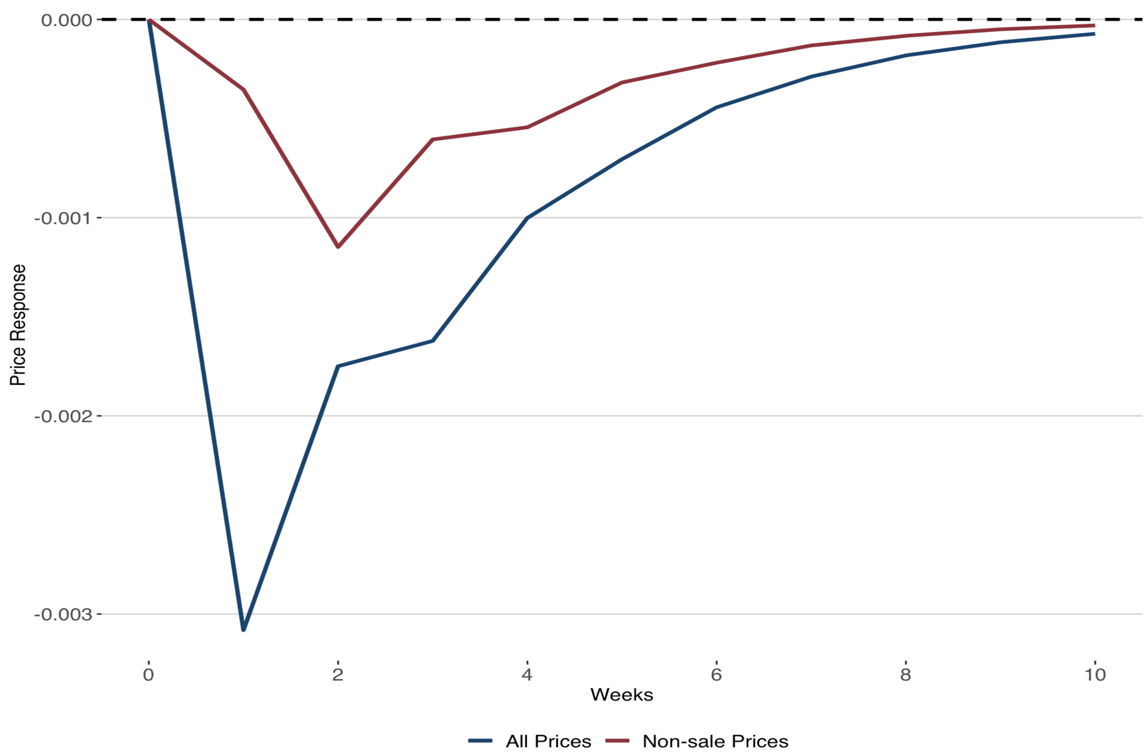

A. price response

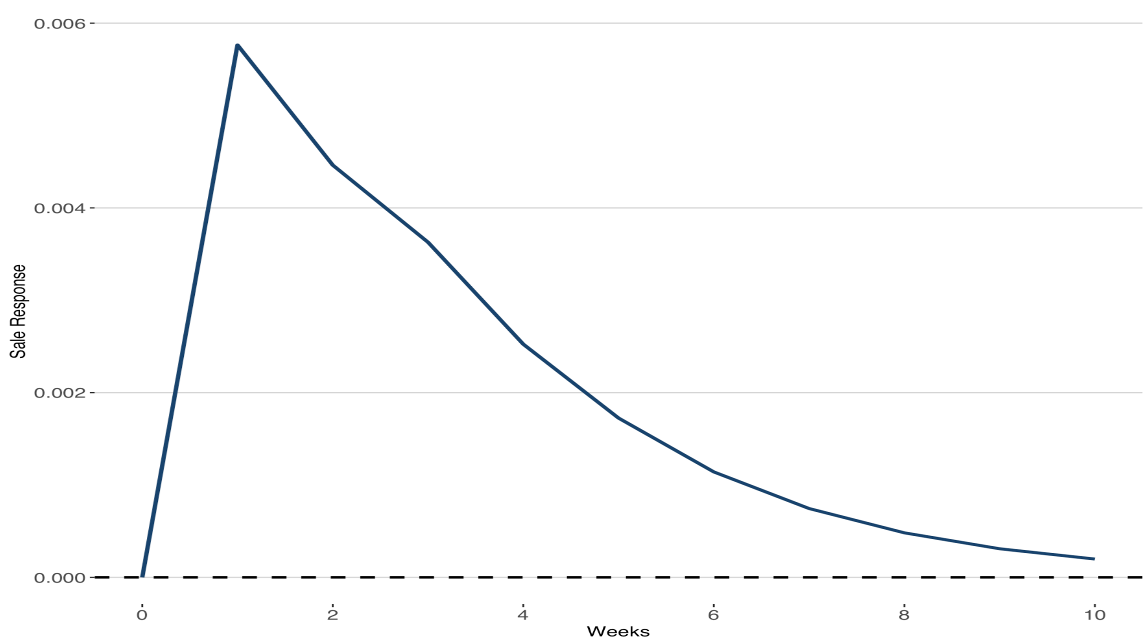

B. Sale response

Figure 4: The effects of a UPC specific demand shock

In Figure 4B, the peak response to an aggregate shock occurs after a week and is 0.6%. The cumulative effect (over all weeks) on sale is 2%. This does not seem to be a large effect and it seems consistent with the view that many sales are planned ahead of time. On the other hand, eliminating sale observations may lead to a gross underestimate of the price response. When using all observations, the peak price response occurs in the first week and it is 0.3%. After eliminating sale observations, the peak response occurs in the second week and it is 0.1%. Thus, the peak response occurs later and its magnitude is about a third of the peak response when using all prices. The cumulative effect when using all observations is 0.9% while the cumulative effect after eliminating sale observations is only 0.4%, which is less than half of the response we get when using all observations. We may add that there is no difference in the size of the shock: After controlling for week fixed effects, the standard deviation of the quantity sold is 0.201 when using all observations and 0.198 when using only regular price observations. Thus, if we eliminate sale observations and focus on regular prices, we may conclude that the price response to a demand shock is much smaller than it actually is.

5. SEQUENTIAL TRADE

We would like to discuss the above findings in terms of the UST model. We start with a simple version and then augment it to allow for storage.

5.1 A simple version

There are many goods and many sellers who can produce the goods at a constant unit cost. We focus on one good with a unit cost of λ. Production occurs at the beginning of the period before the arrival of buyers. Storage is not possible. The number of buyers $\tilde{N}$ is an iid random variable that can take two possible realizations: N with probability $1-q$ and $N+\Delta$ with probability q.

Sellers take prices and the probability of making a sale as given. They know that they can sell at the price P_1 for sure. They may also be able to sell at a higher price, P_2, if demand is high, with probability q. In equilibrium sellers are indifferent between the two price tags: The expected profits are the same for both tags.

It is useful to think of two hypothetical markets. The price in the first market is P_1 and the probability that this market opens is 1. The price in the second market is P_2 and the probability that it opens is q. From the seller's point of view, he can sell any quantity at the price announced in the market, if the market opens but cannot sell anything at that market if the market does not open. A unit with a price tag of P_1 will be sold in the first market. A unit with a price tag of P_2 will be sold in the second market, if this market opens.

Buyers arrive sequentially in batches. The first batch of N buyers buys in the first market at the price P_1. The second market opens only if the second batch of Δ buyers arrives. If this second batch arrives the second market opens at the price P_2.

The demand of each of the active buyer at the price P is: $D(P)$. In equilibrium sellers supply x_1 units to the first market and x_2 units to the second market.

Equilibrium is thus a vector (P_1, P_2, x_1, x_2) such that the expected profits for each unit is zero:

(1) $$P_1 = qP_2 = \lambda$$

And markets that open are cleared:

(2) $$x_1 = ND(P_1) \text{ and } x_2 = \Delta D(P_2)$$

Figure 5 illustrates the equilibrium solution. The demand in market 1 at the price λ, $ND(\lambda)$ is equal to the supply to the first market (x_1). When market 2 opens at the price λ / q, the demand in this market, $\Delta D(\lambda / q)$, is equal to the supply (x_2).

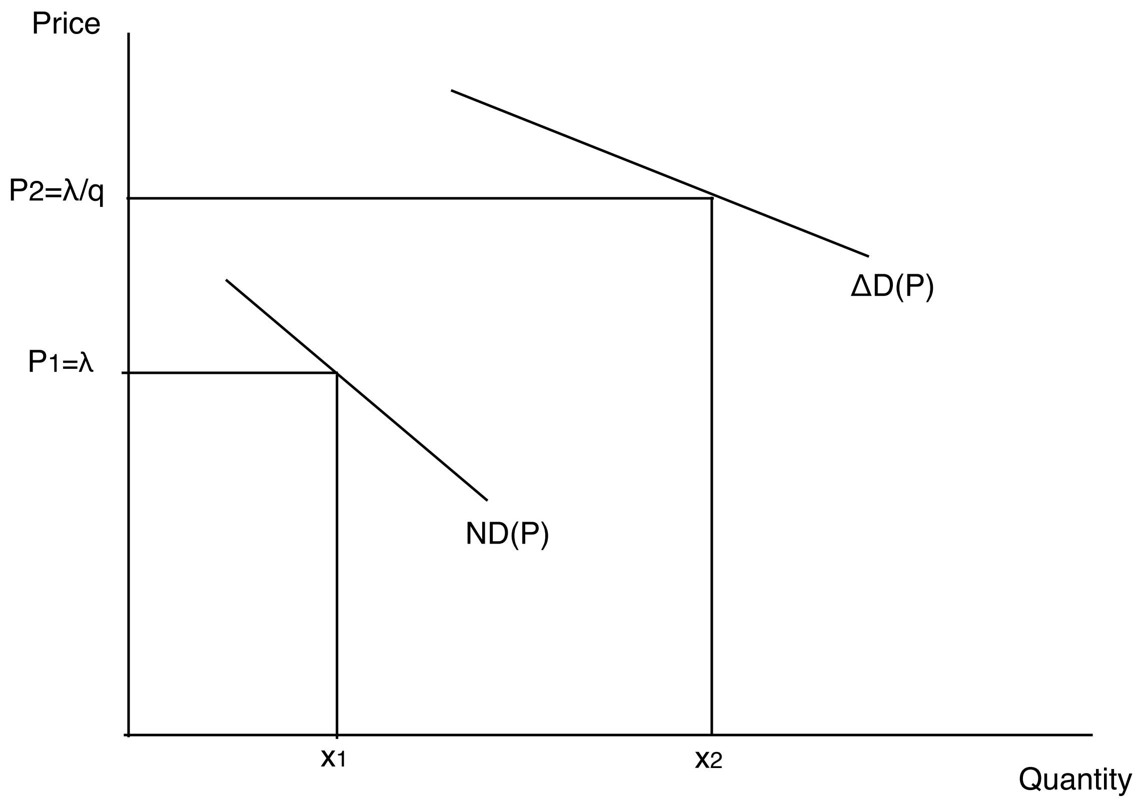

Figure 5: Prices and quantities in the simple version of the UST model

Note that in this simple version posted prices do not change over time. The quantity sold at the low price does not change over time but the quantity sold at the high price fluctuates over time.

3.2 Storage

Bental and Eden (BE, 1993) study a UST model that allows for storage. In their model prices fluctuate as a result of both *iid* demand and supply shocks. A negative demand shock leads to the accumulation of inventories and a reduction in all prices. A temporary reduction in the cost of production has a similar effect. Thus, "temporary sales" may be the result of both demand and supply shocks.

The BE model assumes a convex cost function and exponential depreciation. Here we assume a constant per unit cost and one-hos-shay depreciation. The constant per unit cost simplifies the analysis. The one-hos-shay depreciation is realistic because most supermarket items have an expiration date. It also serves as a tiebreaker and yields predictions about temporary sales that are an important feature of the data.

To simplify, we assume that the good can be stored for one period only. Thus, if a good is not sold in the first period of its life, it can still be sold in the second period but it has no value if it is not sold within the two period.

As before, the number of buyers $\tilde{N}_t$ is iid and can take two possible realizations: $\tilde{N}_t = N$ with probability $1-q$ and $\tilde{N}_t = N+\Delta$ with probability q.

At the beginning of period t the economy can be in one of two states. In state I (I for inventories) the demand in the previous period was low ($\tilde{N}_{t-1} = N$) and the second market did not open. As a result, inventories were carried from the previous period. In state NI (NI for no inventories) demand was high ($\tilde{N}_{t-1} = N+\Delta$) and there are no inventories. The price in the first market is $P(1,I)$ in state I (with inventories) and $P(1,NI)$ in state NI (with no inventories). The quantity offered for sales in market 1 is $x(1,I)$ in state I and $x(1,NI)$ in state NI. The price in the second market (P_2) and the supply (x_2) do not depend on the level of inventories. The quantity sold in the first market is equal to the quantity offered for sale. The quantity sold in the second market is zero if demand is low and x_2 if demand is high. Table 10 describes quantities as a function of last period's demand and this period's demand. Note that the quantity sold depends both on the last period's demand and on this period demand. Production depends only on the last period's demand. With some abuse of notation we write the level of inventories in state I, as $I = x_2$. Production in state I is equal to the demand in the first market while production in state NI is equal to the demand in both markets.

Table 10*: Quantities in period t as a function of the state at $t-1$ and the state at t

	$\tilde{N}_t = N + \Delta$		$\tilde{N}_t = N$	
	Quantity sold	Production	Quantity sold	Production
$\tilde{N}_{t-1} = N + \Delta$	$x(1,NI) + x_2$	$x(1,NI) + x_2$	$x(1,NI)$	$x(1,NI) + x_2$
$\tilde{N}_{t-1} = N$	$x(1,I) + x_2$	$x(1,I)$	$x(1,I)$	$x(1,I)$

A formal analysis and the equilibrium definition is in Eden (2018, Appendix A.3). To make this paper self-contained we repeat here the description of the model. In allocating the available amount of goods (from new production and inventories) across the two markets, the older units receive "priority" in the first market (and the younger units receive "priority" in the second market). Given prices the allocation rule is as follows. If the number of old units (from inventories) is less than the demand in the first market then all old units are supplied to the first market. If the number of old units is greater than the demand in the first market then only old units are supplied to the first market. To motivate this allocation rule, we consider the following example. There are two stores: Store O with old units and store Y with young units. Suppose further that store Y posts the first market low price and store O posts the second market high price. In this case if aggregate demand is low, store O does not sell and the units supplied by store O expire. Alternatively, if store O posts the first market price and store Y posts the second market price, the unsold units supplied by store Y do not expire and can be sold next period. It follows that the joint profits of both stores can be increased if they do not follow our allocation rule. This cannot occur in equilibrium.

A young unit that is not sold in the current period will be sold in the next period at the price $P(1,I)$. The value of a young unit that is not sold in the current period (the value of inventories) is $\beta P(1,I)$, where $0 < \beta < 1$ is a constant that captures discounting, storage costs and depreciation. The value of an old unit that is not sold is zero. Newly produced units are supplied to the second market and in equilibrium the following arbitrage condition must hold.

(3) $$qP_2 + (1-q)\beta P(1,I) = \lambda$$

The left-hand side of (3) is the expected present value of revenues from a newly produced unit allocated to the second market. If the second market opens (with probability q) the seller gets P_2. Otherwise he will get the unit value of inventories, $\beta P(1,I)$. The right hand side of (3) is the unit cost of production. Thus, (3) says that the marginal cost is equal to expected revenues.

We now distinguish between two cases. In the first case, illustrated by Figure 6A, inventories in state I are relatively low and newly produced goods are supplied in state I to both markets. Since newly produced goods are supplied to the first market, the price in the first market is the marginal cost: $P(1,I) = P(1,NI) = \lambda$. Substituting this into (4) yields:

(4) $$P_2 = \frac{\lambda\left(1-(1-q)\beta\right)}{q}$$

In the second case, illustrated by Figure 6B, newly produced goods are supplied to the first market only in state NI. In state I the entire supply to the first market is out of inventories and the supply to the second market is of both newly produced units and old units. Since old units are supplied to both markets, we must have:

(5) $$qP_2 = P(1,I)$$

This says that the expected revenue of supplying an old unit to the second market is the same as the revenue from supplying it to the first market. The solution to (3) and (5) is[8]:

(6) $$P(1,I) = \frac{\lambda}{1+(1-q)\beta} < \lambda \text{ and } P_2 = \frac{\lambda}{q\left(1+(1-q)\beta\right)}$$

[8] It is also possible that all the old units are allocated to the first market and all the new units are allocated to the second market. Also in this case the first market price can be below cost: $P(1,I) \leq \lambda$. See, Eden (2018).

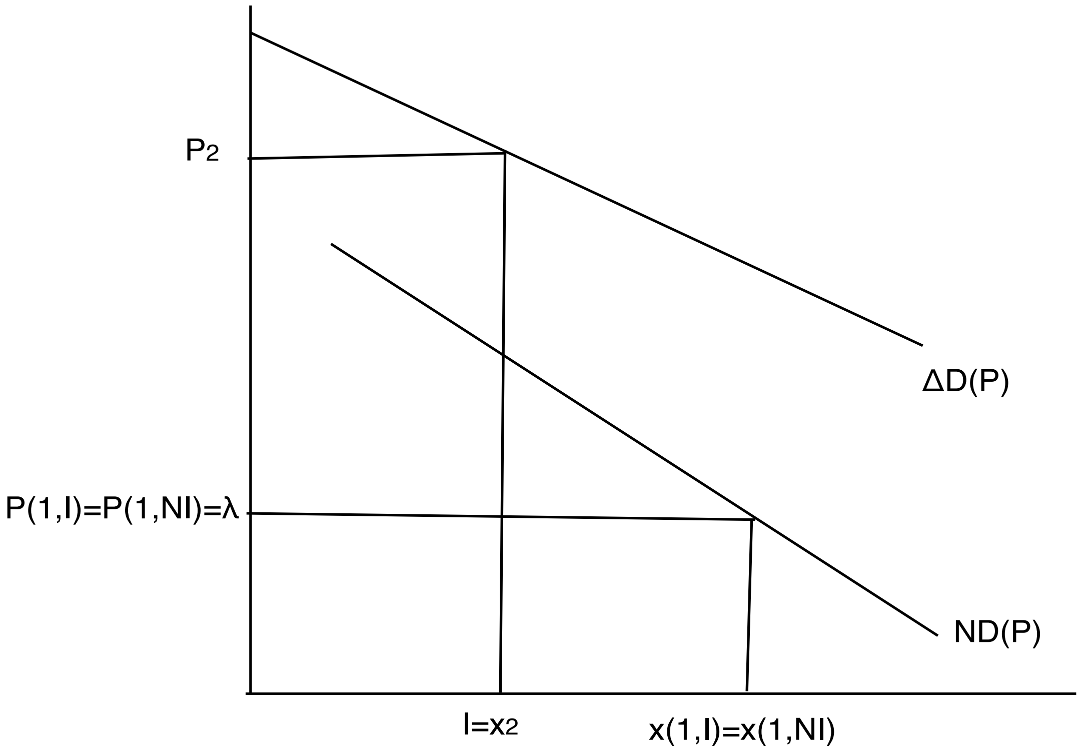

A. In state I, $I=x_2$ "old units" and $x(1,I)-I$ newly produced units are supplied to the first market.

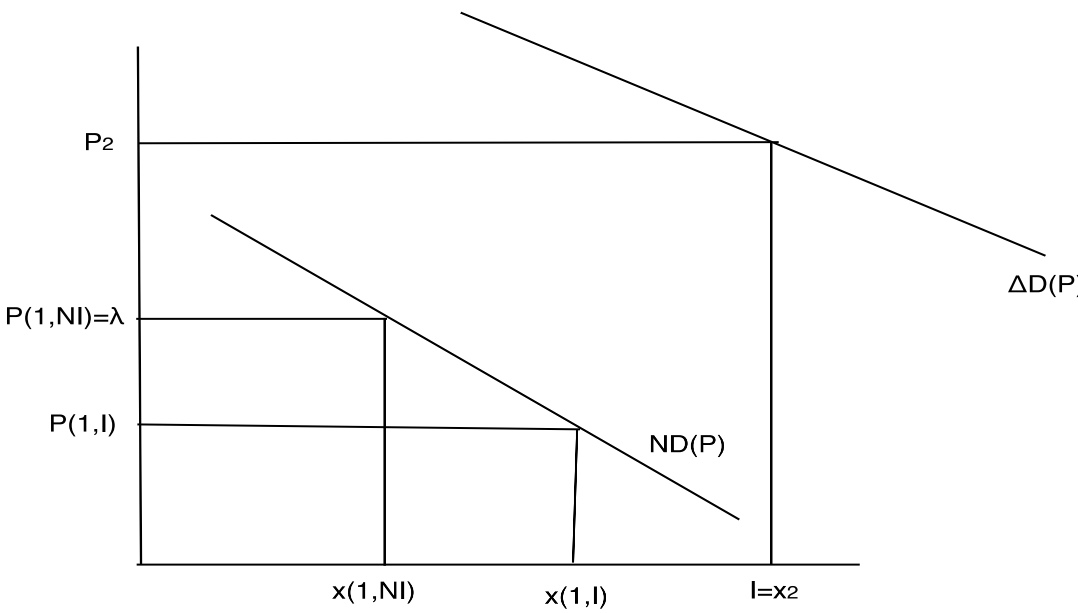

B. In state I, $x(1,I)$ "old units" are supplied to the first market and $\mathrm{I}-x(1,I)$ "old units" are supplied to the second market. No new units are supplied to the first market.

Figure 6: Possible Equilibria

The model described by Figure 6 may account for temporary sales. Some stores offer the newly produced good at the high ("regular") price of market 2. Then if demand is low they accumulate inventories and offer the good for sale at the low price of market 1. We also note that the price and quantity in the first market may change over time. Prices are more volatile than cost. In the example illustrated by Figure 6B, the cost does not change but the price in the first market does.[9]

A 1% change in λ will change prices in both markets (and both states) by 1%. Sales may occur in response to temporary change in cost in addition to sales that occur in response to the accumulation of inventories. There is however a difference between the two types of sales: Sales that occur because of a change in technology will affect the price in both markets while sales that occur because of the accumulation of inventories affect the price in the first market only.

5.3 Accounting for the stylized facts.

We assumed that the life of the good is two periods and there are only two possible realizations of demand. Nevertheless, we think that this rather specialized example can account for some of the observations made above.

1. For the average good, the fraction of weeks in which there is no sale in any of the stores is not small.

This will occur in our example if the probability that demand is at its highest realization is not small.

2. Stores with higher average price tend to have more sales. In our example, if a store posts the low price it will not accumulate unwanted inventories and will not have sales.

---

[9] Eichenbaum, Jaimovich, and Rebelo (EJR, 2011) found that most price changes are associated with a change in cost. Under the constant returns to scale assumption, changes in prices may occur as a result of changes in inventories and without changes in cost. This is not the case in the BE model. In their model, there is no distinction between the producer and the store (the retailer). Once we introduce this distinction, there is no puzzle. The producer set a price that changes in response to both the aggregate level of inventories and technology. Therefore, when we look at the cost from the store's point of view (as EJR did), changes in aggregate inventories are associated with changes in cost and most price changes are associated with cost changes.

3. Removing temporary sales reduces the standard deviation of the average price over weeks. It can be shown that this will occur in the case illustrated by Figure 6B.
4. The standard deviation of quantity does not change much after removing sale observations.

Our model suggests that removing sale observations will increase the standard deviation of the quantity sold. Sales that are planned in advance are likely to work in the opposite direction. It is therefore possible that the effects of the two types of sales cancel each other.

5. Stores with relatively high average price tend to have more sales. In our model stores that advertise the low price are not likely to have sales.
6. A temporary high price may not be cheap relative to the price in other stores. In our example, supply shocks (changes in λ) leads to changes in all prices. If the change in λ lasts for a short time and then reverts to its previous level it may lead to a change in prices that we identify as sales. Otherwise it will lead to a change in regular prices. This is consistent with the observation that temporary sales occur in all bins.
7. The upper part of the cross sectional price distribution is more rigid than the lower part. In our example the high price is constant over periods but the low price may change in response to the level of unwanted inventories. Therefore, the low price fluctuates more than the high price.
8. A negative demand shock increases the fraction of stores that offer the item at a sale price. In our example, stores that post a high price will accumulate unwanted inventories as a result of a negative demand shock and will offer the unsold goods at a sale price.
9. The reaction of prices to a demand shock appears to be much smaller after removing temporary sale observations. In the case illustrated by Figure 6B prices react to the accumulation of unwanted inventories and most of the reaction occurs by stores who offer inventories at sale prices.

5.4 Price rigidity

To better understand the role of temporary sales we consider now the case in which prices cannot be changed. For simplicity we adopt the case illustrated by Figure 6A.

As in Figure 6A, the price in the first market is λ and does not change over time. Initially, stores in the second market choose prices under the constraint that the prices cannot be changed. Whenever demand is low stores in the second market accumulate inventories and since sales are not allowed, they offer the inventories at the high (regular) price and at that price they may not sell and the goods may expire.

The price in the second market must satisfy the following arbitrage condition:

(3') $$qP_2 + (1-q)q\beta P_2 = \lambda$$

This arbitrage condition assumes that the newly produced unit is sold with probability q at the price P_2. If it is not sold in the first period it may still be sold (with probability q) in the second period of its life at the same price. The value of inventories in the case of price rigidity is $q\beta P_2$ and (3') says that the expected discounted revenue must equal the cost.

Solving (3') yields:

(4') $$P_2^* = \frac{\lambda}{q + (1-q)q\beta}$$

Note that (4') is greater than (4). The reason is that the value of inventories is less than in the price flexibility case: $q\beta P_2 = \dfrac{q\beta\lambda}{q+(1-q)q\beta} < \beta\lambda$.

The quantity supplied to the first market is the same as the quantity under price flexibility and is denoted by x_1 . The quantity supplied to the second market under price rigidity satisfies the demand when the second market opens: $x_2^* = \Delta D(P_2^*)$. This quantity is less than the quantity supplied under price flexibility $(x_2^* < x_2)$ because the second market price is lower under price flexibility.

As before there are two states. A state in which there are inventories and a state in which there are no inventories. But here, under price rigidity, the inventory state is defined as a state in which there are usable inventories. This will occur when the accumulated inventories are one period old. When the accumulated inventories are two periods old they cannot be used and we say that there are no inventories. Thus, the probability of an inventory state is lower under price rigidity because sometimes the accumulated inventories cannot be used.

Production under price flexibility is $x_1 + x_2$ in the no inventory state and x_1 in the inventory state. Production under price rigidity is $x_1 + x_2^*$ in the no inventory state and x_1 in the inventory state. In the flexible price case, the inventory state occurs when $\widetilde{N}_{t-1} = N$. In the rigid price case an inventory state occurs when $\widetilde{N}_{t-1} = N$ and $\tilde{N}_{t-3} = \tilde{N}_{t-2} = N + \Delta$. In this case the supply to market 2 at time $t-2$ is of newly produced units and the units that were not sold in the last period are not expired. Thus, as was said before, the probability of the inventory state is lower under price rigidity.

We may thus say that in our model temporary sales cannot be treated as a noise. Eliminating temporary sales will affect average regular prices, average consumption and average production. It may also affect the variance of production over weeks.

5. CONCLUDING REMARKS

We have made the following observations.

1. The fraction of weeks with no sales in any of the stores is much larger than the fraction predicted by the hypothesis that stores use a mixed strategy to choose temporary sales. For example, in the 2005 sample the fraction of weeks in which there is no sale in any of the stores is 45% while the mixed strategy hypothesis predicts 0.5%. Heterogeneity does not fully account for this observation. When looking at UPCs with temporary sales frequencies between 21 and 30 percent, there are no sales in any of the stores in 29% of the weeks while the mixed strategy hypothesis suggests that this fraction is 0.2%.

2. Goods with more fluctuations in regular prices have also more temporary sales. For example, the standard deviation of the average regular price in UPCs with sale frequencies between 20 to 30 percent is more than twice the standard deviation of the

average regular price in UPCs with sale frequencies between 0 to 10 percent. This suggests that some underlying feature of the product determines both regular price variability and sale frequency. In UST models, large demand shocks leads to the accumulation of a relatively large amount of unwanted inventories and it takes longer for inventories to get to the "normal level". It is therefore possible that large demand shocks leads to regular price change and small demand shocks leads to temporary sales. Goods that face more demand uncertainty may experience more large shocks and more small shocks and therefore goods that have more fluctuations in regular prices tend to have more sales.[10]

3. Temporary sales contribute substantially to the weekly variation of the average cross-sectional price of the typical good. In the larger sample of 2005 the standard deviation of the average price is 20% higher than the standard deviation of the average regular price. This is consistent with the hypothesis that in some weeks sales reduce the cross sectional average price more than in other weeks.

4. Stores with relatively high average regular price tend to have more sales. In our model, stores that consistently choose to offer newly produced goods at the low price of market 1 have little or no sales, while stores that choose to offer the newly produced goods at the high price of market 2 have relatively many sales. The model is thus consistent with the observation that stores with relatively high regular price tend to have more sales.

5. A temporary sale price may not necessarily be cheap relative to the price in other stores. But a sale price is more likely to be relatively cheap: The probability that a sale price is in the bottom third of the distribution is around 60%. In our model, temporary change in cost may lead to sales in all stores and this explains why some sale prices are

---

[10] An alternative story may assume that the common underlying feature may be monopoly power. Goods produced by firms with more monopoly power will have more sales because higher monopoly power allows for the use of sales as a discrimination device. According to this alternative story, goods with more monopoly power should also have more fluctuations in the average regular price. The industrial economic literature suggests that this is not the case. See for example, Wu (1979). We therefore think that the correlation between the frequency of sale and the standard deviation of the average regular price is a challenge for the discrimination hypothesis.

not relatively cheap. Sales that occur as a result of demand shocks are relatively cheap in our model.

6. High prices appear to be more rigid than low prices. In the 2 bins division, the standard deviation of the average price in the low-price bin is 40% higher than the standard deviation of the average price in the high-price bin. This is true for the sample of all prices and the sample of regular prices that is obtained after removing temporary sale observations. It is consistent with the hypothesis that price reductions occur as a response to the accumulation of unwanted inventories.

7. Sales play a critical role in the reaction of prices to a demand shock. When running a VAR with three variables: Average log price (average over stores that sale the same UPC), the fraction of stores in which the UPC is on sale and the average log quantity, the peak price response occurs immediately and it is -0.3%. When running the VAR after eliminating sale observations, the peak price response occurs with a one period lag and it is only -0.1% which is a third of the response we get when using all observations. The cumulative response to a negative demand shock is -0.9% when using all observations and -0.4% after eliminating sale observations. Thus, if we eliminate sale observations and look only at regular prices we may conclude that the price response to a demand shock is much smaller than it actually is.

We argue that these stylized facts are consistent with a model in which prices are completely flexible and temporary sales play a role in reacting to demand and supply shocks.

Our model is not consistent with the observation that temporary sales do not contribute to the weekly variation in the quantity sold. The literature suggests that some sales are planned in advance. It is possible that those planned in advance sale increase the weekly variation in the quantity sold. The effects of the two type of sale may thus be in the opposite direction and the total effect may therefore be close to zero.

To better understand the role of temporary sales we analyze in our model the case in which temporary sales are not possible (say because of high menu type costs). In our example, eliminating temporary sales would increase the average regular price and reduce average consumption and average production. These effects on the averages will show up even after aggregating across goods and time.

In the price rigidity version of our model some unwanted inventories expires before they get sold. In more general UST models price rigidity may prolong the effects of a negative demand shock because it will take longer to de-cumulate inventories.

REFERENCES

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

Anderson, Eric & Malin, Benjamin A. & Nakamura, Emi & Simester, Duncan & Steinsson, Jón, 2017. "**Informational rigidities and the stickiness of temporary Sales**," Journal of Monetary Economics, Elsevier, vol. 90(C), pages 64-83.

Bental, Benjamin and Benjamin Eden (1993). "**Inventories in a Competitive Environment**," Journal of Political Economy, University of Chicago Press, vol. 101(5), pages 863-86, October.

________, 1996. "**Money and inventories in an economy with uncertain and sequential trade**," Journal of Monetary Economics, Elsevier, vol. 37(3), pages 445-459, June.

Chevalier, Judith A., and Anil K. Kashyap. *2019. "Best Prices: Price Discrimination and Consumer Substitution." American Economic Journal: Economic Policy, 11 (1): 126-59.*

Coibion Oliver, Yuri Gorodnichenko and Gee Hee Hong., The Cyclicality of Sales, Regular and Effective Prices: Business Cycle and Policy Implications". American Economic Review 2015, 105(3):993-1029.

Dana James D. Jr. "Advance-Purchase Discounts and Price Discrimination in Competitive Markets" Journal of Political Economy, Vol.106, Number 2, April 1998, 395-422.

_______ "Equilibrium Price Dispersion under Demand Uncertainty: The Roles of Costly Capacity and Market Structure" The RAND Journal of Economics, Vol. 30, No. 4 (Winter, 1999), pp. 632-660.

Deneckere Raymond and James Peck., "Dynamic Competition with Random Demand and Costless Search: A Theory of Price Posting". Econometrica, Vol. 80, No. 3 (May, 2012), 1185–1247.

Eden, Benjamin. "Marginal Cost Pricing When Spot Markets are Complete" Journal of Political Economy, Vol. 98, No.6,1293-1306, Dec. 1990.

_______ "**Price Dispersion And Demand Uncertainty: Evidence From U.S. Scanner Data**," International Economic Review, Department of Economics, University of Pennsylvania and Osaka University Institute of Social and Economic Research Association, vol. 59(3), pages 1035-1075, August 2018.

Eichenbaum Martin & Nir Jaimovich & Sergio Rebelo, 2011. "**Reference Prices, Costs, and Nominal Rigidities**," American Economic Review, American Economic Association, vol. 101(1), pages 234-62, February.

Glandon, PJ. 2018. "Sales and the (Mis) Measurement of Price Level Fluctuations." Journal of Macroeconomics 58: 60 - 77, December.

Guimaraes, Bernardo, and Kevin D. Sheedy. *2011.* "Sales and Monetary Policy." American Economic Review*, 101 (2): 844-76.*

Kaplan Greg & Guido Menzio, 2015. "**The Morphology Of Price Dispersion**," International Economic Review, Department of Economics, University of Pennsylvania and Osaka University Institute of Social and Economic Research Association, vol. 56, pages 1165-1206, November.

Kehoe Patrick J. & Virgiliu Midrigan, 2015. "**Prices are Sticky After All"**, Journal of Monetary Economics, Volume 75, October 2015, pages 35-53.

Lach Saul., "Existence and Persistence of Price Dispersion: An Empirical Analysis" The Review of Economics and Statistics, August 2002, 84(3): 433–444.

Nakamura, Emi & Jón Steinsson, 2008. "**Five Facts about Prices: A Reevaluation of Menu Cost Models**," The Quarterly Journal of Economics, MIT Press, vol. 123(4), pages 1415-1464, November.

Prescott, Edward. C., "Efficiency of the Natural Rate" Journal of Political Economy, 83 (Dec. 1975): 1229-1236.

Salop S., and J. Stiglitz, "Bargains and Ripoffs: A Model of Monopolistically Competitive Price Dispersion," Rev. Econ. Stud., Oct. 1977, 44, 493-510.

Sheremirov, Viacheslav, In Press. "**Price dispersion and inflation: new facts and theoretical implications**," Journal of Monetary Economics.

Shilony, Yuval., "Mixed Pricing in Oligopoly," J. Econ. Theory, Apr. 1977, 14, 373-88.

Sorensen, Alan T. "Equilibrium Price Dispersion in Retail Markets for Prescription Drugs" Journal of Political Economy, Vol. 108, No. 4 (August 2000), pp. 833-850.

Varian, H., "A Model of Sales," American Economic Review 70:4 (1980), 651–659.

Wu. S. Y. "An Essay on Monopoly Power and Stable Price Policy". The American Economic Review, Vol. 69, No. 1 (Mar., 1979), pp. 60-72.