Temporary Sales in Response to Demand Shocks

Benjamin Eden[*], Maya Eden[†], Oscar O'Flaherty[‡]

August 4, 2021

Abstract

Temporary sales, defined as large price drops that quickly rebound, are a puzzling phenomenon from a macroeconomic perspective. We argue that temporary sales play an important role in the response of prices to demand shocks. Using data from supermarkets, we find that the average price of a product decreases by almost 1.5% following a negative shock to demand. The standard practice of removing sales from the price distribution overestimates the degree of price rigidity by a factor of 2. We reconcile our empirical findings using a model in which sales occur in response to the accumulation of unwanted inventories.

Keywords: Temporary Sales, Demand Shocks, Pricing Decisions
JEL Classifications: E3, L11

[*]Department of Economics, Vanderbilt University. Email: ben.eden@vanderbilt.edu
[†]Department of Economics, Brandeis University. Email: meden@brandeis.edu
[‡]Department of Economics, Vanderbilt University. Email: oscar.oflaherty@vanderbilt.edu

1 Introduction

Temporary sales – defined as large price drops which quickly rebound – are commonly viewed as a tool for price discrimination that cancels out at the aggregate and contains no information on macroeconomic conditions (Chevalier and Kashyap, 2019; Guimaraes and Sheedy, 2011; Salop and Stiglitz, 1977; Shilony, 1977; Varian, 1980). In light of this, it has become standard practice in macroeconomics to ignore sales and focus on the behavior of regular/non-sale prices (Nakamura and Steinsson, 2008; Eichenbaum et al., 2011; Kehoe and Midrigan, 2015). Here, we examine this common view and argue that temporary sales play an important role in the reaction of stores to demand shocks.

We start by asking whether it is indeed the case that sales cancel out in the aggregate. We find that sales do not cancel out. On average, a product is not on sale in any store in over 40% of weeks. We also find that the frequency of sales increases following negative demand shocks. We then illustrate the importance of temporary sales by comparing the average price reaction to a demand shock with the reaction of the average non-sale price. The average price decreases by almost 1.4% following a negative demand shock. Ignoring sales and analyzing the average non-sale price mitigates this response to 0.7%. Thus, sales account for about half of the average price decline following demand shocks.

We interpret these findings in the context of a model in which temporary sales are reactions to unwanted inventories. In the model, inventories are accumulated after negative demand shocks. 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 (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

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

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". This mechanism is motivated by Aguirregabiria (1999) who found a significant and robust effect of inventories at the beginning of the month on the current price using a unique data set from a chain of supermarket stores in Spain.

In practice, temporary sale decisions are often negotiated between the chain's headquarters and its suppliers. Aguirregabiria (1999) describes that 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 are complex contingent contracts that are determined substantially in advance. However, 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. The flexibility in the timing of sales–as evidenced in our empirical framework–may also reflect the need to respond to inventories that were accumulated as a result of demand shocks.

Chevalier and Kashyap (2019) raise the issue of close substitutes when aggregating prices. We focus on the average posted price at the product level rather than an aggregate price index (Chevalier and Kashyap, 2019) or the average price paid by the consumer (Coibion et al., 2015; Glandon, 2018). As our focus is thus on the observed behavior of the store, this allows us to avoid the concerns of close substitutes in measuring cost-of-living or consumer welfare.

Kryvtsov and Vincent (2021) find a robust relationship between the frequency of sales and the rate of unemployment. We use weekly data to study the question of price rigidity and focus on shocks to the demand for a narrowly defined product. We do not attempt to study aggregate shocks, like shocks to the money supply, which affect the demand for all goods. Nevertheless, we think that our study is relevant for the study of aggregate shocks. For example, Bental and Eden (1996) and Eden (1994) develop monetary versions of the UST model. These monetary models illustrate that 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. Given this, it is likely that shocks which affect the demand for all goods and stores will work in a similar manner as a good-specific demand shock that affects all stores.

The rest of the paper proceeds as follows. Section 2 describes the data and summary statistics of temporary sales. Section 3 presents several stylized facts about store and product-level decision making of temporary sales. Section 4 analyzes the impact of temporary sales on the average posted price in response to demand shocks. Section 5 presents the Uncertain and Sequential Trade (UST) model in light of the empirical results. Section 6 concludes.

2 Data

Our primary analysis covers grocery stores across the US from 2004-2005 using the Information Resources Inc. (IRI) retail scanner dataset.[2] Observations are recorded weekly at the store-product level. Products are defined by their Universal Product Code (UPC) and belong to specific product categories (e.g. Beer, Hot Dogs). Stores in our sample are located in different markets across the US. A market is sometimes classified as a city (Chicago, Los Angeles) or a state (Mississippi). For each store-product-week observation, we see the to-

[2] A complete description of the dataset can be found in Bronnenberg et al. (2008).

tal revenue (Rev_{ijt}) and total quantity sold (Q_{ijt}). We then compute the average price of product i in store j for the week t as the total revenue divided by quantity, $P_{ijt} = \frac{Rev_{ijt}}{Q_{ijt}}$.

2.1 Filtering Process

In order to analyze the effect of temporary sales, we apply the following filter in a sequential manner by market area:

1. We drop all UPC-Store cells that do not have strictly positive quantities in all weeks.
2. We drop all UPCs that are sold by fewer than 11 stores.
3. We drop all categories with less than 10 UPCs.

The first exclusion is applied because we cannot distinguish between product stockouts (product is not carried in the store) and periods of low demand (product is carried by the store but not purchased). These two scenarios carry different supply and demand implications which may obscure our analysis. This step is also required for identifying temporary sale prices. The second exclusion is aimed at reliable measures of the average cross-sectional price distribution. This requirement also leads to a sample of fairly popular brands.[3] The third economizes on the number of category dummies. After applying our filter, we obtain a semi-balanced panel in which the number and identities of stores may vary across UPCs, but does not vary over weeks for a given UPC.

Panel A of Table 1 provides summary statistics for our filtering process. The original sample contains over 1,500 unique products across 56,000 stores for a total of almost 400 million store-product-week observations. Observations span 31 categories and 50 markets across the US in the full sample. Our final sample reduces the total number of observations to under 10 million. The restriction that products are continuously sold accounts for over 90% of this reduction. The filtering process reduced the total number of stores to 546 located

[3]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 of IRI Data

Panel A: Filtering Process

	Stores	Products	Categories	Markets	Observations
Pre-filter	1,589	56,342	31	50	395,756,478
Product Continuously Sold	693	12,114	31	50	35,718,904
Product $\geq$ 11 Stores	546	1,808	25	26	11,440,832
Category $\geq$ 10 UPCs	546	1,686	22	26	9,828,520

Panel B: Final Sample Composition

Market Name	Percent of Sample	Product Category	Percent of Sample
New York, NY	16.7%	Carbonated Beverages	20.6%
Los Angeles, CA	12.0%	Yogurt	15.8%
New England	8.4%	Salted Snacks	11.5%
Seattle/Tacoma, WA	6.9%	Cold Cereal	10.6%
Dallas, TX	5.3%	Milk	10.4%
St. Louis, MO	5.2%	Beer	6.9%
Chicago, IL	5.0%	Soup	6.8%
Philadelphia, PA	4.6%	Margarine/Butter	5.6%
San Francisco, CA	4.4%	Mayonnaise	1.9%
Raleigh/Durham, NC	3.6%	Peanut Butter	1.5%
Portland, OR	3.2%	Mustard/Ketchup	1.4%
Phoenix, AZ	3.2%	Toilet Tissue	1.4%
Washingtion D.C.	2.6%	Frozen Dinner Entrees	1.1%
Houston, TX	2.5%	Spaghetti Sauce	1.1%
Richmond/Norfolk, VA	2.4%	Hot Dogs	0.8%
Roanoke, VA	2.3%	Paper Towels	0.5%
Kansas City	2.0%	Coffee	0.5%
San Diego, CA	1.8%	Cigarettes	0.4%
Knoxville, TN	1.6%	Frozen Pizza	0.4%
Charlotte, NC	1.6%	Laundry Detergent	0.3%
Buffalo/Rochester, NY	1.3%	Household Cleaner	0.2%
Harrisburg/Scranton, PA	1.2%	Facial Tissues	0.2%
Salt Lake City, UT	0.8%		
Boston, MA	0.6%		
Syracuse, NY	0.5%		
Sacramento, CA	0.2%		

Note: This table presents summary statistics for our data filtering process and the final sample composition. Observations are recorded at the store-product-week level. We require that (1) products are continuously sold from 2004-2005, (2) Products are sold by at least 11 stores in a given market, and (3) Categories contain at least 10 products. Totals may not sum to 100% due to rounding error.

across 26 markets. The final sample spans 1,686 unique products which belong to 22 product categories.

Panel B of Table 1 describes the composition of our final sample by market area and product category. We see that New York, NY contains the most observations and comprises 16.7% of our panel. Los Angeles, CA is the only other market to exceed 10% of our sample. Carbonated beverages is our largest product category at 20.6%. Four other categories exceed 10% of observations–Yogurt, Salted Snacks, Cold Cereal, and Milk.

2.2 Temporary Sales

The IRI provides an indicator for sale prices based on a proprietary algorithm. The IRI definition of sales is rather obscure and may include promotion activities in addition to the behavior of the price. We use another definition that focuses on the behavior of prices as our baseline definition and combine it with the IRI definition as a robustness check. 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. This definition provides the extra benefit that it can be used for comparability across datasets such as Nielson's Retail Scanner, and it is similar to other definitions of temporary sales (Coibion et al., 2015; Nakamura and Steinsson, 2008). As a robustness check, we add the requirement that the sales prices also satisfy the IRI definition.

Panel A of Table 2 compares the three measures of temporary sales. We define the sales frequency as the fraction of total prices that are a temporary sale. At the store-product level, this is denoted as $sale\_freq_{ij} = \frac{1}{T}\sum_{t=1}^{T} I(sale_{ijt} = 1)$, where $I(sale_{ijt} = 1) = 1$ if the price (of product i in store j at week t) is a sale price and zero otherwise. We see that the sales frequency for the IRI definition is greater than our definition of a 10% temporary price reduction across the entire distribution of UPC-Store combinations. For example, the average sales frequency across all UPC-Store combinations is 14.4% for our definition of a 10% temporary price reduction. This is about half of the sales frequency of 28% using the IRI definition of a sale. The IRI definition also contains several observations in which a product is always on sale in a given store. It may be the case that the wider definition of sales adopted by the IRI reflects the need to satisfy supplier imposed requirements (Aguirregabiria, 1999;

Table 2: Summary Statistics for Sales

Panel A: Sale Frequency by UPC-Store

	1st Quartile	Median	Mean	3rd Quartile	Maximum
IRI Definition	9.6%	24.0%	28.0%	43.3%	100.0%
10% Temporary Reduction (Baseline)	1.9%	11.5%	14.4%	24.0%	60.6%
Combined Definition (Robustness Checks)	1.0%	10.6%	13.2%	22.1%	58.7%

Panel B: Market and Category Heterogeneity

Market Name	Sales Frequency	Product Category	Sales Frequency
San Francisco, CA	21.6%	Hot Dogs	22.7%
Chicago, IL	20.4%	Salted Snacks	20.0%
San Diego, CA	20.4%	Yogurt	19.1%
Washingtion D.C.	19.0%	Carbonated Beverages	18.7%
Harrisburg/Scranton, PA	18.7%	Cold Cereal	14.1%
Sacramento, CA	18.6%	Coffee	13.6%
Boston, MA	17.1%	Margarine/Butter	12.5%
Los Angeles, CA	16.8%	Frozen Dinner Entrees	12.3%
Seattle/Tacoma, WA	16.8%	Frozen Pizza	11.5%
Portland, OR	16.6%	Facial Tissues	10.7%
Phoenix, AZ	16.5%	Peanut Butter	10.5%
Philadelphia, PA	15.2%	Toilet Tissue	10.4%
St. Louis, MO	14.8%	Spaghetti Sauce	10.3%
Syracuse, NY	13.9%	Household Cleaner	9.6%
Salt Lake City, UT	13.8%	Beer	8.6%
Charlotte, NC	13.8%	Soup	8.5%
Buffalo/Rochester, NY	13.4%	Mustard/Ketchup	7.0%
New York, NY	13.2%	Mayonnaise	6.6%
Dallas, TX	12.9%	Milk	6.4%
Roanoke, VA	12.8%	Paper Towels	6.1%
Richmond/Norfolk, VA	12.5%	Laundry Detergent	5.3%
Houston, TX	11.9%	Cigarettes	0.2%
Raleigh/Durham, NC	10.2%		
Knoxville, TN	9.4%		
Kansas City	8.2%		
New England	5.0%		

Note: This table presents summary statistics for temporary sales. Observations are recorded at the store-product-week level. We define a sale as a temporary reduction in price of at least 10% which returns to the pre-sale price or greater within four weeks. Sales frequency is then defined as the fraction of total prices which are temporary sales. Totals in Panel B may not sum to 100% due to rounding error.

Anderson et al., 2017). We are primarily interested in the behavior of store-level pricing decisions. For this reason, we adopt our definition based on the behavior of prices for our baseline results. We use the combined definition that a sale requires a 10% temporary price reduction and the IRI flag for robustness checks throughout the paper. This added restriction does not significantly affect the ditribution of sales frequency. Overall, the correlation

Figure 1: Variation in Sales Frequency

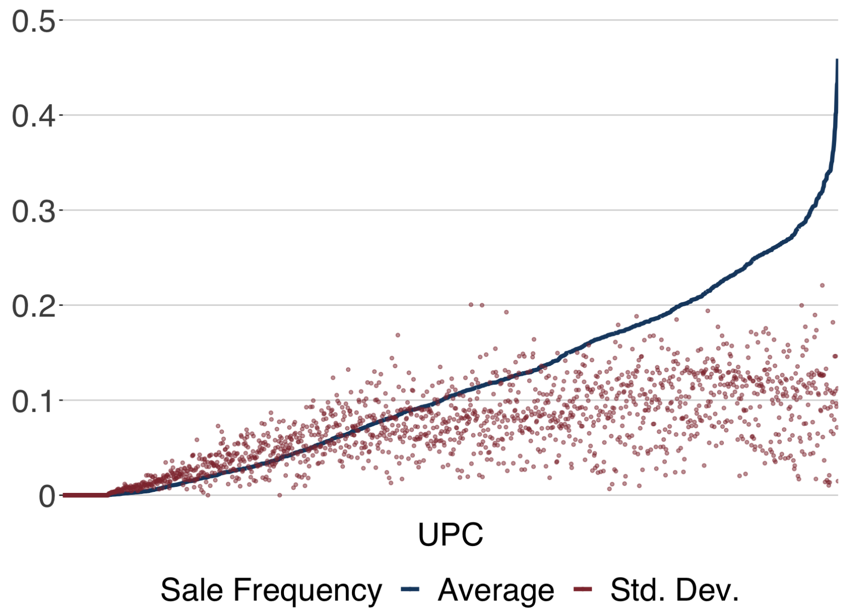

Note: This figure plots the average sale frequency over stores for a given product in blue. The red dots represent the standard deviation of the sales frequency over stores for a given product.

between our baseline and combined definition is 0.99.

Panel B of Table 2 compares the overall sales frequency by market area and product category. We see that San Francisco has the highest sales frequency of 21.6%. The New England market has the lowest sales frequency at 5.0%. Sales frequency also varies across product categories. Cigarettes are the least likely products to have a sale with a frequency of 0.2%. Hot dogs have the highest sales frequency of 22.7%.

Figure 1 provides a more in depth view of the variation in sales frequency across stores and products. The average sales frequency over stores for a given product is plotted in blue, and the standard deviation in red. Similar to product categories, we see that the sales frequency varies widely across products with a range from 0% to almost 50%. The correlation between the mean and the standard deviation is 0.66. Thus, UPCs with higher frequency of sales tend to have more variation across stores. The large standard deviation measures suggest that store managers have a significant role in choosing temporary sales.

3 Stylized Facts

Here we advance the hypothesis that stores set sales in response to cost or demand shocks. An alternative hypothesis assumes that stores use a mixed strategy to determine sales (Sheremirov, 2020; Guimaraes and Sheedy, 2011). The first subsection provides evidence against the alternative hypothesis. The second subsection analyzes the relationship between temporary sales and demand uncertainty, and the third subsection analyzes the dynamic relationship between sales and demand shocks.

3.1 Do stores use a mixed strategy?

The mixed strategy hypothesis is that the price distribution is stable across time, but that individual stores change their place within the distribution. This hypothesis predicts that the fraction of weeks in which a product was not on sale in any of the stores which carried it is small. To test this prediction, we define the fraction of weeks with no sales in any of the stores as follows:

$$num\_sale_{it} = \sum_{j=1}^{J_i} I(sale_{ijt} = 1) \tag{1}$$

$$no\_sale_i = \frac{1}{T} \sum_{t=1}^{T} I(num\_sale_{it} = 0) \tag{2}$$

where J_i represents the total number of stores that sell product i. Thus, no_sale_i is the fraction of weeks in which product i was not on sale in any of the stores.

Panel A of Table 3 presents statistics for these variables. On average, a product is not on sale in any of the stores which sold it for 46% of weeks. This equates to about 48 weeks in our sample. If stores use a mixed strategy with the same sale frequency, the implied probability of no_sale_i can be written as $(1 - sale\_freq_i)^{J_i}$. In our sample, 56 stores carry a product on average. Using the average sale frequency in the data, the implied mixed strategy estimate of no_sale_i is then $(1-.144)^{56} = 0.016\%$. This differs significantly from the percentage found

Table 3: Behavior of Temporary Sales

Panel A: Sale Behavior by UPC

	1st Quartile	Median	Mean	3rd Quartile	Maximum
Sale in No Store (no_sale_i)	12.50%	45.19%	46.43%	75.96%	100%

Panel B: Heterogeneity by Frequency of Sale

	0-5%	5-10%	10-15%	15-20%	>20%
Percent of Sample	32.1%	16.8%	16.3%	14.4%	20.4%
Sale Frequency ($\overline{sale\_freq}$)	1.7%	7.6%	12.3%	17.4%	25.8%
Number of Stores ($\bar{J}$)	34.9	50.7	56.9	71.7	82.1
Sale in No Store ($\overline{no\_sale}$)	83.6%	47.1%	31.5%	22.8%	15.8%
Mixed Strategy $(1 - \overline{sale\_freq})^{\bar{J}}$	54.7%	1.8%	0.1%	0.0001%	2.2e-9%

Note: This table presents statistics about the fraction of weeks in which a product is not on sale in any of the stores (no_sale_i). Panel A presents summary statistics for this variable. Panel B presents similar statistics grouping products by their sale frequency. Means are denoted by bars.

in the data. This result may be driven by heterogeneity in the sales frequency over products. For example, the phenomenon of no sale in any store may occur primarily in products with low sale probabilities.

Panel B of Table 3 examines the possibility that our results are driven by this heterogeneity. We divide UPCs into five bins according to their frequency of sale: [0, 5%); [5, 10%); [10, 15%); [15, 20%); [20, 100%]. We present results using the mean of variables in their respective bin which are denoted by bars. We see that a mixed strategy may provide a close approximation for products with a low probability of sale. For products in the [0, 5%) bin, a mixed strategy implies that no stores list a sale price in 54.7% of weeks. The observed frequency in the data is 83.6%. However, these products only comprise 32% of our sample. A mixed strategy provides a poor approximation for all other bins. Products experience no sale prices in all stores in 47% of weeks on average in the [5, 10%) bin. The estimated probability is 1.8% for a mixed strategy. The mixed strategy probability converges towards zero in the higher frequency bins. Although the probability of no sale in the data does decrease with the frequency of sale, this probability remains significantly different than zero. In the > 20% bin, the fraction of weeks with no sales in any of the stores is 15.8%.

3.2 Temporary sales and demand uncertainty

The main point of the paper is that temporary sales are a reaction to demand shocks and not merely a discrimination device. If temporary sales are a reaction to demand shocks, then stores and products that face more demand uncertainty should have more sales. We start by establishing these correlations between demand uncertainty and the frequency of temporary sales.[4]

3.2.1 Empirical Strategy

Variable Definitions We use the standard deviation of log units over time as a proxy for demand uncertainty for each product-store combination, SDU_{ij}. However, sale prices can both respond to and cause changes in demand uncertainty which may bias our estimates. To help mitigate the issue of reverse causality, we calculate SDU using only observations in which a non-sale price is listed. Representing the quantity sold of product i at store j in week t as Q_{ijt}, this can be calculated as follows:

$$q_{ijt} = \log(Q_{ijt}) \tag{3}$$

$$q_{ij} = \frac{1}{T_{ij}^{reg}} \sum_{t=1}^{T} q_{ijt} I(sale_{ijt} = 0) \tag{4}$$

$$SDU_{ij} = \sqrt{\frac{\sum_{t=1}^{T} \left((q_{ijt} - q_{ij}) I(sale_{ijt} = 0)\right)^2}{T_{ij}^{reg} - 1}} \tag{5}$$

where T_{ij}^{reg} represents the number of periods in which product i in store j is sold at a regular (non-sale) price. $I(sale_{ijt} = 0)$ is an indicator function equal to one if the price of product i

[4] Using only the Chicago market, Eden (2018) provides evidence about the correlation between demand uncertainty and the frequency of temporary sales only at the product level.

in store j at week t is a regular (non-sale) price and zero otherwise.

We follow a similar procedure for prices to define the average log price, p_{ij}, and the standard deviation of log price, SDP_{ij}. We then define store-level observations as the average over UPCs:

$$sale\_freq_j = \frac{1}{I_j}\sum_{i=1}^{I_j} sale\_freq_{ij} \tag{6}$$

$$SDU_j = \frac{1}{I_j}\sum_{i=1}^{I_j} SDU_{ij} \tag{7}$$

where I_j is the number of products sold by store j in our sample. Using the subsample of regular prices, we calculate the average standard deviation of log prices (SDP_j), average log units (q_j), and the average log price (p_j)) in a similar manner.

Specification After computing our proxy for demand uncertainty, we estimate:

$$sale\_freq_k = \alpha + \beta SDU_k + \Gamma X_k + \epsilon_k \tag{8}$$

where $X_k = (SDP_k, q_k, p_k)$ is a vector of controls. We estimate Equation (8) for the store level ($k = j$) and the product level ($k = i$).[5] Our coefficient of interest, β, represents the effect of a 1 percentage point increase in the standard deviation of units on the sales frequency.

3.2.2 Results

Figure 2 presents a scatterplot of temporary sales frequency and the standard deviation of log units at the store level. The blue line plots an estimate of Equation (8) without other covariates. There is a clear positive relationship between the sale frequency and standard deviation of units. Specification (1) of Table 4 shows that a 1 percentage point increase in the standard deviation of units is associated with a 0.3 percentage point increase in the sale

[5]Product-level variables are computed in a similar manner as the store-level variables.

Figure 2: Demand Uncertainty and Temporary Sales

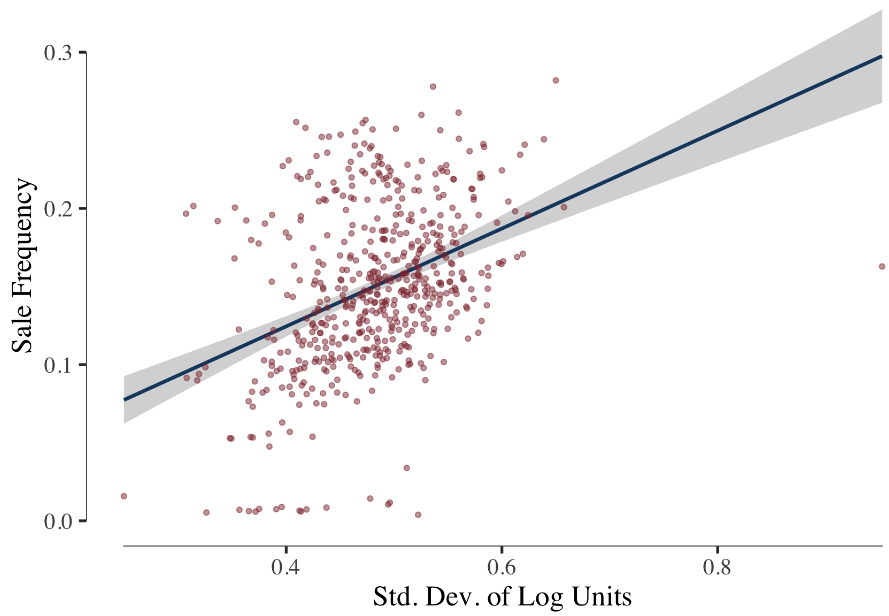

Note: This figure presents a scatterplot of temporary sales frequency (y-axis) and the standard deviation of log units (x-axis) at the store level. A regression fit on the data points is plotted in blue.

frequency for stores on average. The estimate of β reduces to 0.07 after controlling for the average of log units sold, the average of log price and the standard deviation of log price. This estimate remains statistically significant at the 95% level. Much of this effect appears to be captured by the standard deviation of prices. This suggests a correlation between SDU and SDP. Thus, stores may react to demand shocks by changing regular prices and by temporary sales.

Columns (3) and (4) present the estimates of Equation (8) at the product level. The product-level results also suggest a positive relationship between temporary sales and demand uncertainty. The coefficient for SDU is larger at the product level when compared to the store-level estimates. The product-level results suggest that stores issue more sales for products that experience more demand uncertainty.

One concern is that products may be on sale more often for reasons that are not related to demand uncertainty. Another concern is that differences in the characteristics and

Table 4: Demand Uncertainty and Temporary Sales

	Dependent variable:					
	Store-level ($sale\_freq_j$)		Product level ($sale\_freq_i$)		Store-product level ($sale\_freq_{ij}$)	
	(1)	(2)	(3)	(4)	(5)	(6)
SDU	0.313***	0.070**	0.422***	0.245***	0.138***	0.081***
	(0.032)	(0.033)	(0.017)	(0.019)	(0.005)	(0.006)
SDP		1.952***		1.162***		0.468***
		(0.097)		(0.052)		(0.020)
Ave. ln(Units)		−0.001		0.010***		0.012***
		(0.005)		(0.003)		(0.001)
Ave. ln(Price)		−0.006		−0.005**		0.180***
		(0.010)		(0.002)		(0.007)
UPC Fixed Effects					X	X
Store Fixed Effects					X	X
Observations	546	546	1,686	1,686	94,505	94,505
Adjusted R^2	0.148	0.525	0.271	0.497	0.619	0.650

Note: This table presents regression estimates of Equations (8)-(9). SDU and SDP represent the standard deviations of log units and log prices, respectively. Standard errors are clustered at the store level for the product-store regressions in columns (5)-(6).
*p<0.1; **p<0.05; ***p<0.01

demographics of stores may lead to this result. To account for these possibilities, we estimate

$$sale\_freq_{ij} = \alpha_i + \gamma_j + \beta SDU_{ij} + \Gamma X_{ij} + \epsilon_{ij} \tag{9}$$

where α_i are product fixed effects and γ_j are store fixed effects. The product fixed effects control for differences in the means across products. The inclusion of store fixed effects allows for different pricing strategies across stores. Columns (5) and (6) of Table 4 present the results of Equation (9). Standard errors are clustered at the store level. The estimate of β is now 0.138 without controls and 0.081 with controls, and both remain statistically significant.

To address potential endogeneity problems, we re-estimate Equations (8) and (9) using the 2005 sample to measure the dependent variable and the 2004 sample to measure the independent variables. These results are presented in Appendix Table A.1. The results are very similar to those in Table 4 suggesting that endogeneity is not a problem. Overall, our results suggest that demand uncertainty is positively correlated with temporary sales.

3.3 Temporary sales and demand shocks

The previous section provided evidence that stores with more demand uncertainty also issue more temporary sales. This suggests that stores use sales to react to demand shocks. We examine this hypothesis by exploiting the dynamic nature of the data to estimate several panel vector autoregressions (PVAR).

3.3.1 Empirical Strategy

Specification We estimate a panel vector auto-regression of order two:

$$y_{ijt} = \alpha_{ij} + \gamma_{it} + A_1 y_{ijt-1} + A_2 y_{ijt-2} + \epsilon_{ijt} \tag{10}$$

where y_{ijt} is a vector consisting of log price (p_{ijt}), a sale indicator ($sale_{ijt}$), and log quantity (q_{ijt}) in that order for a store-product-week combination. The inclusion of product-store fixed effects, α_{ij}, controls for differences in the average price, sales frequency, and units sold across stores for a given product. The product-time fixed effects, γ_{it}, control for product-specific seasonality and shocks.

Interpretation We use the estimated coefficients of Equation (10) to compute orthogonalized impulse response functions (IRFs) and simulate the effect of a quantity shock on temporary sales. The inclusion of the product-time fixed effects implies that the error term in (10) does not reflect shocks to the aggregated demand for the product and is due to store-specific shocks. Our IRFs therefore describe the response to a store-specific shock (that may vary across products). We will discuss the effects of shocks to the aggregated demand for products in the next section. In this context, the impulse response function k periods after a shock at time t (IRF_{t+k}) represents how much *more* likely a given store is to issue a sale. For

Figure 3: Impulse Response of Temporary Sales to Demand Shocks

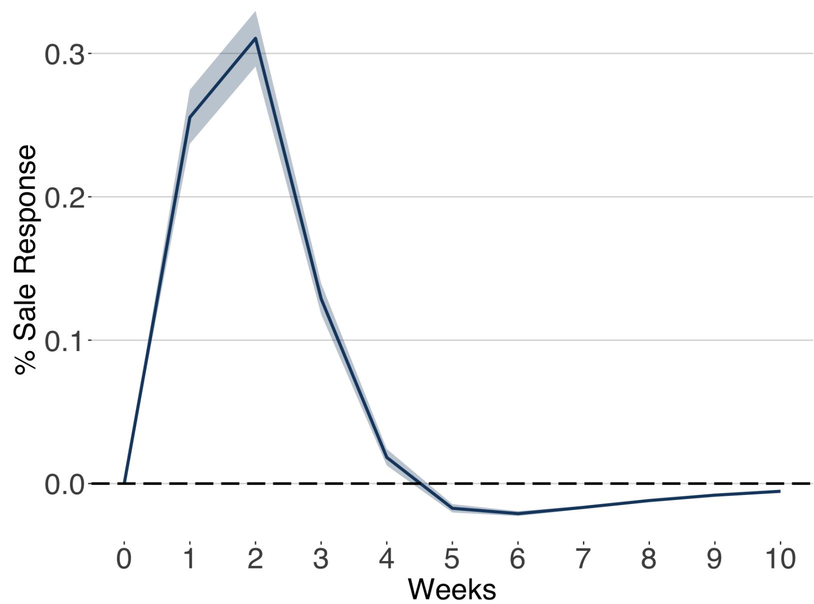

Note: This figure plots the orthogonalized impulse response function of temporary sales in response to a one standard deviation negative quantity shock. Confidence intervals are bootstrapped at the 95% confidence level.

example, a store that experiences a one standard deviation negative demand shock in week t is IRF_{t+k} percentage points more likely to have a sale in period $t + k$.

3.3.2 Results

Figure 3 plots the orthogonalized impulse response function of temporary sales in response to a one standard deviation negative quantity shock. It describes the effect of a store-specific demand shock on the probability that a store will issue a temporary sale. Confidence intervals are bootstrapped at the 95% confidence interval. We see positive effects on the probability of sale in the first three weeks after a demand shock. The impulse response function shows that a store is 0.25 percentage points more likely to place a product on sale one week after a store-specific demand shock. This response peaks at 0.3pp two weeks after the shock, and begins to subside to about 0.1pp at three weeks. After three weeks, the effect converges to zero.

Table 5: Impulse Response of Temporary Sales to Demand Shocks

	Weeks after Shock						Max. Cumulative Effect
	1	2	3	4	5	6	
Baseline							
2 Lags	0.26 [0.24, 0.27]	0.31 [0.29, 0.33]	0.13 [0.12, 0.14]	0.02 [0.01, 0.02]	-0.02 [-0.02, -0.01]	-0.02 [-0.02, -0.02]	0.71 [0.67, 0.76]
4 Lags	0.26 [0.24, 0.29]	0.39 [0.37, 0.41]	0.34 [0.32, 0.36]	0.05 [0.03, 0.07]	-0.02 [-0.03, 0]	-0.03 [-0.03, -0.02]	1.04 [1.00, 1.09]
Market Fixed Effects	0.22 [0.2, 0.24]	0.26 [0.24, 0.27]	0.10 [0.09, 0.11]	0.01 [0.01, 0.02]	-0.02 [-0.02, -0.01]	-0.02 [-0.02, -0.02]	0.59 [0.55, 0.63]
Market FEs & 4 Lags	0.23 [0.21, 0.25]	0.31 [0.29, 0.33]	0.28 [0.27, 0.30]	0.04 [0.02, 0.06]	-0.01 [-0.02, -0.01]	-0.03 [-0.03, -0.02]	0.87 [0.83, 0.92]
10% Reduction & IRI Flag							
2 Lags	0.35 [0.33, 0.37]	0.33 [0.31, 0.35]	0.12 [0.12, 0.13]	0.02 [0.01, 0.02]	-0.02 [-0.02, -0.01]	-0.02 [-0.02, -0.02]	0.82 [0.78, 0.86]
4 Lags	0.36 [0.34, 0.38]	0.40 [0.38, 0.42]	0.32 [0.3, 0.34]	0.02 [0, 0.04]	-0.02 [-0.03, -0.01]	-0.03 [-0.03, -0.02]	1.10 [1.05, 1.14]
Market Fixed Effects	0.30 [0.29, 0.32]	0.27 [0.25, 0.29]	0.10 [0.09, 0.1]	0.01 [0.01, 0.01]	-0.01 [-0.02, -0.01]	-0.02 [-0.02, -0.01]	0.68 [0.65, 0.72]
Market FEs & 4 Lags	0.32 [0.30, 0.34]	0.32 [0.31, 0.34]	0.27 [0.25, 0.28]	0.02 [0, 0.03]	-0.02 [-0.03, -0.01]	-0.03 [-0.03, -0.02]	0.92 [0.88, 0.96]

Note: This table plots the orthogonalized impulse response functions of temporary sales in response to a one standard deviation negative quantity shock. Rows represent the specification, and columns specify the period after the shock. The last column presents the maximum cumulative effect for each specification. Forecast confidence intervals bootstrapped at the 95% confidence level are in brackets. The first row plots our baseline results. The second row plots the estimates using a VAR(4). The third and fourth rows are similar to first and second rows, but replace Product-Week fixed effects in Equation (10) with Product-Market-Week fixed effects. The bottom panels adds the restriction that a sale is also indicated by the IRI flag.

We provide several robustness checks for these results. We estimate Equation (10) with four lags rather than two. We change γ_{it} in Equation (10) to product-market-time fixed effects (γ_{imt}). Lastly, we provide estimates for the alternative definition of a sale as a temporary price reduction that coincides with the IRI flag. In Table 5, rows represent the specification, and columns the period after shock. The last column presents the maximum cumulative effect for each specification.[6] Upper and lower forecast confidence bands bootstrapped at the 95% confidence level are presented in brackets.

[6] This is typically the fourth week after a demand shock, but may vary across specifications. The cumulative effect can be found be summing the individual response estimates.

We find similar estimates in the first two weeks after a demand shock using two and four lags when comparing the first and second rows of Table 5. However, at three weeks, we estimate a 0.29 percentage point increased probability of a sale using four lags compared to the baseline result of 0.14pp. As in the baseline results, impulse responses converge to zero after three weeks. This results in a maximum cumulative effect of 1.04pp with four lags compared to the baseline of 0.71pp. Including market fixed effects does not change the overall dynamics or relative effects in the impulse response functions. However, the magnitude of the effects are slightly mitigated. This results in a maximum cumulative effect of 0.59pp and 0.87pp for the estimates with two and four lag lengths, respectively.

The added restriction that a sale must coincide with the IRI flag increases the probability of a sale price in the first week after a demand shock across all specifications. Without the IRI flag, the peak response always occurred in the second week. With the IRI flag, the peak response varies across specifications between the first and second week. The increased response in the first week also leads to larger cumulative effects compared to our baseline specifications.

4 Shocks to the Aggregated Demand for a Product

The previous section showed that sales respond to store-level demand shocks. We now show that this is also the case when the shock occurs to the aggregated demand for a product. We then assess the role of temporary sales in the response of the average price (across stores) to a demand shock.

Specification We estimate a panel vector autoregression of order two:

$$y_{it} = \alpha_{iw} + A_1 y_{it-1} + A_2 y_{it-2} + \epsilon_{it} \tag{11}$$

where t indexes times and w represents the week of the year that period t belongs to. In our sample we have 104 weeks, so $t \in \{1, ..., 104\}$ and $w \in \{1, ..., 52\}$ since there are 52 weeks in

Figure 4: Impulse Response of Aggregate Sales Frequency to Demand Shocks

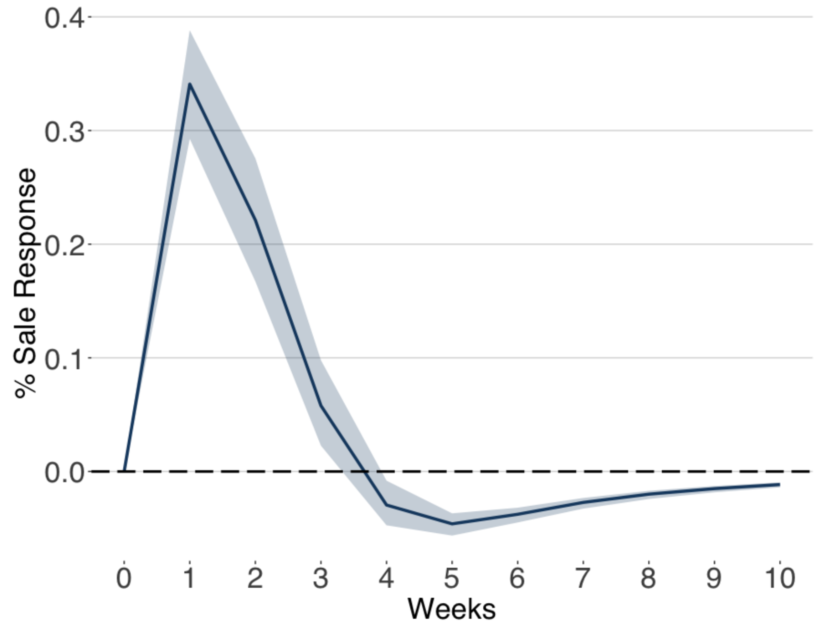

Note: This figure plots the orthogonalized impulse response function of temporary sales in response to a one standard deviation negative quantity shock from Equation 11. Confidence intervals are bootstrapped at the 95% confidence level.

a year. Thus, $w = 1$ for weeks with time index $t = 1$ and $t = 53$, $w = 2$ for $t = 2$ and $t = 54$ and so on. The product/week-of-year fixed effects, α_{iw}, control for seasonality. The vector y_{it} consists of the log average price over stores (p_{it}), the average sales frequency over stores ($sale_{it}$), and log of the total quantity sold over stores (q_{it}) in that order.[7]

4.1 Sales Response

Figure 4 presents the response of the aggregate sales frequency on average in response to a one standard deviation shock to the total units sold over all stores. The aggregate sales frequency increases by 0.3 percentage points one week after the shock. This effect remains greater than 0 percentage points for three weeks. After this, the response begins to converge to zero. Appendix Table A.2 shows that a shock to the total demand for a product (over stores) has a maximum cumulative effect of 0.62 percentage points on the average sales

[7] To control for trending factors such as inflation, a product-specific linear trend is removed from p_{it} and q_{it} before estimating Equation (11).

Figure 5: Impulse Response of the Average Price to Demand Shocks

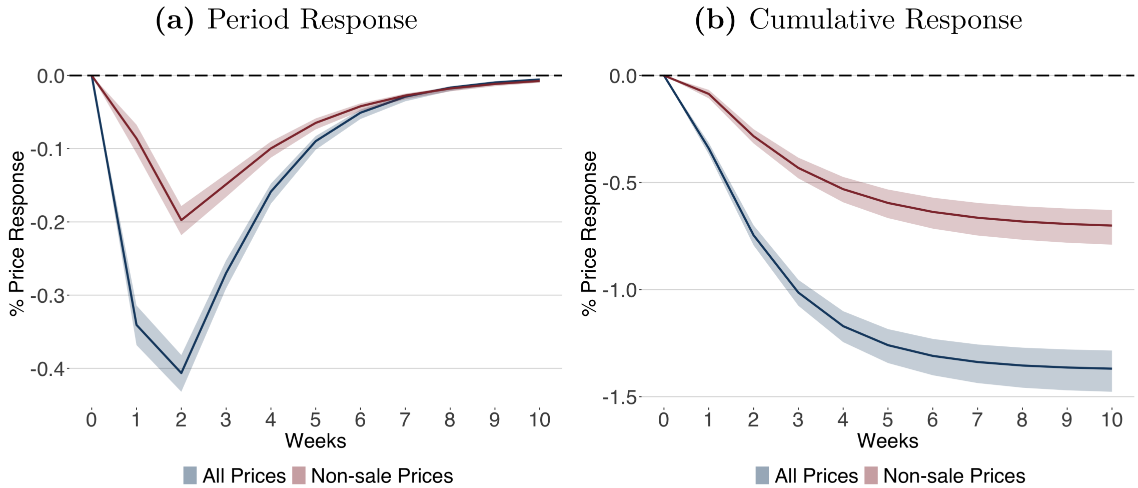

Note: This figure plots the orthogonalized impulse response function of all prices (blue) and non-sale prices (red) in response to a one standard deviation negative quantity shock from Equation 11. Confidence intervals are bootstrapped at the 95% confidence level.

probability (over stores). Robustness checks for alternative lag lengths and the definition of sales are also included in this table.

4.2 Contribution to Price Dynamics

We now explore how this increase in temporary sales affects the response of the average price (across stores) to demand shocks. To do this, we re-estimate Equation (11) using the log of the average non-sale/regular price over stores. The resulting impulse response function is then compared to the IRF when using all prices to compute the average price.[8]

Figure 5 plots the response of all prices (blue) and non-sale prices (red) to a one standard deviation negative shock to the quantity sold over all stores. Confidence intervals are bootstrapped at the 95% confidence level. As expected, prices decrease after a negative demand shock. However, this response is mitigated for non-sale prices. Non-sale prices experience a trough decline of 0.20% two weeks after a demand shock. Including sale prices estimates an effect of 0.41%. The response of all prices is larger in magnitude than non-sale prices for the first six weeks following the shock. After six weeks, both impulse response functions

[8]For these regressions, the sales vector is omitted from the specification.

converge to zero.

Panel (b) shows that the cumulative price response for all prices remains greater (in absolute value) than non-sale prices up to ten weeks following the shock. After 10 weeks, the cumulative response of all prices is 1.37% while the response of regular prices is 0.70%. The relative contribution of sales to the cumulative price response after 10 weeks is then $\frac{1.37-0.70}{1.37} = 49\%$. Thus, sales account for almost half of the average price decline for a product following a demand shock.

Table 6 presents several robustness checks for these results. We estimate Equation (11) with four lags rather than two. We provide estimates for the alternative definition of a sale as a temporary price reduction that coincides with the IRI flag. In Table 6, rows represent the specification, and columns the period after shock. The last column presents the maximum cumulative effect for each specification. Upper and lower forecast confidence bands bootstrapped at the 95% confidence level are presented in brackets.

We see that using four lags in the VAR leads to larger cumulative decreases in both the all price and non-sale price response functions. This leads to a cumulative response after ten weeks of -1.88% and -1.04% for all prices and regular prices, respectively. However, this has a small effect on the relative contribution of sales to the total price response which has changed from 48.8% to 44.9%. Similarly, using the alternative definition of sales does not significantly affect the contribution to the cumulative price response.

Overall, the results suggest that temporary sales play a significant role in the response to demand shocks.

5 A Model

To account for our empirical findings, we present an Uncertain and Sequential Trade (UST) model in which sales are determined endogenously in response to demand shocks. The first subsection presents the baseline UST model. The second subsection extends the UST model to allow for storage of goods. The final subsection uses the model to account for the stylized

Table 6: Impulse Response of Average Prices to Demand Shocks

	Weeks after Shock						Cumulative Effect (10 Weeks)	Contribution of Sales
	1	2	3	4	5	6		
Baseline								
Two Lags								
All Prices	-0.34	-0.41	-0.27	-0.16	-0.09	-0.05	-1.37	48.8%
	[-0.37, -0.31]	[-0.43, -0.38]	[-0.29, -0.25]	[-0.18, -0.15]	[-0.1, -0.08]	[-0.06, -0.05]	[-1.48, -1.28]	
Regular Prices	-0.09	-0.20	-0.15	-0.10	-0.06	-0.04	-0.70	
	[-0.11, -0.07]	[-0.22, -0.18]	[-0.17, -0.13]	[-0.11, -0.09]	[-0.07, -0.06]	[-0.05, -0.04]	[-0.79, -0.63]	
Four Lags								
All Prices	-0.30	-0.30	-0.31	-0.24	-0.19	-0.16	-1.88	44.9%
	[-0.33, -0.27]	[-0.32, -0.27]	[-0.34, -0.29]	[-0.26, -0.21]	[-0.21, -0.18]	[-0.17, -0.14]	[-2.02, -1.76]	
Regular Prices	-0.07	-0.15	-0.20	-0.16	-0.11	-0.09	-1.04	
	[-0.09, -0.05]	[-0.18, -0.13]	[-0.22, -0.18]	[-0.18, -0.14]	[-0.13, -0.1]	[-0.11, -0.08]	[-1.16, -0.93]	
10% Reduction & IRI Flag								
Two Lags								
All Prices	-0.34	-0.41	-0.27	-0.16	-0.09	-0.05	-1.37	47.5%
	[-0.37, -0.31]	[-0.43, -0.38]	[-0.29, -0.25]	[-0.18, -0.15]	[-0.1, -0.08]	[-0.06, -0.05]	[-1.48, -1.28]	
Regular Prices	-0.08	-0.20	-0.15	-0.10	-0.07	-0.05	-0.72	
	[-0.1, -0.06]	[-0.22, -0.18]	[-0.17, -0.14]	[-0.12, -0.09]	[-0.08, -0.06]	[-0.05, -0.04]	[-0.81, -0.65]	
4 Lags								
All Prices	-0.30	-0.30	-0.31	-0.24	-0.19	-0.16	-1.88	44.4%
	[-0.33, -0.27]	[-0.32, -0.27]	[-0.34, -0.29]	[-0.26, -0.21]	[-0.21, -0.18]	[-0.17, -0.14]	[-2.02, -1.76]	
Regular Prices	-0.06	-0.16	-0.20	-0.16	-0.12	-0.10	-1.05	
	[-0.08, -0.04]	[-0.18, -0.14]	[-0.22, -0.18]	[-0.18, -0.14]	[-0.13, -0.1]	[-0.11, -0.09]	[-1.15, -0.95]	

Note: This table plots the orthogonalized impulse response functions of the average price for a product in response to a one standard deviation negative quantity shock. Rows represent the specification, and columns specify the period after the shock. The last two columns present the cumulative price response and the contribution of sales to the cumulative response after ten weeks, respectively. Forecast confidence intervals bootstrapped at the 95% confidence level are in brackets. The first row plots our baseline results. The second row plots the estimates using a VAR(4). The third and fourth rows are similar to first and second rows, but add the restriction that a sale is also indicated by the IRI flag.

facts.

5.1 Uncertain and Sequential Trade Model

Uncertain In the UST model, the economy is comprised of many consumers and firms. Each firm makes their production decision at the beginning of the period *before* the number of consumers is known. We assume the number of consumers for good i in period t is given by N_{it} where N_{it} is an iid random variable. For simplicity of the exposition, we describe the scenario in which the number of consumers can take two possible realizations:

$$N_{it} = \begin{cases} x_i & \text{with prob. } \pi_i, \\ x_i + \Delta_i & \text{with prob. } 1 - \pi_i. \end{cases}$$

All consumers have the same demand function. An individual consumer who buys good i in store j at price P_{ijt} demands $D_i(P_{ijt})$ units of the good.

Sequential Consumers arrive sequentially in the model. The first group of x_i consumers arrive early with certainty. The second group of Δ_i consumers arrive later with probability $1 - \pi_i$. It is useful to think of two hypothetical markets. The first market (*Early*) opens with certainty and serves the group of x_i consumers. The second market (*Late*) opens with probability $1 - \pi_i$ and serves the second batch of Δ_i consumers if they arrive. Firms take the price in each of the hypothetical markets as given. Thus, firms can sell good i at the price $P_{i,Early,t}$ with certainty or at the price $P_{i,Late,t}$ if the additional consumers arrive. Firms that sell at the early price supply the market with $Q_{i,Early,t}$ units of good i while firms that sell at the late price supply $Q_{i,Late,t}$ units.

Equilibrium The cost of production is λ_i per unit. Equilibrium is then defined as a vector of prices and quantities $(P_{i,Early,t}, P_{i,Late,t}, Q_{i,Early,t}, Q_{i,Late,t})$ such that the expected profits for each unit is 0:

$$P_{i,Early,t} = (1 - \pi_i) P_{i,Late,t} = \lambda_i \tag{12}$$

And markets that open are cleared:

$$Q_{i,Early,t} = x_i D_i(P_{i,Early,t}) \tag{13}$$

$$Q_{i,Late,t} = \Delta_i D_i(P_{i,Late,t}) \tag{14}$$

Figure 6: Prices and Quantities in the Baseline UST Model

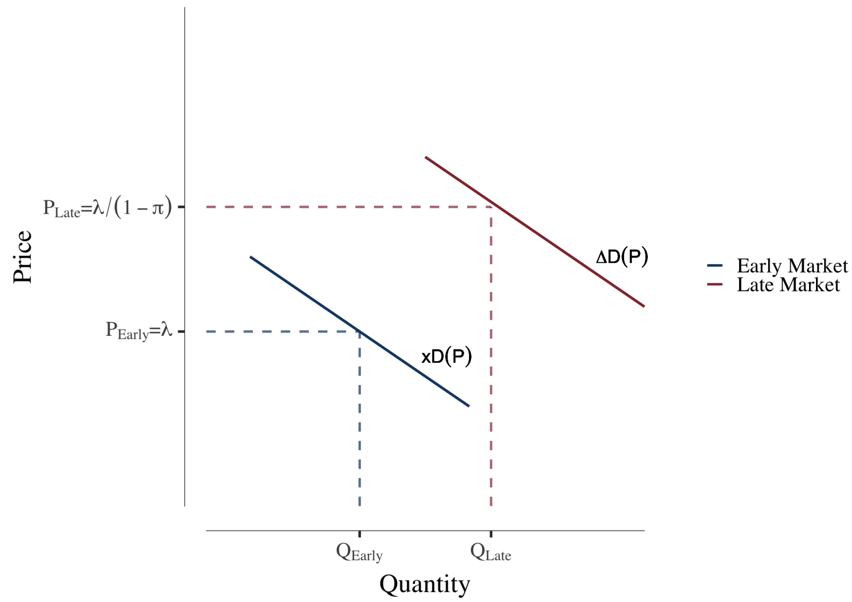

Note: This figure plots a possible equilibrium in the baseline Uncertain and Sequential Trade model. Prices and quantities are denoted by P and Q respectively. Subscripts denote *Late* and *Early* markets.

Thus, in equilibrium, firms are indifferent between the two prices as the expected profits are the same.

Figure 6 plots a possible equilibrium solution for a given good. The product and time subscripts are omitted for simplicity. At the price λ, the total demand of the first group of consumers, $xD(\lambda)$, is equal to the quantity supplied, Q_{Early}. If the second group of consumers arrives to the market, they purchase the good at price $\frac{\lambda}{1-\pi}$. Again, markets clear and quantity demanded is equal to quantity supplied, $\Delta D(\frac{\lambda}{1-\pi}) = Q_{Late}$. Note that, in this simple version of the UST model, prices do not change over time (Equation 12). The quantity sold at the low price also does not change over time. However, the quantity sold at the high price fluctuates between Q_{Late} and 0 depending on the realization of demand.

5.2 Extended Model with Temporary Sales

Bental and Eden (1993) develop a UST model that allows for storage. Including storage allows prices to fluctuate as a result of demand shocks. In the model, a negative demand shock leads to the accumulation of inventories and a reduction in future prices.[9] In what follows, we describe the one-good model under the assumptions of a constant per unit cost (λ) and two states of demand as in the previous subsection. Following Eden (2018) we assume that goods have an expiration date and follow one-hoss shay depreciation which is a natural assumption for our analysis of supermarket goods. For simplicity, 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. After these two periods, the good then has no value.

With the inclusion of storage, the economy at the beginning of period t can now be in one of two states: inventories (Inv) or no inventories ($NoInv$). In state Inv, the demand in the previous period was low ($N_{t-1} = x$) and the additional Δ consumers did not arrive. Thus, inventories are carried from the previous period as a result. In state $NoInv$, demand was high ($N_{t-1} = x + \Delta$) and there are no inventories. Prices are now a function of last period's demand, $P_j(N_{t-1})$. This extends the equilibrium vector of prices in the baseline UST model to $(P_{Early}^{Inv}, P_{Early}^{NoInv}, P_{Late}^{Inv}, P_{Late}^{NoInv})$. Similarly, the equilibrium vector of quantities is now extended to $(Q_{Early}^{Inv}, Q_{Early}^{NoInv}, Q_{Late}^{Inv}, Q_{Late}^{NoInv})$. Inventories at the beginning of the period are given by:

$$I = \begin{cases} \Delta D(P_{Late}) & \text{if } N_{t-1} = x, \\ 0 & \text{otherwise.} \end{cases} \tag{15}$$

These inventories are allocated over the two hypothetical markets:

$$I = q_{Early}^{Stored} + q_{Late}^{Stored} \tag{16}$$

The supply to each of the hypothetical markets is the sum of newly produced and stored

[9] This is consistent with Aguirregabiria (1999) who finds that markups are negatively correlated with inventory levels.

units:

$$Q^{Inv}_{Early} = q^{Stored}_{Early} + q^{New}_{Early} \tag{17}$$

$$Q^{Inv}_{Late} = q^{Stored}_{Late} + q^{New}_{Late} \tag{18}$$

A newly produced unit that was not sold in the current period can be sold in the next period. Since stored units will expire if they are not sold, we assume that stored units are supplied to the *Early* market first. They are supplied to the *Late* market only if the quantity of inventories exceeds the demand in the *Early* market. Thus, the value of an individual stored unit can be given by βP^{Inv}_{Early} where β is a discount factor that captures discounting, storage costs, and depreciation with the restriction that $\beta \in (0, 1)$.

Equilibrium The full equilibrium vector is $(I, P^{Inv}_{Early}, P^{NoInv}_{Early}, P^{Inv}_{Late}, P^{NoInv}_{Late}, Q^{Inv}_{Early}, Q^{NoInv}_{Early}, Q^{Inv}_{Late}, Q^{NoInv}_{Late}, q^{Stored}_{Early}, q^{New}_{Early}, q^{Stored}_{Late}, q^{New}_{Late})$ which must satisfy Equations (15)-(18) in addition to the following conditions:

1. Prices and quantities in the *Late* consumer market do not depend on the state of inventories:

$$P_{Late} = P^{Inv}_{Late} = P^{NoInv}_{Late} \tag{19}$$

$$Q_{Late} = Q^{Inv}_{Late} = Q^{NoInv}_{Late} \tag{20}$$

2. If the level of inventories is positive, stored units are supplied to the *Early* consumer market first:

$$q^{Stored}_{Early} = \min\left\{Q^{Inv}_{Early}, I\right\} \tag{21}$$

$$P^{Inv}_{Early} \geq (1-\pi)P_{Late} \qquad \text{with equality if } q^{Stored}_{Late} > 0 \tag{22}$$

Equation (21) states that stored units are supplied to the *Late* consumer market only if the level of inventories is larger than the demand of the *Early* consumer market. Equation (22) requires that supplying stored goods to the *Early* consumer market is optimal for firms. When stored goods are supplied to both markets ($q_{Late}^{Stored} > 0$), the price in the *Early* market must equal the expected price in the *Late* market. Otherwise, it is optimal to allocate all stored units to the *Early* market.

3. Price in the *Early* consumer market must be less than or equal to λ with equality if new units are supplied:

$$P_{Early}^{Inv} \leq \lambda \qquad \text{with equality if } q_{Early}^{New} > 0 \tag{23}$$

$$P_{Early}^{NoInv} = \lambda \tag{24}$$

4. The expected marginal revenue of supplying a newly produced good to the *Late* consumer market is equal to marginal cost:

$$(1-\pi)P_{Late} + \pi\beta P_{Early}^{Inv} = \lambda \tag{25}$$

To gain the intuition, this was simply $(1-\pi)P_{Late} = \lambda$ in the UST model without storage. With storage, the firm may still sell the good in the next period if the *Late* market does not open in the current period. Thus, expected revenues also include the value of inventories mentioned earlier, $\pi\beta P_{Early}^{Inv}$.

5. Markets clear:

$$Q_{Early}^{Inv} = xD(P_{Early}^{Inv}) \tag{26}$$

$$Q_{Early}^{NoInv} = xD(P_{Early}^{NoInv}) \tag{27}$$

$$Q_{Late} = \Delta D(P_{Late}) \tag{28}$$

Figure 7 plots two possible equilibria. As in the previous subsection, the blue and red

Figure 7: Possible Equilibria in the UST Model with Storage

(a) Low Inventory **(b)** High Inventory

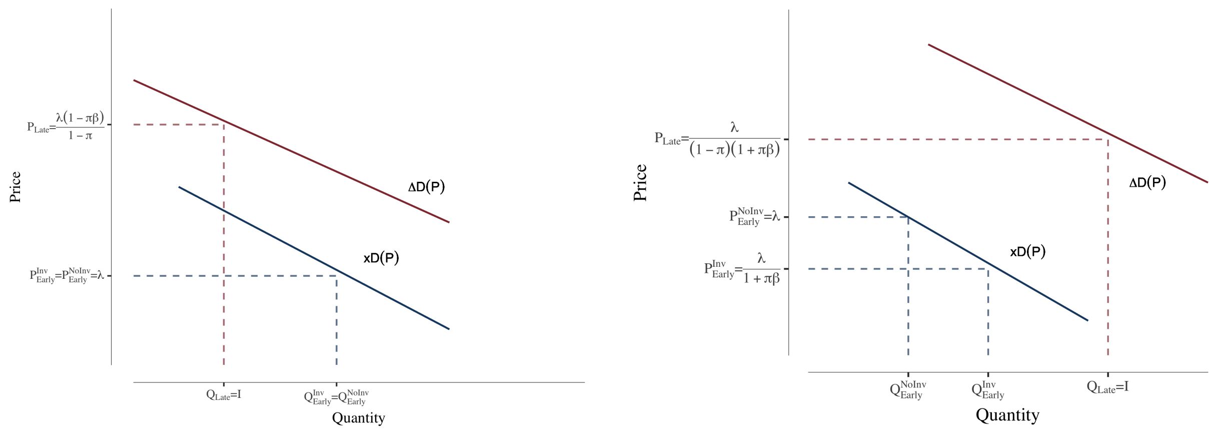

Note: This figure plots possible equilibria in the extended Uncertain and Sequential Trade model with storage. Panels A and B represent scenarios in which the inventory level is lower/higher than the quantity demanded in the *Early* market, respectively. Prices and quantities are denoted by P and Q respectively. Subscripts denote *Late* and *Early* markets. Superscripts denote the state of inventories where Inv represents a positive amount of inventories and $NoInv$ represents no inventories.

lines depict the *Late* and *Early* consumer markets, respectively. Panel A depicts the case in which both newly produced goods and stored goods are supplied to the *Early* market if inventories are carried. In this case, the price in the *Early* market is equal to marginal cost regardless of the state of inventories, $P_{Early}^{Inv} = P_{Early}^{NoInv} = \lambda$. As a result, the quantities in the *Early* market do not depend on the level of inventories, $Q_{Early}^{Inv} = Q_{Early}^{NoInv}$. We can determine the price in the *Late* market by solving the equilibrium condition (25), $P_{Late} = \frac{\lambda(1-\pi\beta)}{1-\pi}$. Lastly, the equilibrium quantity in the *Late* market is less than the quantity in the *Early* market.

Panel B of Figure 7 depicts the case in which inventories are greater than the quantity demanded in the *Early* market. Since stored units are supplied to both markets, we must have $(1-\pi)P_{Late} = P_{Early}^{Inv}$. Substituting into Equation 25 and solving for prices yields:

$$P_{Early}^{Inv} = \frac{\lambda}{1+\pi\beta} \tag{29}$$

$$P_{Late} = \frac{\lambda}{(1-\pi)(1+\pi\beta)} \tag{30}$$

The price in the *Early* market is equal to marginal cost if inventories are not carried, $P_{Early}^{NoInv} = \lambda$. Thus, unlike Panel A, this equilibrium represents a situation in which the price and quantity in the *Early* market vary with the state of inventories.[10]

5.3 Accounting for the stylized facts

Sections 3 and 4 presented several stylized facts about sales. First, the fraction of weeks in which there is no sale in any store is larger than predicted by using a mixed strategy. Second, stores and products that face more demand uncertainty have more sales. Third, sales play an important role in the response to demand shocks. In this section we use the UST model to discuss these findings.

In the UST model, stores that post a high price accumulate unwanted inventories when the realization of demand is low. The accumulated inventories are typically offered at a low price, so that they will be sold before reaching the expiration date. After the store sells all the "unwanted" inventories it may switch to the high price, so the reduction in price may be temporary and may satisfy our definition of temporary sale.

In our model, there are periods of high demand. After a period of high demand, there are no "unwanted" inventories and stores do not reduce prices. This is consistent with the observation that in many weeks there are no temporary sales in any of the stores.

Price dispersion in the UST model arises in response to demand uncertainty. When the distribution of demand is degenerate we get the standard Walrasian equilibrium with no accumulation of "unwanted inventories" and no temporary sales. This is the intuition behind the result in Eden (2018) who showed that when the distribution of aggregate demand is close to uniform, more demand uncertainty leads to more temporary sales.

Here we extend the results in Eden (2018) to the store level. In our model, newly produced goods are typically allocated to both (high and low price) markets. In the no inventories state that occurs after the realization of high demand, stores are indifferent between the two

[10]It is also possible that only stored units are allocated to the *Early* market and only new units are supplied to the *Late* market. This a special scenario of the equilibrium depicted in Panel B.

markets and some stores may choose the low price and some may choose the high price. If this choice is correlated over time and some stores consistently choose the high price in the no inventories state and some consistently choose the low price, we may find a correlation between demand uncertainty and the frequency of temporary sales also on the store level. Stores that choose the high price for newly produced goods will experience more fluctuations in the amount they sell because the high price market opens only when the realization of demand is high. These stores will also accumulate unwanted inventories when demand is low and will offer the stored goods at a temporary sale price.

In our regressions, average price does not seem to play a significant role in explaining the variation of the frequency of sales over stores. This suggests that variation in demand uncertainty across stores is not fully explained by differences in pricing strategies. We may assume that some buyers do not search over stores, so that each store has a group of non-shoppers that are loyal to it. If the demand of non-shoppers fluctuates in some stores more than in other stores, the stores that experience more fluctuations in the demand of non-shoppers will tend to accumulate unwanted inventories more often and will tend to have temporary sales more often.

The model implies that after a negative demand shock stores will have more unwanted inventories and more sales. Importantly, sales do not cancel out when aggregating over stores. Thus, the average sales frequency over stores also increases following a demand shock. This increase in temporary sales leads to a reduction in the average price. This is consistent with the impulse response functions we estimate.

6 Conclusion

It has become increasingly standard in macroeconomics to remove sale observations and analyze regular/non-sale prices. We argue that temporary sales play an important role in the response to demand shocks. Using scanner data, we find that the average price of a product decreases by almost 1.5% following a shock to the total quantity sold across stores.

If sale observations are removed and the average regular price is analyzed instead, then the cumulative response of prices is cut in half. Thus, ignoring sales underestimates the price reaction to a demand shock and overestimates the degree of price rigidity.

We then develop a model to account for this finding guided by several stylized facts seen in the data. We find that in about 45% of the weeks, an average item is not on sale in any store. In our model there are no temporary sales in any of the stores after a period of high demand. We also find a correlation between demand uncertainty and the frequency of sales at the store and product levels. In the model, products with more demand uncertainty should have more sales. Furthermore, stores that adopt a strategy of charging a high price for newly produced goods will experience more fluctuations in the quantity sold and will have more sales. We also find that the correlation between the frequency of sales and demand uncertainty holds even after controlling for the average price of the store. This suggests that demand uncertainty by non-shoppers plays an important role.

References

Aguirregabiria, Victor, "The Dynamics of Markups and Inventories in Retailing Firms," *Review of Economic Studies*, 1999.

Anderson, Eric, Benjamin A. Malin, Emi Nakamura, Duncan Simester, and Jón Steinsson, "Informational rigidities and the stickiness of temporary Sales," *Journal of Monetary Economics*, 10 2017, *90*, 64–83.

Bental, Benjamin and Benjamin Eden, "Inventories in a Competitive Environment," *Journal of Political Economy*, 1993.

_ **and** _ , "Money and inventories in an economy with uncertain and sequential trade," *Journal of Monetary Economics*, 1996, *37*, 445–459.

Bronnenberg, Bart J., Michael W. Kruger, and Carl F. Mela, "The IRI marketing data set," *Marketing Science*, 7 2008, *27* (4), 745–748.

Chevalier, Judith A. and Anil K. Kashyap, "Best prices: Price discrimination and consumer substitution," *American Economic Journal: Economic Policy*, 2 2019, *11* (1), 126–159.

Coibion, Olivier, Yuriy Gorodnichenko, and Gee Hee Hong, "The cyclically of sales, regular and effective prices: Business cycle and policy implications," *American Economic Review*, 3 2015, *105* (3), 993–1029.

Dana, James D., "Advance-purchase discounts and price discrimination in competitive markets," *Journal of Political Economy*, 1998, *106* (2), 395–422.

_ , "Equilibrium Price Dispersion under Demand Uncertainty: The Roles of Costly Capacity and Market Structure," *The RAND Journal of Economics*, 1999, *30* (4), 632–660.

Deneckere, Raymond and James Peck, "Dynamic Competition With Random Demand and Costless Search: A Theory of Price Posting," *Econometrica*, 2012, *80* (3), 1185–1247.

Eden, Benjamin, "Marginal Cost Pricing When Spot Markets Are Complete," *Journal of Political Economy*, 1990, *98* (6), 1293–1306.

_ , "The Adjustment of Prices to Monetary Shocks when Trade is Uncertain and Sequential," *Journal of Political Economy*, 1994.

_ , "Price Dispersion and Demand Uncertainty: Evidence from U.S. scanner data," *International Economic Review*, 8 2018, *59* (3), 1035–1075.

Eichenbaum, Martin, Nir Jaimovich, and Sergio Rebelo, "Reference prices, costs, and nominal rigidities," *American Economic Review*, 2 2011, *101* (1), 234–262.

Glandon, P. J., "Sales and the (Mis)measurement of price level fluctuations," *Journal of Macroeconomics*, 12 2018, *58*, 60–77.

Guimaraes, Bernardo and Kevin D. Sheedy, "Sales and monetary policy," *American Economic Review*, 4 2011, *101* (2), 844–876.

Kehoe, Patrick and Virgiliu Midrigan, "Prices are sticky after all," *Journal of Monetary Economics*, 10 2015, *75*, 35–53.

Kryvtsov, Oleksiy and Nicolas Vincent, "The Cyclicality of Sales and Aggregate Price Flexibility," *The Review of Economic Studies*, 2 2021, *88* (1), 334–377.

Nakamura, Emi and Jon Steinsson, "Five Facts About Prices: A reevaluation of menu cost models," *Quarterly Journal of Economics*, 2008.

Prescott, Edward C, "Efficiency of the Natural Rate," *The Journal of Political Economy*, 1975, *83* (6), 1229–1236.

Salop, Steven and Joseph Stiglitz, "Bargains and Ripoffs: A Model of Monopolistically Competitive Price Dispersion," *Review of Economic Studies*, 1977.

Sheremirov, Viacheslav, "Price dispersion and inflation: New facts and theoretical implications," *Journal of Monetary Economics*, 2020.

Shilony, Yuval, "Mixed pricing in oligopoly," *Journal of Economic Theory*, 1977.

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

A Appendix

A.1 Empirics

Table A.1: Demand Uncertainty and Temporary Sales (Regressions that address Potential Endogeneity Issues)

	Dependent variable:					
	Store-level ($sale\_freq_j$)		Product level ($sale\_freq_i$)		Store-product level ($sale\_freq_{ij}$)	
	(1)	(2)	(3)	(4)	(5)	(6)
SDU	0.294***	0.110***	0.432***	0.260***	0.114***	0.072***
	(0.029)	(0.033)	(0.018)	(0.020)	(0.005)	(0.005)
SDP		1.392***		1.149***		0.361***
		(0.092)		(0.055)		(0.017)
Ave. ln(Units)		−0.001		0.004		0.013***
		(0.006)		(0.004)		(0.001)
Ave. ln(Price)		−0.007		−0.010***		0.161***
		(0.011)		(0.002)		(0.009)
UPC Fixed Effects					X	X
Store Fixed Effects					X	X
Observations	546	546	1,686	1,686	94,505	94,505
Adjusted R^2	0.159	0.421	0.256	0.463	0.556	0.577

Note: This table presents regression estimates of Equations (8)-(9) that address potential endogeneity issues. This table repeats the regressions of Table 4, but here we use the 2005 sample to measure the dependent variable and the 2004 sample to measure the independent variables. SDU and SDP represent the standard deviations of log units and prices, respectively. Standard errors are clustered at the store level for the product-store regressions in columns (5)-(6).
*p<0.1; **p<0.05; ***p<0.01

Table A.2: Aggregate Impulse Response of Temporary Sales to Demand Shocks

	Weeks after Shock						Max. Cumulative Effect
	1	2	3	4	5	6	
Baseline							
2 Lags	0.34	0.22	0.06	-0.03	-0.05	-0.04	0.62
	[0.29, 0.39]	[0.17, 0.28]	[0.02, 0.09]	[-0.05, -0.01]	[-0.05, -0.04]	[-0.04, -0.03]	[0.52, 0.74]
4 Lags	0.34	0.22	0.14	-0.03	-0.03	-0.02	0.70
	[0.29, 0.39]	[0.17, 0.27]	[0.09, 0.2]	[-0.08, 0.02]	[-0.07, 0]	[-0.04, 0]	[0.58, 0.82]
10% Reduction & IRI Flag							
2 Lags	0.41	0.22	0.05	-0.03	-0.04	-0.03	0.68
	[0.35, 0.46]	[0.17, 0.28]	[0.02, 0.09]	[-0.05, -0.02]	[-0.05, -0.03]	[-0.04, -0.03]	[0.58, 0.80]
4 Lags	0.40	0.23	0.13	-0.06	-0.04	-0.02	0.77
	[0.36, 0.47]	[0.18, 0.29]	[0.08, 0.19]	[-0.12, -0.01]	[-0.07, -0.01]	[-0.04, -0.01]	[0.66, 0.89]

Note: This table plots the orthogonalized impulse response functions of the aggregated sales frequency for a product in response to a one standard deviation negative shock to the total quantity sold. Rows represent the specification, and columns specify the period after the shock. The last column presents the maximum cumulative effect for each specification. Forecast confidence intervals bootstrapped at the 95% confidence level are in brackets. The first row plots our baseline results. The second row plots the estimates using a VAR(4). The third and fourth rows are similar to first and second rows, but add the restriction that a sale is also indicated by the IRI flag.