

Rigid Prices: Evidence from U.S. Scanner Data*

Jeffrey R. Campbell[†] Benjamin Eden[‡]

April, 2007

Abstract

This paper uses over two years of weekly scanner data from two small US cities to characterize time and state dependence of grocers' pricing decisions. In these data, the probability of a nominal adjustment *declines* with the time since the last price change even after controlling for heterogeneity across store-product cells and replacing sale prices with regular prices. We also detect state dependence: The probability of a nominal adjustment is highest when a store's price substantially differs from the average of other stores' prices. However, extreme prices typically reflect the selling store's recent nominal adjustments rather than changes in other stores' prices.

JEL Classification: E31, L16, L81.

Keywords: Nominal Price Changes, Time-Dependent Pricing, State-Dependent Pricing, Hazard Function.

*The authors gratefully acknowledge research support from the National Science Foundation through grant 013048 to the NBER. Gadi Barlevy, Marco Bassetto, Mario Crucini, Szabolcs Lorincz, and participants at the 2004 SED meetings made many helpful suggestions. The views expressed in this paper are those of the authors and do not reflect those of the Federal Reserve Bank of Chicago, the Federal Reserve System, or its Board of Governors. A replication file is available at <http://www.nber.org/~jrc/scanner>

[†]Federal Reserve Bank of Chicago and NBER. E-mail: jcampbell@frbchi.org

[‡]Vanderbilt University and University of Haifa. E-mail: ben.eden@vanderbilt.edu

1 Introduction

This paper measures time and state dependence in grocers’ pricing decisions using scanner data. The observations cover most transactions of items in five product categories in two Midwestern cities. These weekly records allow accurate calculation of the age of an item’s price and of the corresponding prices paid at other stores.

Twenty-three percent of the prices in these data change in an average week. If the probability of a nominal adjustment remains the same as a price ages, then this implies that the average price lasts about four weeks. However, a randomly chosen price remains unchanged for more than eight weeks. The discrepancy between these duration estimates reflects a fact at odds with familiar macroeconomic models of pricing: the frequency of nominal adjustment *declines* with the time since the last price change. We find this counterintuitive time dependence even after controlling for heterogeneity in price flexibility across products and stores. That is, occasional spells of flexibility punctuate otherwise rigid prices.

Models of state-dependent pricing – such as those of Barro (1972); Sheshinski and Weiss (1977); Caplin and Spulber (1987); Caplin and Leahy (1991); and Dotsey, King, and Wolman (1999) – imply that the benefit of a nominal adjustment is highest when the price differs substantially from other sellers’ prices. Our findings reproduce this qualitatively: Increasing the difference between an item’s price and the average price for the same item at other stores substantially raises the probability of a price change. However, the probability of changing a price close to average far exceeds zero, and most price changes occur with the original price close to average. In simple menu-cost models, extreme prices arise from the erosion of a fixed nominal price by other sellers’ adjustments. Therefore, they are older than average. We find that most extreme prices are relatively young (less than a month old). That is, grocers deliberately select extreme prices which they then quickly abandon. Taken together these results suggest to us that sellers extensively experiment with their prices.

Many papers that examine price data collected to construct the CPI precede this work. Examples from the Euro zone, Israel, Poland, and the United States include Dhyne et al.

(2006), Lach and Tsiddon (1992), Konieczny and Skrzpacz (2005), and Klenow and Kryvtsov (2005). These micro-CPI data record the prices for many more items than do the scanner data we employ. We therefore view this paper as complementing earlier studies by looking at a narrower but richer set of data. Our data are weekly. This is an advantage over the monthly CPI data because stores typically change prices more than once in a single month. Furthermore, our data are more detailed than CPI data. In Bils and Klenow (2004) for example, there is one product category called “margarine”. We examine the pricing of 62 specific margarine products. This allows us to construct the relevant comparisons of prices across stores required for measuring state dependence. Finally, scanners directly measure transaction prices with little human intervention; unlike BLS enumerators.

Dutta, Bergen, and Levy (2002) and Chevalier, Kashyap, and Rossi (2003) examined scanner data of prices at a single Chicago supermarket chain. These observations share the high frequency of the data we employ, and in addition they record the supermarket’s markup over wholesale cost. The advantage of this paper’s data arises from their coverage of multiple sellers. In leveraging this feature, we follow Kashyap (1995). He compared price adjustments of three retailers selling a few identical items. Our work examines 135 grocery products each sold in five or more stores.

The remainder of this paper proceeds as follows. The next section discusses the source of the data and how we use it to detect nominal price changes. It also presents summary statistics and foretells our results with the behavior of the price for a specific item at one store. Section 3 measures time dependence of pricing decisions, and Section 4 studies the dependence of price changes on stores’ relative prices. Section 5 unifies the study of time and state dependence by estimating linear regression models of the decision to change a nominal price. Section 6 discusses the robustness of our results to different measurement strategies. Section 7 examines how wholesale price dynamics and price synchronization impact our results, and Section 8 offers concluding remarks.

2 Data

Our data source is the ERIM scanner data set collected by A.C. Nielsen. The James M. Kilts Center for Marketing at the University of Chicago’s Graduate School of Business graciously makes these data available on its web site.¹ Nielsen collected these data from two small Midwestern cities - Springfield, Missouri and Sioux Falls, South Dakota - from the fifth week of 1985 through the twenty-third week of 1987. The data come from the checkout scanners of these cities’ supermarkets and drug stores. The sample includes observations from 19 stores in Springfield and 23 stores in Sioux Falls. Together, they account for about 80 percent of the two markets’ grocery and drug retail sales. We identify a product with a Universal Product Code (UPC). These differ across different packagings of the same good (e.g. 8 oz. and 16 oz. sizes) and across different varieties of that good (e.g. flavored and unflavored margarine). For each product in six categories - ketchup, margarine, peanut butter, sugar, toilet tissue, and tuna - the data record the revenues from the sales of that product as well as the quantity sold at each store. Nielsen also issued identification cards to approximately 10 percent of each city’s households. These customers presented their cards at stores’ checkout counters, and Nielsen used the resulting observations to construct household-level purchase histories. These allow us to observe the exact transaction prices and locations for goods purchased by these households.

Our baseline measure of a good’s price in a particular store is average revenue per unit, because records of individual stores’ prices constructed from households’ purchase histories are incomplete for all but the most popular products and stores. The measure of revenues equals the amount the store would have received at the register if customers had redeemed no coupons, so changes in coupon redemption do not directly influence these price measures.

We refer to a specific product sold in a given store as a store-product cell (or simply cell whenever this does not lead to confusion). The ERIM data contain 271,028 price observations from cells with sales in both the first and last sample weeks. We call these continuing

¹The data can be found at <http://gsbwww.uchicago.edu/kilts/research/db/erim/>.

cells.² A store might not sell a given product to any household during a particular week. In that case, we do not observe a transaction price. The sample of store-product cells with positive sales in every sample week contains only 89,175 prices and accounts for 52 percent of continuing cells' sales. We believe that this is too exclusive a sample for our purposes, so we also use data from all cells with spells of missing data that last no longer than two weeks per spell. For a store-product-week with no sales, we assume that the price equals the most recently observed price. The resulting sample includes 203,811 observations, of which 3,009 are imputed with the most recently observed price. They account for 84 percent of all continuing cells' sales. The final criterion for inclusion of a product in our sample is that no fewer than five stores sold it in any week. This eliminates 24,846 prices from the data. The comparison prices allow us to determine whether a store's price is close to other stores' prices for the same good, and they need not appear in our final data set.³

2.1 Measuring Price Changes

The final balanced panel has 178,965 prices, and they account for 75 percent of continuing cells' sales. The division of revenues by units sold yielded 39,919 prices which cannot be expressed in whole cents (e.g. \$1.3529). Such a fractional price could arise either from technical mistakes in price setting or from time aggregation. Technical errors occur when the price displayed in the store differs from that in the computer. For example, suppose that a store manager decided to change the price on Monday from \$1.29 to \$1.40. The price must be changed both on the computer and on the shelf. Suppose that erroneously it changes

²These data also contain 436,357 prices from store-product cells without sales in both the first and last sample weeks. The pricing and sales dynamics of such cells is of substantial independent interest, but the complete analysis of the decision to introduce or retire a product offering lies well beyond the scope of this paper.

³We have also created versions of every table and figure in this paper using only the 89,175 prices from complete store-product cells. Although those estimates are somewhat less precise than those reported here, they lead us to the same conclusions.

only on the computer. Then those customers who notice the lower price on the shelf will complain and receive the item for \$1.29, while others who are less attentive will pay \$1.40. At some point, the store will correct the shelf price, but the earlier mixing can nevertheless result in an average price of \$1.3529.⁴


The data weeks begin with Monday, so a fractional price can also arise from time aggregation when a store manager changes an item’s price during the middle of a week. To illustrate this possibility, change the previous example to suppose that the manager changes the price from \$1.29 to \$1.40 successfully on a Wednesday. Those customers buying the good on Monday or Tuesday would pay \$1.29 while those buying later would pay \$1.40. We might therefore observe in the second week a fractional price like \$1.3529 as a result of time aggregation. The average revenue price changes twice in three weeks, while the daily price in the computer changes only once. The spurious price change arises because the second week’s average revenue embodies two prices.

To address these problems we replaced fractional prices by the minimum price in the individual purchase history data whenever this was possible. In our example, if we observe three consumers purchasing the item in the second week for \$1.29, \$1.29, and \$1.40 we change the second week’s price to \$1.29. This was done for 26,051 prices out of the 39,919 fractional prices in the data.⁵ For the remaining prices that we could not replace, we checked whether they are part of a descending or increasing sequence of prices (as in the example). If the fractional price was a part of a monotone sequence we concluded that it is likely to

⁴The Federal Trade Commission executed two studies of scanner accuracy in 1996 and 1998. These defined an error as we have here: a failure of the scanned price to match the lowest posted price in the store. In 1996, the error rate from a sample of randomly selected items was 4.82 percent. In 1998 this dropped to 3.35 percent. The original studies are available at <http://www.ftc.gov/reports/scanner1/scanners.htm> and <http://www.ftc.gov/reports/scanner2/scanner2.htm>.

⁵The purchase history gives us multiple prices for 5,982 of the available 26,051 prices, and the median range among those that are different is 17 cents. Therefore, we have also experimented with using the maximum available purchase-history price. This changes none of our results substantially.

Figure 1: The Price of Fleischmann's Margarine⁽ⁱ⁾



Note: (i) Weekly observations of the price of Fleischmann's Margarine at a store in Sioux Falls, South Dakota and the average of all other stores' prices for the identical product. Dates are the final days of the given week.

be the result of time aggregation and replaced the fractional price by the following week's price. Applying this rule to our example changes the price of \$1.3529 in week 2 to \$1.40; and the number of price changes in the corrected data is accurate. We changed 3,126 prices in this manner. Average-revenue prices that were not in whole cents but were either greater than or less than both the previous and following week's prices were rounded to the nearest whole cent but otherwise left unchanged. If we change our example so that the store lowers the price from \$1.40 to \$1.26 on the beginning of week 3 and maintains this price through week 3, then the three weeks' average prices are \$1.29, \$1.3529, and \$1.26. We count two price changes after rounding the second week's price down to \$1.35, as actually occurred.

Figure 1 plots one store's price of a single product (Fleischmann's Margarine) along with

the average of all other store’s prices for the same item.⁶ This price changed 42 times during the 123 week sample. It begins at \$1.06, and it rises to \$1.09 in April of 1985, and remains roughly constant for the rest of the year. The average of other firms’ prices approximately equaled \$1.10 throughout these eleven months. The price changed much more often in 1986. After two very modest price increases in January, the price dropped to \$0.92 for four weeks and then returned to \$1.09. Throughout these changes, the average price at other stores fluctuated little. The return to \$1.09 lasted only seven weeks. In April, the price entered a period with very frequent changes that ended only in October. At that point, it approximately settled at \$1.15. The year ended with a dramatic temporary price increase. In 1987, the price returned to a pattern of much less frequent price changes. It ended the sample period at \$1.15.

The Figure shows seven price increases that last exactly one week and are followed by a return to the original price. Those are fractional prices that were not corrected because there were no purchase history data to replace them and they are not within a monotone sequence of prices. However we suspect that these observations might not be real price changes. If these are indeed “mistakes” they would affect our calculations of the weekly frequency. It would change from $42/123 = 0.34$ to $28/123 = 0.23$. We thought of replacing all prices which follow and precede exactly the same price with that common bracketing price, but doing so might be imposing too much of our prior on the data: A manager could decide to increase the price, realize that this leads to a loss of revenues, and revert to the previous price after a week. With this in mind, we leave such prices unchanged. However, we check our results’ robustness to treating such one-week price increases as mistakes in Section 6.

⁶Here and throughout the paper, we construct the average of all other stores’ prices for the same item by dividing total sales of the item across all other stores by the number of units sold by those stores.

Table 1: Summary Statistics


Category	Number of <i>UPCs</i>	Number of Observed Prices	Frequency of Sale Prices	Annualized Rate of Price Change
All Products	135	178,965	3.4	-2.4
Ketchup	6	7,626	4.7	0.4
Margarine	62	99,753	2.4	-3.6
Peanut Butter	13	9,840	4.0	10.9
Sugar	14	14,268	4.0	-0.5
Tissue	12	13,407	5.0	-4.4
Tuna	28	34,071	4.8	-3.1

2.2 Summary Statistics

Before proceeding to examine price changes, we document some of the scanner data set’s most salient characteristics. Table 1 provides summary statistics for the sample we use. Its first two columns report the number of products observed by category and the number of prices recorded. The sample includes observations of 135 products, and most of these are either margarine or tuna.

The third column of Table 1 reports the fraction of prices that we identify as sale prices. We wish to ensure that this paper’s results do not merely reflect firms’ switching between sale and “regular” prices, because some authors discount these as variation arising from a simple pricing rule which ignores available macroeconomic information rather than a conscious change in that rule. To identify sales, we look for price declines of 10 percent or more in a given week that the store completely reverses within 2 weeks. All prices between the initial decline and the reversal are sale prices. With this criterion, only 3.4 percent of the observations are sale prices. Tissue’s frequency of sale prices, 5.0 percent, exceeds that of any other category. [Klenow and Kryvtsov \(2005\)](#) report that the BLS identifies 15 percent of food prices collected to produce the CPI as sale prices. Either the stores in our sample

Figure 2: Inflation Rates for Margarine⁽ⁱ⁾



Note: (i) Annualized monthly inflation rates.

use sale prices relatively infrequently or the BLS uses a more liberal definition of a sale.⁷

Table 1's final column reports the annualized average rates of price change in percentage points. We find this of interest because inflation expectations impact firms' price choices. The prices in all categories but Peanut Butter declined over the sample period. The corresponding average annual growth rate of the consumer price index for margarine is -1.7 percent. The matching *CPI*'s for the other categories all display price growth, so the deflation in Springfield and Sioux Falls did not typify the national experience.

Aggregate fluctuations in inflation also concern price-setting producers. To illustrate the sort of aggregate fluctuations facing the sample's stores, Figure 2 plots two monthly measures of annualized inflation for margarine over the scanner data's sample period. The first measure uses a geometric average fixed-weight price index conceptually similar to the

⁷This procedure identifies none of the prices in Figure 1 as sale prices.

Table 2: Standard Deviations of Log Prices^{(i),(ii)}

Category	{UPC \times Market \times Date} \cup				
	{UPC}	{UPC \times Market}	\emptyset	{UPC \times Store}	{UPC \times Store \times Sale}
All Products	14.6	14.2	10.8	9.6	9.0
Ketchup	14.7	13.2	10.1	9.2	8.8
Margarine	13.7	13.4	9.9	8.8	8.3
Peanut Butter	13.8	13.7	8.0	7.5	7.0
Sugar	11.9	11.1	7.5	7.3	6.2
Tissue	12.9	12.4	8.1	7.3	6.7
Tuna	18.4	18.0	15.4	13.6	12.6

Notes: (i) The table reports residual standard deviations in percentage points from the regressions of the price's logarithm against the given set of regressors. (ii) In the table, {X,Y} indicates a set of dummy variables that span the unique combinations of the values of X and Y.

CPI. We built this with the scanner data following the procedure of [Richardson \(2003\)](#). The second measure is the margarine *CPI* itself. The *CPI*-based inflation varies much less than the scanner-based inflation. Their standard deviations are 11.0 and 21.5 percent. Their sample correlation is 0.35, which suggests that location-specific shocks dominate the scanner-data based price index. We could not locate a national-level *CPI* for toilet tissue. The other categories' scanner-based inflation rates also display considerably greater variance than their corresponding *CPI*-based rates.

Next, consider the variability of prices across stores and time. The first column of Table 2 reports residuals' standard deviations from a regression of the price's logarithm against a set of 135 UPC dummy variables. The first row reports the residuals' standard deviation for all of the observations, and the remaining rows report the standard deviations for each category's residuals. These range from 11.9 percent for Sugar to 18.4 percent for Tuna. The table's remaining columns report residual standard deviations from regressions that include progressively richer sets of dummy variables. The regression underlying the second

column’s results includes two sets of 135 UPC dummy variables, one set for each market. This accounting for systematic differences between prices in Springfield and Sioux Falls lowers the standard deviations little. The regression for the third column includes one set of UPC dummy variables for each market and week for a total of $135 \times 2 \times 123$ dummies. As Figure 2 suggests, removing date-specific means substantially lowers variation. For example, margarine’s standard deviation drops from 13.4 percent to 9.9 percent. Because $(9.9/13.4)^2$ approximately equals one half, the cross-sectional variance of prices at a given date and the time-series variance of the average price across dates roughly equal each other.

Table 2’s final two columns further decompose the cross-sectional dispersion of prices. The regression used for the fourth column adds store-specific UPC dummies that are invariant across time to the regression in the third column. The average UPC has approximately 10.8 stores, so this adds $1455 \approx 135 \times 10.8$ dummy variables. We expect these dummies to substantially reduce the standard deviation if stores consistently follow “low-price” or “high-price” strategies.⁸ In fact, adding store-specific UPC dummies lowers the standard deviation at most 1.8 percentage point. This indicates that there are few *systematic* differences in the prices for a given product across either markets or stores. The final column quantifies the contribution of stores switching between sale prices and regular prices to price dispersion. For this, we added two sets of store-specific UPC dummies to the regression from the third column, one for regular prices and another for sale prices (a total of 2×1455 dummies). Accounting for the differences between sale and regular prices lowers the standard deviations somewhat, but overall Table 2 indicates that substantial price dispersion remains even after controlling for heterogeneity across markets, stores, products, and weeks.

Next we consider the frequency of price changes, which Table 3 reports. The first column gives the weekly average fraction of prices that changed for the whole sample and each of the six categories. Overall, 23 percent of prices change in a given week. The second column

⁸Because the store-specific dummies also vary across UPC’s, this regression will account for persistent heterogeneity across stores in the pricing of particular items that does not reflect store-wide pricing strategies.

Table 3: The Frequency of Price Changes⁽ⁱ⁾

Weekly Data

Category	Original Data	Sales Replaced ⁽ⁱⁱ⁾	Monthly Data ⁽ⁱⁱⁱ⁾	Bils-Klenow ^(iv)
All Products	23	20	39	26
Ketchup	25	20	46	20
Margarine	23	21	39	28
Peanut Butter	27	22	46	31
Sugar	19	14	33	23
Tissue	21	16	41	24
Tuna	24	19	38	27

Notes (i) The table’s entries are frequencies expressed in percentage points. (ii) The observations used to calculate this column’s results have sale prices replaced with the most recent non-sale (or “regular”) price. (iii) Monthly observations are constructed by using the price of each store-product cell in the first week of each calendar month. (iv) The entries in this column are the frequencies of price changes reported in Table 1 of Bils and Klenow (2004) for the respective categories. See Footnote 9 in the text for more information regarding the mapping of the ERIM categories into BLS categories. The first frequency in this column is the simple average of those for the six reported categories.

reports the frequencies after first replacing sale prices with the most recent non-sale (or “regular”) price. This equals 20 percent for All Products, so most price changes in these data remain even after accounting for sales. The final columns compare the price changes in these data with those tabulated by the BLS while constructing the *CPI*. The third column computes the average monthly frequency that is obtained by “visiting” a store during the first week of each month, and the last column reports the BLS estimates as reported in [Bils and Klenow \(2004\)](#).⁹ If the probability of a price changing did not depend on the price’s age,

⁹Two of the scanner data’s categories have identically named BLS item categories, Margarine and Peanut Butter. We matched Ketchup with “Other condiments (excl olives, pickles, and relishes),” Sugar with “Sugar and artificial sweeteners,” Tissue with “Cleaning and toilet tissue, paper towels, napkins,” and Tuna with

Table 4: Estimates of Average Price Durations⁽ⁱ⁾

Category	Average Durations ⁽ⁱⁱ⁾		Inverse Frequency Estimates	
	All Prices	New Prices	Inverse of Average ⁽ⁱⁱⁱ⁾	Average of Inverses ^(iv)
All Products	8.3	4.3	4.3	5.5
Ketchup	6.7	3.8	4.0	4.7
Margarine	8.2	4.4	4.3	5.5
Peanut Butter	7.0	3.8	3.7	4.3
Sugar	10.1	4.8	5.2	6.7
Tissue	7.0	4.5	4.7	6.0
Tuna	9.2	3.9	4.1	5.1

Note: (i) Table entries measured in weeks. (ii) These columns report the average duration of prices set in the sample’s first 85 weeks. (iii) This column reports the inverses of the weekly frequencies from the first column of Table 3. (iv) This column reports the average of inverse price frequencies calculated for each store-product cell.

as in the [Calvo \(1983\)](#) model of price adjustment, then the weekly frequencies we observe would imply that 65 percent of prices change in a given month. However, we find instead that approximately 39 percent of prices change when sampled monthly. This suggests that the [Calvo](#) assumption of a constant weekly probability of price adjustment does not hold good in our data. In any case, the prices in our data appear to be somewhat more flexible than those in the BLS sample.

Table 4 provides direct and inverse-frequency type estimates of the average price duration. The first two columns report direct measures using average realized durations of prices before the sample’s last year. Average duration can only be calculated using completed price spells without employing strong distributional assumptions, but because the typical price duration is short relative to the sample period’s length, we observe the completion of all but 2 spells that begin before the sample’s last year. The first column reports the average duration for

“Canned fish or seafood.”

all prices charged. This equals 8.3 weeks for all products. Across categories, it varies from 7.0 weeks for Peanut Butter to 10.1 weeks for Sugar. The second column reports the average durations of the sample’s newly-set prices. This is about half of the average duration for all prices.

The last two columns of Table 4 provides estimates based on the inverse of the price-change frequency. The first of the two is the standard estimate. This is very close to the average duration for newly set prices.¹⁰ Let λ_i denote sample frequency of price changes for store-product cell i . The inverse-frequency estimate in the third column equals the inverse of the average of λ_i across cells. Jensen’s inequality implies

$$(1) \quad \frac{1}{\frac{1}{N} \sum_{i=1}^N \lambda_i} < \frac{1}{N} \sum_{i=1}^N \frac{1}{\lambda_i}$$


That is, the average of the inverse frequencies exceeds the inverse of the average frequency. Baharad and Eden (2004) argue that if the frequency of price change is cell-specific and there is no within-cell heterogeneity across time, then the correct measure of duration is the right-hand side of (1). We report this in Table 4 final column. In fact, the average of the inverses lies about half-way between the inverse of the average in the third column and the direct measure of price duration in the first column.

3 Time Dependence

We now turn to the measurement of time dependence in grocers’ decisions to change nominal prices. For this we use the unconditional hazard function, which plots the adjustment frequency as a function of the price’s age. In the Calvo (1983) model of stochastic price setting, this function does not vary with the price’s age, whereas in the Taylor (1980) model

¹⁰In an infinite sample, the two estimates equal each other by construction. They differ in Table 4 because the first spell (with unknown duration) contributes to the inverse-frequency estimate but not the direct estimate in the second column and because the direct estimate ignores spells that begin in the sample’s last year.

Figure 3: Sample Hazard Function for Price Changes




of staggered pricing it equals zero until the interval of price rigidity passes, at which point it jumps to one. Standard models of state dependent pricing clearly imply the unconditional hazard function increases with the price’s age for very young prices, because a producer gains nothing from changing a newly-set and hence optimal price.

Figure 3 plots the unconditional hazard function estimated using the observations from all product categories. The chance of a newly-set price changing equals 47 percent. As the price ages, this probability drops precipitously. It equals 32 percent for a two-week-old price and 21 percent for a three-week-old price. As the price ages further, the hazard function continues its decline at a more gradual pace. For very old prices, the probability of a price change equals only 8 percent. The unreported hazard functions calculated separately for each product category all resemble Figure 3.

If occasional measurement errors (like those we suspect in Figure 1) infect our data, then they can make the measured hazard function drop, because measuring a single price with

Figure 4: Example Hazard for Price Changes with Rigid and Flexible Prices



error induces two sequential “observed” price changes. This only causes the hazard to drop from one to two weeks, but in the data the drop from one to two weeks only slightly exceeds that from two to three weeks. Apparently, measurement error alone cannot rationalize the estimated hazard function’s shape.

Heterogeneity provides a simple explanation for the decreasing hazard function. We illustrate this here with an example adapted from [Darby, Haltiwanger, and Plant \(1985\)](#). Suppose a price is either flexible or rigid. Within each type there is a constant probability of changing the price, but flexible prices change more frequently. The hazard function initially reflects the average probability across the two groups. As a cohort of prices set on a given date ages, the fraction of rigid prices among the survivors increases. The hazard function declines (as in [Figure 3](#)) and asymptotes to the probability of a rigid price changing. [Figure 4](#) plots the implied raw hazard function from this example. It assumes that 70 percent of all new prices are flexible and that the weekly probability of a price change is 65 percent for a

flexible price and 8 percent for a rigid price. This simple example reproduces Figure 3 well.

Standard models can accommodate heterogeneity in price duration between product-store cells, but they do not produce heterogeneity across time within a given cell. To check whether within-cell heterogeneity contributes to the decreasing hazard, we calculated the hazard within each product-store cell for young prices (with ages less than or equal to 3 weeks) and old prices (with ages greater than 3 weeks). Whenever the hazard for young prices was higher than the hazard for old prices we said that it is declining. Table 5 displays the fraction of cells that exhibit decreasing hazard function. Overall, 88.9 percent of cells have a declining hazard. Under the null hypothesis of a constant hazard, these two estimates have a known asymptotic distribution (as the number of weeks per cell grows) which allows us to measure the statistical significance of their difference.¹¹ About 46.5 percent of cells have slopes that are negative and statistically significant at the 5 percent level. Only 1.1 percent of the cells have positive and statistically significant slopes. The bottom panel of Table 5 reports the same estimates after first replacing sale prices with the most recent non-sale prices. This has no dramatic impact on the results.¹²

In summary, these data display a counterintuitive form of time dependence in price setting. Older prices are *less* likely to change than newly-set prices. This does not merely reflect heterogeneity across stores or products in the frequency of price setting. Instead, hazard functions calculated using only one store-product cell’s observations typically decrease with the price’s age. That is, the price for a given product at a given store switches between periods of apparent rigidity and flexibility, as did the price for margarine in Figure 1.

¹¹Under the null hypothesis that the hazard function is constant and equal to λ , the estimated hazards for young and old prices from a sample of T weeks’ prices are asymptotically normally distributed with common mean λ , standard deviations $\sqrt{\lambda(1-\lambda)/((1-(1-\lambda)^3)\sqrt{T})}$ and $\sqrt{\lambda(1-\lambda)/((1-\lambda)^3\sqrt{T})}$, and zero covariance. We base our test on this joint distribution.


¹²Because a product-store combination defines a cell, the results of Table 5 account for systematic differences across stores’ price change frequencies.

Table 5: Slopes of Within-Cell Hazard Function⁽ⁱ⁾

Category	< 0	Significantly < 0	Significantly > 0
Original Data			
All Products	88.9	46.5	1.1
Ketchup	82.3	25.8	0.0
Margarine	87.2	45.9	1.6
Peanut Butter	85.0	37.5	1.3
Sugar	94.0	45.7	0.0
Tissue	89.9	33.0	0.0
Tuna	93.9	61.0	0.7
Sales Replaced with Regular Prices			
All Products	88.9	47.8	1.0
Ketchup	83.9	37.1	0.0
Margarine	87.7	45.4	1.2
Peanut Butter	88.8	41.3	1.3
Sugar	95.7	56.9	0.0
Tissue	79.8	35.8	0.9
Tuna	94.2	59.9	1.1

Note: (i) The table's top panel gives the percentage of cells with decreasing sample hazard functions, the second and third give the percentages with sample hazard functions with statistically significant negative and positive slopes. The significance level used was 5%. The bottom panel mimics the top after first replacing sale prices with the most recent non-sale (or "regular") price. A "cell" refers to a product-store combination.

Figure 5: Hazard for Price Changes as a Function of the Relative Price⁽ⁱ⁾



Note: (i) The relative price equals the store's price in the previous week divided by the average of all other prices for the same UPC charged in the current week.

4 State Dependence

Existing models of state-dependent pricing cannot easily generate a decreasing hazard function like that we observe, because the benefit of changing a newly-set price is small and grows as the price ages. Nevertheless, the insights of state-dependent pricing models might yet improve our understanding of stores' nominal adjustments. In this section, we examine this possibility. We use a noisy indicator of the benefit of a nominal adjustment, the price relative to the average of other stores' prices for the same product.¹³

Figure 5 plots the observed frequency of price changes as a function of the relative price's logarithm. The relative price equals the ratio of the store's nominal price in the *previous*

¹³One possible objection to this measure of relative prices is that its denominator might include prices charged at other stores owned by the same firm. We address this below in Section 6.


week divided by the sales-weighted average of prices for the same good at all other stores in the current week. This measures the real gap that a price adjustment in the current week could close. Before estimation, we accounted for some stores systematically following high-price or low-price rules by normalizing the mean of each store-product cell’s log relative price to zero. On the horizontal axis, zero indicates a relative price equal to the average for this store-product cell. We divided the interval $[-1/2, 1/2]$ into twenty equally sized bins and calculated the price change frequency for each of them. The dark solid curve in Figure 5 gives the frequencies for all store-item-week observations. The grey solid gives the analogous frequencies calculated after first replacing sale prices with the most recent non-sale price. For visual reference, the light horizontal line gives the unconditional frequency of a price change, 23 percent; and the dashed curve plots the sample’s distribution of relative prices.

There are four notable features of Figure 5. First, the minimum frequency substantially exceeds zero. For both samples, it approximately equals 15 percent. Thus, even a store with an “average” price might change it. Second, firms’ prices cluster around the average of *other* firms’ prices. Together, these two observations strongly suggest that the relative price cannot substantially improve forecasts of the occurrence of nominal adjustment. Third, moving the relative price away from its average substantially increases the probability of a nominal adjustment. The estimated probability of a nominal adjustment is 60 percent when the price is 35 to 40 percent below the average of others’ prices and 53 percent when it is 35 to 40 percent above that average. In this sense, these observations display a basic feature of menu-cost pricing models. Fourth and finally, replacing sale prices lowers the average frequency of price changes.¹⁴

In light of the negative association of a price’s age with the probability that it changes, Figure 6 plots the frequency of price changes against the mean-adjusted logarithmic relative price for samples of young prices – those with ages three weeks or less – and for old prices

¹⁴Category-specific versions of Figure 5 all display the same features, but the estimated frequencies are considerably noisier.

Figure 6: Young and Old Prices' Hazards as Functions of the Relative Price⁽ⁱ⁾




Note: (i) The relative price equals the store's price in the previous week divided by the average of all other prices for the same UPC charged in the current week. The calculations exclude observations from the initial (left-censored) price spells. Young prices are defined to be those less than four weeks old.

– those with ages of four weeks or more. As Figure 3 suggests, the adjustment frequencies of new prices exceed those of old prices substantially. Furthermore, extreme relative prices increase the nominal adjustment frequencies of both young and old prices. Figure 6 also plots the estimated relative-price distributions for both young and old prices. Unsurprisingly, both of their modes are near the mean of zero. What is surprising is that young prices display *more* dispersion than do old prices. Comparing the two distributions' peaks makes this excess dispersion particularly clear. To quantify this, we calculated the standard deviations of both price distributions. These are 16.6 percent and 8.7 percent for young and old prices.

The finding that young prices are more dispersed than old prices comes nowhere near the standard assumption that all stores changing their prices choose the same price. To examine

Figure 7: The Fraction of Young Prices by Relative Price⁽ⁱ⁾



Note: (i) Young prices are those with ages less than four weeks. The plotted fractions exclude prices one week old from both the numerator and denominator.

this surprising result further we plot in Figure 7 the fraction of young prices as a function of the relative price. The horizontal line plots the overall fraction of young prices, 46 percent, for reference. We see that the minimum of 34 percent occurs when the relative price is close to zero and the fraction of young prices increases with the absolute value of the relative price. Young prices constitute 71 percent of the prices between 35 and 40 percent below the average of others' prices and 58 percent of the prices between 35 and 40 percent above that average. The results are similar if we exclude one-week-old prices (which might embody measurement errors) from the analysis. In a standard model, a fixed nominal price becomes extreme as inflation erodes it. In these data, stores set many extreme relative prices.

5 Forecasting Price Changes

Figure 6 goes some distance towards unifying the consideration of time and state dependence. This section continues in that direction by presenting forecasting models of the decision to change a store’s nominal price. The estimated models reinforce the findings above. Quantitatively, the price’s age contributes more to the models’ forecasts than does the relative price.

All of the models we estimate have the simple linear-in-probabilities form,

$$\Pr[p_{i,t} \neq p_{i,t-1}] = \beta' x_{i,t},$$

where $x_{i,t}$ is a vector of variables known at the time that $p_{i,t}$ is chosen.¹⁵ It includes three sets of dummy variables spanning the sets of stores, products, and calendar dates, the mean-adjusted relative price, the inverse of the price’s age, and their squares. Finally, it contains the percentage deviation of the number of units sold by the store in the previous period from its mean as well as its square. We add these to $x_{i,t}$ because Golosov and Lucas (2007) emphasize that firms’ with high sales have a greater incentive to change prices in their state-dependent model. We have also experimented with Probit and Logit specifications for this forecasting model and obtained similar results. We use the linear-in-probabilities model to take advantage of its ease of interpretation.

We estimated the linear-in-probabilities model using ordinary least squares separately for each category and for the sample as a whole. Table 6 reports the estimated coefficients for the models’ regressors of interest, their heteroskedasticity-corrected standard errors, and each model’s R^2 . To demonstrate that the key results from this exercise do not merely reflect sale prices, Table A.1 in the appendix reports the analogous estimates after first replacing each sale price with the most recent regular price.

Consider first the model estimated with all products’ data. The regressors together

¹⁵This includes the logarithmic relative price defined above, which includes the prices all other stores charge at time t . We use this for conformity with the analysis in Section 4.

Table 6: Linear-in-Probabilities Estimates⁽ⁱ⁾

Category	Relative Price		Lagged Units Sold		Inverse Age		R^2
	Linear	Squared	Linear	Squared	Linear	Squared	
All Products	6.5*** (1.2)	46.4*** (3.4)	0.1 (0.2)	-0.1 ** (0.0)	46.4*** (1.4)	-11.1*** (1.3)	15.4
Ketchup	-15.7*** (5.3)	70.6*** (23.2)	-2.0 ** (0.9)	0.0 (0.2)	54.6*** (7.5)	-24.2*** (6.6)	12.2
Margarine	7.3*** (1.7)	45.5*** (4.0)	-0.1 (0.3)	0.0 (0.0)	33.2*** (1.9)	1.8 (1.7)	16.4
Peanut Butter	-12.8 ** (5.2)	122.2*** (17.5)	-0.6 (1.0)	-0.3*** (0.1)	46.9*** (6.8)	-14.4 ** (5.9)	16.6
Sugar	9.7*** (3.8)	132.5*** (10.7)	2.3*** (0.7)	-0.4*** (0.1)	36.8*** (4.8)	-5.9 (4.5)	20.9
Tissue	-7.1 * (4.3)	104.5*** (15.1)	0.2 (0.6)	-0.1 (0.1)	39.6*** (5.1)	-14.8*** (4.6)	14
Tuna	5.9*** (1.8)	30.7*** (4.0)	-0.1 (0.4)	-0.1 ** (0.0)	76.9*** (3.4)	-36.6*** (3.1)	20

Note: (i) Each column reports estimated coefficients (in percentage points) multiplying the indicated variable. Heteroskedasticity-consistent standard errors are below each coefficient in parentheses. The superscripts *, **, and *** indicate statistical significance at the 10, 5, and 1 percent levels. Table A.1 in the appendix reports analogous estimates from data where all sale prices were replaced with the most recent non-sale price.

explain 15.4 percent of the variation in the decision to change the nominal price. As Figure 5 suggests, the regression function is convex in the relative price.¹⁶ The coefficients multiplying linear and squared terms in lagged units have approximately equal magnitude and opposite signs. The coefficient multiplying the squared inverse price age is negative but less than half the magnitude of that multiplying the corresponding linear term. Thus, the regression function strictly decreases in the price’s age.


Because the sample size varies greatly across the categories, so does the precision of the estimated coefficients. The standard errors for the two categories with the fewest price observations, Ketchup and Peanut Butter, are particularly large. For all categories, a joint exclusion test for the two relative price terms rejects the null hypothesis at the one percent level. All of the estimated coefficients on the squared term are positive, as the convex hazard in Figure 5 suggests.

The influence of lagged units sold on nominal adjustments varies more across the categories. For Peanut Butter and Tuna, only the squared term is individually significant, and it appears with a negative sign. Thus, prices that produce either extremely large or small sales tend to be left alone. For Sugar, the coefficient on the linear term is positive and significant while that on the quadratic term is small. This category’s results conform to the intuition from Golosov and Lucas (2007): A firm selling a large number of units facing a given menu cost adjusts price more frequently than a rival selling fewer units because the return to the adjustment is larger. For Ketchup though, the coefficient on the linear term is negative and significant. Overall, only the results for Sugar conform to the intuition of Golosov and Lucas (2007).

Finally, consider the coefficients multiplying the inverse of the price’s age. The coefficients on the linear term are all positive and statistically significant at the one percent level. Those

¹⁶Figure 5 might suggest that the probability of a price change is a function of relative price’s absolute value or another piecewise linear function. We have experimented with such specifications and obtained very similar results to those reported here.

Figure 8: Implied Hazard Functions from the Linear-in-Probabilities Estimation



multiplying the squared term are negative and statistically significant except for Margarine. The two coefficients are jointly statistically significant at the 1 percent level for all categories. Figure 8 plots the estimated hazard functions. We choose the intercepts so that the hazard for a newly set price equals 47 percent as in Figure 3. The hazards all decline with age.

A variable's statistical significance indicates that it has some forecasting value, but it does not show that it matters quantitatively. To assess each variable's contribution to the forecasts, Table 7 reports root mean-squared errors (in percentage points) from several specifications of the linear-in-probabilities model. The first column reports the in-sample *rmse*'s from forecasting price changes with only a constant. Their maximum possible value is 50 percent. The remaining columns report the *rmse*'s from models with progressively richer specifications for $x_{i,t}$. The second column corresponds to a model which includes only the dummy variables for the store, product, and calendar date. These variables lower the *rmse*'s from 1.0 to 2.3 percentage points. The third column gives the results from models

Table 7: Root Mean-Squared Errors⁽ⁱ⁾

Category	Store, UPC, & Date Indicators plus Quadratic in				
	Constant	\emptyset	Relative Price	& Units Sold ⁽ⁱⁱ⁾	& Inverse Age ⁽ⁱⁱⁱ⁾
All Products	43.2	41.7	41.6	41.5	39.9
Ketchup	44.4	43.4	43.3	43.2	42.2
Margarine	43.2	41.6	41.4	41.4	39.6
Peanut Butter	45.2	43.5	43.2	43.1	41.8
Sugar	40.4	38.7	37.7	37.6	36.3
Tissue	41.3	39.6	39.4	39.4	38.7
Tuna	43.8	41.5	41.5	41.4	39.6

Notes: (i) The table’s entries are in-sample root mean-squared errors (in percentage points) from forecasts of $I\{p_{i,t} \neq p_{i,t-1}\}$ based on linear-in-probabilities models that include the specified set of regressors. The maximum possible value for these is 50. (ii) The linear-in-probabilities models underlying these results include quadratic terms in the relative price and units sold. (iii) The linear-in-probabilities models underlying these results include quadratic terms in the relative price, units sold, and the inverse duration.

that add the two relative price terms, and the fourth column has results from adding the terms in lagged units sold to that specification. For all categories but Sugar, adding these variables lowers the *rmse*’s very little. For the final column, we added the two terms in the price’s inverse age to the fourth column’s specification. This yields the same definition of $x_{i,t}$ used in the original regression analysis. The *rmse*’s drop from 0.7 to 1.9 percentage points. For all categories, the price’s age is the most quantitatively useful forecaster of nominal adjustments.

6 Robustness

In this section, we document robustness to changes in measurement strategy. We begin by showing that the results hinge on neither the correction for time aggregation nor the specific

definition of a sale price. Firms ultimately control stores' prices, so we also explore the consequences of using a firm's average price for an item (across its stores) as the unit of analysis. We conclude with a brief discussion of other empirical characterizations of time dependence, which indicate that the declining hazard function manifests itself in data sets collected to create the CPI.

6.1 Measurement Choices

Three measurement decisions permeate the results: the correction for time aggregation and measurement error based on fractional prices, the definition of a sale, the decision not to replace prices that differ from a common preceding and following price. To examine the importance of the time aggregation and measurement error correction, we recomputed every table and figure after rounding all prices to the nearest whole cent and without the correction as described in Section 2. Unsurprisingly, removing the time aggregation correction increases the frequency of price adjustment and decreases the average price duration, from 23 to 34 percent and from 8.3 to 6.4 weeks. However, this modification leads to no other substantial changes. We noted in Section 2 the possibility that prices which follow and precede exactly the same price are errors. Replacing such prices with their bracketing price decreases the frequency of price adjustment to 17 percent and increases the average price duration to 13.1 weeks, but it leaves the paper's central results intact. Given the scarcity of sale prices as we define them, one might speculate that our definition is too conservative. We investigated this by adopting an alternative which required a price to drop at least 5 percent (instead of 10 percent) and fully recover within four weeks (instead of two weeks). Here also, the *only* substantial change to the results was unsurprising: the frequency of sale prices rose from 3.4 to 7.8 percent.

6.2 Firms

Chains need not charge identical prices at their stores, but they frequently do so. To examine whether our results are robust to changing the unit of observation from the store to the firm, we wish to repeat our analysis after aggregating the store-level observations into firms. Unfortunately, the ERIM scanner data does not indicate which firms own which stores, so this is not straightforward. Nevertheless some groups of stores seem to belong to a common firm, because their prices *exactly* equal each other more often than not. We use this to assign stores to individual firms with a simple rule. First, we compare the median transaction prices from the household data across all of the sample's stores. Second, we divide stores into firms to maximize the number of firms given the requirement that two stores must belong to a common firm if 50% or more of available price comparisons between them result in an exact match.

The rule divides the 42 sample stores into 9 firms in Sioux Falls and 6 firms in Springfield. The three largest firms in Sioux Falls have 6, 4, and 3 stores. The first and third firms own only groceries, while the second owns only drug stores.¹⁷ The remaining firms in Sioux Falls each have a single store. The rule identifies four multiple-store grocers in Springfield, which have 9, 5, 4, and 3 stores. We cross-checked the results for Springfield using the Southwestern Bell Yellow Pages listings for grocers from 1987, and we gained additional information for some of the listed firms from the 1987 *Chain Store Guide Directory: Supermarkets and Grocery Chains*. With these, we identified four multiple-store grocers serving Springfield: Consumers Market (10 stores), Ramey Super Markets (7 stores), Dillon Food Stores (4 stores), and Smitty's (3 stores). Apparently, either our data is missing three stores from the largest two firms or assigns them to single-store firms. In either case, we believe that this simple rule has grouped stores into firms with sufficient accuracy for the purpose of assessing robustness.

After assigning the stores to firms, we constructed firm-level average prices, corrected

¹⁷Our simple rule did *not* use information about stores' types when grouping them into firms.

them for time aggregation, and selected a balanced panel of firm-item pairs that were always sold by at least five firms. This results in a somewhat smaller sample of 67 products. In a given week, 39 percent of the firm-level prices change. This increase is unsurprising given that a change will occur if the price changes at any store. The switch to firm-level prices changes no other result meaningfully.

6.3 Other Data Sets

One might wonder whether this paper’s results (particularly the declining hazard function) only arise in the ERIM data. We conclude our robustness checks with a discussion of the existing evidence on this point. [Dhyne et al. \(2006\)](#) report that the hazard function for nominal adjustment decreases when measured with monthly *CPI* data in nearly every Euro zone country. These authors did not examine the hazard function after first controlling for heterogeneity across store-product cells, so all we can say is that existing evidence from the Euro zone is consistent with a decreasing hazard like that we document. The results of [Nakamura and Steinsson \(2007\)](#) are more supportive. They have confirmed that this paper’s finding of a decreasing hazard function holds good in U.S. data used to construct the CPI even after controlling for unobserved heterogeneity across product-store cells. Overall, the available evidence indicates that the declining hazard is not an idiosyncratic feature of the ERIM data.

7 Two Conjectures

We have used no economic theory to derive this paper’s empirical results. In this section, we explore two ex-ante plausible conjectures about how our results can be reconciled with standard models of price setting. We consider the possibility that the decreasing hazard reflects grocers’ optimal response to an unobservable feature of producer price dynamics, and we examine the hypothesis that it arises from grocers synchronizing price changes for

multiple goods. We find that neither of these conjectures accounts for our the observed decreasing hazard.

7.1 Producer Prices

This paper characterizes retail price dynamics. Suppose that producer prices also have a decreasing hazard. The producer’s price equals the retailer’s marginal cost, so periods of frequent adjustment in producer prices could lead a retailer facing menu costs to also bunch price changes together. In this case, the decreasing hazard in retail-level prices would switch to an increasing hazard if producer prices were constant; and the menu-cost model works as it should given the evolution of retailer’s marginal cost. This interpretation implies that we have documented a potentially puzzling feature of *producer* prices.


The ERIM data contain no information about producer prices, but we can test this hypothesis indirectly using the linear-in-probabilities model. The version presented in Section 5 includes dummy variables for the calendar date. These control for the effects of any category-wide changes in marginal cost, but they assign the effects of UPC-specific producer price changes to the residual. To address this, we have estimated a second version of the model with UPC-specific calendar-date dummies that can control for variation of individual products’ prices.

Figure 9 plots the resulting estimated hazard functions. These were computed following the procedure used for those in Figure 8. Just as before, the linear and quadratic terms in the inverse age are jointly significant in every category. The hazards decline less over the 12 week horizon plotted, but they still decline substantially. We cannot credibly attribute the declining hazard to producer price dynamics.

7.2 Synchronized Price Changes

Although macroeconomic models of price setting typically consider the decision to adjust a single good’s price, Lach and Tsiddon (1996) document that stores selling multiple products

Figure 9: Hazard Functions from Estimation with UPC-Specific Weekly Dummy Variables



typically change their prices simultaneously. This observation led [Midrigan \(2006\)](#) to investigate a model of price setting in which paying a single menu cost allows a firm to change multiple prices. He finds that this modification of the standard menu-cost model explains the anomalous observation of many *small* price changes that appear not to benefit the seller greatly. Small price changes always accompany a complementary large price change. [Lach and Tsiddon \(2007\)](#) verified that this prediction holds good in Israeli data.

A single-good monopolist facing a menu cost will not change a newly set price, because the benefits of changing an optimal price are zero. This implies an increasing hazard function for young prices. However, the previous success of adding economies of scope to the menu-cost model leads us to wonder if they could bring its predictions in line with this paper's observation of a downward sloping hazard. The answer clearly determines how one should interpret our results, so we address it empirically and theoretically.

To explore the relevance of price synchronization for our observations, we estimated a

version of the linear-in-probabilities model from Section 5 which includes the fraction of *other* prices at the same store changing in the same week. We expect that changing more prices in the store lowers the cost of changing a price, so its associated coefficient should be positive. Table 8 reports the estimates, and they conform to these expectations. The fraction of other prices being changed in the current week always has a positive and significant coefficient. Adding this variable substantially increases the regressions' R^2 values, but it changes the point estimates and statistical significance of the other coefficient estimates only little. The implied hazard functions (not plotted to conserve space) are very similar to those in Figure 9. We conclude that price synchronization contributes to the understanding of price changes in these data, but it does not change the declining hazard result.

We now examine whether the leading theoretical explanation for price synchronization – returns to scope in price setting – can also explain the declining hazard. For this, we calculate the solution to a simple version of the Midrigan (2006) model of a monopolist choosing multiple prices with a single menu cost. In it, a single producer sets a vector of two prices, p_t . Their optimal values are collected in p_t^* , and the firm's payoff at time t includes a quadratic loss from the deviation between them, $-(p_t - p_t^*)'(p_t - p_t^*)$. For simplicity, we assumed that $p_t^* = p_{t-1}^* + u_t$, where the elements of u_t are independent from each other and across time. Price adjustment is an all-or-nothing affair: The cost of price adjustment is $F \geq 0$ regardless of how many price change. The firm discounts future profits with the constant rate δ .

For a particular parameterization, Figure 10 plots the optimal price-adjustment rule as a function of both elements of $p_t - p_t^*$.¹⁸ It is similar to Figure 3 in Midrigan (2006). The grey dots indicate points at which price adjustment is optimal. Because p_t^* has no drift, price

¹⁸For this, we assume that the support of u_t is an equispaced grid of $2M + 1$ points with step size σ centered around zero. The probability distribution of u_t on this support is chosen to approximate a normal distribution with mean 0 and variance γ^2 . We set $M = 25$, discount profits with an annual rate of five percent applied at the weekly frequency ($\delta = 0.999$), use a step size (σ) of 0.01, set the standard deviation of u_t (γ) to 0.10, and set fixed costs to 0.025.

Table 8: Linear-in-Probabilities Estimates Accounting for Price Synchronization⁽ⁱ⁾


Category	Relative Price		Lagged Units Sold		Inverse Age		Fraction of Other	R^2
	Linear	Squared	Linear	Squared	Linear	Squared	Prices Changed	
All Products	-0.8 (1.1)	39.6*** (3.1)	0.3 (0.2)	-0.1*** (0.0)	42.6*** (1.4)	-18.3*** (1.2)	85.5*** (0.6)	27
Ketchup	-17.4*** (5.1)	56.3*** (21.2)	-1.7* (0.9)	0.0 (0.2)	49.1*** (7.2)	-25.7*** (6.4)	73.6*** (3.2)	19.4
Margarine	-5.0*** (1.6)	37.8*** (3.3)	0.0 (0.3)	0.0 (0.0)	33.5*** (1.8)	-10.2*** (1.6)	89.9*** (0.8)	29.4
Peanut Butter	-18.5*** (5.1)	109.7*** (16.7)	-0.3 (1.0)	-0.3** (0.1)	45.9*** (6.5)	-22.5*** (5.6)	80.1*** (2.7)	25.8
Sugar	6.7* (3.8)	127.7*** (10.6)	2.9*** (0.7)	-0.5*** (0.1)	25.5*** (4.6)	-4.2 (4.3)	66.9*** (2.3)	28.2
Tissue	-2.9 (4.1)	101.1*** (13.4)	0.6 (0.6)	-0.1 (0.1)	35.3*** (5.0)	-16.3*** (4.5)	62.1*** (2.2)	20.8
Tuna	0.0 (1.8)	24.7*** (3.7)	0.1 (0.3)	-0.1*** (0.0)	62.1*** (3.1)	-37.2*** (2.8)	95.1*** (1.4)	33.2

Note: (i) Each column reports estimated coefficients (in percentage points) multiplying the indicated variable.

Heteroskedasticity-consistent standard errors are below each coefficient in parentheses. The superscripts *,

, and * indicate statistical significance at the 10, 5, and 1 percent levels.

Figure 10: Optimal Price Setting in a **Midrigan**-Style Model



The grey dots mark states with price adjustment, which sets the vector of gaps to the origin.

adjustment returns the vector of gaps to the origin. There is a region of inaction around this return point that is roughly circular.

There are two things worth noting about this figure. First, small price changes are possible. If the accumulated shocks place the firm outside of the region of inaction but close to either axis, then one of the prices will adjust only little. This was the insight of **Lach and Tsiddon (2007)** further developed by **Midrigan (2006)**. Second, the vector of prices immediately after adjustment lies in the region of inaction's center, so if desired prices do not frequently jump the probability of changing a very young price should be small. The probability of a price change can only increase as the price ages if it starts at zero, so this

strongly suggests that the probability of price adjustment increases with the price’s age for very young prices. This is indeed the case. The shock to p_t^* can be large enough to induce two price changes in a row, but the hazard still increases. The model’s hazard jumps from 36 percent after one week to 50 percent and then asymptotes at approximately 52 percent. We conclude from this that returns to scope in price setting hold little promise for explaining the declining hazard.

8 Conclusion

Price setting in the ERIM scanner data displays state dependence: Increasing the difference between the current price and the average of other stores prices for the same good raises the probability of that nominal price changing. Our other two results present a greater challenge to conventional theories of costly price adjustment. Prices in the tails of the distribution are *younger* than average, and the hazard for nominal price adjustment *decreases* with time.

What do these results mean? One heuristic description is that they combine experimentation and inattention. A producer unsure of the profit-maximizing price tries several and eventually settles on one. This price remains in place until trying to improve upon it becomes worthwhile. With this in mind, we have worked with a toy model of a monopolist learning her demand curve. She sells a single product to a high-value consumer who always purchases the good and to a low-value consumer whose store visits are governed by a Markov chain. The monopolist only observes the price and the number of units sold. Her optimal behavior requires her to punctuate long stretches of high prices with occasional sales to detect the low-value consumer’s presence. If such an experiment succeeds, the low price sticks until the low-value consumer leaves. Otherwise, the price returns to its high value. It is not hard to see that this model’s hazard function initially decreases. This qualitative success leads us to believe that further development of similar models will improve our understanding of this paper’s evidence.

References

- Baharad, E. and B. Eden (2004). Price rigidity and price dispersion: Evidence from micro data. *Review of Economic Dynamics* 7(3), 613–641. [14](#)
- Barro, R. J. (1972). A theory of monopolistic price adjustment. *Review of Economic Studies* 34(1), 17–26. [1](#)
- Bils, M. and P. J. Klenow (2004). Some evidence on the importance of sticky prices. *Journal of Political Economy* 112(5), 947–985. [2](#), [12](#)
- Calvo, G. A. (1983). Staggered prices in a utility-maximizing framework. *Journal of Monetary Economics* 12(3), 383–398. [13](#), [14](#)
- Caplin, A. S. and J. Leahy (1991). State-dependent pricing and the dynamics of money and output. *Quarterly Journal of Economics* 106(3), 683–708. [1](#)
- Caplin, A. S. and D. F. Spulber (1987). Menu costs and the neutrality of money. *Quarterly Journal of Economics* 102(4), 703–726. [1](#)
- Chevalier, J. A., A. K. Kashyap, and P. E. Rossi (2003). Why don’t prices rise during periods of peak demand? *American Economic Review* 93(1), 15–37. [2](#)
- Darby, M. R., J. Haltiwanger, and M. Plant (1985). Unemployment rate dynamics and persistent unemployment under rational expectations. *American Economic Review* 75(4), 614–637. [16](#)
- Dhyne, E., L. J. Álvarez, H. Le Bihan, G. Veronese, D. Dias, J. Hoffman, N. Jonker, P. Lünneken, F. Rumier, and J. Viilmunen (2006, Spring). Price setting in the Euro area: Some stylized facts from individual consumer price data. *Journal of Economic Perspectives* 20(2), 171–192. [1](#), [30](#)

- Dotsey, M., R. G. King, and A. L. Wolman (1999). State-dependent pricing and the general equilibrium dynamics of money and output. *Quarterly Journal of Economics* 114(2), 655–690. [1](#)
- Dutta, S., M. Bergen, and D. Levy (2002). Price flexibility in channels of distribution: Evidence from scanner data. *Journal of Economic Dynamics and Control* 26, 1845–1900. [2](#)
- Golosov, M. and R. E. Lucas, Jr. (2007). Menu costs and Phillips curves. *Journal of Political Economy* 115(2), 171–199. [23](#), [25](#)
- Kashyap, A. K. (1995). Sticky prices: New evidence from retail catalogs. *Quarterly Journal of Economics* 110(1), 245–274. [2](#)
- Klenow, P. J. and O. Kryvtsov (2005). State-dependent or time-dependent pricing: Does it matter for recent U.S. inflation? NBER Working Paper # 11043. [2](#), [8](#)
- Konieczny, J. D. and A. Skrzpacz (2005). Inflation and price setting in a natural experiment. *Journal of Monetary Economics* 52(3), 621–632. [2](#)
- Lach, S. and D. Tsiddon (1992). The behavior of prices and inflation: An empirical analysis of disaggregated price data. *Journal of Political Economy* 100(2), 349–389. [2](#)
- Lach, S. and D. Tsiddon (1996, December). Staggering and synchronization in price-setting: Evidence from multiproduct firms. *American Economic Review* 86(5), 1175–1196. [31](#)
- Lach, S. and D. Tsiddon (2007). Small price changes and menu costs. *Managerial and Decision Economics Forthcoming*. [32](#), [35](#)
- Midrigan, V. (2006, August). Menu costs, multi-product firms, and aggregate fluctuations. Unpublished Working Paper, Federal Reserve Bank of Minneapolis. [32](#), [33](#), [35](#)
- Nakamura, E. and J. Steinsson (2007, January). Five facts about prices: A reevaluation of menu cost models. Harvard University. [30](#)

- Richardson, D. H. (2003). Scanner indexes for the consumer price index. In R. C. Feenstra and M. D. Shapiro (Eds.), *Scanner Data and Price Indices*, Volume 64 of *Studies in Income and Wealth*, Chapter 2, pp. 39–65. University of Chicago Press. [10](#)
- Sheshinski, E. and Y. Weiss (1977). Inflation and costs of price adjustment. *Review of Economic Studies* 44(2), 287–303. [1](#)
- Taylor, J. B. (1980). Aggregate dynamics and staggered contracts. *Journal of Political Economy* 88(1), 1–23. [14](#)

Table A.1: Linear-in-Probabilities Estimates using Sale-Replacement Records⁽ⁱ⁾

Category	Relative Price		Lagged Units Sold		Inverse Age		R^2
	Linear	Squared	Linear	Squared	Linear	Squared	
All Products	6.4*** (1.1)	36.4*** (2.5)	-0.7*** (0.2)	0.0 (0.0)	27.2*** (1.4)	7.0*** (1.3)	15.1
Ketchup	-5.6 (5.2)	49.5*** (18.6)	-1.9*** (0.7)	0.1 (0.1)	30.4*** (6.7)	-0.2 (6.2)	13.1
Margarine	9.9*** (1.6)	33.3*** (3.6)	-0.7*** (0.2)	0.0 (0.0)	22.1*** (1.8)	11.3*** (1.7)	15.2
Peanut Butter	-15.8*** (5.1)	90.7*** (14.4)	-0.9 (0.7)	-0.2*** (0.1)	29.7*** (6.3)	1.7 (5.7)	15.6
Sugar	-14.3*** (4.9)	70.7*** (13.7)	1.4*** (0.5)	-0.2*** (0.1)	12.1*** (4.6)	21.1*** (4.5)	18.7
Tissue	-7.2* (4.0)	80.2*** (18.4)	-1.0** (0.4)	0.1 (0.0)	18.0*** (4.5)	7.8* (4.3)	13.7
Tuna	8.0*** (1.9)	31.2*** (3.3)	-0.8*** (0.3)	0.0 (0.0)	40.4*** (3.2)	-2.1 (3.0)	20.6

Note: (i) Each column reports estimated coefficients (in percentage points) multiplying the indicated variable.

Heteroskedasticity-consistent standard errors are below each coefficient in parentheses. The superscripts *,

, and * indicate statistical significance at the 10, 5, and 1 percent levels.

Table A.2: Root Mean-Squared Errors using Sale-Replacement Records⁽ⁱ⁾
Store, UPC, Date, & Sale Indicators plus Quadratic in

Category	Constant	\emptyset	Relative Price	& Units Sold ⁽ⁱⁱ⁾	& Inverse Age ⁽ⁱⁱⁱ⁾
All Products	40.7	39.5	39.5	39.5	37.7
Ketchup	41.2	40.4	40.4	40.2	38.9
Margarine	41.4	40.1	40.1	40.0	38.3
Peanut Butter	42.9	41.5	41.4	41.3	39.8
Sugar	36.4	35.0	34.9	34.8	33.1
Tissue	36.6	35.3	35.3	35.2	34.3
Tuna	40.8	38.9	38.9	38.9	36.7

Notes: (i) The table's entries are in-sample root mean-squared errors (in percentage points) from forecasts of $I\{p_{i,t} \neq p_{i,t-1}\}$ based on linear-in-probabilities models that include the specified set of regressors. The maximum possible value for these is 50. (ii) The linear-in-probabilities models underlying these results include quadratic terms in the relative price and units sold. (iii) The linear-in-probabilities models underlying these results include quadratic terms in the relative price, units sold, and the inverse duration.