Menu Costs and Firm Size: Evidence from Consumer Packaged Goods*

Tsolmon Otgon†
Georgetown University

September 25, 2025

Click here for the latest version.

Abstract

I explore the magnitude and heterogeneous effects of menu costs on pricing dynamics. To do so, I develop a dynamic discrete choice model of a firm facing menu costs when adjusting prices, using methods from industrial organization. I obtain descriptive evidence that is consistent with the model using scanner data from grocery stores in the US. I then estimate the model and find that the menu cost is substantial, amounting to 7.2% of overall retail profit at the average-sized retailer. When applied to a high-inflation scenario, the model generates an increase in both frequency and magnitude of price changes. In addition, as the menu cost varies little with the retailers' size, smaller retailers face a higher burden of the menu costs. Being unable to adjust prices frequently, they exhibit earlier and more pronounced responses to anticipated cost increases.

JEL codes: L11, L14, E31
Keywords: menu cost, price rigidity, retailers

*I am grateful to my advisor, Nathan Miller, for invaluable guidance and support. I thank Anita Rao, John Rust, and Louise Laage for their advice and encouragement. All errors are my own. Researcher's own analyses calculated (or derived) based in part on data from Nielsen Consumer LLC and marketing databases provided through the NielsenIQ Datasets at the Kilts Center for Marketing Data Center at The University of Chicago Booth School of Business. The conclusions drawn from the NielsenIQ data are those of the researcher(s) and do not reflect the views of NielsenIQ. NielsenIQ is not responsible for, had no role in, and was not involved in analyzing and preparing the results reported herein.

†Email: to292@georgetown.edu

1 Introduction

Menu costs—the frictions firms face when adjusting prices—are a leading explanation for the widely observed price rigidity. Such rigidity is a powerful mechanism that amplifies demand shocks, including monetary policy shocks, affecting output (Nakamura and Steinsson 2013, Blanco et al. 2024b). Beyond the general prevalence of sticky prices, research has documented substantial heterogeneity in the degree of rigidity across firms, depending on their size and market power (Gertler and Gilchrist 1994, Ottonello and Winberry 2020). Understanding the magnitude of menu costs and how these costs differentially shape the behavior of firms of varying sizes is therefore essential not only for effective price stabilization but also for broader economic policy design.

In this paper, I examine the trade-offs that firms face when adjusting prices in response to demand and cost conditions and derive implications for price dynamics. I employ industrial organization research methods to explore the magnitude and heterogeneous effects of menu costs in a specific setting, the consumer packaged goods (CPG) retail sector. Leveraging institutional knowledge of this sector and high-frequency detailed price and sales data, I estimate the value of menu costs and how these costs vary with retailer size.

I develop a dynamic discrete choice model of retail price setting with menu costs. The model has a finite dynamic structure in which a firm chooses an initial retail price.[1] Over time, the chosen price may become suboptimal due to stochastically evolving demand and marginal cost states, both governed by first-order Markov processes. In each period, the firm decides whether to maintain the existing price or adjust it. Although price updates are permitted in any period, doing so requires paying the menu costs, which generates price rigidity while the state variables are changing.

I apply the model to the consumer packaged goods (CPG) retail sector, a large industry that supplies households with thousands of daily-use products. This sector is particularly well-suited for the analysis because it provides rich, high-frequency scanner data on prices and sales. Moreover, the existing research literature offers valuable insights into its institutional background (e.g., Anderson and Fox 2019), which helps in structuring the theoretical model. For example, annual planning cycles are common in the sector, with contracts taking effect at the beginning of the year.

From the broad set of consumer packaged goods, I focus on prices and sales of

1. The firm maximizes the joint surplus, which is allocated between the manufacturer and the retailer according to a contractual rule. The specific mechanism of surplus division lies beyond the scope of this study.

oatmeal. This category is particularly suitable for the analysis because it is dominated by a single manufacturer, the Quaker Oats Company, which allows me to abstract from complex dynamic oligopolistic interactions and focus on the magnitude of menu costs. Moreover, Quaker's widespread presence across retail chains enables me to study how these costs vary with retailer size. The analysis concentrates on changes in non-promotional prices, which I recover from observed weekly average price data.[2] I find that in 73% of cases, prices are adjusted at most once per year. The timing of price changes across different geographic markets within the same retail chain is highly correlated, consistent with centralized decision-making at the chain level. Finally, I document that retailers with higher sales volumes update prices more frequently.

The estimation of the theoretical model parameters proceeds in two steps. First, I estimate a logit demand model for Quaker products as a function of the company's flagship products and calculate demand shifters. Identification of menu cost relies on the difference in expected discounted profits of the firm with the existing price vector and with the optimal price vector. I exploit variation in oatmeal sales across different retail chains to analyze the relationship between the retailers' size and the menu costs, as it might be more costly to update the prices at larger retailers. The marginal cost, on the other hand, is imputed from the value of the price chosen in a market when a price is updated. The model is then estimated using a maximum likelihood approach following Rust (1987).

Estimation results indicate that the menu costs are substantial, amounting to 7.2% of the annual profit generated at an average-sized retail chain. However, I do not find evidence that the magnitude of menu costs differs based on the retail chain size. As a result, the ratio of menu cost to payoff is smaller for larger retail chains, which leads to more opportunities to recover the menu cost. Therefore, the prices are updated more often at larger retail chains, leading to less rigid pricing behavior compared to smaller chains. Additionally, menu costs result in a lower probability of observing price adjustments close to the retail contract renewal time.

In a counterfactual simulation, I analyze a high-inflation scenario in which the firm's marginal cost increases by a constant monthly rate. The firm can observe the marginal cost at the start of each month, but does not expect further cost increases. Compared to a low-inflation scenario, this experiment generates an increase in both the magnitude and frequency of price changes. This result is relevant for macroeconomic research, as the standard menu cost models cannot generate simultaneous

2. As promotional schedules are set at the start of the year and vary little from year to year, this variation of prices is not as informative about the market conditions (Anderson and Fox 2019).

increases in both the magnitude and frequency of price changes.

Next, I simulate a one-time permanent increase in marginal cost, comparing scenarios in which the shock is unexpected versus anticipated. When the shock is anticipated, the timing of price adjustments varies systematically with retailer size. On average, smaller retailers raise prices two to three months before the cost increase takes effect, whereas larger retailers typically raise prices after the shock occurs. Because smaller firms have limited opportunities to adjust prices, they use available chances to increase prices, even in advance of the shock. They also exhibit larger price increases relative to the baseline (no cost shock) than larger retailers. These earlier and stronger responses to anticipated cost shocks are consistent with empirical evidence in the macroeconomic literature, beginning with Gertler and Gilchrist (1994).

Finally, I examine the welfare implications of reducing menu costs. Lower menu costs increase overall welfare, with firms' profits benefiting the most due to more frequent price adjustments. More frequent adjustments allow prices to better match demand, which also generates a modest improvement in consumer welfare, even though the direct burden of menu costs falls on firms.

This research contributes to several strands of the economic literature. First, it adds to the work on estimating menu costs using micro-level data. Early studies relied largely on direct observations of the price-setting process within stores (e.g., Levy et al. 1997, Zbaracki et al. 2004). A smaller literature has employed dynamic discrete choice models with (S, s)-type rules[3] to estimate menu costs. For instance, Slade (1998) uses retail price and sales data on saltine crackers in a small town, while Aguirregabiria (1999) estimates a dynamic model incorporating inventory costs and menu costs for a Spanish retail chain. However, both studies treat promotional pricing as equivalent to regular price changes, which may lead to underestimating menu costs. Stella (2014) instead estimates bounds on menu costs using multiple specifications and inequality restrictions. The present paper advances this literature in several directions. I incorporate institutional features of the consumer packaged goods sector when developing the model. I restrict attention to changes in regular (non-promotional) prices, since promotional schedules are pre-planned and vary little across years. Moreover, because price changes across markets belonging to the same retail chain exhibit a strong positive correlation, I model pricing decisions at the retail-chain level. Finally, by exploiting data from multiple retail chains, I provide point

3. First introduced in Barro (1972), this rule implies that adjustments occur when a key metric—such as the real price—falls below s or rises above S.

estimates of menu costs and document how their magnitude varies systematically with retailer size.

Second, this paper relates to the literature on micro-level price adjustment, which has documented the prevalence of "uniform" prices across broad geographical areas (e.g., DellaVigna and Gentzkow 2019, Gagnon and López-Salido 2020, Arcidiacono et al. 2016, Tabanakov, Goli, and Chintagunta 2024, Hitsch, Hortacsu, and Lin 2019). Non-promotional retail prices are highly sticky and show limited responsiveness to both demand and cost shocks. For example, Gagnon and López-Salido (2020) and Basker and Noel (2009) find that retail prices react only modestly to demand shocks such as natural disasters or competitor entry. Similarly, Butters, Sacks, and Seo (2022) show that local cost shocks often trigger delayed adjustments that occur at the chain level rather than in the directly affected region. Other studies highlight systematic pricing patterns, such as the prevalence of price endings in 9 cents (Knotek II 2024) and adjustments in round-number increments (Conlon and Rao 2020). This project contributes to this literature by showing that the probability of price changes varies systematically with the timing of retail contract renewal.

Third, this research connects to the literature on heterogeneous firm responses to economic shocks. Beginning with Gertler and Gilchrist (1994), a long line of studies has documented that firms' reactions vary systematically with characteristics such as size, market power, and age (e.g., Crouzet and Mehrotra (2020), Duval et al. 2024). A large body of work seeks to rationalize this heterogeneity. For example, Ottonello and Winberry (2020) show that firms with higher default risk tend to respond more strongly to demand shocks, while Hsieh and Rossi-Hansberg (2023) and Ganapati (2025) demonstrate how large fixed investments in technology and capacity shape firms' differential reactions. This paper contributes to this literature by highlighting how the relative importance of menu costs influences pricing decisions across firm sizes. Since menu costs vary little with retailer size, smaller retailers—who face fewer opportunities to adjust prices—respond earlier and more sharply to cost shocks than their larger counterparts.

The rest of the paper is organized as follows: Section 2 describes the empirical context, data processing, and motivating evidence; Section 3 presents the theoretical model featuring menu costs; and the following section explains the estimation strategy and results. The counterfactual exercises are reported in Section 5, and the final section concludes.

2 Data and Background

2.1 Empirical Context

2.1.1 Institutional Features

I apply a dynamic discrete choice model of pricing to the consumer packaged goods (CPG) sector. The CPG sector is a major part of the economy, supplying most of the products that households consume on a daily basis, from food to personal care items. Beyond its economic significance, it offers exceptionally rich, high-frequency data on prices and sales at the product level.

A common practice in this sector is joint annual planning between the manufacturers and retailers (Anderson and Fox 2019). The retail plans are finalized at the beginning of each year, specifying agreed-upon retail prices, promotional calendars, and other terms and conditions. While firms retain the option to adjust retail prices as market conditions evolve during the year, doing so entails substantial costs—including the replacement of shelf tags, as well as administrative and logistical expenses. As a result, mid-year price adjustments are relatively rare.

Prices of individual products in the CPG sector typically consist of a stable “base" price combined with temporary price promotions. The frequency and depth of promotions reflect the retailer’s pricing strategy—for example, “hi-lo pricing” versus “everyday-low-price”, and are determined at the start of the year. Promotional schedules vary little from year to year and are finalized well in advance; for instance, promotional plans are typically locked in 12–16 weeks before the promotion starts (Anderson and Fox 2019). Although economic research has often regarded promotions as equivalent to price adjustments, in practice, they represent pre-planned activities rather than flexible responses to contemporaneous market conditions.

2.1.2 Quaker Oats Company

As the analysis in this paper centers on the magnitude of menu costs, I abstract from modeling dynamic oligopolistic competition. Instead, I focus on a specific product market—oatmeal—which is characterized by a dominant manufacturing firm. As of 2018, Quaker was the largest oatmeal producer in the United States, accounting for approximately 70% of oatmeal volume sold in grocery stores. The company has a long history: the Quaker trademark was first registered in 1877, and in 1888, Quaker joined six other American oat millers to form the American Cereal Company, which was renamed the Quaker Oats Company in 1901 (The Quaker Oats Company 2025).

Over the subsequent decades, Quaker expanded its product portfolio to include breakfast cereals, snacks, and grits, while continuing to hold a dominant position in the U.S. oatmeal market.

Understanding the oatmeal production process is essential for identifying instrumental variables for oatmeal prices. The production begins with the purchase of raw oats from farmers, followed by cleaning and sifting using specialized machinery (How Products Are Made 2025). The oats are then softened through steaming and separated from their hulls. Next, the oatmeal is cut or rolled into flakes and roasted to enhance flavor and texture. The final steps involve measuring, vacuum-sealing, and packaging the product for distribution. Most of Quaker's oatmeal is processed in Cedar Rapids, Iowa, which hosts the largest oat mill in the world (Cedar Rapids Tourism Office 2025).

2.2 Data Processing

The empirical analysis in this paper is based on the NielsenIQ Retail Scanner (RMS) data for oatmeal, specifically the "Cereal – Hot" product module within the "Grocery" department. The RMS data provides weekly store-level observations on average transaction prices and quantities sold for each Universal Product Code (UPC). Combining the sales of each UPC to calculate the total sales in each market is straightforward. However, identifying the prices and markets involves several challenges.

Price data at the UPC level is essential for analyzing price rigidity, as it enables precise tracking of changes in "base" (non-discounted) prices. However, the sheer number of oatmeal products in the sample renders structural model estimation computationally infeasible. Aggregating UPCs to the brand level is not a suitable alternative, since doing so risks conflating genuine price changes with shifts in the mix of UPCs sold within a brand. To address this, I focus on the Quaker Oats Company's flagship products—the 42-oz Old Fashioned and 42-oz Quick Oats. These two items are the company's top sellers nationally, together accounting for more than half of total oatmeal sales. Moreover, the prices of other Quaker products tend to move proportionally with those of these flagship products.

Second, a single retail chain can operate thousands of stores and employ an area-based pricing strategy, which is in line with Tabanakov, Goli, and Chintagunta (2024). The prices at multiple stores belonging to the same retail chain in a given area tend to be uniform, and identifying these areas helps to significantly increase computational efficiency. To do this, for each retail chain, I find the stores with high price similarity and group them into clusters. To measure similarity in price movements, I compute the Euclidean distance between the weekly prices of the Quaker's flagship products at each pair of stores, a and b:

$$\text{Dist}_{ab} = \left[\sum_{t=1}^{T} (p_{at} - p_{bt})^2 \right]^{\frac{1}{2}} \tag{1}$$

Using these pairwise distance values, I apply Ward's minimum variance method (Kaufman and Rousseeuw 2009) to perform hierarchical clustering. This approach iteratively merges the pairs of stores that result in the smallest increase in within-cluster variance.[4] Although the clusters are identified solely based on price data, the stores within each cluster tend to be concentrated in distinct geographical areas. I define each of these areas as a separate market for the retail chain. As the stores are clustered for each retail chain separately, the geographical market borders of different retail chains do not necessarily coincide.

After forming store clusters for each retail chain, I observe that prices within a cluster are highly similar but exhibit some small variation. This variation arises because the Retail Scanner data reports weekly average prices, which can fluctuate depending on the timing of promotional events. For example, if a promotion begins midweek, the share of products sold at the discounted price can vary from store to store, leading to small differences in weekly average prices. Since the objective is to identify changes in non-promotional prices, it is important to filter out this noise. Therefore, within each market (retailer–area pair), I select the store with the highest total oatmeal sales to represent the price in the area.[5]

Finally, I combine weekly data into a monthly frequency, and for standardization, item sizes are converted into the number of oatmeal portions based on the standard serving size indicated on Quaker Oats packaging: one portion is defined as 1.41 ounces (equivalent to 0.5 cup) of dry oatmeal.[6] With these filters, I obtain weekly price and sales data from 2011 to 2018 for 21 different retail chains.

4. I use the linkage function with the "ward" method from the `scipy.cluster.hierarchy` module in Python.

5. To further ensure consistency, I impose several additional restrictions on the selected stores. First, the store must have continuous data coverage throughout the sample period. Second, I exclude stores that underwent ownership changes, as such transitions lead to significant shifts in pricing strategy. Finally, I omit retail chains that rely heavily on frequent or coupon-based discounts, as their pricing patterns tend to be particularly noisy.

6. Unit prices are calculated by dividing the reported price by `prmult`, which records the number of units in a pack.

2.2.1 Identifying Base Price Changes

The NielsenIQ RMS database reports the weekly average transaction price of each product without distinguishing between regular and promotional periods. As a result, the time series graph of a product price at a given store exhibits a stable baseline (base or regular price) and periodic drops due to price promotions. After each promotion, the prices generally return to the baseline levels, and the base prices often remain stable for longer than a year.

As shifts in the underlying regular prices are the main focus of this paper, it is essential to recover them from the transactional price data. For this, I first identify periods with price promotions following Mansley (2022). A period is labeled as a promotional period if it meets the following criteria:

$$p_{nt} < \max_{\tau \in \{1,\dots,T\}} \left\{ k p_{nt-\tau} \right\} \text{ AND } p_{nt} < \max_{\tau \in \{1,\dots,T\}} \left\{ k p_{nt+\tau} \right\} \tag{2}$$

Here, p is the weekly average transactional price of the product in market n at period t. T stands for the maximum allowed discount length in weeks, and $k \in [0, 1]$ is a multiplier that controls the sensitivity of the filter. For instance, if $k = 0.95$, changes with a magnitude higher than 5% of the price will be considered as discounts. In the following analysis, I set $T = 9$ to identify the price promotions up to 9 weeks in length.[7] I also set $k = 1$ to identify any temporary price drop as a price promotion.

I assume that prices observed during non-promotional periods are equal to the base prices and that they remain unchanged during the promotional events.[8] With this assumption, I recover the base prices from the observed store-level weekly average pricing data (green line in Figure 1). For the remainder of the paper, I define a price update as a change in the base (non-discounted) price, excluding temporary discounts from the analysis of price changes for the reasons outlined in Section 2.1.1. I also create an indicator variable showing the periods when the base price shifts occur, a_{it}^{obs}.[9] I aggregate the weekly prices into a monthly frequency by choosing the last

7. The majority of observed price promotions last no more than five weeks. However, some promotions are relatively shallow in depth but persist for longer durations. These extended promotions are likely driven by coupon-based discounts or membership-specific pricing.

8. The NielsenIQ RMS data occasionally records isolated price spikes that last for a single week. These are not treated as changes in base prices, as they are not accompanied by corresponding drops in sales volumes. Such anomalies are likely due to recording or processing errors and are excluded from the analysis.

9. Although the price changes are implemented at the end of the calendar year, such adjustments are unlikely to happen precisely on December 31st. In practice, price changes can occur several days before or after the end of the year. To account for this, I treat any price change observed after December 20th as an "end-of-year price change" and move it to January of the following year.

Figure 1: Example of Recovered Base Prices

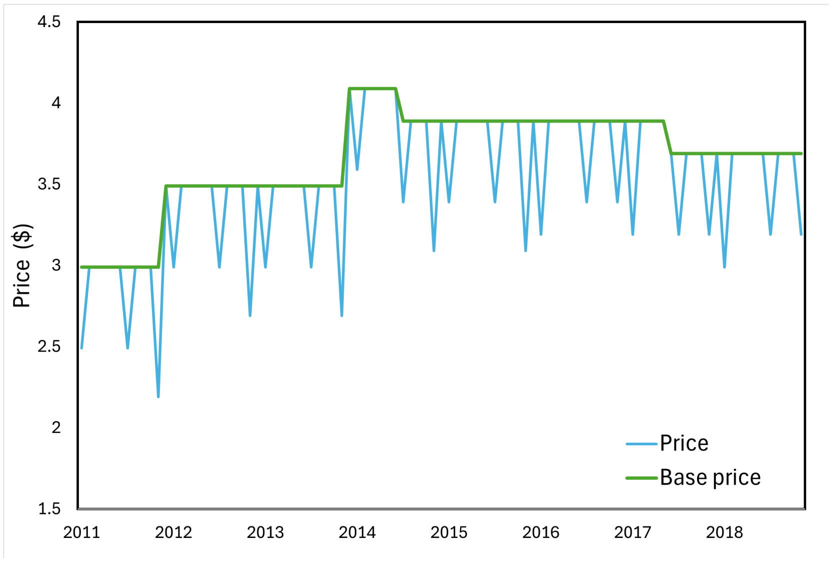

Notes: The graph visualizes the base prices (green) recovered from the weekly average transactional prices (blue) with simulated product data in a hypothetical market. The blue line shows that the average weekly price can have short-term drops due to price promotions but exhibits a steady level reflecting the "base" price that can remain unchanged for several years.

observation of each month.

2.3 Demand Estimation

The utility function of consumer i at market n at time t:

$$u_{qint} = \begin{cases} \alpha_n + \beta p_{qnt} + \gamma Dcnt_{qnt} + \xi_{qnt} + \epsilon_{qint}, & \text{if } q = 1 \\ \epsilon_{qint}, & \text{if } q = 0 \end{cases} \tag{3}$$

Here, $q = 1$ denotes Quaker's oatmeal and $q = 0$ is the outside option not to purchase Quaker's oatmeal, p_{qnt} is the price of Quaker's flagship product per portion, and $Dcnt_{nt}$ is the weighted average (by volume of oatmeal sold) percentage price promotion of the Quaker products, and α_n captures market-level fixed effects.

The market sizes are calculated following the procedure employed in Miller and Weinberg (2017) with population data obtained from the U.S. Census Bureau[10], I find the share of all Quaker's oatmeal products, s_{1nt}, in market n at time t. The total volume of oatmeal sold by Quaker is represented by the volume of the flagship

10. Available at https://www.census.gov/data/datasets/time-series/demo/popest/2010s-counties-total.html

Table 1: Demand Regression Results

	OLS	2SLS
Constant	-1.4	
	(0.07)	
Price	-0.09	-9.34
	(0.28)	(1.1)
% Discount	3.65	4.53
	(0.15)	(0.2)
N	4884	4884
Adjusted R^2	0.422	0.433
Market-level Effects		Yes

Notes: The table shows results of logit estimation of the relative market share of Quaker based on $ base price (per portion) of the flagship product and average percentage price promotions observed for Quaker products. The difference between the OLS and 2SLS estimates suggests endogeneity issues mitigated by the instrumental variables approach. Table A.1 in Appendix A shows the first-stage estimation results.

products scaled up to match the sample average of Quaker's market share in each market. Assuming that ϵ_{qint} follows Type I Extreme Value distribution, I estimate the following model:

$$ln\left(\frac{s_{1nt}}{s_{0nt}}\right) = \alpha_n + \beta p_{1nt} + \gamma Dcnt_{1nt} + \xi_{1nt} \tag{4}$$

Table 1 shows the results of OLS and 2SLS regressions of Equation (4). The price parameter estimated by 2SLS suggests that the price elasticity of Quaker's oatmeal demand is around -1.9 in a median-sized market. The results of the first-stage regression of base prices on the instrumental variables, along with descriptions of the data sources, are presented in Table A.1 in the Appendix.

As noted in Section 2.1.1, price promotions are determined at the beginning of the year and cannot be adjusted frequently in response to short-term fluctuations. Based on this, I treat the discounts as an exogenously determined variable. In contrast, since prices are endogenously determined, I employ an instrumental variable approach to address potential endogeneity. I use cost shifters as instruments, including the price of oats; retail electricity price in the industrial sector in Iowa; wages in the food manufacturing sector in Iowa, where Quaker processes the majority of its oatmeal; commercial sector wages in the geographic areas where the retail store is located, and prices of diesel and cardboard. In addition, I incorporate Hausman-type instrumental variables (Hausman 1994), which are valid under the assumption that marginal cost shocks are common across markets. This is a reasonable assumption in this context with a focus on a single manufacturer. Using the results of 2SLS estimation, the demand shifter is calculated as:

$$\hat{d}_{nt} = ln\left(\frac{s_{1nt}}{1 - s_{1nt}}\right) - \hat{\beta}p_{1nt} \quad (5)$$

For the structural estimation of the supply-side model, the demand shifter $\hat{d}_{nt}$ in (5) is discretized into two bins for each market: low (below the mean base price observed in the market) and high (above the market average). Each bin is assigned a value that equals the median of $\hat{d}_{nt}$ within the bin. Based on the observed transitions in demand states over time, I compute the corresponding demand transition probabilities for each market.

2.4 Motivating Evidence

This section presents descriptive statistics from the data sample. I find that price changes are infrequent and are more commonly observed at the beginning or middle of the year. I then employ a reduced-form approach to document a positive relationship between the frequency of price changes and sales.

2.4.1 Descriptive Statistics

Figure 2 displays the distribution of annual price changes implemented by retail chains. In most cases, prices are adjusted at most once per year: 37% of observations involve a single price change, while 36.1% show no price change at all within the year. The prices at a given retailer tend to remain stable for extended periods, often spanning several months to multiple years. On average, a retailer makes 1.14 price changes per year, and the monthly probability of a price change varies considerably across retailers, ranging from 1% to 19.1%.

Although price levels may vary across different geographical areas that belong to the same retailer, the timing of price adjustments appears to be centralized at the whole retail chain level. For a given chain, I find a strong positive correlation in the timing of price changes across the different areas identified in Section 2.2. This pattern does not hold, however, for retail chains with distinct banner names that are owned by the same parent company, suggesting that pricing decisions are made independently for each retail chain.

Figure 3 presents the distribution of dollar amounts for price changes, conditional on a price update. Price adjustments commonly occur in increments of 10

Figure 2: Number of Price Changes per Year

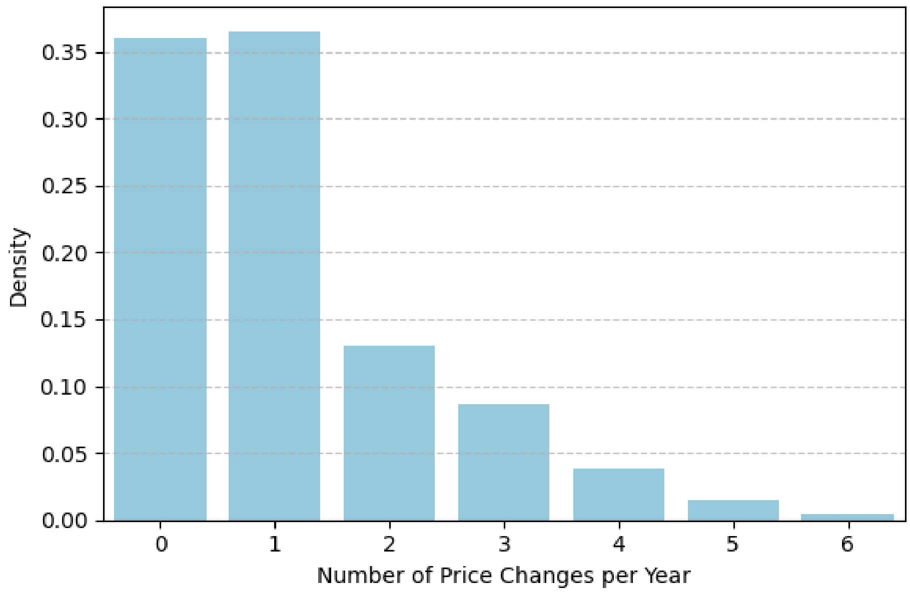

Notes: The figure shows the number of price changes in a year at a single retail chain. The numbers are based on the 9-year sample of 21 retail chains.

Figure 3: Distribution of Observed Price Changes

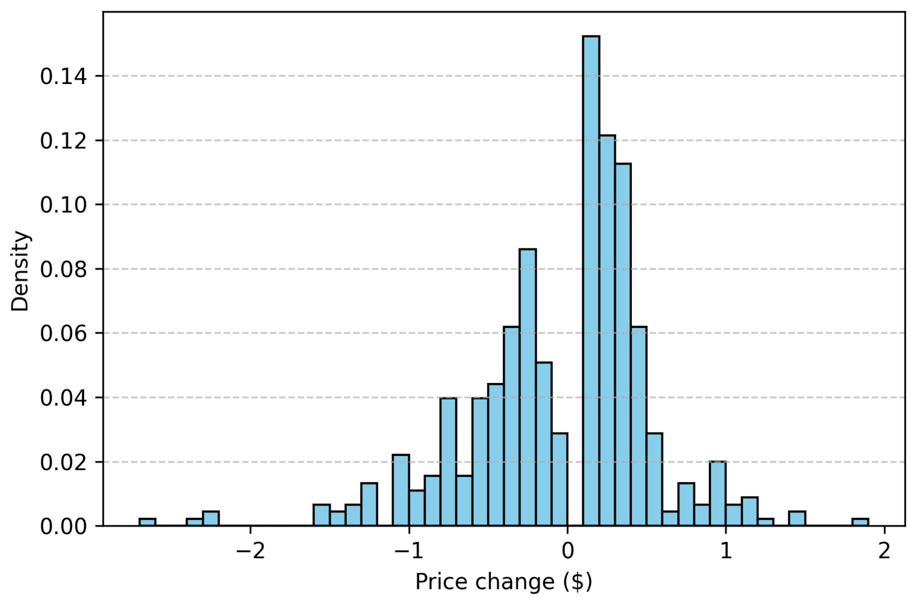

Notes: The figure presents the distribution of price change magnitudes, conditional on a price adjustment taking place. The observations are at the retail chain and geographic area level.

cents, consistent with findings in the retail pricing literature such as Conlon and Rao (2020).

Figure 4: Price Change Probability by Months

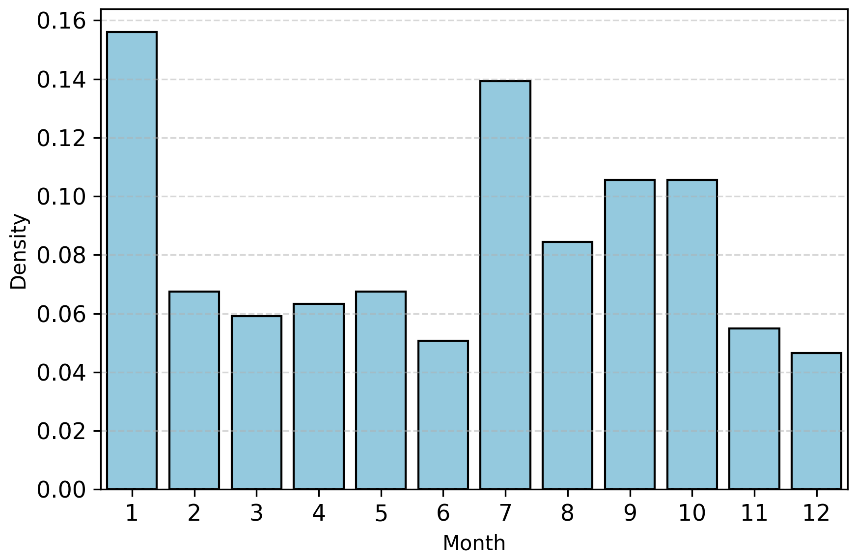

Notes: The chart displays the monthly frequency of base price changes throughout the year. Each observation is at the retail chain level, where a price update in any geographic market is treated as a pricing decision made by the entire chain. The numbers are based on the data of 21 different retail chains over the period of 9 years.

Figure 4 shows the monthly frequency of price changes observed in the sample. The beginning of the year is the most common period for price adjustments, with approximately 15.8 percent of all observed price changes occurring in January. This pronounced spike is consistent with findings from the marketing literature on pricing practices in the consumer packaged goods retail sector (Anderson and Fox 2019). Additionally, price updates are more likely to occur in the middle of the year, particularly between July and October.

Another notable observation is that price changes appear to be more elaborate depending on the size of the retail chain. At a small local retailer, prices of various Quaker product groups tend to be updated simultaneously by the same magnitude. For instance, prices of all Quaker products are increased by 20 cents. In contrast, at a larger national retail chain, although the prices are updated simultaneously, the price changes can vary across products: while prices of some products are reduced, others remain unchanged or are increased. Although a detailed investigation of the mechanisms is beyond the scope of the present study, it represents an intriguing direction for future research.

Table 2: Regression Results

	Number of price changes per year		Number of non-January price changes	
	(1)	(2)	(3)	(4)
Constant	1.02	1.03	0.85	0.86
	(0.15)	(0.16)	(0.14)	(0.15)
Volume (mln.)	0.006	0.005	0.005	0.005
	(0.001)	(0.001)	(0.001)	(0.001)
N	189	189	189	189
R^2 (Overall)	0.02	0.02	0.02	0.02
Time Effects	-	Yes	-	Yes

Notes: The table reports the results of a panel regression examining the relationship between the number of price changes for Quaker's flagship product at a retail chain (per year) and the annual volume of Quaker products sold at the retailer (measured in millions of portions sold per month). I find a positive and significant relationship between the sales volume and the total number of price changes observed in a year (Columns 1-2) and the number of mid-year price changes (Columns 3-4). Standard errors are clustered at the retail chain level.

2.4.2 Motivating Regressions

In this section, I examine whether the varying frequency of price updates is associated with the size of retail chains. It is possible that adjusting the price at a larger retailer where the manufacturer generates higher revenue yields more gains. Using panel data covering 21 retail chains over a nine-year period between 2011 and 2018, I regress the number of observed annual price changes for Quaker Oats Company's flagship oatmeal products on the sales volume of Quaker products.

Table 2 shows a statistically significant positive relationship between the number of portions sold and the frequency of price changes in a year (Columns 1-2). I then focus specifically on price changes occurring outside of January (Columns 3-4), which reflect the mid-year pricing decisions. I also find a statistically significant positive correlation between sales volume and the frequency of mid-year price changes. These results support the claim that decisions to update prices depend on the revenue Quaker receives from a retail chain.

3 The Model

I develop a formal model of price setting in the presence of menu costs.[11] A manufacturer and retailer are choosing a retail price to maximize the joint surplus from the sales at a retail chain.[12] As the retail contracts are created and renewed on an annual basis, according to research literature on the consumer-packaged goods retail sector, the model has a finite dynamic structure with 12 periods, each representing a month.

At the start of the year, a vector of retail prices is chosen for the different markets the retail chain operates in. The market is defined as a retail chain and geographical area combination. As the year progresses, the chosen retail prices can diverge from the optimal level due to stochastic demand and marginal cost. Each month, the firm has an option to update the price, but to do so, it must pay a fixed menu cost.[13] I assume that the cost of price updates is fixed across time and does not depend on the magnitude of the price changes. However, menu costs may vary across different retailers depending on characteristics of the retail chain, such as the total volume of products sold. This formulation allows for heterogeneity in contractual frictions across retailers, which can influence the frequency and timing of within-year price adjustments.

Each retail market is characterized by distinct demand states and their associated transition probabilities. The marginal cost reflects both the manufacturer's and the retailer's costs and is assumed to be identical across all markets. This assumption is reasonable given the presence of a single manufacturer. The realizations of the stochastic states are observed by the firms at the beginning of each period, and any price updates take effect immediately.

3.1 Dynamic Problem

The retailer operates as a monopolist in N different markets, each with a different size, demand states, and associated transition probabilities. At the start of every period, a manufacturer and retailer pair can observe the vector of demand states in

11. Menu costs in the model include various types of costs that can occur when adjusting retail prices: physical cost of changing price tags, as well as administrative and logistical costs.

12. The joint surplus is then split between the manufacturer and retailer following a pre-determined rule, which is beyond the scope of this paper.

13. The model assumes that price adjustments made at the time of contract renewal, at the beginning of the year, do not incur additional costs. However, in practice, changing prices, even at the start of the planning period, can involve some costs, such as updating price tags and revising retail plans. This cost can be included in an extended version of this model and estimated by leveraging the observed frequency of price updates in January.

each market $\mathbf{d}_t = (d_{1t}, ..., d_{Nt})$ and the marginal cost mc_t, which is the same in all the markets. The joint payoff from sales π_t is given as:

$$\pi(\mathbf{d}_t, mc_t, \mathbf{p}_t) = \sum_n (p_{nt} - mc_t) s(p_{nt}, d_{nt}) M_n \tag{6}$$

where p_{nt} is the price in market n, mc is the marginal cost, which I assume is constant across markets, and quantity sold is decomposed into market share of the firm, $s(p_{nt}, d_{nt})$, and the market size, M_n. Following the consumer utility function in (3), the market share of the manufacturer $s(p_{nt}, d_{nt})$ is

$$s(p_{nt}, d_{nt}) = \frac{d_{nt} + \beta p_{nt}}{1 + d_{nt} + \beta p_{nt}} \tag{7}$$

The retail contract must be renewed and initial retail prices selected at the start of a year ($t = 1$). During this time, the firms solve the following problem to set the initial retail prices $\mathbf{p}_1 = (p_{11}, ..., p_{N1})$:

$$\max_{\mathbf{p}_1} V_1(\mathbf{d}_1, mc_1, \mathbf{p}_1) \tag{8}$$

Starting from February, the firms, need to decide whether to keep the current prices or update them by comparing the current retail price vector $\mathbf{p}_t = (p_{1t}, ..., p_{Nt})$ to the optimal vector $\mathbf{p}_t^* = (p_{1t}^*, ..., p_{Nt}^*)$. For each price update, the firms are required to pay the menu cost FC that depends on a vector of retailers' characteristics X_r such as the volume of product sold, number of stores, or number of employees. The value function can be written as follows:

$$V_t(\mathbf{d}_t, mc_t, \mathbf{p}_t) = max\Big\{ \underbrace{\pi(\mathbf{d}_t, mc_t, \mathbf{p}_t) + \beta EV_{t+1}(\mathbf{d}_{t+1}, mc_{t+1}, \mathbf{p}_t)}_{\text{If keeping the current price}} \\ \underbrace{\max_{p^*}\big[\pi(\mathbf{d}_t, mc_t, \mathbf{p}^*) + \beta EV_{t+1}(\mathbf{d}_{t+1}, mc_{t+1}, \mathbf{p}^*)\big] - FC(X_r)}_{\text{If updating the price}} + \epsilon_t \Big\} \tag{9}$$

Here, ϵ_t is the difference between two Type I extreme value shocks $\epsilon_t(1)$ for price adjustment case and $\epsilon_t(0)$ for keeping the old price vector: $\epsilon_t = \epsilon_t(1) - \epsilon_t(0)$.

This decision-making is repeated every month until the year ends in December.

3.1.1 Backward iteration

This dynamic problem is solved using backward iteration starting from the last period ($t = 12$). At this time, the firms have only one month's worth of profit to consider:

$$V_t(\mathbf{d}_t, mc_t, \mathbf{p}_t) = \max\Big\{\underbrace{\pi(\mathbf{d}_t, mc_t, \mathbf{p}_t)}_{V_t(a=0;\mathbf{d}_t,mc_t,\mathbf{p}_t)}, \underbrace{\pi(\mathbf{d}_t, mc_t, \mathbf{p}^*) - RC(X_r)}_{V_t(a=1;\mathbf{d}_t,mc_t,\mathbf{p}_t)} + \epsilon_t\Big\} \tag{10}$$

Here, $V_t(a = j; \mathbf{d}_t, mc_t, \mathbf{p}_t)$ is the deterministic part of the value function conditional on choice of action $j \in 0, 1$. The value function for each month from February to November can be rewritten as:

$$V_t(\mathbf{d}_t, mc_t, \mathbf{p}_t) = \max\Big\{V_t(a = 0; \mathbf{d}_t, mc_t, \mathbf{p}_t), V_t(a = 1; \mathbf{d}_t, mc_t, \mathbf{p}_t) + \epsilon_t\Big\} \tag{11}$$

4 Estimation

There are 5 parameters to be estimated: marginal cost transition probabilities, $Prob_{mc,ll}$, $Prob_{mc,hh}$, and r_0, r_1, r_2 from the following function of menu cost:

$$FC = r_0 + r_1 X_r + R_2 X_r^2 \tag{12}$$

Here, X_r is the volume of oatmeal sold at a given retail chain, measured in 100 thousands. I set the discount rate to $\beta = 0.997$ which corresponds to 0.96 annual discount rate. The sample for the estimation of the supply model consists of monthly data of 21 retail chains, covering the period from January 2011 to January 2018.

4.1 Methodology

The prices in each market are split into three bins: low, medium, and high. The values of price and demand grids can differ from market to market. The marginal cost grid is fixed with two values: low and high (\$0.08 and \$0.1 per portion, with an average price per portion observed of \$0.16). Demand states d_t^{obs} and transition probabilities for each market are obtained from the demand estimation in Section 2.3.

The parameters of the supply model are estimated using a Maximum Likelihood approach following Rust (1987). Starting with a guess of the parameters $\tilde{\theta} = \{P\tilde{r}ob_{mc,ll}, P\tilde{r}ob_{mc,hh}, \tilde{r}_0, \tilde{r}_1\,\tilde{r}_2\}$, I solve the dynamic program for each retailer using backward iteration described in Section 3.1.1. Next, employing the approach in Section 4.2, I impute mc for periods when a price change is observed. For the rest of the periods, I find B marginal cost paths with the highest probability of occurring, conditional on the current guesses $P\hat{r}ob_{mc,ll}$ and $P\hat{r}ob_{mc,hh}$. With the imputed marginal cost, I calculate the conditional choice probabilities:

$$CCP_j(\tilde{x}_t;\tilde{\theta}) = \frac{V_t(a=j;\tilde{x}_t)}{\sum_{l=\{0,1\}} V_t(a=l;\tilde{x}_t)} \tag{13}$$

Here, $\tilde{x}_t$ contains the state variables for the retail-level dynamic problem and ϵ_t in (10) and (11) is the difference between two Type I extreme value shocks $\epsilon_t(1)$ for the price adjustment case and $\epsilon_t(0)$ for keeping the old price vector. Next, the log-likelihood is calculated as:

$$\begin{aligned}\log L = &\sum_t \sum_f a_{ft}^{obs} \sum_b Prob_b \log\Big(CCP_{1ft}(d_t^{obs}, mc_{bt})\Big) \\ &+ \sum_t \sum_f \big(1 - a_{ft}^{obs}\big) \sum_b Prob_b \log\Big(1 - CCP_{1ft}(d_t^{obs}, mc_{bt})\Big)\end{aligned} \tag{14}$$

Here, a_{ft}^{obs} is the action observed at retailer f at time t, with $a_{ft}^{obs} = 1$ corresponding to the choice of adjusting the prices, and $a_{ft}^{obs} = 0$ means the old price vector was kept. Additionally, $b \in B$ is the probability rank of a marginal cost path at the periods when the price changes were not observed, $Prob_b$ is the probability of this marginal cost path. Then, the parameter guesses are updated, and this process is iterated to maximize the log-likelihood until convergence.

4.2 Simplification for Marginal Cost Estimation

To pin down the marginal cost given the model parameters, I compute the retailer's optimal price vector at each point on the marginal cost grid. I then identify the marginal cost value that generates a predicted price closest to the observed price. During this step, a fine price grid is essential for distinguishing between different marginal cost values. However, increasing the size of the price grid for the retailer-level problem would make the estimation computationally infeasible. Since the retailer operates across multiple markets, a large price grid increases exponentially with the number of markets. For instance, if each market has a grid of length H, the size of the price state space of the retailer would become H^N.

To address this challenge, I introduce a simplifying assumption that preserves computational efficiency. When a price change is observed, I assume the new price reflects a forward-looking decision, taking into account that any further within-year adjustment would incur an additional menu cost. I assume that this menu cost is sufficiently large such that the observed price is chosen to maximize the expected discounted profit until the remainder of the year. This assumption is supported by the data, which shows that multiple mid-year price changes for the same product-market are extremely rare. This may reflect the substantial burden associated with the price adjustment process and suggests that firms set their prices with the intention of maintaining them throughout the year.

This assumption allows us to consider the market-level decision separately, as the new price p^*_{nt} is the optimal price for the market until the end of the year. The state variables in each market are the current price p_{nt}, demand state d_{nt}, and marginal cost mc_t. The value function in market n at period t:

$$V_{nt}(p_{n,t-1}, d_{nt}, mc_t) = \max_{p^*_{nt}} \pi_{nt} + \beta E\big[V_{n,t+1}(p^*_{nt}, d_{n,t+1}, mc_{t+1})\big] \tag{15}$$

The return function in the market is:

$$\pi_{nt} = s(p^*_{nt}, d_{nt})\big(p^*_{nt} - mc_t\big) \tag{16}$$

Here, $s(p^*_{nt}, d_{nt})$ denoting the market share of the product. As each market-level problem is considered separately, a large price grid does not overly burden the computation. I chose a grid with 200 points between the minimum and maximum observed per-portion prices. At this level of grid detail, I can map each observed price change to a point on the marginal cost grid. I assume that the observed p^{*obs}_{nt} was chosen as:

$$p^{*obs}_{nt} = \arg\max s(p^*_{nt}, d^{obs}_{nt})\big(p^*_{nt} - mc_t\big) + \beta E\big[V_{n,t+1}(p^*_{nt}, d_{n,t+1}, mc_{t+1})\big] \tag{17}$$

Then, I can find the marginal cost value that produces p^*_{nt} closest to p^{*obs}_{nt}. If prices have changed in several markets simultaneously, I average across them to impute the marginal cost: $mc^{imp}_t = \frac{\sum_k mc_{kt}}{K}$. Here $k \in K$ and K is the total number of markets with price changes at t.

Table 3: Dynamic Model Parameter Estimates

	Estimate	Std. Error
$Prob_{mc,ll}$	0.77	0.37
$Prob_{mc,hh}$	0.78	0.35
r_0	2.75	0.115
r_1	9.35×10^{-4}	12.1×10^{-4}
r_2	4.61×10^{-6}	8.34×10^{-6}

Notes: The table presents the parameters of the structural model estimated using monthly data for 21 retail chains for the period between January 2010 and January 2018. The Nelder-Mead algorithm was used to minimize the negative of the log-likelihood. The Hessian matrix was approximated numerically for the calculation of standard errors.

4.3 Identification

Marginal costs and the menu costs are identified from distinct sources of variation in the data.

The menu cost parameter r_0 are identified from the difference in expected profits with current and optimal price vectors. The menu costs can vary across retailers, but are constant over time. I drop the first months of each year as the prices can be adjusted during these periods at no additional cost. In other months, I compute the difference between (i) the expected discounted profit from retaining the existing price vector and (ii) the expected profit with the optimal price vector that maximizes returns throughout the end of the year. The profit differences can be associated with the observed decisions to adjust prices or not. Additionally, I use variation in oatmeal sales volume across retail chains to identify r_1 and r_2.

The marginal cost transition parameters are identified using the observed chosen price level. Given a parameter guess, I impute marginal cost when a price change is observed assuming the price was chosen as profit maximizing until the end of year. For the periods without price changes observed, I find B number of marginal cost paths with the highest probabilities of occurrence and use the probabilities of these paths to weight the likelihood.

4.4 Estimates and Implications

The results of the estimation of parameters of the structural dynamic model are presented in Table 3. The estimates of r_0, r_1, and r_2 imply that the magnitude of menu costs for changing prices at an average-sized retailer is $275.5 thousand, which is around 7.2% of the variable profit generated at this retailer. However, due to high standard errors of r_1 and r_2, the null hypothesis that the menu costs do not depend on the volume of product sold cannot be rejected.

4.5 Discussion

4.5.1 Price Adjustment Timing

The model with menu costs captures the variation in the likelihood of price adjustments based on the time relative to the contract renewal. Figure 4 shows that in the data sample, price adjustments occur more frequently either in January or around the middle of the year, during the summer. This model can explain the higher probability of price updates in the middle of the year through two distinct factors.

Figure 5: Monthly Profit and Value Function

(a) Monthly Profit

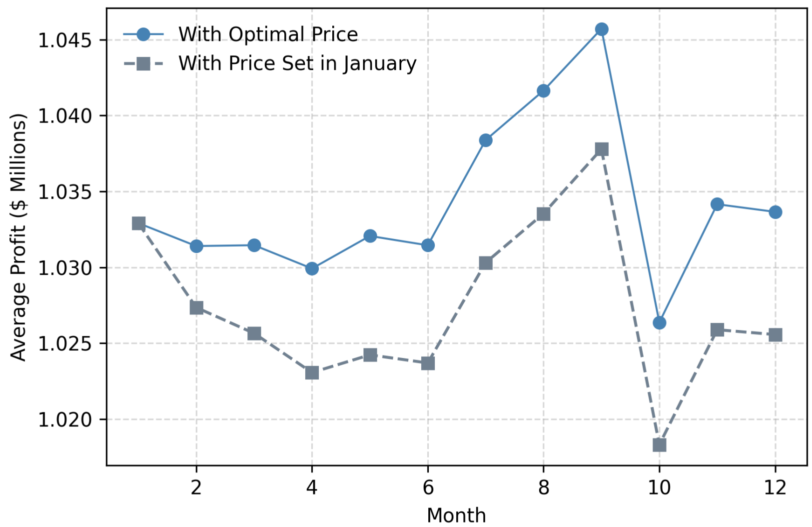

(b) Difference of Profit

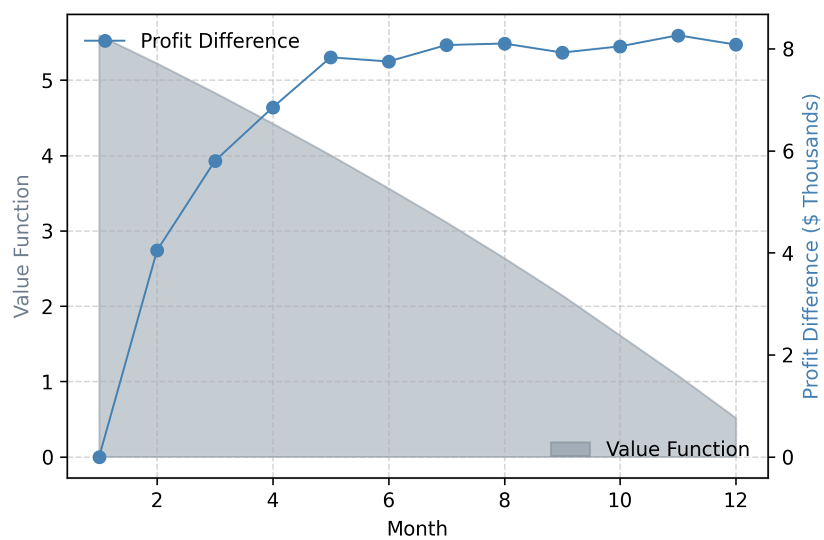

Notes: Panel (a) displays the monthly joint profit (in $ millions) of the manufacturer and retailer under two different pricing scenarios. The blue line with circular markers represents the profit obtained using prices that are optimal given the current month's stochastic state realizations. The dashed gray line depicts the joint profit calculated with the price set at the beginning of the year. Panel (b) shows the increasing difference in monthly profits (in $ thousands) with the optimal price and the January-set price. The gray shaded area shows the decreasing value function. Both panels are based on simulated averages from 1,000 replications of the dynamic model (12,000 months in total). These simulations use stochastic state transitions and retailer characteristics calibrated to resemble those of an average retailer.

First, due to the Markovian nature of demand and marginal cost states, selected

Figure 6: Distribution of Price Changes by Months

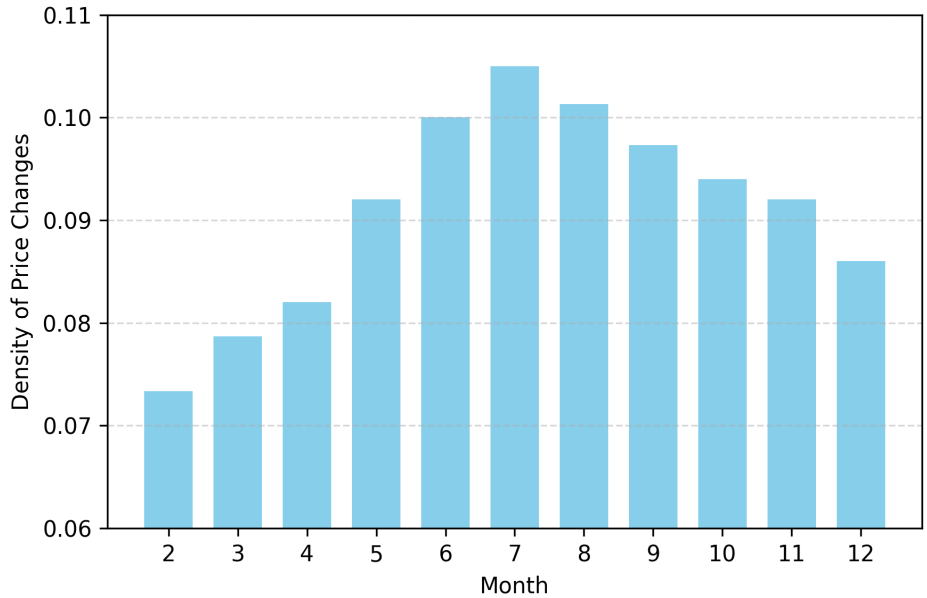

Notes: The graph shows the distribution of price changes by months (excluding January). The probability of observing a price change has an inverted-U shape, with higher probabilities in the middle of the year. The values are based on simulated averages from 1,000 replications of the dynamic model (12,000 months in total). The simulations use stochastic state transitions and retailer characteristics calibrated to resemble those of an average retailer.

prices are relatively close to the optimal prices shortly after the retail contract renewal in January. Therefore, the firms have less incentive to update prices at this time. As the year progresses, however, the chosen prices begin to diverge more from the evolving optimal price levels. Firms choose to update the prices when the deviation becomes sufficiently large to justify the menu costs. Panel (a) of Figure 5 illustrates this by comparing profits under two scenarios: with optimal prices (solid blue line) and with prices chosen in January (dashed gray line). Panel (b) shows the difference in profits between these two pricing strategies. This difference is zero in January and increases throughout the year, exceeding $5,000 by December. [14]

On the other hand, since this dynamic problem is finite and ends in December, the expected discounted value of future profit is highest in January and declines as the year progresses, as shown by the gray area in Panel (b) of Figure 5. This implies that the remaining time to recover the menu cost becomes shorter, reducing the incentive to update the price. Instead, the firms can decide to wait until the start of the following year to reset the prices without the penalty of menu cost. Due to this effect,

14. To protect the confidentiality of retailers, the demand states, transition probabilities, and retailer sizes used in the simulations are calibrated to approximate those of an average retailer.

there are fewer price adjustments close to the year-end.

Figure 6 shows the simulated monthly distribution of price changes. In line with the incentives explained above, there are fewer price changes observed after January and towards December.

4.5.2 Retailer Size and Frequency of Price Changes

The estimation results in Section 4.4 indicate that the menu costs vary little with the volume of product sold at a retail chain. Consequently, the ratio of menu cost to revenue declines with retail chain size. Manufacturers are more likely to recover this cost when adjusting the prices at larger retail chains where they generate higher profits. The decrease in the relative importance of menu costs leads to a positive correlation between the market size of the retailer and the number of price updates observed in a year.

Figure 7 shows that the ratio of the menu cost to variable profit declines non-linearly as market size increases. This arises because the menu costs increase little with market size, while profits increase linearly, holding demand parameters constant. As a result, in sufficiently large markets, the share of the menu costs in total profit approaches zero.

However, the ratio of the price adjustment cost to profit close to 0 does not necessarily lead to price updates occurring in every period. When the frequency of the price updates increases, the marginal gain from each update falls. As both demand and marginal cost states follow Markov processes, when the price is optimized by the forward-looking firm, it tends to remain close to the optimal level in the following periods. Reoptimizing prices that are already close to optimal will incur the same magnitude of menu costs but yield fewer gains.

Figure 8 shows a concave relationship between the variable profit and the frequency of price changes. This implies diminishing marginal profit gains from additional price updates. Taken together, the declining importance of menu costs with retailer size and the diminishing marginal benefits of price adjustments produce an approximately linear relationship between retailer size and the frequency of price changes.

Figure 9 illustrates the near-linear relationship between retail chain market size and the simulated number of price updates per year. Both the total number of price changes (blue line) and the number of mid-year price changes (gray dashed line) increase with market size. The rise in total price changes is primarily driven by more frequent mid-year updates, indicating that larger retailers tend to update prices more

Figure 7: Menu Costs and Profit.

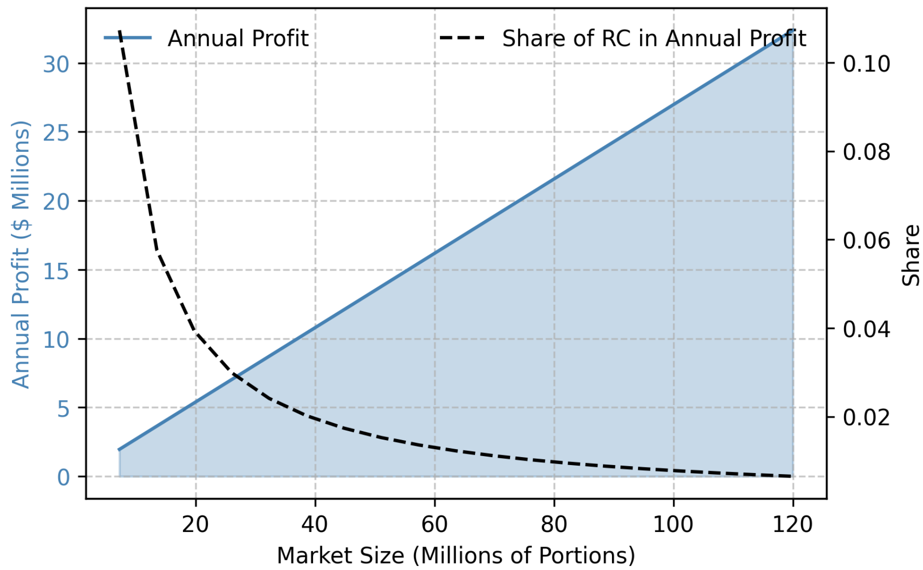

Notes: The graph shows the relationship between the retailer's market size and the manufacturer and retailer's joint surplus on the left vertical axis. The black dashed line shows the share of the estimated menu costs in the joint surplus (the right axis). The values and transition probabilities of the demand shifter were chosen to match those of the average retailer.

often during the year. In contrast, the number of price changes occurring at the beginning of the year, depicted by the vertical gap between the two lines, does not appear to depend on the retailers' size.

5 Counterfactual Exercises

5.1 Low and High Inflation

Economic literature has been documenting changes in both magnitude and frequency of price changes during high and low inflation periods (e.g. Gagnon 2009, Nakamura and Steinsson 2008). However, the standard menu cost models used in macroeconomic research struggle to generate these two changes simultaneously (Blanco et al. 2024a).

I run a simplified inflation experiment with an exogenously increasing marginal cost. In this simple scenario, the marginal cost grows at a constant rate each month. The firm can observe the marginal cost at the start of the period, but does not expect further increases.

Figure 10 shows the number of price changes per year and the magnitude of

Figure 8: Frequency of Price Changes and Profit

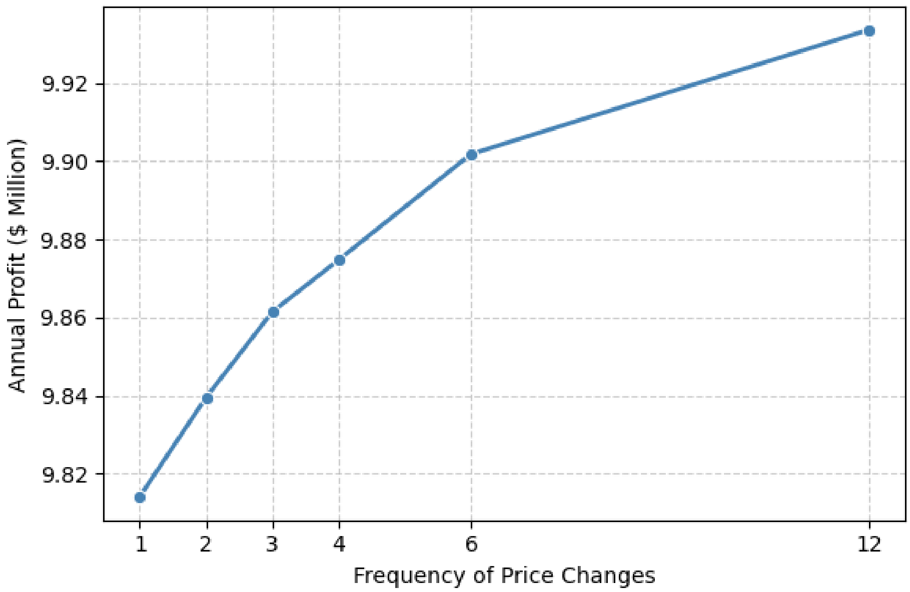

Notes: The figure presents the average annual simulated profit of the manufacturer and the frequency of price changes per year. Price changes are spaced evenly throughout the year. For example, in the case of two annual price updates, they are deterministically scheduled in January and July. The values and transition probabilities of the demand shifter are calibrated to match those observed for the average retailer. Profit values are based on 1,000 simulations of the dynamic model.

these changes increases as the growth rate of marginal cost is raised. The marginal cost growth rate of 2% per month (annual growth rate 26.8%) results in 20% more frequent price adjustments per year. At the same time, the magnitude of price changes also increases from $0.43 to $0.56. This means that although the prices are changed more frequently, the changes are also larger in magnitude. This result is relevant to macroeconomic literature as it can generate results not available in standard menu cost models.

5.2 Anticipated Cost Shocks

I explore the effect of cost shock anticipation for small and large retail chains. I consider a scenario in which the marginal cost grid permanently increases by 30 cents (equivalent to 1 cent per portion of oatmeal) in the middle of the year.

The anticipation of the shock leads both retailers to set initial prices higher in January. However, the difference is significantly larger for the smaller retailers. Moreover, smaller retailers increase their prices further 2-3 months before the shock occurs. Panel (a) of Figure 11 shows that when anticipating a permanent marginal cost in-

Figure 9: Number of Price Changes per Year

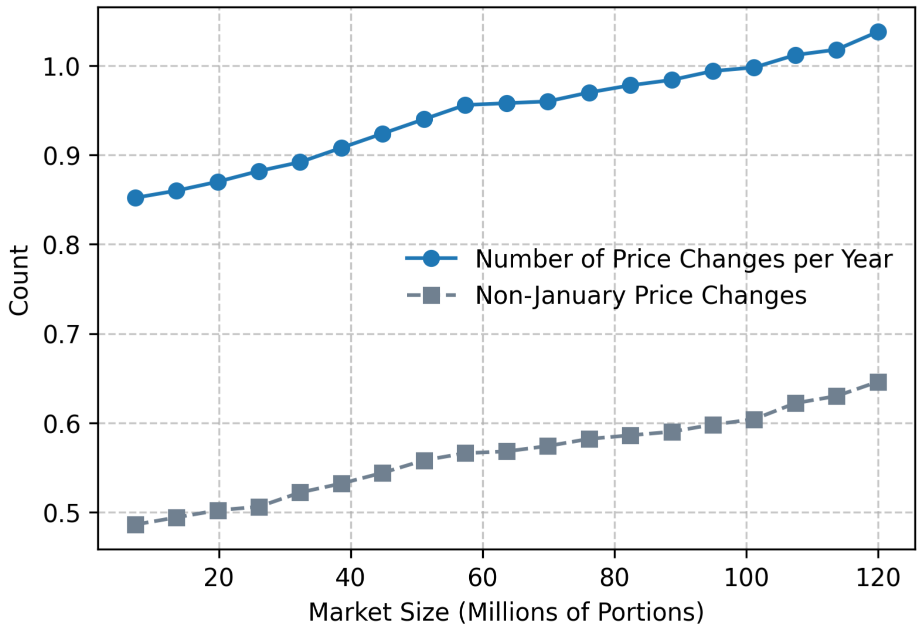

Notes: The graph depicts the average simulated number of price changes per year as a function of retail chain market size. The blue line with circular markers represents the average number of all observed price changes, while the dashed gray line captures only mid-year price changes, excluding updates made at the beginning of the year. The values and transition probabilities of the demand shifter are calibrated to match those observed for the average retailer. Price change frequencies are computed from 1,000 simulations of the dynamic model.

crease, a smaller retailer starts raising their prices well before the shock. As they are constrained in their capacity to change the prices often, they adjust their price if they receive a positive ϵ shock from Equation (9) even if it happens before the shock takes place.

Even after the permanent cost shock occurs, the difference in average prices between the anticipated and unanticipated cases remains more pronounced for the smaller retailer. This pattern reflects the limited flexibility of smaller firms, which have fewer opportunities to update their prices. As a result, in most cases, the small retailer is unable to respond to an unanticipated cost increase before the end of the calendar year.

This is consistent with empirical evidence of the heterogeneous reactions of firms of different sizes, beginning with Gertler and Gilchrist 1994, which highlight that smaller firms exhibit larger price reactions compared to large firms in response to cost increases, resulting in them being squeezed out of the market (Baqaee, Farhi, and Sangani 2024).

Figure 10: Cost Increase and Price Changes

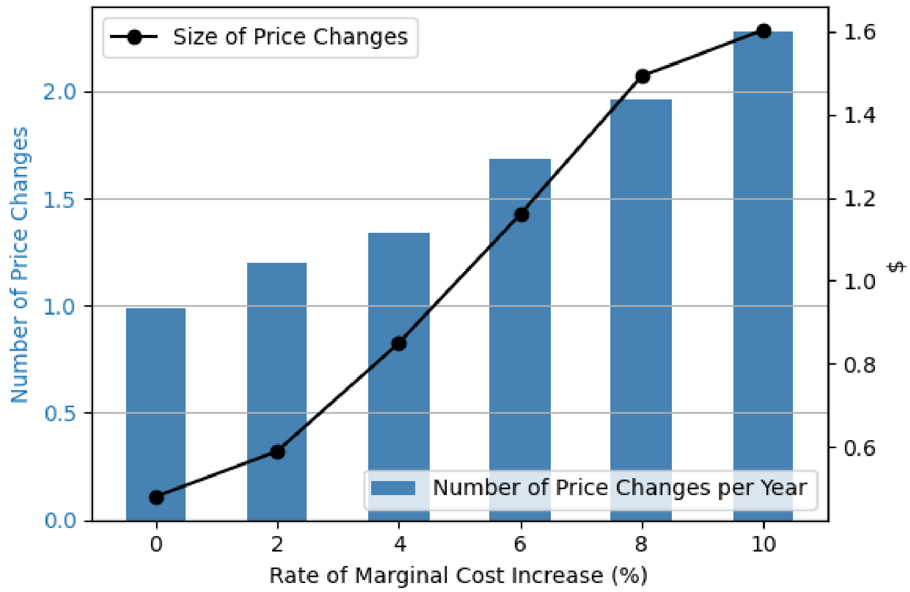

Notes: The chart shows the number of price changes per year (left vertical axis) and the standard deviation of price changes (right vertical axis) when the marginal cost is increasing at a constant rate each month. For instance, in the baseline case, the marginal cost growth rate is 0, and the average-sized retailer adjusts the prices once per year on average, and the standard error of prices is \$0.43. Both the magnitude and frequency of price changes increase as the rate of marginal cost growth is faster. The graph shows the average of 200 simulations of the dynamic problem (2400 periods).

5.3 Welfare Effects of Menu Costs

I examine how a reduction in price adjustment costs affects social welfare. I lower the menu costs for an average-sized retail chain from the estimated baseline level and examine the implications for firms' profit and consumer surplus.

As menu costs decrease, firms have more financial capacity to optimize their prices more frequently. Joint surplus of the manufacturer and retailer is calculated as the difference between the retail price and marginal cost, multiplied by quantity sold, and summed up across all markets the retailer operates in:

$$\pi_{ft} = \sum_{n=1}^{M} (p_{nt} - mc_t) q_{nt} \tag{18}$$

Column (6) of Table 4 shows that the firms are changing prices more frequently. As a result, firms' net surplus (after paying the menu costs) is increasing, as shown in Column (7). With greater flexibility to update prices, firms can capture a larger share of the surplus when menu costs are lower.

Figure 11: Effect of Marginal Cost Increase

(a) Prices at a Small Retailer

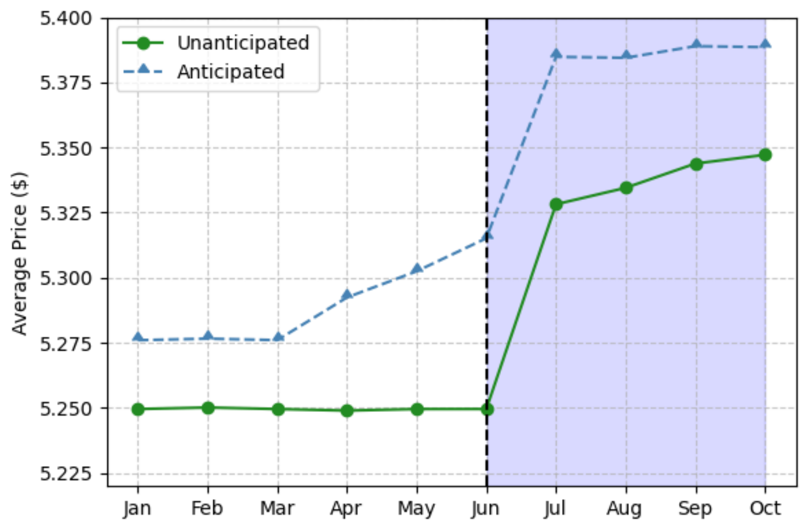

(b) Prices at a Large Retailer

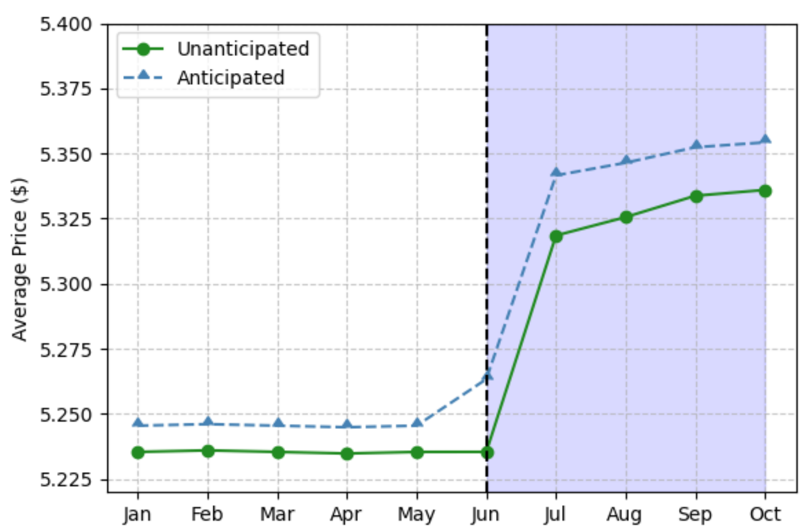

Notes: The graphs show retail prices at retailers of different sizes with anticipated and unanticipated permanent cost shocks. The marginal cost grid increases by $0.3 per package (approximately 1 cent per portion) in June. Panel (a) shows that for the smaller retailer (with a market size equal to 1 million portions of oatmeal per month), Panel (b) shows the prices set by a large retailer with a market size of 20 million portions per month. The graphs show the average price of the flagship product across 250 (3000 months) Monte-Carlo simulations of the dynamic model.

Consumer surplus is calculated following Small and Rosen (1981) and is reported in Column (3) of Table 4. At first, consumer welfare increases as retail prices are adjusted more frequently and better match demand conditions. However, lowering the menu costs further is not beneficial for consumers, as the firms can extract more consumer surplus through more flexible pricing.

Table 4: Welfare Effects of Lowering Menu Costs

Menu Cost (1)	% of Est. Menu Cost (2)	Consumer Surplus (3)	Δ CS (4)	Profit (5)	Num. of Price Changes (6)	Net Profit (7)	Δ Net Profit (8)
275.5	1.0	2934.2		3343.8	0.5	3202.8	
206.6	0.75	2949.4	15.1	3344.8	0.6	3210.5	7.7
137.7	0.5	2991.2	57.0	3347.0	0.8	3235.5	32.7
68.9	0.25	2969.4	35.1	3350.2	1.2	3264.8	62.0
0	0	2884.5	-49.8	3354.6	1.8	3354.6	151.8

Notes: The table reports welfare outcomes under varying levels of menu costs. The top row corresponds to the baseline case, where the menu cost is set to its estimated value of $275.5 thousand for an average-sized retail chain. Columns (1) and (2) present the absolute value of the menu cost and its proportion relative to the baseline, respectively. Column (3) reports the average annual consumer surplus (in thousands of dollars), calculated following Small and Rosen (1981) and summed across the markets the retail chain operates in. Column (5) shows the joint annual profit of the manufacturer and retailer (also in $ thousands), while Column (6) reports the average number of price adjustments per year. Column (7) presents the joint profit net of menu costs. Columns (4) and (8) display the differences in consumer surplus and net profit compared to the baseline case. All figures are based on 250 simulations of the dynamic model (3000 periods).

6 Conclusion

In this paper, I develop a model of pricing in the presence of menu costs. I apply the model to the U.S. consumer packaged goods (CPG) retail sector and estimate the parameters using a detailed dataset on oatmeal prices and sales from grocery store chains. The results indicate that menu costs are substantial but vary little with retailer size. Consequently, prices are adjusted more frequently at larger retail chains. The model also explains the lower probability of price changes close to the retail contract renewal time.

Counterfactual exercises reveal how menu costs can lead to more frequent and larger price adjustments during periods of high inflation. Prices at smaller retail chains exhibit greater sensitivity to cost increases. Furthermore, when cost increases are anticipated, smaller retailers tend to start adjusting their prices earlier than larger ones. These findings are relevant to macroeconomic literature on price rigidity and heterogeneous responses of firms of different sizes to shocks.

Lowering the menu costs enables firms to update prices more frequently, increasing their profitability. However, low menu costs are not beneficial for consumers as firms can extract more consumer surplus through more flexible pricing.

References

Aguirregabiria, Victor. 1999. "The dynamics of markups and inventories in retailing firms." *The review of economic studies* 66 (2): 275–308.

Anderson, Eric T, and Edward J Fox. 2019. "How price promotions work: A review of practice and theory." *Handbook of the Economics of Marketing* 1:497–552.

Arcidiacono, Peter, Patrick Bayer, Jason R Blevins, and Paul B Ellickson. 2016. "Estimation of dynamic discrete choice models in continuous time with an application to retail competition." *The Review of Economic Studies* 83 (3): 889–931.

Backus, Matthew, Christopher Conlon, and Michael Sinkinson. 2021. *Common ownership and competition in the ready-to-eat cereal industry.* Technical report. National Bureau of Economic Research.

Baqaee, David R, Emmanuel Farhi, and Kunal Sangani. 2024. "The supply-side effects of monetary policy." *Journal of Political Economy* 132 (4): 1065–1112.

Barro, Robert J. 1972. "A theory of monopolistic price adjustment." *The Review of Economic Studies* 39 (1): 17–26.

Basker, Emek, and Michael Noel. 2009. "The evolving food chain: competitive effects of Wal-Mart's entry into the supermarket industry." *Journal of Economics & Management Strategy* 18 (4): 977–1009.

Berry, Steven, James Levinsohn, and Ariel Pakes. 1995. "Automobile prices in market equilibrium." *Econometrica: Journal of the Econometric Society,* 841–890.

Berry, Steven T, and Philip A Haile. 2014. "Identification in differentiated products markets using market level data." *Econometrica* 82 (5): 1749–1797.

Blanco, Andrés, Corina Boar, Callum J Jones, and Virgiliu Midrigan. 2024a. *Nonlinear dynamics in menu cost economies? evidence from us data.* Technical report. National Bureau of Economic Research.

———. 2024b. *The inflation accelerator.* Technical report. National Bureau of Economic Research.

Butters, R Andrew, Daniel W Sacks, and Boyoung Seo. 2022. "How Do National Firms Respond to Local Cost Shocks?" *American Economic Review* 112 (5): 1737–72.

Cedar Rapids Tourism Office. 2025. *Plan Your Trip: Cedar Rapids.* Accessed: 2025-05-25. https://www.tourismcedarrapids.com/plan-your-trip/cedar-rapids/.

Conlon, Christopher T, and Nirupama L Rao. 2020. "Discrete prices and the incidence and efficiency of excise taxes." *American Economic Journal: Economic Policy* 12 (4): 111–143.

Crouzet, Nicolas, and Neil R Mehrotra. 2020. "Small and large firms over the business cycle." *American Economic Review* 110 (11): 3549–3601.

DellaVigna, Stefano, and Matthew Gentzkow. 2019. "Uniform pricing in us retail chains." *The Quarterly Journal of Economics* 134 (4): 2011–2084.

Duval, Romain, Davide Furceri, Raphael Lee, and Marina M Tavares. 2024. "Market power and monetary policy transmission." *Economica* 91 (362): 669–700.

Gagnon, Etienne. 2009. "Price Setting during Low and High Inflation: Evidence from Mexico." *The Quarterly Journal of Economics* 124 (3): 1221–1263. ISSN: 00335533, 15314650, accessed August 31, 2025. http://www.jstor.org/stable/40506256.

Gagnon, Etienne, and David López-Salido. 2020. "Small price responses to large demand shocks." *Journal of the European Economic Association* 18 (2): 792–828.

Ganapati, Sharat. 2025. "The modern wholesaler: Global sourcing, domestic distribution, and scale economies." *American Economic Journal: Microeconomics* 17 (1): 1–40.

Gertler, Mark, and Simon Gilchrist. 1994. "Monetary policy, business cycles, and the behavior of small manufacturing firms." *The quarterly journal of economics* 109 (2): 309–340.

Hausman, Jerry A. 1994. *Valuation of new goods under perfect and imperfect competition.* National Bureau of Economic Research Cambridge, Mass., USA.

Hitsch, Günter J, Ali Hortacsu, and Xiliang Lin. 2019. *Prices and promotions in us retail markets: Evidence from big data.* Technical report. National Bureau of Economic Research.

How Products Are Made. 2025. *How Oatmeal is Made.* Accessed: 2025-05-25. https://www.madehow.com/Volume-5/Oatmeal.html.

Hsieh, Chang-Tai, and Esteban Rossi-Hansberg. 2023. "The industrial revolution in services." *Journal of Political Economy Macroeconomics* 1 (1): 3–42.

Kaufman, Leonard, and Peter J Rousseeuw. 2009. *Finding groups in data: an introduction to cluster analysis.* John Wiley & Sons.

Knotek II, Edward S. 2024. "The roles of price points and menu costs in price rigidity." *Journal of Monetary Economics* 145:103559.

Levy, Daniel, Mark Bergen, Shantanu Dutta, and Robert Venable. 1997. "The magnitude of menu costs: direct evidence from large US supermarket chains." *The Quarterly Journal of Economics* 112 (3): 791–824.

Mansley, Ryan. 2022. "Intertemporal Price Discrimination in Consumer Packaged Goods." *Available at SSRN 4244900.*

Miller, Nathan H, and Matthew C Weinberg. 2017. "Understanding the price effects of the MillerCoors joint venture." *Econometrica* 85 (6): 1763–1791.

Nakamura, Emi, and Jón Steinsson. 2008. "Five facts about prices: A reevaluation of menu cost models." *The Quarterly Journal of Economics* 123 (4): 1415–1464.

———. 2013. "Price rigidity: Microeconomic evidence and macroeconomic implications." *Annu. Rev. Econ.* 5 (1): 133–163.

Ottonello, Pablo, and Thomas Winberry. 2020. "Financial heterogeneity and the investment channel of monetary policy." *Econometrica* 88 (6): 2473–2502.

Rivers, Douglas, and Quang Vuong. 2002. "Model selection tests for nonlinear dynamic models." *The Econometrics Journal* 5 (1): 1–39.

Rust, John. 1987. "Optimal replacement of GMC bus engines: An empirical model of Harold Zurcher." *Econometrica: Journal of the Econometric Society,* 999–1033.

Slade, Margaret E. 1998. "Optimal pricing with costly adjustment: evidence from retail-grocery prices." *The Review of Economic Studies* 65 (1): 87–107.

Small, Kenneth A, and Harvey S Rosen. 1981. "Applied welfare economics with discrete choice models." *Econometrica: Journal of the Econometric Society,* 105–130.

Stella, Andrea. 2014. "The magnitude of menu costs: A structural estimation." In *2014 Meeting Papers,* vol. 436. Society for Economic Dynamics.

Tabanakov, Semyon, Ali Goli, and Pradeep K. Chintagunta. 2024. *Retail Pricing and Organizational Structure.* Kilts Center at Chicago Booth Marketing Data Center Paper. Available at SSRN: https://ssrn.com/abstract=4870276 or http://dx.doi.org/10.2139/ssrn.4870276, June.

The Quaker Oats Company. 2025. *Quaker Oats History.* Accessed: 2025-05-25. https://www.quakeroats.com/about-quaker-oats/quaker-history.

U.S. Bureau of Labor Statistics. 2025. *Producer Price Index by Commodity: Pulp, Paper, and Allied Products: Corrugated Paperboard in Sheets and Rolls, Lined and Unlined [WPU091405].* https://fred.stlouisfed.org/series/WPU091405. Retrieved from FRED, Federal Reserve Bank of St. Louis, July 17, 2025.

Zbaracki, Mark J, Mark Ritson, Daniel Levy, Shantanu Dutta, and Mark Bergen. 2004. "Managerial and customer costs of price adjustment: direct evidence from industrial markets." *Review of Economics and statistics* 86 (2): 514–533.

Appendix

A First-stage demand regression

Table A.1: First-stage estimation results

	Stage 1 of 2SLS
Oats price	-1.95
	(8.92)
Electricity price	0.05
	(0.02)
Diesel price	0.001
	(0.153)
Cardboard price	-0.001
	(0.002)
Wages (commercial)	-0.01
	(0.03)
Wages (food production)	-0.006
	(0.1)
Hausman IV	0.81
	(0.08)
% Discount	0.0329
	(0.01)
N	4884
F	37.2
Market-level Effect	Yes

Notes: The table shows OLS regression results of $ price of oatmeal per portion on instrumental variables. Data on oats prices ($ per gallon) were obtained from Macrotrends commodity price series. Monthly retail electricity prices (in $ per kWh), in the industrial sector in Iowa, were sourced from the U.S. Energy Information Administration (Link to recreate the dataset: Electricity Data Browser). Historical weekly U.S. No. 2 diesel retail prices (in $ per gallon) were obtained from the U.S. Energy Information Administration website. Diesel prices, reported by the Petroleum Administration for Defense Districts (PADD), were merged with the designated market areas (DMA) defined by Nielsen. Weekly price data were converted to a monthly frequency by retaining the final weekly observation of each month. Producer price index for pulp, paper, and allied products was used as a cardboard price proxy and sourced from the FRED database (U.S. Bureau of Labor Statistics 2025). Average weekly wage data (in $ thousands) for private food manufacturing in Iowa and retail trade for all U.S. states were obtained from the Quarterly Census of Employment and Wages published by the Bureau of Labor Statistics (Available at the BLS website).

B Test of Conduct

I compare the performance of the model with menu costs against a more conventional model with flexible pricing using a non-nested testing approach. The testing procedure follows Backus, Conlon, and Sinkinson (2021), which is based on Berry and Haile (2014), and is used to differentiate between a pair of structural models. The null hypothesis that both models of supply side perform equally well is tested against the alternative hypothesis that one of the models performs better than the other. The test checks which model produces an unobserved marginal cost that better satisfies an orthogonality condition between the exogenous product characteristics z_t. The exogenous characteristics include variables that affect demand but are excluded from marginal costs v_t, as well as those affecting both demand and marginal costs x_t, and that affect marginal cost but not demand w_t.

First, Berry and Haile (2014) assumes that marginal cost equals marginal revenue (or the generalized residual):

$$mc_{jt} = \psi_j\big(\mathbf{s}_t, \mathbf{p}_t, D(\mathbf{z}_t)\big) \tag{B.1}$$

Here, $\mathbf{s}_t$ denotes market shares, $\mathbf{p}_t$ observed prices, and $D(z_t)$ the demand system. This is extended by Backus, Conlon, and Sinkinson (2021) to marginal cost equals the difference between the prices and markup $\eta_j^m(\mathbf{s}_t, \mathbf{p}_t, D(\mathbf{z}_t))$ implied by the model:

$$mc_{jt} = p_{jt} - \eta_j^{model}\big(\mathbf{s}_t, \mathbf{p}_t, D(\mathbf{z}_t)\big) \tag{B.2}$$

Then, the marginal cost is characterized by the exogenous product characteristics, assuming an additively separable unobservable term:

$$mc_{jt} = h_s(x_{jt}, w_{jt}) + \omega_{jt} \tag{B.3}$$

The conditional moment restriction is as follows:

$$E[\omega_{jt}|z_t] = 0 \tag{B.4}$$

Then, the test statistic is calculated as follows:

$$T = \frac{\sqrt{n}}{\sigma}\Big(Q(\eta^1) - Q(\eta^2)\Big) \tag{B.5}$$

Here, η is the firm markups implied by a given model, and $Q(\eta)$ is a GMM objective of the test defined as a function of markups. $\frac{\sigma}{\sqrt{n}}$ is the standard error of $Q(\eta^1) - Q(\eta^2)$. Rivers and Vuong (2002) shows that T follows a standard normal distribution. Therefore, at the 5% significance level, the null hypothesis is rejected if $|T| > 1.96$. If $T < -1.96$ we conclude that Model 1 is preferred, and if $T > 1.96$, we conclude that Model 2 is closer to the truth.

Backus, Conlon, and Sinkinson (2021) choose the following form of the GMM objective:

$$Q(\eta) = \Big(\frac{1}{n}\sum_t \hat{\omega}_t \widehat{\Delta\eta^{1,2}}\Big)^2 \tag{B.6}$$

Here, $\Delta\eta^{1,2} = \eta^1 - \eta^2$ is the difference between markups implied by Model 1 and Model 2. $\hat{\omega}_t$ is a part of the marginal cost which is not explained by the exogenous product characteristics and $\widehat{\Delta\eta^{1,2}}$ are the fitted values of the following regression on exogenous characteristics of all products in the market $\mathbf{z}_t$:

$$\Delta\eta^{1,2} = g(\mathbf{z}_t) + e_t \tag{B.7}$$

Next, I use the base prices and demand shifters to impute marginal cost for the model with flexible pricing (Model 2). In this model, the firm can update the price without incurring a menu cost. I assume that the manufacturer uses the demand specification described in Section 2.3. The manufacturer's problem in Model 2 would be:

$$\max_{p^*} \lambda s(p_{m,t}, d_m)(p_{m,t} - mc_t) \tag{B.8}$$

Then, the marginal cost is imputed based on the first-order condition:

$$mc_{m,t} = p_{m,t} - \frac{1}{\beta(1 - s_{m,t})} \tag{B.9}$$

For empirical implementation of the test, I use data for the variables v_t that affect demand but are excluded from the marginal cost: BLP-style instrumental variables (Berry, Levinsohn, and Pakes 1995) that are the sum of the price of all products not produced by Quaker, and the number of all oatmeal products (distinct UPC) sold in the market. I also include the average percentage of price promotions observed in the market. For the variables that affect the marginal cost but are excluded from demand w_t, I use the input cost data described in Table A.1, such as price of oats, electricity, wages, and diesel, as well as Hausman-style instrumental variables.

I first regress the marginal cost implied by the two models on all exogenous product characteristics that affect marginal cost to get the unobserved portion of the marginal cost. I find the standard error of $Q(\eta^1) - Q(\eta^2)$ by block bootstrapping with 1000 blocks with 12 elements (months) each. The test statistic is calculated as $T = -9.08$. Therefore, I conclude that Model 1 with menu costs performs better than the model with flexible pricing at 1% significance level.

Figure C.1: Lag Between Marginal Cost Increase and Price Response

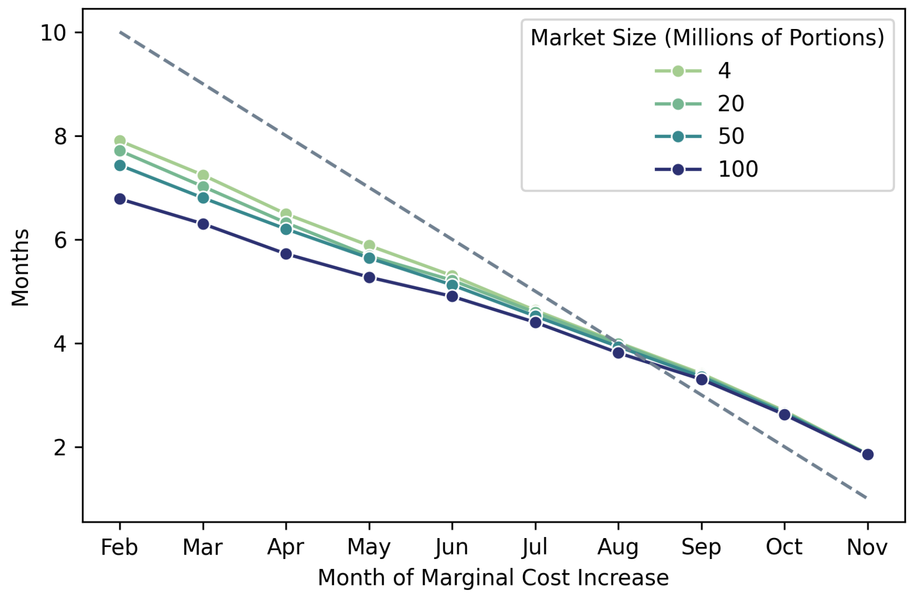

Notes: The graph shows the lag (in months) between the time of the marginal cost increase (from the low to high state) and the price increase by the retail chains' size. The numbers are based on 500 Monte-Carlo simulations of the dynamic model for each value of the market size.

C Marginal Cost Pass-Through Timing

In this counterfactual exercise, I examine how the timing of marginal cost changes can lead to different speeds of pass-through to the retail prices. Suppose that a year starts and the retail prices are set when the marginal cost is in a "low" state. Keeping the demand shock unchanged, I deterministically increase the marginal cost to the "high" state at different months of the year, starting from February.

Figure C.1 shows that if the marginal cost increases in February, on average, it will take a smaller retailer 8 months to adjust the price, while it will take around 6.7 months for the large retailer to update the price. As larger retailers can update their prices more frequently, as explained in Section 4.5.2, it is not surprising that they require less time to adjust the price. The 45-degree line in the figure shows the number of months a retailer would have to wait until the beginning of next year, when the retail contract would be renewed and the prices can be reset without incurring the menu costs.

Figure C.1 also shows that the marginal cost rises occurring later in the year are followed by increasingly prompt price adjustments. Two underlying mechanisms lead to this effect. First, marginal cost increases during summer have a higher probability

Figure C.2: Probability of Waiting Until Next Year

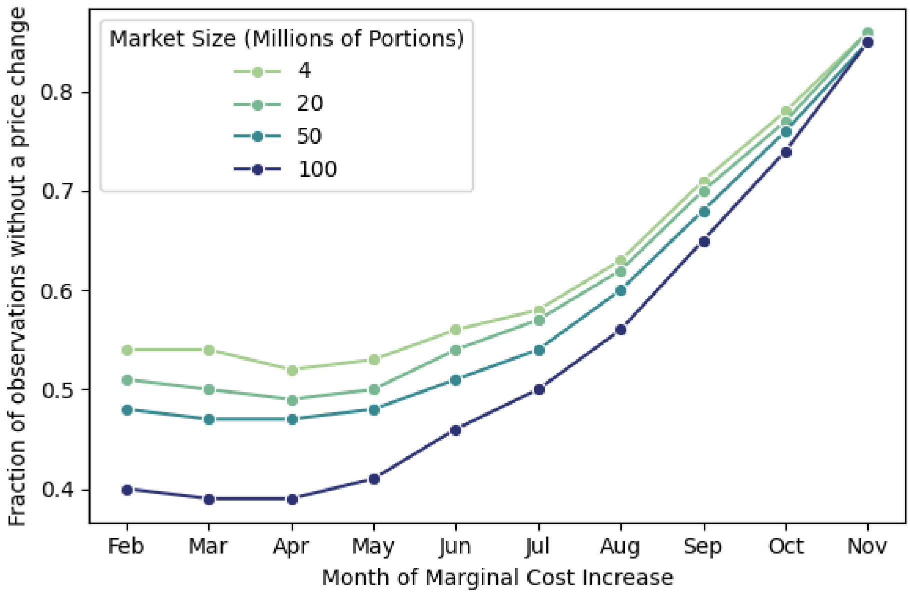

Notes: The graph shows the share of simulations when the marginal cost increase is not followed by the same-year price update. The data is based on 500 Monte-Carlo simulations of the dynamic model for each market size.

of triggering a price update. Second, as cost shocks occur closer to year-end, the firms often defer price changes until the start of the following year, when retail contracts are renewed. For example, when a marginal cost increase occurs after August, retail prices are, on average, not adjusted within the same calendar year.

Figure C.2 illustrates the behavior described above and shows the probability of postponing a price change to the following year. The probability of deferring the price update increases sharply for shocks occurring later in the year. For instance, when a cost shock occurs in November, 86% of cases result in the retailer waiting until the new year to adjust prices. This tendency to delay is less pronounced among larger retailers, particularly when cost shocks arise in the first half of the year.