Price Stickiness:
Empirical Evidence of the Menu Cost Channel

Eric Anderson
Northwestern University

Nir Jaimovich
Duke University and NBER

Duncan Simester
MIT

May 2012

This paper has benefited from comments by Jeff Campbell, Alberto Cavallo, Nathan Fong, Robert Hall, Pete Klenow, Roy Mill, Emi Nakamura, Nicolas Vincent, Ivan Werning, and participants at the 2011 NBER Pricing Dynamics workshop, 2011 Milton Friedman Institute Pricing Dynamics Conference, the 2011 Yale Customer Insights Conference and seminar participants at the University of Alberta and the University of Wisconsin.

Price Stickiness: Empirical Evidence of the Menu Cost Channel

A leading explanation in the economic literature is that monetary policy has real effects on the economy because firms must incur a cost when changing prices. Yet, empirical validation and quantification of the effects of menu costs on pricing is scant. Using a 55-month unique database of cost and price changes at a large retailer we find that absent these menu costs, cost changes would result in up to 18% more price changes. We confirm that a decision to forgo a price change when costs change is not merely a short temporal delay. We also show that the effects we measure are allocative in the sense that they have a persistent impact on both prices and unit sales. Finally, we provide evidence that the menu cost channel only operates when cost shocks are small in magnitude. This is consistent with theory, and provides the first empirical evidence of boundary conditions in the menu cost channel.

1. Introduction

Why does monetary policy have real effects on the economy? A leading explanation in the macroeconomics literature is that firms incur a cost (a "menu cost") when changing their prices, which results in less frequent price adjustments. This framework of "state-dependent" pricing has been studied extensively since the work of Barro (1972) and Sheshinski and Weiss (1977).[1]

Unfortunately, there is almost no empirical evidence that validates and quantifies the effects of menu costs on price stickiness (in what follows we refer to this as the "menu cost channel"). The reason is that such an analysis requires a unique data set that is hard to obtain. Specifically, the appropriate data set has at least two prerequisites. First, it needs to include measures (or proxies) of menu costs that vary across products, time or geography. Second, it needs to include accurate measures of cost shocks and other covariates. These are required to rule out the possibility that an interaction between the magnitudes of the menu costs and other factors contribute to the decision to change prices.

We report the findings from a large-scale empirical study with a national U.S. retailer that enables us to construct such a data set. A key to our identification of the menu cost channel relies on a pricing rule that the retailer enforces. Like many retailers, the firm links the prices of different color and flavor variants of a product. If the price of one variant changes, then the prices of all other variants need to change as well. As we later show, there is significant variation in the number of links across products.

The menu costs we study are attributed to in-store labor costs. Consider the task of changing the price of Cheerios (which has one variant) with changing the price of nail polish (for which one item has 62 different color variants). To change the price, a store employee must locate the product in the store. Once Cheerios is located, the employee simply changes a single on-shelf price sticker. However, when changing the price of nail polish the employee must match each sticker with the shelf location of each of the 65 colors. The retailer conducts time and motion studies that clearly show the time to change the price of an item increases with the number of variants.[2]

---

[1] While important details differ across the work that followed, a central and common assumption is that a fixed cost must be incurred upon a price change. See for example other prominent examples that built on this work: Akerlof and Yellen (1985), Mankiw (1985), Caplin and Spulber (1987), Caplin and Leahy (1991, 1997), Bertola and Caballero (1990), Danziger (1999), Dotsey, King, and Wolman (1999), Burstein (2006), Golosov and Lucas (2007), and Gertler and Leahy (2008).

[2] Midrigan (2012) assumes in his theoretical work that there exists a fixed cost of changing a single price but that the cost of changing additional prices is zero. This could lead to outcomes in which items with more variants have a lower menu cost, which leads to a higher likelihood of price changes in response to cost shocks. We address his work later in this Introduction.

Like most retailers, in-store labor costs represent the largest expense after cost of goods sold and therefore these expenses are carefully monitored and budgeted. In the short-run labor capacity is fixed, and so a small number of additional price changes on any single day do not result in the firm hiring an extra employee or paying overtime. But, it does create an opportunity cost as less labor is allocated to valuable activities such as stocking shelves, answering customers' questions, or completing transactions. Returning to our example, changing the price of nail polish requires a larger time allocation than Cheerios and this imposes a larger opportunity cost by reducing the time available for these other activities. In Section 2 we provide additional details on the specific policies this retailer adopts to manage these opportunity costs.

The firm's policy of linking prices across variants provides an ideal opportunity to measure whether these costs affect price stickiness. If these types of menu costs play no part in the decision to change prices then we would expect more price variation on items with more variants, as these items tend to have higher unit sales volumes. However, if menu costs increase with the number of variants and these costs contribute to price stickiness, this will tend to reduce price variation on items with more variants. This identification strategy, together with our access to clean measures of cost shocks and a rich set of covariates for each product, allows us to validate the menu cost channel.[3]

We find that among products that have a single variant, a cost increase leads to an immediate price increase 71.2% of the time. But, if a product has seven or more variants, then the probability of a price increase is just 59.8%. We calculate that 18% more price changes would have been observed if all items had only a single variant.

Importantly, we show that these effects are persistent. Price changes that do not occur at the time of a cost shock are not merely delayed; among items for which prices are not initially increased, only 5.8% have a price increase within the next 90 days. When we look over longer horizons (e.g. 360 days), there is no evidence of a delayed price change. Finally, we confirm that the effects we measure have long-run impacts on both prices and unit sales.

We also identify boundary conditions on the menu cost channel. We anticipate that menu costs will be weighed against the cost of not adjusting a price. For this reason, when the cost increase is large we do not expect that the menu cost channel will impact pricing decisions. This is what we observe in our data. When cost increases are large, the probability of a price increase is also large. It is also invariant to the number of product linkages, indicating that menu costs do not play an important role. It is only

[3] These cost change events are difficult to infer from other data sources, such as the widely utilized Dominick's data. For example, the Dominick's data does not report the regular price of an item and the cost metric may capture a weighted average cost of inventory (Peltzman 2000). These limitations make it difficult to identify the timing or magnitude of cost or regular price changes.

when the cost increase is small that we observe the menu cost channel influencing pricing decisions.

The contribution of this paper is to provide clear empirical evidence of the menu cost channel. We identify institutional features that create variation in menu costs across products. We then show that when a retailer is faced with a cost increase these menu costs have a meaningful impact on the decision to raise the price. Indeed, menu costs play an important role in contributing to price stickiness.

Related Literature

There have been surprisingly few attempts to directly measure the link between menu costs and the frequency of price changes. A search of the literature reveals one important predecessor. Levy, Bergen, Dutta, and Venable (1997) begin by de-composing the cost of changing prices in a sample of US supermarkets, and then examining how item-pricing laws (which require separate pricing stickers on each unit) affect the frequency of price changes. It is this second portion of their paper that is most closely related to this paper. The four retailers in their study that are not subject to item-pricing laws change prices on 15.6% of products each week. In contrast, a different retailer that is subject to item pricing laws changes prices on just 6.3% of products weekly. They also show that for the retailer subject to item pricing laws, price changes occur three times more frequently on the items exempt from item pricing than on items for which item pricing is required. This is perhaps the first study to directly measure a link between menu costs and price stickiness. The key differences with respect to our work is that their data is at an aggregate level using either store-level data or aggregating across large groups of products (those subject to item pricing and those that are exempt).[4] As a result they do not have access to detailed controls describing differences across stores and products. In particular, they do not have access to cost data. Controlling for cost shocks is crucial as without it we cannot refute the hypothesis that the probability of price adjustment differs across products merely because of different cost shocks (or other product differences).

The same research team also has a series of studies in which they document the magnitude of menu costs in different markets. For example, Levy et al. (1998) and Dutta et al. (1999) document the price change process and provide direct measurements of menu costs at large US supermarket retailers and drugstores (respectively). The menu costs that they document are comprised primarily of in-store labor costs.[5] While these studies provide valuable documentation of the importance of in-store labor costs when adjusting retail prices, neither of the studies measure how

[4] The unit of analysis in this study is a cost change on an individual product.

[5] In contrast, Zbaracki et al. (2004) measure the magnitude of the costs of price adjustments in industrial markets and highlight the importance of managerial costs (information gathering, decision-making and communication) and customer costs (communication and negotiation).

these costs influence the frequency of price changes, which is the primary focus of this paper.

Other empirical research investigating menu costs has relied upon indirect inferences of menu costs through the frequency and magnitude of price adjustments.[6] A notable recent example is Midrigan (2012). Citing evidence in Lach and Tsiddon (2007) and Levy et al. (1997), he considers a model in which there are economies of scope in changing prices. Midrigan assumes in his theoretical work that there exists a fixed cost of changing a single price but that the cost of changing additional prices is zero. This could lead to outcomes in which items with more variants have a higher likelihood of price changes in response to cost shocks. In our work we find the opposite result: the more variants a product has the lower the probability its price will change. We interpret this finding as evidence that the cost function increases with the number of variants. Therefore, under a strict interpretation of Midrigan's model, our results may appear inconsistent with Midrigan. However, a more general interpretation of the central thesis in Midrigan's paper is that there are economies of scope, so that the cost of changing prices does increase with the number of price changes but at a decreasing rate. Our empirical results provide support for this more general interpretation; we find evidence that the marginal cost of changing the price of an additional variant increases at a decreasing rate with the number of variants.

The paper proceeds as follows. In Section 2 we describe our data together with the institutional processes used by the retailer that provided the data. In Section 3 we estimate the menu cost channel. Specifically, we investigate whether the additional opportunity cost of changing prices on items with more variants is of sufficient magnitude to influence the firm's pricing decisions. The finding that the firm is less likely to raise prices on items that have more variants provides empirical evidence that menu costs contribute to pricing decisions. In Section 4 we investigate whether the effect is temporary or enduring effects by investigating how quickly the firm changes prices in future periods. We also ask whether the effects are allocative by evaluating how they affect quantities sold in subsequent periods. The paper concludes in Section 5 with a review of the findings.

2. Description of the Data and Institutional Background

The analysis in this paper uses data provided by a large United States retailer. The retailer operates a large number of stores that sell items in grocery, health and beauty and general merchandise product categories. We begin by describing this retailer's policy of linking prices across items. We then review details of the firm's institutional

[6] Examples include Rotemberg (1982), Cecchetti (1986), Carlton (1986), Danziger (1987), Ball, Mankiw and Romer (1988), Lach and Tsiddon (1992 and 1996), Ball and Mankiw (1995), Kashyap (1995), Warner and Barsky (1995), Golosov and Lucas (2007), Midrigan (2007), Bils and Klenow (2004), Nakamura and Steinsson (2008), Klenow and Kryvtsov (2008).

processes that illustrate the importance of in-store labor costs. Finally, we conclude the section by describing the three datasets that are used in our analysis.

Uniform Pricing Rules (linkages)

Like many other retailers the firm follows a "uniform pricing" rule, which requires that all variants of a product have the same price. For example, a product such as Stacy's Pita chips has two variants (cinnamon and parmesan), while Gold Em Spices have 25 variants. This retailer assigns a common "Primary Stock Keeping Unit" (hereafter "PrimarySKU") to a family of variants, and then a Stock Keeping Unit or "SKU" number to every individual variant. A key feature of the data is that the retail price is the same for every SKU under the same PrimarySKU. Thus, if the retailer decides to change the price of the product then prices for all of its variants must change.

The use of uniform pricing policies is common across grocery retailers, and this has led to a growing academic literature focused on explaining why retailers adopt this practice. Explanations for "uniform pricing" have focused on simplifying the purchasing decision (Hauser and Wernerfelt 1990; Iyengar and Lepper 2000; and Draganska and Jain 2001), avoiding an adverse quality signal for the lower-priced item (Anderson and Simester 2001; and Orbach and Einav 2007), the managerial cost of setting different prices for different variants (Leslie 2004; and McMillan 2005), homogeneity in consumer preferences across different flavors (Draganska and Jain 2006; Anderson and Dana 2008), demand uncertainty (Orbach and Einav 2007), and customer fairness (Andersen and Simester 2008).[7]

In-Store Labor Costs

In-store labor costs represent a large portion of this retailer's cost structure. The retailer establishes budgets for labor expenses at each store and complying with these budgets plays an important role in determining both bonuses and promotions. As part of the monitoring of labor expenses, the number of regular price changes allowed is 100 SKUs per day, five days per week (Tuesday through Saturday). As a basis for comparison, this retailer's stores typically stock approximately 20,000 SKUs.

The number of price changes is calculated at the variant level, so that changing the price of two different colors of the same item is counted as two price changes. The policy is enforced by a reporting system that counts how many days each month there are more than 100 daily price changes. The same report also tracks how many items receive more than one price change within a 32-day period, and how many price changes are smaller than 4-cents. Part of the annual bonuses of specific employees depend upon these

[7] It is important to clarify that the paper does not address why this retailer has adopted uniform pricing. More generally, the optimality of the firm's pricing decisions is beyond the scope of the paper. We are also unable to speculate on how the firm's policies would change if there were a dramatic increase in the inflation rate.

measures; smaller bonuses are received when there are too many daily price changes, prices of individual items are changed too frequently, and/or there are too many small price changes. While compliance with the 4-cent policy is very high (averaging over 99%), there are many instances in which the retailer does not comply with the other two policies. In particular, compliance with the daily limit on price changes averages 91.8%, indicating that the restriction on the frequency of price changes is not trivially satisfied. As we discussed in the Introduction, if the number of price changes exceeds this capacity then labor is re-allocated within the store. The re-allocated labor is substituted from critical activities such as, inventory management, re-stocking and merchandising shelves, and serving or assisting customers.

The decision to measure the frequency of price changes at the variant level is informative. Price changes on items that have multiple variants are interpreted as multiple price changes. Thus, the uniform pricing rule that the retailer follows implies that changing the price of an item with more variants more quickly exhausts the planned capacity.

It is the uniform pricing rule that gives rise to our identification of the menu cost channel. If the cost of changing prices increases with the number of variants, then this rule induces a natural variation in the cost of changing prices across items. Holding everything else constant, products with more variants invoke a larger cost.

Description of the Data

We obtained the following three datasets from the retailer:

1. A record of every wholesale cost and regular retail price change over a 55-month period.
2. A product hierarchy mapping individual SKUs to PrimarySKUs.
3. Two hundred weeks of transaction data for each item sold at a sample of 102 stores.

The first dataset describes every wholesale cost change and every change to the regular retail price during the period between March 2005 and September 2009. This data is compiled into monthly reports. The monthly reports are used by senior management to monitor variation in profit margins in each product category, together with the frequency of price and cost changes. Moreover, the cost data is interpreted by the firm as the effective marginal cost of an item when conducting analysis to support managerial decisions. The price and cost change reports also include the total unit volume for the item over the prior 12-months. In Appendix 1 we provide formal definitions and summary statistics for each of these variables.

Notice that the price and cost change database focuses solely on regular price changes and does not consider price changes due to temporary sales. Therefore, like Golosov and Lucas (2007), and Nakamura and Steinsson (2008), we exclude temporary sales from our analysis.[8] For ease of exposition we will use the term "prices" to denote regular retail prices.

An important feature of the data is that when the wholesale cost changes the firm makes an explicit decision about whether to change the retail price at the same time. In particular, when the wholesale cost changes the record of the cost change indicates whether the retail price was changed at the same time. Discussions with the company confirm that when the category manager communicates the cost change to the operations team the category manager is required to also provide a decision on whether to change the retail price. We will exploit the simultaneity of these cost and retail price events in the next section by investigating whether a cost increase was less likely to lead to a coinciding price increase for items that have a large number of variants. In Section 4 we extend the time horizon by investigating whether cost increases lead to price increases in subsequent (or prior) periods.

Our second source of data is a hierarchy mapping individual SKUs to PrimarySKUs. This data is used to calculate the number of variants for each PrimarySKU. Throughout the paper (unless noted) we use a product hierarchy dated July 2010. The intersection of the cost data and the product hierarchy yields 11,368 cost increases and 4,194 cost decreases.

The third dataset describes weekly transactions at a sample of 102 of the firm's stores. This transaction data extends from the first week in 2006 through October 2009 (200 weeks). We will use this data in Section 4 to investigate whether the firm's initial decision to increase prices has an enduring impact on prices and quantities sold.

A shortcoming of our paper is that the results rely on a single (albeit very large) retailer. This introduces a risk that findings using data from this firm are not representative of other firms. In response, we begin by noting that while uniform pricing is a common practice among retailers, the generalizability of our findings does not rely on other retailers enforcing the same policy. We merely use the policy at this retailer to identify the variation in menu costs across items. Instead what we require is that other retailers place the same focus on managing in-store labor costs.

The existence and magnitude of in-store labor costs as a source of menu costs is now well documented in the literature (see the discussion of related literature in the Introduction). Moreover, this retailer's focus on managing in-store labor costs is standard industry practice. Because a small reduction in labor costs can have a large

[8] At this retailer temporary sales are implemented through a different operational process and have separate labor budgets allocated to them. They are not counted in the 100 per day price change capacity and in almost all cases are funded by manufacturers through a separate funding channel that does not affect the wholesale costs that we observe. Moreover, in separate research we find that the depth and frequency of temporary sales offered by this retailer are unaffected by the timing of wholesale cost changes.

impact on net profit margin, retailers carefully monitor labor costs and labor capacity is carefully planned for all standard operational tasks.[9] Time and motion studies are commonly used to identify opportunities to reduce the time spent on standard, in-store activities. For example, an apparel retailer found that saving one second from the checkout process for each customer would produce savings of $15,000 in annual labor costs across its 34 stores (O'Connell 2008). Similarly, another retailer reported that flipping a box of bananas over prior to stocking could enable employees to grasp more bananas at a time and speed up the stocking process, yielding annual savings of $100,000.

This issue of generalizability to other firms frequently arises with empirical work, particularly for studies relying on unusually detailed data. Obtaining data of this nature from a single firm is a major undertaking. Yet absent such a unique micro data set, the empirical validation and quantification of the menu cost channel is not possible. While we recognize that replication would be reassuring, this is the first study to provide a direct empirical validation that menu costs contribute to price stickiness using detailed micro-level data.

3. Measuring Menu Costs Using the Number of Variants

In this section we investigate how the number of variants influences the retailer's decision to increase prices at the time of a cost increase. We begin by describing the frequency and magnitude of cost and price changes in our data set. We then consider the distribution of price and cost changes by their absolute value. The motivation for this stems from the common prediction in menu cost models that if the cost of changing prices is meaningful then we should not observe many small price changes, particularly for higher valued items. Indeed we show that our data is characterized by almost no small price changes. We then move to the main part of the section and describe the number of variants in our sample of PrimarySKUs. We conclude this section by evaluating how the probability of a price change depends on the number of variants.

Recall that our data includes 11,368 observations in which the cost increased and 4,194 observations in which it decreased (where the unit of analysis is a PrimarySKU). In Table 1 we report how often these cost changes resulted in price changes (at the time of the cost change).

[9] Labor costs are typically the second largest source of retailers' costs and are directly controllable by the retailer (Atlanta Retail Consulting 2011). For a typical grocery store, the Food Marketing Institute estimates that labor costs are 14.8% of total sales and that net profit margins are only 3% of total sales (FMI Research 2008).

Table 1. Frequency of Cost and Price Changes

	Cost Increases	Cost Decreases
Price Increased	70.6%	5.7%
Price Decreased	0.9%	9.2%
No Price Change	28.5%	85.2%
Sample Size	11,368	4,194

The table reports the percentage of times that the retail price changed when the cost changed.

While cost increases often resulted in a price increase, cost decreases rarely led to price decreases. The potential for asymmetries in the frequency of price increases versus price decreases have been discussed elsewhere in the literature (Peltzman 2000).[10] They were also acknowledged by the retailer's management, who confirmed that the firm uses different criteria when deciding whether to change prices in response to a cost increase versus a cost decrease. For this reason we will initially restrict attention to cost increases, which represent almost 75% of the data. In Appendix 3 we turn attention to cost decreases and highlight additional asymmetries in the retailer's response.

The Size of the Cost and Retail Price Changes

The size of the retail price changes offers a preliminary check on the role of menu costs at this retailer. If menu costs are meaningful then we should not observe many small price changes, particularly for higher valued items. In Table 2 we report the distribution of the absolute magnitude of the retail price changes grouped by the prior retail price of the item. As a basis for comparison we also report the distribution of the absolute size of the cost changes.

Small cost changes are common in the data, but small price changes are rare. For example, 35% of the 15,562 cost changes are less than 10-cents in absolute magnitude, but just 4% of the 12,824 price changes fall within this range. This scarcity of small price changes is what we would expect in the presence of menu costs.

[10] Other references include Karrenbrock 1991; Neumark and Sharpe 1992; Borenstein, Cameron and Gilbert 1997; Jackson 1997; Noel 2009; Hofstetter and Tover 2010; and Green, Li and Schurhoff 2010. Peltzman (2000) does not find any evidence of this asymmetry when studying price changes at Dominick's Finer Foods in Chicago. He attributes this null finding to a distinction between an individual firm decisions and market outcomes. Our findings could be considered a counter-example to Peltzman's supermarket example. Notably, the retailer in this study and the supermarket in Peltzman's study compete in similar retail markets.

Table 2. Absolute Size of Cost and Price Changes By the Prior Retail Price

Absolute Size of Price Changes

Prior Retail Price	Under 10-cents	Under 20-cents	Under 50-cents	Sample Size
Under $5	6%	19%	72%	5,326
$5 to $10	3%	6%	36%	3,834
$10 to $15	3%	4%	13%	1,654
$15 to $20	1%	3%	10%	812
$20 to $30	1%	2%	4%	647
$30 to $40	4%	5%	9%	215
$40 to $50	0%	0%	5%	124
Over $50	0%	0%	2%	212
Total	4%	11%	44%	12,824

Absolute Size of Cost Changes

Prior Retail Price	Under 10-cents	Under 20-cents	Under 50-cents	Sample Size
Under $5	59%	83%	97%	7,018
$5 to $10	22%	48%	85%	4,363
$10 to $15	8%	14%	50%	1,746
$15 to $20	5%	8%	33%	888
$20 to $30	4%	7%	18%	764
$30 to $40	2%	4%	15%	329
$40 to $50	4%	5%	13%	217
Over $50	5%	6%	7%	237
Total	35%	54%	76%	15,562

The table reports the absolute size of cost and retail price changes grouped by the item's prior retail price.

The Number of Variants

In Table 3 we report frequency distributions of the number of SKUs under each PrimarySKU (*NUMBER OF SKUS*). The first column is a distribution of the number of PrimarySKUs, while the second column is a distribution of the number of SKUs. The last two columns report the distribution of revenue and units sold in the previous 12 months. The 8,073 PrimarySKUs include a total of 12,932 individual SKUs with an average of 1.60 SKUs per PrimarySKUs and a maximum of 62 (the brand of nail polish referred to in the opening paragraphs). The frequency distribution reveals that, while PrimarySKUs with a single variant represent 79.9% of all PrimarySKUs, they only represent 59.3% of revenue and 49.9% of individual SKUs.

Table 3. Frequency Distribution of the Number of SKUs under Each PrimarySKU

NUMBER OF SKUS	PrimarySKU Frequency	SKU Frequency	Revenue Weighted	Units Weighted
1	79.9%	49.9%	59.3%	50.4%
2	9.3%	11.6%	13.4%	14.0%
3	4.1%	7.8%	8.2%	8.6%
4	2.3%	5.7%	5.3%	6.6%
5	1.2%	3.9%	3.4%	4.0%
6	0.9%	3.3%	2.6%	3.7%
7	0.5%	2.2%	1.6%	1.8%
8	0.4%	1.8%	0.8%	1.5%
9	0.3%	1.5%	0.6%	0.7%
10	0.1%	0.9%	0.3%	0.9%
11	0.1%	0.6%	0.3%	1.0%
12	0.1%	0.9%	0.4%	0.6%
13	0.1%	1.2%	0.7%	1.4%
14	0.1%	0.9%	0.3%	0.6%
15	0.1%	0.9%	0.3%	0.5%
Over 15	0.4%	7.0%	2.6%	3.6%

The table reports a frequency distribution of the *NUMBER OF SKUS* by PrimarySKU, by SKU, by revenue and by units. The revenue and units measures are calculated using the prior 12-months of sales data (reported in the cost and price change reports). The sample includes 8,073 PrimarySKUs for which there was either a cost change or retail price change in our 55-month data period. In all four distributions we exclude items for which the *NUMBER OF SKUS* or either of the weighting variables is missing.

The Number of Variants and the Probability of a Price Change

In Figure 1 we report how the probability of a price increase (following a cost increase) changes according to the *NUMBER OF SKUS*. We report both weighted and unweighted averages, where the weighting uses the previous 12-months of revenue for each PrimarySKU. The findings reveal a strong negative relationship: items with more SKUs were less likely to receive a price increase following a cost increase. This is a key finding in the paper, and is consistent with the firm forgoing price increases in order to avoid larger menu costs on items with additional variants.

Figure 1a. The Probability Prices Increase Following a Cost Increase (Unweighted)

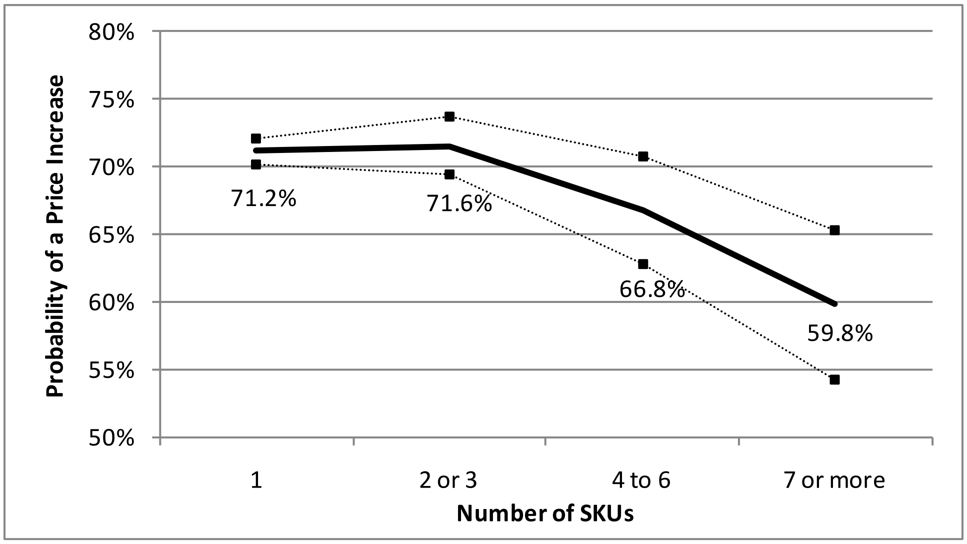

Figure 1b. Weighted by Prior Revenue

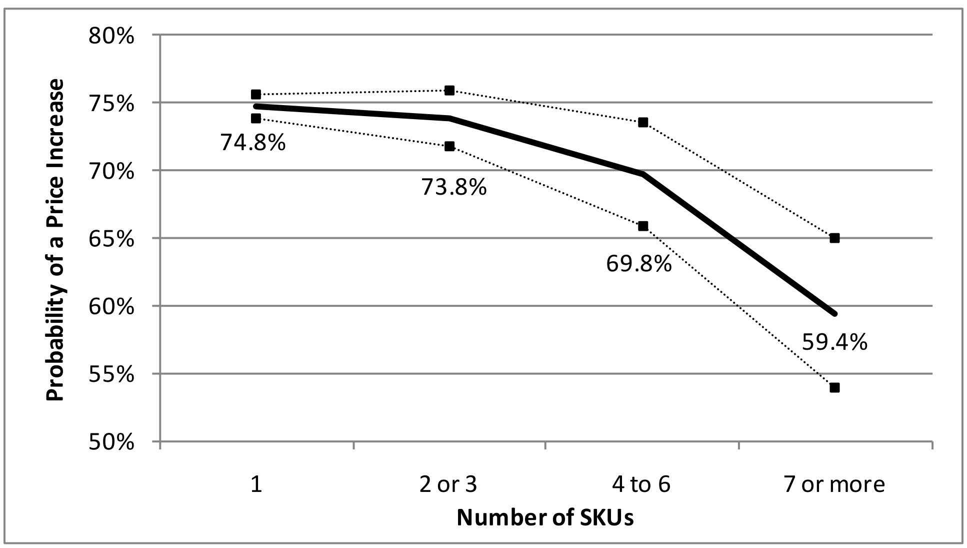

The figures report the probability of a price increase following a cost increase. The square markers indicate the 95% confidence interval. We report both weighted and unweighted averages, where the weighting uses total revenue over the prior 12-months. There are a small number of missing observations for this weighting variable. To maintain a consistent comparison we omit these observations from both the weighted and unweighted averages.

To evaluate the importance of the effects in Figure 1 it is helpful to understand how the reluctance to raise prices on items with more variants affects the overall frequency of price changes at this firm. We address this issue by asking the following question: if the probability of a price increase was the same for items with multiple variants as it is for items with a single variant how many more price increases would we see? To answer this question, we calculate the probability of a price increase following a cost increase for items that had only a single variant. We restrict attention to items that had at least one cost increase in our 55-month data period.

There were 8,701 cost increases on items with a single variant, and these cost increases resulted in 6,191 price increases. Therefore, the probability of a price increase following a cost increase on an item with a single variant is 71.2% (see Figure 1a). Using this probability we calculate how many "projected" price increases we would expect to observe on items with multiple variants if price increases on these items occurred at the same rate. The findings are reported in Table 4.[11] Cost increases on items with multiple variants represented a total of 10,491 increases on individual SKUs. If price increases occurred at the same rate as on items with a single variant then we would have observed 7,465 price increases, which is 541 (8%) more than we actually observed. When weighting by revenue, there would have been 1,191 additional price increases, or 18% more.

Table 4. Overall Frequency of Price Increases on Items With Multiple Variants

	Unweighted	Weighted (by Revenue)
Cost Increases	10,491	10,491
Actual Price Increases	6,924	6,654
Projected Price Increases	7,465	7,845
Projected minus Actual	541	1,191

The table reports the actual number of cost and price increases on items with at least two variants. The table also projects how many price increases would occur if price increases on these items occurred at the same rate as they occur on items with a single variant. We report both weighted and unweighted results, where the weighting uses total revenue over the prior 12-months.

These initial findings are consistent with our interpretation that the menu costs associated with changing the retail price are larger when an item has more variants, and that this leads to price stickiness. However, this is the not the only explanation for these univariate results. Notably, it is possible that the relationship may reflect an interaction between *NUMBER OF SKUS* and other factors that contribute to the decision to increase prices. Next, we estimate a series of models that control for this possibility.

[11] When weighting by revenue the probability of a price increase on items with a single variant increases to 74.8%. We use this probability in the weighted analysis.

Other Factors That Contribute to the Decision to Raise the Price

There are several factors in addition to menu costs that could contribute to the decision to raise the retail price following a cost increase. For example, we would expect the size of the cost change, the purchase volume, and the prior profit margin to influence the decision to change prices. The larger the cost change, the more likely we will observe a price increase. Larger purchase volumes increase the profit implications of changing prices and so we would expect retailers to prioritize price increases on higher volume items. Similarly, discussions with the retailers' pricing managers confirm that they focus on maintaining profit margins within a targeted range. This suggests that if prior to the cost increase the profit margin was low, then the retailer is more likely to respond to cost increases that push the profit margin further outside the targeted range. Collectively these arguments suggest that price increases will be more likely when the cost change and unit volume are larger and the prior profit margin was lower.

There is also now an extensive literature establishing that there is a kink in the demand curve around 99-digit price endings (for example $2.99). Levy et al. (2010) present evidence that retailers seek to preserve these price endings, and are less likely to increase prices that currently end with 99-cents (see also Knotek 2008 and 2010). The retailer's pricing policy suggests that this retailer recognizes the kink in the demand function; approximately 45% of the retailer's prices end with 99-cents. Therefore we construct a binary variable indicating whether the prior retail price ended in 99-cents (*Prior 99-cent Price Ending*).[12]

It is also possible that a cost increase is more likely to lead to a price increase in larger product categories, in which there are more substitutes. Raising the price in these categories is likely to result in a smaller effect on category sales, as customers are more likely to substitute purchases to other items. To measure product category size we count the number of PrimarySKUs in each item's product category (*Category Size*).

In Table 5 we report the marginal effects from a logistic model in which the unit of observation is a cost increase on a PrimarySKU, and the dependent variable is a binary variable indicating whether the price increased. The independent variables include the *NUMBER OF SKUS* for that PrimarySKU. We also report an alternative specification, including the log of *NUMBER OF SKUS* (Model 2). For completeness the models include fixed year and month effects. The marginal effects for the different specifications are reported in Table 5 (to ease exposition we omit the year and month fixed effects). In all of the models standard errors are clustered by the month of the observation.[13]

[12] We also investigated prices that end with 9-cents (such as $1.49). However, almost all of the prices at this retailer have a 9-cent ending (over 95%), making it difficult to reliably estimate the impact of a 9-cent ending versus other single-digit endings.

[13] We also considered clustering by the product category. However, there are too many categories for clustering to be meaningful.

The findings in Table 5 confirm that the relationship between *NUMBER OF SKUS* and the probability of a price increase survives controlling for all of these explanatory variables. The larger the *NUMBER OF SKUS* the lower the probability of a price increase following a cost increase.

To help interpret the magnitude of this relationship we also estimated a linear probability model using OLS. We use binary indicator variables to identify items with 2 or 3 variants, 4 to 6 variants, or 7 or more variants.[14] These findings are reported as Model 3 in Table 5. The findings reveal that if there are 2 or 3 variants then the probability of a price increase following a cost increase is 4.3% lower than when the item has only a single variant. If there are 4, 5 or 6 variants, the probability is 12% lower compared to a single variant, and if there are 7 or more variants the probability difference is 20%.

We note that Model 2, which uses a log transformation of *NUMBER OF SKUS*, yields a small improvement in model fit compared to Model 1. This could be interpreted as evidence supporting Midrigan's (2012) argument that the cost of changing prices exhibits economies of scope. This outcome seems reasonable in this setting; changing the price of the 20th nail polish color would seem to be easier than initially finding the nail polish in the store and changing the price of the first variant. In Model 3 we observe further evidence that the marginal cost of changing the price of an additional variant is decreasing in the number of variants. The average *NUMBER OF SKUS* for items with 2 or 3 variants is 2.32; for items with 4 to 6 variants is 4.67; and for items with 7 or more variants is 12.74. In Figure 2 we plot the implied probability of a price increase against these averages. The resulting plot exhibits a decreasing effect as the number of variants increase, which is also consistent with economies of scope.

[14] This is the grouping of NUMBER OF SKUS that we used in the univariate analysis (see Figure 1).

Table 5. Factors that Contribute to the Decision to Raise the Price

	Model 1	Model 2	Model 3
NUMBER OF SKUS	-0.0122** (0.0030)		
Log NUMBER OF SKUS		-0.0723** (0.0126)	
NUMBER OF SKUS = 2 or 3			-0.0435** (0.0134)
NUMBER OF SKUS = 4 to 6			-0.1249** (0.0296)
NUMBER OF SKUS = 7 or more			-0.1987** (0.0369)
Prior 99-cent Price Ending	-0.1752** (0.0187)	-0.1774** (0.0184)	-0.1898** (0.0200)
Size of Cost Change (%)	0.3643** (0.0972)	0.3649** (0.0966)	0.2080** (0.0617)
Prior Profit Margin (%)	-0.7262** (0.0530)	-0.7449** (0.0529)	-0.7785** (0.0601)
Purchase Volume (log)	0.0091* (0.0039)	0.0116** (0.0038)	0.0116** (0.0042)
Category Size (00's)	0.0164* (0.0082)	0.0162* (0.0080)	0.0125* (0.0056)
Model	logistic	logistic	OLS
Log pseudolikelihood	-5,745	-5,725	
R^2 or pseudo R^2	0.1421	0.1450	0.1635
Sample Size	10,998	10,998	10,998

The table reports marginal effects from logistic models (Models 1 and 2) and coefficients from an OLS model (Model 3). In all three models the dependent variable is a binary variable indicating whether the retailer increased its price following a cost increase. Fixed year and month effects were included but are omitted from this table. Standard errors are in parentheses. The standard errors are clustered by the month of the observation (month*year).
*significantly different from zero, $p < 0.05$; **significantly different from zero, $p < 0.01$.

Figure 2. Estimated Probabilities of Price Increase Following a Cost Increase

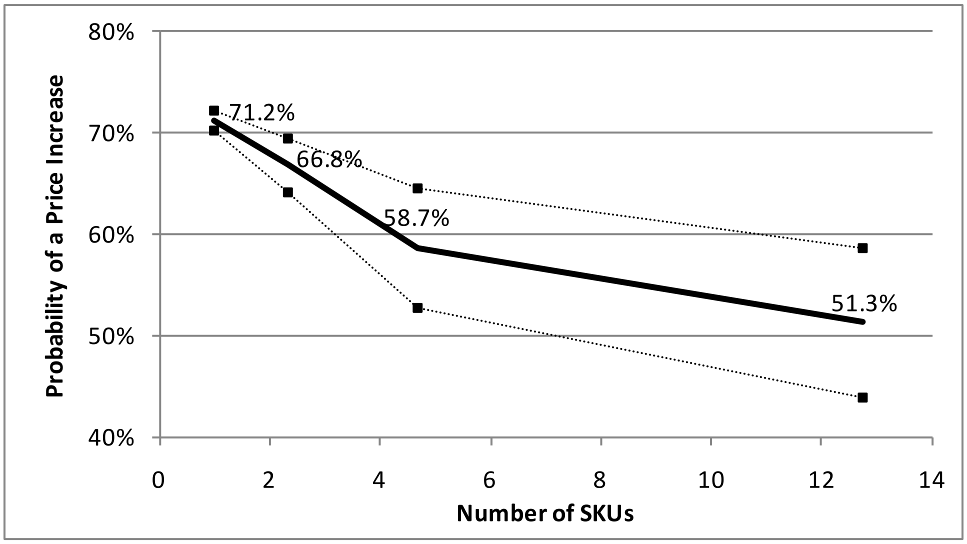

Figure 2 interprets the coefficients from Table 5 by reporting the implied probability of a price increase following a cost increase. We index the findings by setting the probability of a price increase for an item with just one variant at 71.2% (the actual probability). The x-axis uses the average NUMBER OF SKUS in each of the four product groupings.

Beyond the *NUMBER OF SKUS*, the coefficients for the other variables reveal several additional findings of interest. First, as expected, the retailer is more likely to increase the price when the cost increase is larger. Second, there is a significant effect of the prior profit margin on the probability of a price change. When the initial profit margin is lower the firm is more likely to respond to a cost increase with a price increase. Third, if the prior price ended with 99-cents there is a lower probability of a price change. This is consistent with the findings previously reported by Levy et al. (2010), and suggests that the firm finds it profitable to maintain prices just below the kink in the demand curve. Fourth, the firm is more likely to increase prices on items with larger purchase volumes. Finally, we also see evidence that the firm is more likely to increase prices on items that are in larger product categories.

Robustness Checks

In what follows we describe four different robustness checks.

Differences in profit margins and quantities sold

If items with more variants have lower profit margins and/or sell fewer total units then the benefit of changing prices on these items would be lower. This may explain why the firm is less likely to raise prices on these items. However, there is strong evidence to refute this explanation. First, we explicitly control for the profit margins and unit volumes in our analysis. Second, we compare the median unit volume, revenue and gross profits at the item level. This comparison reveals that items with multiple variants generally sell more units and have profit margins that are just as high as products with a single variant.

Category fixed effects

Thus far we have exploited variation in *NUMBER OF SKUS* both within and across categories (an example of a category is “soda beverages,” which includes multiple PrimarySKUs). In practice, much of the variation in *NUMBER OF SKUS* occurs across categories and so to investigate the robustness of the results we re-estimated the models with category fixed effects. The coefficients of interest again remain statistically significant, though smaller in magnitude than those reported those in Table 5. We conclude that our main empirical findings are not driven solely by variation between categories.

Frequency of cost shocks

Our discussion of the institutional details in Section 2 suggests that the relationship between *NUMBER OF SKUS* and the probability of a price increase reflects the retailer’s focus on minimizing in-store labor costs. Notice that the in-store labor costs associated with changing prices are only relevant for the retailer and do not extend to the manufacturer. Therefore, if our interpretation is correct we would not expect to see a relationship between *NUMBER OF SKUS* and the frequency of cost changes. To investigate this we regressed the number of cost changes in our data period on the same set of explanatory variables that we used in Table 5. The findings confirm that there is no evidence that the *NUMBER OF SKUS* contributed to systematic variation in the frequency of cost changes.

Endogeneity of the number of variants

A final possible concern is that there are other unobserved factors that influence both the probability of a price increase and the number of variants. In Appendix 2 we explore the sources of variation in the *NUMBER of SKUS* and address the potential endogeneity with an instrumental variables (IV) model. We show that all the results reported in Table 5 survive in the IV model.

Isolating When Menu Costs Play a More Prominent Role

We would expect that the contribution of menu costs to price stickiness would depend upon the firm’s other motivations for changing the price. In particular, we know from Table 5 that when the cost increase is large enough the firm is highly motivated to increase the price. In these cases we might expect a price change irrespective of the menu costs. In contrast, when the cost change is small the motivation to increase prices is weaker. This is when we would expect menu costs to play a more prominent role.

To investigate this prediction we used the median sized cost increase (5.98%) to split the sample into two sub-samples of equal size. In Figure 3 we illustrate how the probability of a price increase varied with the number of variants for each of these sub-samples. For large cost increases we see that the number of variants has relatively little impact on the probability of a price increase. When we move from an item with a single variant to an item with 7 or more variants the probability of a price increase only varies from 72% to 69%.

Figure 3. Estimated Probabilities of Price Increase Following a Cost Increase by Size of Cost Increase

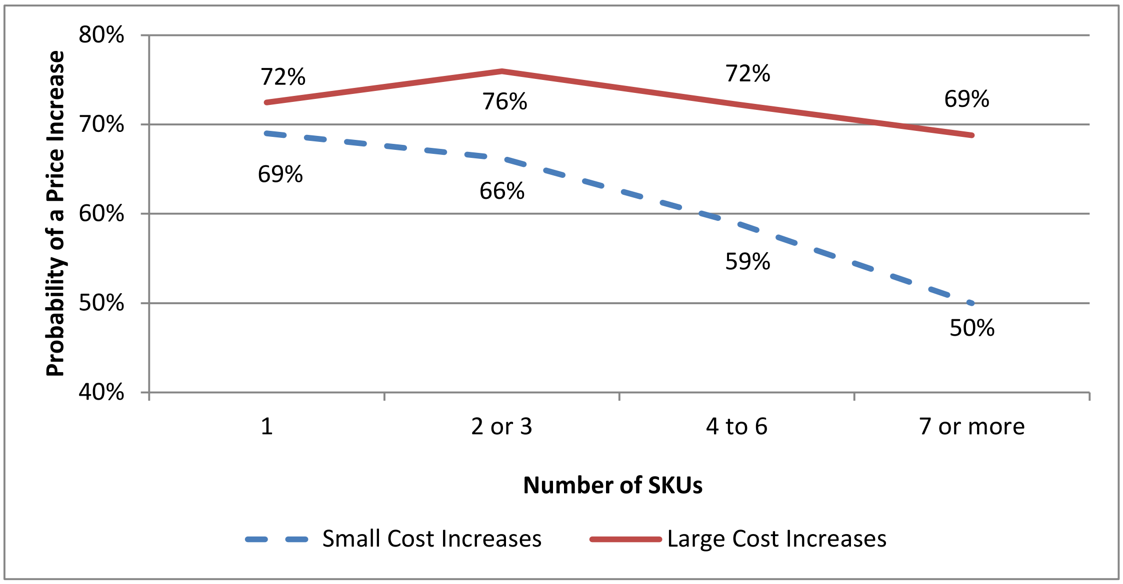

Figure 3 report the probability of a price increase following a cost increase. The figure uses the same sample of observations as those used in Table 5. Small cost increases include the 5,496 observations with cost increases less than the median (5.98%), and the large cost increases include the 5,499 observations with cost increases larger than the median.

However, when the cost increase is small menu costs appear to play a much larger role. For items with a single variant the probability of a price change is relatively high (comparable to the probabilities for large cost increases). This is consistent with relatively small menu costs providing little disincentive to changing prices. However, as the number of variants grows, the probability of a price increase (after a small cost increase) falls from 69% to just 50%. These findings can be interpreted as some of the first empirical evidence of boundary conditions for the menu cost channel. They suggest that the menu cost channel plays a more central role when cost changes are relatively small and the firm has a weaker motivation for changing prices.

In Appendix 1 we formally estimate how the interaction between *NUMBER OF SKUS* and *Size of Cost Change* affects the probability of a price increase in our multivariate model. The findings replicate the pattern observed in Figure 3. The larger the number of variants the lower the probability of a price increase, but this effect is attenuated when the cost increase is larger.[15]

Summary

We have shown that the firm is less likely to respond to a cost increase by increasing the price when the item has more variants. While cost increases lead to price increases 71.2% of the time on items with a single variant, this falls to 59.8% on items with seven or more variants. We attributed this increased stickiness in the prices of items with more variants to the additional in-store labor costs of changing prices on items with multiple variants. In particular, this behavior is consistent with the firm's internal management policies, which are designed to deter frequent price changes that would exceed the firm's in-store labor capacity. Reassuringly, the findings survive controlling for a range of other factors that contribute to the decision to increase prices following a cost increase.

We have exploited a unique feature of the data in that it documents each unique cost change event, and the firm's explicit pricing decision at that time. For this reason our analysis has focused on *immediate* price changes that coincide with the cost change. However, it is possible that the pricing response to a cost change is merely delayed to allow the firm to smooth out price changes and operate within its in-store labor capacity constraint. We investigate this possibility in the next section by asking: how sticky is sticky? In particular, if the firm forgoes a price increase when a cost increases, does this accelerate the timing of the next price increase?

4. How Sticky is Sticky?

The analysis in this section proceeds in three steps. First, we focus on the items that had cost increases and report how many of them had a price increase within the next 30, 90, 180 or 360 days. In doing so we exclude any initial price increases that occurred at the time of the cost increase (these are studied in the previous section). Second, we use our sample of store transaction data and report the weekly trends in the prices and profit margins in the weeks after the cost increase. Finally, we conclude the section by investigating whether the effects are allocative, in the sense that they had long-term impacts on the quantities purchased.

Frequency of Subsequent Price Increases

To investigate the possibility that not raising the price at the time of a cost increase merely results in a short delay in the timing of the price increase, we compare the incidence of *future* price increases. In particular, for each observation (cost increase) we report whether there was a price increase in the next 30 days, 90 days, 180 days and

[15] The interaction coefficient is significantly different from zero ($p < 0.01$) when using the log transformation of NUMBER OF SKUS (Model 2), but only marginally significant ($p < 0.10$) without the log transformation (Model 1).

360 days. For cost increases that had an immediate price increase we do not count this initial price increase when evaluating whether there was a price increase in the subsequent periods.[16] The results are reported in Table 6, where we distinguish between items with and without an immediate price increase.

Table 6. Proportion of Items with a Future Price Increase

	Items WITHOUT Immediate Price Increase	Items WITH Immediate Price Increase	Difference
Future Price Increases			
Within 30 days	1.77%	1.31%	0.47%*
	(0.21%)	(0.11%)	(0.22%)
Within 90 days	5.82%	4.52%	1.30%**
	(0.37%)	(0.21%)	(0.41%)
Within 180 days	12.48%	10.75%	1.73%**
	(0.55%)	(0.33%)	(0.62%)
Within 360 days	24.31%	29.75%	-5.44%**
	(0.83%)	(0.57%)	(1.04%)
Sample Sizes			
Within 30 days	4,113	9,791	13,904
Within 90 days	3,951	9,561	13,512
Within 180 days	3,630	8,971	12,601
Within 360 days	2,670	6,373	9,043

The table reports the proportion of items that had a future price increase within the indicated periods. For example, there were 4,113 cost increases on items that did not have a price increase at the time of the cost increase and the cost increase occurred at least 30-days before the end of the data period. Among these items 1.77% had a price increase in the next 30-days. The difference in sample sizes reflects the omission of cost increases for which the data ended (or the item was discontinued) after the shorter evaluation period but before the end of the longer evaluation period. *significantly different from zero, $p < 0.05$; **significantly different from zero, $p < 0.01$.

Of the cost increase events without an immediate price increase, there were 4,113 events that occurred at least 30 days before the end of the 55-month data period.

[16] We also exclude any items that were discontinued within the specified evaluation period or cost increases that occurred too close to the end of the data period to observe whether there were subsequent price increases.

Among these, only 1.77% had a price increase in the 30 days after the cost increase. This is only a very small proportion and indicates that for the vast majority of these events forgoing a price increase at the time of the cost increase is not merely a short delay in the timing of the price increase. After 180 days, we only observe a subsequent price increase for 12.48% of these events.

The results in Table 6 reveal that items that do not have an immediate price increase were slightly more likely to have a future price increase in the next 30-days, 90-days and 180-days. However, the difference in these probabilities between the two groups of items is relatively small. The differences are not large enough to compensate for the difference in the initial decision whether to increase the price at the time of the cost increase.

Because this retailer requires that manufacturers provide advance notice of impending cost changes, it is possible that the retailer changes the retail price before the cost changes take effect. In particular, it is possible that the retailer increased prices in anticipation of future cost increases. To investigate this possibility we also evaluated the incidence of prior price increases. This analysis reveals that only 2% of the items had a prior price increase within 30-days of the cost increase. This proportion is slightly higher for items on which the firm did not raise prices at the time of the cost increase (3.7% versus 1.4%). While this may reflect the firm raising prices in anticipation of the cost increase (rather than raising prices at the time of the cost increase), this difference is again too small to fully compensate for the different pricing decisions at the time of the cost increase.

We have shown that there is little evidence that the firm "smoothes" the frequency of price changes by delaying price changes to future periods (or accelerating them to prior periods). The incidence of future (and prior) price increases is only slightly higher if the price did not increase when the cost increased. However, simply comparing the incidence of future price changes does not capture any differences in the magnitude of those future price increases. Thus in our next analysis we compare the prices and profit margins of the two groups of items over the subsequent 12 months. This reveals the net impact of the (small) difference in the frequency of future price changes, together with any differences in the magnitude of those price changes.

Do the Prices and Profit Margins Recover?

Recall that we documented in Section 3 that the firm was more likely to raise prices on items that had lower initial profit margins. We interpreted this as evidence that the firm's pricing decisions were designed to maintain profit margins within a "target range". While the initial decision to increase prices on some items and not on others will lead to initial divergence in the profit margins of these items, we would expect that this divergence will eventually be mitigated if the firm's goal is to maintain margins within a target range.

To investigate the trend in the prices and profit margins we turn to our sample of weekly store transactions. Recall that our transaction data reports aggregate weekly store-level transactions for every item in a sample of 102 stores for the period between January 2006 and October 2009. We use the 52 weeks before the cost change to calculate baseline averages for the *Retail Price* and *Profit Margin* for each item.[17] In the 52-weeks after the baseline we then calculate the percentage change from the baseline to yield a weekly *Retail Price Index* and a weekly *Profit Margin Index* for each item. To compare how the initial pricing decision at the time of the cost increase affected these indexes we then calculate the difference in the weekly indexes between items that had an initial price increase and those that did not. The findings are reported in Figure 4.

Figure 4. Difference in Price and Profit Margin Indexes for Items With and Without a Price Increase at the time of the Cost Increase

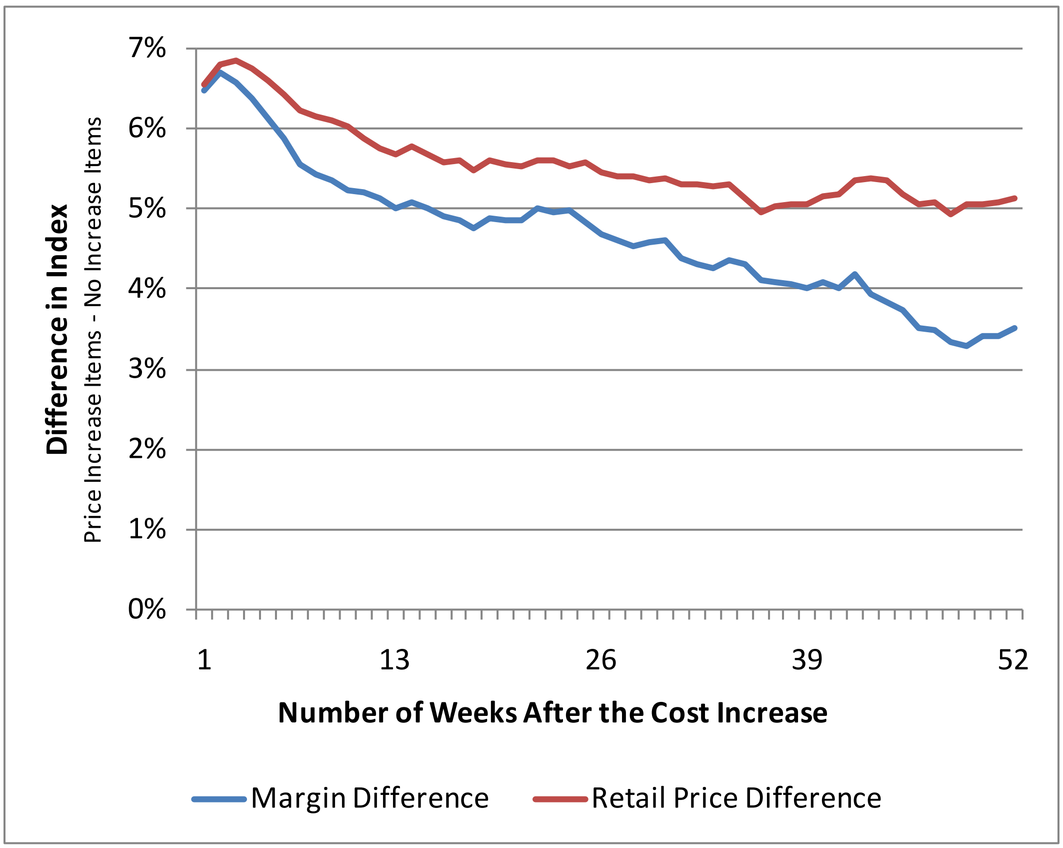

The x-axis identifies the number of weeks after a cost increase. The y-axis describes the difference in the average *Profit Margin* and *Retail Price* indexes for items that had a price increase in Week 0 (at the time of the cost increase) and items that did not. Positive (negative)

[17] In the transaction data we only observe prices and profit margins in weeks for which there was a transaction and so we restrict attention to items for which there were 52 weeks of consecutive transactions after the cost increase and 52 weeks of consecutive transactions before the cost increase. This ensures that the weekly outcomes are calculated using the same sample of items. For items with multiple qualifying cost increases (105 weeks of consecutive sales) we focus on the first cost increase.

values indicate that items that had a price increase had a higher (lower) index than items that did not. The indexes represent percentage changes from the prior 52-week baseline period. The analysis only includes items that had sales in each of the 52-weeks before and after the cost increase. For items that had multiple qualifying cost increases we focus on the first cost increase. The sample sizes include 1,701 items that had a price increase at the time of the cost increase and 633 items that did not.

As we would expect, the initial price increase at the time of the cost increase resulted in higher prices and higher margins for these items. This is reflected in Figure 4, where the difference in the indexes indicates that in the weeks immediately after the cost increase, prices and profit margins were almost 7% higher on items that had an initial price increase (compared to items without an initial price increase). In subsequent weeks the differences steadily decreased, indicating that the price and margin indexes for the two groups of items began to converge. After 52-weeks approximately half of the initial difference in the profit margins remained.

The trends in Figure 4 confirm that the effects of the initial pricing decision are enduring rather than transitory. Persistent differences in both the indexed retail prices and the profit margins remain even a year later. We next ask whether these differences were allocative.

Is the Decision to Increase Prices Allocative?

To evaluate the importance of the effects that we report it is helpful to understand the extent to which they influenced the number of items that were purchased in subsequent periods. We address this issue by comparing sales in the 52-weeks after the cost increase between our two groups of items. In Table 7 we report these averages as the percentage difference compared to the 52-week baseline period.[18] As a basis of comparison we also report the average retail price and cost in the year after the cost increase. When averaging across items the items are weighted using the quantity sold in the baseline period.

[18] To control for changes in product lines at different stores we restrict attention to stores in which the item was introduced prior to the 52-week baseline period and was continued through the 52-week evaluation period. This yields a larger sample size than the analysis in Figure 5, where we required that there were sales in every week of the baseline and evaluation periods.

Table 7. Transaction Outcomes in the 52-Weeks After the Cost Increase

	Items WITHOUT Immediate Price Increase	Items WITH Immediate Price Increase	Difference
Quantity Sold	1.91%	-3.21%	-5.12%**
	(0.70%)	(0.42%)	(0.77%)
Retail Price	6.53%	12.49%	5.96%**
	(0.25%)	(0.15%)	(0.27%)
Cost	7.21%	9.41%	2.19%**
	(0.14%)	(0.31%)	(0.29%)
Sample Size	1,033	2,674	

This table reports the percentage change in each measure in the 52-weeks after the cost increase compared to the previous 52 weeks. The measures for each item are weighted using the quantity sold in the baseline period. The table distinguishes between observations in which the cost increase resulted in an immediate price increase, and those that did not. Standard errors are in parentheses. *significantly different from zero, $p < 0.05$; **significantly different from zero, $p < 0.01$.

The findings confirm that the decision to increase the price in response to a cost increase is allocative; there is a significant impact on the quantity sold over the next 52-weeks. Without a price increase there was a 1.91% increase in quantity sold compared to the Baseline Period. In contrast, items with an initial price increase had a 3.21% *decrease* in units sold over the same period. The difference between these two outcomes reveals a net loss of 5.12% in unit sales growth for items that had a price increase at the time of the cost increase. This 5.12% lower sales growth can be compared with the 5.96% larger increase in retail prices over the same period. This corresponds to an average price elasticity of approximately -1.

Summary

We have presented evidence that the decision to forgo a price increase when the cost increases is not just a decision to delay the price increase. Moreover, a comparison of indexed prices and profit margins in the period after the initial price increase confirms that the effects of the initial pricing decision are persistent. The retail prices and profit margins converge slowly over the next 52-weeks, with enduring differences even at the end of the period.

We also investigated whether the outcome is allocative by evaluating how the decision to increase prices at the time of the cost increase influences the prices and quantities sold over the next 52-weeks. The findings confirm that the initial pricing decision is not just sticky, it also affects transaction outcomes. Items that had an initial price increase experienced a net 5.12% drop in units sold (relative to items without an initial price increase).

5. Conclusions

Starting with the seminal work of Barro (1972) and Sheshinksi and Weiss (1977), much of the analysis of monetary policy effects has relied on models with fixed costs of price adjustment. Yet, there has been little micro evidence validating and quantifying the effects these costs have on the probability of price adjustments.

Building on a 55-month database of cost and price changes at a large retailer this paper helps to fill this gap. We find that absent menu costs, the number of price changes would increase by up to 18%. The identification of this effect stems from the retailer's pricing rule that requires all variants of a product to have the same price. Since different products have a different number of variants, this pricing rule leads to variation in the opportunity cost of changing prices across products.

In addition to documenting existence of the menu cost channel, we also identify boundary conditions. When cost increases are very large, we find that the decision to raise price is independent of menu costs. But, for smaller cost increases we find that the menu cost channel plays a central role.

A limitation of our research is that we study a single retail chain. However, this allows us to understand important institutional facts and acquire unique data. Both are essential to understanding the menu cost channel. Further, while we study a single retailer we analyze a census of products offered by the retailer. In this sense, our empirical study is extremely large as we analyze thousands of brands and hundreds of manufacturers. Future work is needed to investigate the menu cost channel in other empirical settings.

References

Akerlof, George A., and Janet L. Yellen (1985). "Can Small Deviations from Rationality Make Significant Differences to Economic Equilibria?" *American Economic Review*, 75 (September 1985), pp. 708–20.

Akerlof, George A., and Janet L. Yellen (1985). "A Near-rational Model of the Business Cycle, with Wage and Price Inertia," *Quarterly Journal of Economics*, vol. 100 (Supplement 1985), pp. 823–38.

Atlanta Retail Consulting (2011). *The Retail Store Labor Management Improvement Process*, http://www.atlantaretailconsulting.com/documents/ARC_Labor_Management_Whitepaper.pdf.

Anderson, Eric T. and James Dana (2009). "When is Price Discrimination Profitable?," *Management Science*, 55(6), June, 980–989.

Anderson, Eric T. and Duncan I. Simester (2001). "Does Demand Fall When Customers Perceive That Prices Are Unfair? The Case of Premium Pricing for Large Sizes," *Marketing Science*, 27(3), 492-500.

Anderson, Eric T. and Duncan I. Simester (2008). "Price Discrimination as a Signal: Why an Offer to Spread Payments May Hurt Demand," *Marketing Science*, 20(3), 315-327.

Barro, Robert J. (1972). "A Theory of Monopolistic Price Adjustment," *The Review of Economic Studies*, 39 (January 1972), 17–26.

Bertola, Giuseppe, and Ricardo Caballero (1990). "Kinked Adjustment Costs and Aggregate Dynamics," in *NBER Macroeconomics Annual*, Cambridge: The MIT Press.

Bijmolt, Tammo H. A., Herald J. Van Heerde, and Rik G. M. Pieters (2005), "New Empirical Generalizations on the Determinants of Price Elasticity," *Journal of Marketing Research*, Vol. XLII, 141-156.

Borenstein, Severin; Cameron, A. Colin; and Gilbert, Richard, (1997). "Do Gasoline Prices Respond Asymmetrically to Crude Oil Price Changes?" *Quarterly Journal of Economics*, 112, (February), 305-39.

Burstein, Ariel (2006). "Inflation and Output Dynamics with State Dependent Pricing Decisions," *Journal of Monetary Economics*, 53(7) (October), 1235-1257.

Caplin, Andrew, and Daniel Spulber (1987). "Menu Costs and the Neutrality of Money," *Quarterly Journal of Economics*, 102 (November 1987), 703–25.

Caplin, Andrew, and John Leahy (1997). “Aggregation and Optimization with State-Dependent Pricing,” *Econometrica*, 65 (May 1997), 601–25.

Caplin, Andrew, and John Leahy (1991). “State-Dependent Pricing and the Dynamics of Money and Output,” *Quarterly Journal of Economics*, 106 (August 1991), 683–708.

Danziger, Leif. (1999). “A Dynamic Economy with Costly Price Adjustments,” *American Economic Review*, 89 (September 1999), 878–901.

Dotsey, Michael, Robert G. King, and Alexander L. Wolman. “State-Dependent Pricing and the General Equilibrium Dynamics of Money and Output,” *Quarterly Journal of Economics*, 114 (May 1999), 655–90.

Draganska, Michaela and Dipak C. Jain (2001). “Product-Line Length Decisions in a Competitive Environment.” Working paper, Graduate School of Business, Stanford University, Palo Alto CA.

Draganska, Michaela and Dipak C. Jain (2006). “Consumer Preferences and Product-Line Pricing Strategies.” *Marketing Science*, 25(2), (March-April), 164-174.

FMI Research (2008). *The Food Retailing Industry Speaks, 2008.* Food Marketing Institute, Arlington VA.

Gertler, Mark and John Leahy (208). “A Phillips Curve with an Ss Foundation,” *Journal of Political Economy*, 116(3), 533-572.

Golosov Mikail, Robert E. Lucas Jr. (2007). “Menu Costs and Phillips Curves”, *Journal of Political Economy*, 115(2), 171-199.

Green, Richard C., Dan Li and Norman Schuerhoff (2010). “Price Discovery in Illiquid Markets: Do Financial Asset Prices Rise Faster Than They Fall?” *Journal of Finance*, 65(5), 1669-1702.

Hauser, John R. and Birger Wernerfelt (1990). “An Evaluation Cost Model of Consideration Sets.” *Journal of Consumer Research*, 16(4), 393–408.

Hofstetter, Marc and Jorge Tover (2010). “Common Knowledge Reference Price and Asymmetric Price Adjustments,” *Review of Industrial Organization*, 37(2), (September), 141-159.

Iyengar, Sheena S. and Mark R. Lepper (2000). “When Choice is Demotivating: Can One Desire Too Much of a Good Thing?” *Journal of Personality and Social Psychology*, 79(6), 995–1006.

Jackson, William E., III. (1997). "Market Structure and the Speed of Price Adjustments: Evidence of Non-monotonicity." *Review of Industrial Organization*, 12, (February), 37-57.

Karrenbrock, Jeffrey D. (1991). "The Behavior of Retail Gasoline Prices: Symmetric or Not?" Federal Reserve Bank St. Louis Review, 73, (July/August), 19-29.

Knotek, Edward S. II (2008). "Convenient Prices, Currency and Nominal Rigidity: Theory with Evidence from Newspaper Prices." *Journal of Monetary Economics*, 55, 1303–1316.

Knotek, Edward S., II (2010). "Convenient Prices and Price Rigidity: Cross-Sectional Evidence." *Review of Economics and Statistics*, forthcoming.

Leslie, Phillip (2004). "Price Discrimination in Broadway Theater," *The RAND Journal of Economics*, 35(3), 520-541.

Levy, Daniel, Mark Bergen, Shantanu Dutta and Robert Venable (1997). "The Magnitude of Menu Costs: Direct Evidence from Large U. S. Supermarket Chains." *Quarterly Journal of Economics*, 112(3), 791-825.

Levy, Daniel, Dongwon Lee, Haipeng (Allan) Chen, Robert J. Kauffman and Mark Bergen (2010). "Price Points and Price Rigidity." *Review of Economics and Statistics*, forthcoming.

Mankiw, N. Gregory (1985). "Small Menu Costs and Large Business Cycles: A Macroeconomic Model," *Quarterly Journal of Economics*, 100 (May 1985), 529–38.

McAlister, Leigh and Edgar Pessemier (1982). "Variety Seeking Behavior – An Interdisciplinary Review." *Journal of Consumer Research*, 9(3), 311-322.

McMillan, Robert S. (2005). *Estimating Demand for Differentiated Products With Continuous Choice and Variety-Seeking: An Application To The Puzzle Of Uniform Pricing*, Ph.D. Dissertation, Graduate School of Business, Stanford University, Palo Alto CA.

Midrigan, Virgiliu (2012), "*Menu-Costs, Multi-Product Firms and Aggregate Fluctuations*." forthcoming, Econometrica

Neumark, David, and Sharpe, Steven A. (1992). "Market Structure and the Nature of Price Rigidity: Evidence from the Market for Consumer Deposits." *Quarterly Journal of Economics*, (May), 657-80.

Noel, Michael (2009). "Do Retail Gasoline Prices Respond Asymmetrically to Cost Shocks? The Influence of Edgeworth Cycles." *RAND Journal of Economics*, 40(3), 582-595.

O'Connell, Vanessa (2008). "Stores Count Seconds to Trim Labor Costs," *Wall Street Journal*, November 13, http://online.wsj.com/article/SB122651745876821483.html.

Orbach, Barak Y. and Liran Einav (2001). "Uniform Prices for Differentiated Goods: The Case of the Movie-Theater Industry." *International Review of Law and Economics*, 27(2), 129-153.

Peltzman, Sam (2000). "Prices Rise Faster Than They Fall," *Journal of Political Economy*, 108(3), 466-502.

Ratner, Rebecca K. and Barbara E. Kahn (2002). "The Impact of Private Versus Public Consumption on Variety-Seeking Behavior." *Journal of Consumer Research*, 29(2), (September), 1246-257.

Sheshinski, Eytan, and Yoram Weiss (1977. "Inflation and Costs of Price Adjustment," *Review of Economic Studies*, 44 (June 1977), 287–303.

Simonson, Itamar (1990). "The Effect of Purchase Quantity and Timing on Variety-Seeking Behavior," *Journal of Marketing Research*, 27(2), (May), 150-162.

Tellis, Gerard J. (1988), "The Price Elasticity of Selective Demand: A Meta-Analysis of Econometric Models of Sales," *Journal of Marketing research*, Vol. XXV, 331-41.

Zbaracki , Mark, 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* (May 2004), 86(2), (May), 514-533.

Appendix 1

Price and Cost Change Reports: Variable Definitions

Variable	Definition
NUMBER OF SKUS	The number of SKUs associated with that PrimarySKU.
Prior 99-cent Price Ending	1 if prior retail price ended in 99-cents; 0 otherwise.
Prior Profit Margin	The % profit margin prior to the cost change.
Size of the Cost Change	The size of the cost change in %.
Log Purchase Volume	The log of the number of units sold in the prior 12 months.

This table provides formal definitions for the variables constructed from the price and cost change reports.

Summary Statistics

Variable	Average	Standard Error	Sample Size
Prior 99-cent Price Ending	39.63%	0.40%	15,200
Size of the Cost Change	9.10%	0.15%	15,200
Log Purchase Volume	10.15	0.02	15,089
Category Size	127.96	2.16	13,597

This table reports summary statistics for the 15,200 cost increases in the cost and price change database. Missing observations reflects missing data in the database.

Interaction Between NUMBER OF SKUS and Size of Cost Change

	Model 1	**Model 2**
NUMBER OF SKUS	-0.0226** (0.0066)	
NUMBER OF SKUS x Size of Cost Change	0.1339 (0.0807)	
Log NUMBER OF SKUS		-0.1229** (0.0195)
Log NUMBER OF SKUS x Size of Cost Change		0.7228** (0.2337)
Prior 99-cent Price Ending	-0.1754** (0.0187)	-0.1767** (0.0183)
Size of Cost Change (%)	0.1981 (0.1533)	0.2896** (0.1016)
Prior Profit Margin (%)	-0.7268** (0.0524)	-0.7440** (0.0523)
Purchase Volume (log)	0.0091* (0.0039)	0.0114** (0.0039)
Category Size (00's)	0.0163* (0.0081)	0.0160* (0.0079)
Model	logistic	logistic
Log pseudolikelihood	-5,735	-5,707
R^2 or pseudo R^2	0.1435	0.1477
Sample Size	10,998	10,998

The table reports marginal effects from logistic models (Models 1 and 2) in which the dependent variable is a binary variable indicating whether the retailer increased its price following a cost increase. Fixed year and month effects were included but are omitted from this table. Standard errors are in parentheses. The standard errors are clustered by the month of the observation (month*year). *significantly different from zero, $p < 0.05$; **significantly different from zero, $p < 0.01$.

Appendix 2: Sources of Variation in the Number of SKUs

We have treated the *NUMBER OF SKUS* as an exogenous variable and investigated how this variable is associated with the firm's response to a cost increase. In this appendix we investigate sources of variation in the *NUMBER OF SKUS*. In particular we study how heterogeneity in customer preferences contributes to the decision to increase the number of variants. These findings allow us to link characteristics of consumer behavior to price stickiness.

As we will discuss, we expect an item to have more variants when customers' preferences are more heterogeneous and/or when individual customers have a greater preference for variety. We introduce metrics for: (1) heterogeneity in preferences across customers; and (2) individual customer's preference for purchasing different flavor and color variants ("variety seeking"). These metrics are constructed from two years of individual transactions by almost 780,000 customers using the retailer's frequent shopping card. The data includes a 24-month purchasing history (August 1, 2004 through August 10, 2006) for a random sample of 779,734 customers including approximately 17 million transactions. Each transaction is a shopping basket on a single visit to a store. These 17 million transactions included an average of 4.47 items, representing a total purchase volume of almost 75 million items. The transaction histories are complete within the 24-month period, though they only include purchases on occasions that customers used their frequent shopping card. Each record is an item in an order, and the record includes a unique customer identifier, an order number, the order date, the item number, the quantity purchased and the price of the item.

Motivating Example

We motivate our analysis using a simple example to illustrate the factors that may contribute to differences across items in the number of variants. Assume that there are two possible colors for a PrimarySKU: red and blue. Without loss of generality we will assume that red is more popular among the mass of M customers in the market, but that a minority of customers will only buy blue if it is offered or nothing at all'. We denote the proportion of customers who will only buy blue as α (where $\alpha < 0.5$). Notice that we can use α as a measure of heterogeneity in customers' preferences: higher values of α (up to 0.5) indicate more heterogeneity. The retail price and variable cost of both variants is fixed at p and c respectively, and so the only decision the retailer makes is which variants to sell. We will assume that there is a fixed cost of selling each variant, which we denote by k. We also rule out degenerate solutions by assuming that the firm always sells the red color, otherwise we will not observe the product at all.[19] The question of interest is: will the firm also want to sell the blue color? If customers buy at most a single unit then the firm will sell the less popular blue variant iff $k < \alpha M(p-c)$.

[19] This implies that $k < (1-\alpha)M(p-c)$.

It is also possible that the retailer may want to introduce additional variants because individual customers prefer variety. There is now an extensive psychological literature documenting customers' preference for variety and evaluating alternative explanations for this phenomenon (see for example McAlister and Pessemier 1982; Simonson 1990; and Ratner and Kahn 2002). To illustrate the role of variety-seeking we can introduce a third segment of customers who will buy up to two units, but only one of each color. The addition of this third segment results in βM customers who will buy both variants, αM customers who only buy the blue variant, and $(1-\alpha-\beta)M$ customers who only buy the red variant. The incremental profit the firm expects to earn from selling the blue variant is contributed by the first two segments and is equal to: $(\alpha+\beta)M(p-c)$. The firm will introduce the less popular blue variant iff this incremental profit exceeds k.

We can summarize this example by recognizing that the expected number of variants is larger when:

1. The cost of introducing an additional variant is lower (k is smaller).
2. There is more heterogeneity between customers in their preferences (α is larger).
3. Individual customers have a greater preference for variety (β is larger).
4. There are more customers in the market (M is larger).
5. There is a higher profit margin ($p - c$ is larger).

Measuring the profit margin (p-c) is straightforward as we have data describing the unit profit margins. As a proxy for the size of the market (M) we use the unit sales volumes in the prior twelve months. Measuring the cost of introducing a variant is less straightforward as we do not have detailed data describing the cost to the manufacturer of supplying an additional variant. However, we do have a measure of the partial cost to the retailer of merchandising an additional variant. In particular, we have the physical dimensions of the product (measured in inches). Larger products take up more shelf space, suggesting that the opportunity cost of introducing an additional variant is larger on products with larger physical dimensions.

Unfortunately there is no standard measure to describe heterogeneity in preferences across customers or the preference for variety for an individual customer. However, inspection of our large sample of individual transaction data suggested some possible metrics. It is again helpful to use an example. Let us consider a (hypothetical) PrimarySKU that has at least two variants. We will label the most popular variant "SKU A" and the second most popular variant "SKU B" and assume that SKU A sells 1,000 units in our historical transaction data, while SKU B sells a total of 800 units. We can further identify whether the 800 units were purchased by the same customers who purchased SKU A, or different customers. For the sake of our illustration we will assume that 600 of the units were sold to customers who only buy SKU B (they do not buy SKU A); and 200 units were sold to customers who buy both SKU A and SKU B.

The first metric measures customers' preference for variety and describes how many customers purchased both variants. In particular, we calculate the following measure:

$$Variety\ Seeking\ =\ \frac{Units\ of\ SKU\ B\ by\ customers\ who\ also\ purchased\ SKU\ A}{Total\ units\ of\ SKU\ A}$$

This measure is bounded by 0 and 1 and can be interpreted as a proportion (recall that SKU A is the more popular SKU). Higher values of this measure indicate that sales of both variants were more similar because many customers purchased both variants. In our example this measure would have a value of 0.2. The second measure focuses on the heterogeneity in preferences across customers:

$$Heterogeneity\ =\ \frac{Units\ of\ SKU\ B\ by\ customers\ who\ did\ not\ purchase\ SKU\ A}{Total\ units\ of\ SKU\ A}$$

This measure is also bounded by 0 and 1, with higher values indicating that sales of both variants were more similar because different customers prefer different variants. In our example *Heterogeneity* would have a value of 0.6. The third measure measures the overall parity in sales of the two most popular variants:

$$Overall\ Sales\ Parity\ =\ \frac{Units\ sold\ of\ SKU\ B}{Total\ units\ sold\ of\ SKU\ A}$$

This measure is again bounded between 0 and 1, with higher values indicating that sales are distributed more equivalently across the two most popular variants. Intuitively, *Overall Sales Parity* represents the additional sales contributed by the second most popular SKU, with *Variety Seeking* and *Heterogeneity* diagnosing the source of those sales. The *Overall Sales Parity* measure can be calculated by adding the other two measures together, and in our example, *Overall Sales Parity* has a value of 0.8.

These measures can only be calculated for PrimarySKUs that have at least two variants (*NUMBER OF SKUS* > 1). Because we will want to use these measures to evaluate whether cost increases led to a price increase, we also restrict attention to PrimarySKUs for which we observed a price or cost change in our five year sample of cost and price change data.[20] This yields an intersection of 934 PrimarySKUs.

In the table below we report the pair-wise correlation between *NUMBER OF SKUS* and each measure. There are several findings of interest. First, and most importantly, there

[20] To avoid truncation errors we must also restrict attention to PrimarySKUs for which the two most popular SKUs were introduced before the start of our individual transaction period (August 1, 2004) and were not discontinued before the end of the transaction period (August 10, 2006). We also omit any PrimarySKUs for which the most popular variant sells fewer than 100 units over the two years of data.

is a strong positive correlation between the *NUMBER OF SKUS* and our measures of preference heterogeneity and variety seeking. The more evenly sales are distributed across the two most popular variants, the more likely that the PrimarySKU has a large number of variants. The correlations are stronger when using the log of *NUMBER OF SKUS*, and in that case the positive correlations extend across both Heterogeneity and Variety Seeking.

***Number of SKUS*: Pair-Wise Correlations**

	NUMBER OF SKUS	Log of *NUMBER OF SKUS*
Heterogeneity and Variety Seeking		
Overall Sales Parity	0.1890**	0.2347**
Heterogeneity	0.1899**	0.1907**
Variety Seeking	0.0242	0.1146**
Profit Margin and Purchase Volume		
Profit Margin	-0.0568	-0.0993**
Purchase Volume	0.2412**	0.2887**
Physical SKU Size		
Width (inches)	-0.0612	-0.0322
Height (inches)	-0.0832*	-0.0437
Depth (inches)	-0.1318**	-0.0941**

The table reports pair-wise Pearson correlation coefficients between *NUMBER OF SKUS* and the explanatory variables. The sample size for each correlation is 934. *significantly different from zero, $p < 0.05$; **significantly different from zero, $p < 0.01$.

Second, there is a statistically significant relation between the *Purchase Volume* and the *NUMBER OF SKUS*. This is consistent with our prediction that in larger markets firms will be more willing to invest in additional variants. It also amplifies the importance of the phenomenon; although not all items have multiple variants, items that have multiple variants contribute disproportionately to the volume of overall transactions. A reluctance to change prices on higher volume items will tend to have a greater impact on the level of price adjustments in the overall economy.

The correlations between our measures of physical SKU size and the *NUMBER OF SKUS* are consistently negative. Recall that we interpreted physical SKU size as a measure of the opportunity cost of introducing additional variants. The strongest correlation is for SKU depth. This may reflect the need for multiple rows of products (facings) when the depth of a SKU prevents the retailer from carrying sufficient stock in a single product facing. Surprisingly, there is no evidence that retailers are more likely to introduce additional variants when the items have higher profit margins. The results suggest the relationship operates in the opposite direction, so that items with lower profit margins have more variants. However, we caution that items with higher prices (and higher profit margins) tend to have lower purchase quantities, and so this simple pair-wise correlation may be influenced by the relationship between *NUMBER OF SKUS* and *Purchase Volume*.

In separate analysis available from the authors we used OLS to regress *NUMBER OF SKUS* (and its log transformation) on this set of independent variables. The findings replicate the results of the pair-wise correlations. We also validated our *Heterogeneity* and *Variety Seeking* measures by confirming that their construction does not appear to mechanically introduce correlation with *NUMBER OF SKUS* and showing that temporal changes in these measures are predictive of changes in *NUMBER OF SKUS*. In particular, there is strong evidence that the retailer discontinues the second variant when sales in prior periods are low compared to the most popular variant.

In the next sub-section we focus on the component of *NUMBER OF SKUS* that can be attributed to preference heterogeneity and variety seeking. We will investigate how variation in this component relates to the retailers' decision to increase the retail price following a cost increase.

Heterogeneity, Variety Seeking and Price Changes

Our analysis uses a 2-stage GMM estimator that is analogous to 2-stage least squares but accommodates clustering of the standard errors. In particular, we estimate the following system of linear models:

$$1^{st}\text{ Stage: } \textit{NUMBER OF SKUS}_i = a + b\ \textit{Heterogeneity}_i + c\ \textit{Variety Seeking}_i + \mathbf{dX}_i + \varepsilon_i$$

$$2^{nd}\text{ Stage: Retail Price Increase} = \alpha + \beta\ \textit{Predicted NUMBER OF SKUS}_i + \mathbf{BX}_i + \eta_i$$

The unit of analysis is a cost increase event, and *Retail Price Increase* is a binary variable indicating whether the retailer increased its price (the same dependent variable that we used in the findings presented in Table 5). The *Predicted NUMBER OF SKUS* variable is the predicted values from Model 1; **B** is a vector of coefficients; and $\mathbf{X}_i$ describes the matrix of other explanatory variables (the same variables that were included in our earlier models reported in Table 5). In the first model a, b, c and **d** are all estimated coefficients.

If heterogeneity and variety seeking are valid instruments for the *NUMBER OF SKUS* this system of equations can be interpreted as an instrumental variable regression. This is of particular interest if there is concern that the results in Table 5 suffer from an omitted variables problem.[21] Before presenting estimates of these coefficients we will first discuss how we will treat items with only a single variant, for which it is not possible to calculate the *Heterogeneity* and *Variety Seeking* measures.

Recall that we can only calculate *Heterogeneity* and *Variety Seeking* for items with at least two variants. This is less than half of the data (see Table 3) and so omitting observations for items with a single variant would result in the loss of most of the data (together with systematic truncation of the variable of interest). Our solution is to calculate an average of these measures for each product category. The logic is that heterogeneity in preferences and variety seeking are likely to be similar across different items in the same category. For example, we would expect that heterogeneity in customers' preferences for whitening toothpaste should be relatively similar irrespective of the toothpaste brand. By using a common measure within a product category we obtain *Heterogeneity* and *Variety Seeking* measures even for items that only have a single variant (as long as other items in the category have at least two variants).[22]

Results

Below we report the GMM estimates when using either *NUMBER OF SKUS* or log of *NUMBER OF SKUS* as the endogenous variable. Fixed month and year effects were included in each model but are omitted from the table. The standard errors are clustered using the month of the decision.

The coefficient of interest is β, which is the coefficient for the predicted *NUMBER OF SKUS* (or log of the *NUMBER OF SKUS* in Model 2). We see that this coefficient is significantly less than zero in both models. This is consistent with our prediction that the retailer is less likely to increase prices following a cost increase if an item has a more variants.

---

[21] We controlled for a range of observable factors that are likely to affect the retailers' decision in Table 5. However, there may be other unobservable factors that are correlated with both *NUMBER OF SKUS* and the retailer's decision.

[22] As support for our claim that heterogeneity in preferences and variety seeking are likely to be similar across different items in the same category we compared the correlation in our measures of *Heterogeneity* and *Variety Seeking* both within and between categories. In particular we randomly paired items by selecting either within the same category or from across the entire pool of items. We then calculated the correlation across these pairs. When randomizing across the entire sample of items we do not observe any correlation between randomly selected pairs of items in either measure. However, randomly assigning pairs within a product category yields a significant positive correlation for both measures. We interpret this as evidence that heterogeneity in preferences and variety seeking are similar across different items in the same category.

GMM Results for 2^{nd} Stage

	Model 1	Model 2
Predicted NUMBER OF SKUS	-0.0630^{**} (0.0173)	
Predicted Log NUMBER OF SKUS		-0.2827^{**} (0.0738)
Prior 99-cent Price Ending	-0.1672^{**} (0.0233)	-0.1670^{**} (0.0227)
Size of Cost Change (%)	0.1973^{**} (0.0700)	0.2012^{**} (0.0692)
Prior Profit Margin (%)	-0.7032^{**} (0.0732)	-0.7835^{**} (0.0803)
Purchase Volume (log)	0.0274^{**} (0.0075)	0.0333^{**} (0.0083)
Category Size (00's)	0.0202^{**} (0.0054)	0.0155^{**} (0.0043)
Wald Chi^2 (20 d.f.)	624.17	625.67
First Stage R^2	0.1063	0.1754
Second Stage R^2	0.0386	0.0741
Sample Size	7,142	7,142

The table reports coefficients from a 2-stage system of linear models estimated using GMM. The endogenous variable is *NUMBER OF SKUS* in Model 1 and the log of this measure in Model 2. The exogenous instruments are *Heterogeneity* and *Variety Seeking*. Fixed year and month effects were included, but are omitted from this table. The standard errors are clustered by the month of the observation (month*year).
*significantly different from zero, $p < 0.05$; **significantly different from zero, $p < 0.01$.

These findings also offer an additional source of reassurance about the analysis in Section 3. Recall that we directly estimated the relationship between *NUMBER OF SKUS* and the probability that a cost increase leads to a price increase. While we included explicit controls for other observable factors that are likely to influence the probability of a price increase, we did not control for unobservable factors. As a result, it is possible that the findings in Section 3 reflect the omission of unobservable factors that are jointly correlated with both variables. The findings at least partially address this concern. The estimation restricts attention to the component of *NUMBER OF SKUS* that is associated with *Heterogeneity* and *Variety Seeking*. It is difficult to identify alternative explanations for why these instruments would affect the retailer's decision to increase prices. This provides greater confidence that the retailers' apparent reluctance to increase prices on items with more variants is not due to an omitted variable that is correlated with both measures.[23]

Summary

We have investigated why some items have more variants than others. Two factors that appear to contribute to this decision is heterogeneity in preferences across customers, and a desire for variety by individual customers. We then show that the variation in *NUMBER OF SKUS* that can be attributed to these factors helps to explain when the firm increases prices in response to a cost increase.

[23] The construction of our *Heterogeneity* and *Variety Seeking* measures suggest two reasons that the findings may be conservative. First, the reliance on category-level measures of preference heterogeneity and variety seeking excludes any within-category variation in *NUMBER OF SKUS*. The findings prevail despite (not because of) the absence of this within-category variation. Second, the adjusted R^2 values in the first stage model indicate that the instruments only explain a relatively small amount of variation in *NUMBER OF SKUS*. We caution that *Heterogeneity* and *Variety Seeking* are merely metrics for heterogeneity in customers' preference and variety seeking. They are not the only possible measures of these phenomena, and it is possible that alternative measures would explain more of the variation in *NUMBER OF SKUS*.

Appendix 3: Cost Decreases

While we might expect a similar pattern of findings if we study the response to cost increases and cost decreases, the literature suggests otherwise. There is a growing body of evidence suggesting that firms use different criteria for deciding when to increase versus decrease the price, and that this leads to asymmetries.[24] In this appendix we search for further evidence of asymmetries by measuring how often the retailer decreases prices in response to a cost decrease.

The price and cost change database includes 5,793 examples of cost decreases. In Section 3 we reported that just 11.9% of these cost decreases resulted in price decreases (5.7% led to price increases). In comparison, recall that cost increases resulted in price increases 69.8% of the time. While cost increases often lead to a price increase, cost decreases rarely result in price decreases. Discussions with the managers at the retailer confirmed that the decision to lower a price depends on different factors than decisions to increase prices. In particular, price decreases are often made in response to competitive price comparisons. To investigate factors that affected the probability of a price decrease we re-estimated our logistic and OLS models using a new dependent measure. The new (binary) dependent variable indicates whether the price decreased following a cost decrease. The findings are reported below.

The results in the next page reveal no evidence that the decision to decrease the price is related to the *NUMBER OF SKUS*. This is somewhat surprising, as the argument that it is more costly to change prices on items with multiple variants applies equally to price increases and decreases. The results also contrast sharply with the evidence that the number of variants influences the retailers' willingness to increase prices. They represent further evidence of asymmetries in the way that retailers evaluate price increases and price decreases.

There are several other examples of asymmetries. While retailers appear to be very reluctant to increase a price that ends in 99-cents, they do not exhibit the same reluctance when deciding whether to decrease the price. This is consistent with the evidence that the kink in the demand curve occurs above and not below the 99-cent level (Levy et al. 2010).

While we reported that the firm is more likely to raise prices following a cost increase on higher volume items, the reverse is true of price decreases when costs decrease. The firm is less likely to lower prices on higher volume items.

[24] As we acknowledged in Section 3, asymmetries between price increases and price decreases have been recognized elsewhere in the literature.

Factors that Contribute to the Decision to Lower the Price

	Model 1	**Model 2**	**Model 3 OLS**
NUMBER OF SKUS	-0.0004 (0.0024)		
Log NUMBER OF SKUS		0.0103 (0.0108)	
NUMBER OF SKUS = 2 or 3			0.0245 (0.0195)
NUMBER OF SKUS = 4 to 6			0.0235 (0.0315)
NUMBER OF SKUS = 7 or more			-0.0225 (0.0286)
Prior 99-cent Price Ending	0.0189 (0.0127)	0.0198 (0.0126)	0.0156 (0.0115)
Absolute Size of Cost Decrease (%)	0.3796** (0.0437)	0.3823** (0.0439)	0.5519** (0.0740)
Prior Profit Margin (%)	-0.2320** (0.0432)	-0.2275** (0.0427)	-0.1991** (0.0402)
Purchase Volume (log)	-0.0083** (0.0030)	-0.0088** (0.0030)	-0.0093* (0.0045)
Category Size (00's)	-0.0015 (0.0051)	-0.0013 (0.0051)	0.0028 (0.0118)
Model	logistic	logistic	OLS
Log pseudolikelihood	-1,030	-1,029	
R^2 or pseudo R^2	0.1684	0.1689	0.1098
Sample Sizes	4,099	4,099	4,099

The table reports marginal effects from logistic models (Models 1 and 2) and coefficients from an OLS model (Model 3). In all three models the dependent variable is a binary variable indicating whether the retailer decreased the price following a cost decrease. Fixed year and month effects were included but are omitted from this table. Standard errors are in parentheses. The standard errors are clustered by the month of the observation (month*year). *significantly different from zero, $p < 0.05$; **significantly different from zero, $p < 0.01$.

In order to maintain the profit margins within a targeted range we would expect that the firm would be more likely to lower the price in response to a cost decrease when the previous profit margin was larger. The results do not support this prediction.

Instead the firm is less likely to lower the price after a cost decrease on items that previously had higher profit margin. It appears that the firm prefers to maintain higher profit margins on items that previously had higher margins.

The one covariate that continues to play a consistent model between the two models is the size of the cost change. As we would expect, firms are more likely to lower the price when the cost decrease is larger (notice that we measure the absolute size of the cost decrease).