Menu Costs and Price Rigidities: Micro Evidence

Eric Anderson
Northwestern University

Nir Jaimovich
Duke University

Duncan Simester
MIT

January 2011

Menu Costs and Price Rigidities: Micro Evidence

A leading explanation in the economic literature is that monetary policy has real effects on the economy because firms must incur a fixed-cost when changing prices. Yet, empirical support regarding the nature of menu costs and their effect on pricing is scant. In this paper, we use a 55-month database of cost and price changes at a large retailer that allows us to document the impact of menu costs on pricing policy. The identification of menu costs stems from the retailer's adoption of a "uniform pricing" rule that requires all variants of a product to have the same price. Differences in the number of variants across products lead to variation in the cost of changing the prices of these products. We show that the retailer is significantly less likely to raise prices following a cost increase on items that have more variants. Absent these menu costs price increases would be 8% to 18% more frequent. Finally, we show that variation in consumer preferences partially explains why firms offer multiple product variants, which identifies a link between consumer heterogeneity and price stickiness.

1. Introduction

Why does monetary policy have real effects on the economy? A leading explanation in the literature is that firms incur a cost (a "menu cost") when changing their prices and this contributes to 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), who were the first to consider the impact of fixed costs on price adjustments. While important details differ across the work that followed, a central assumption in all of them is that a fixed cost must be incurred upon a price change.[1] However, empirical support regarding the nature of menu costs is scant. We have limited understanding of what factors may create menu costs. We also have little evidence that relates these menu costs to price stickiness.[2]

To address these issues we conduct a large-scale empirical study with a national U.S. retailer. This retailer is concerned about the cost of changing prices and uses an incentive system to minimize the number of price changes that can occur each day. This incentive system discourages a large number of price changes on any day.

We begin by identifying a new source of variation in menu costs. Like many other retailers, this retailer has rules that link price changes across different variants of an item. For example, two flavors of Stacy's Pita chips are linked by a rule so that if the retailer decides to change the price of pita chips then two shelf tags must be changed. These "uniform pricing" rules exist for several reasons. First they avoid customer confusion or perceptions of unfairness. Second, in our data there are no cost differences between linked items so there are no supply side factors to justify a price difference. Third, the rules greatly simplify managerial decision-making as they reduce multiple pricing decisions to a single decision. These rules allow us to measure the impact of menu costs on pricing policy. In our data there is significant variation in the number of linkages that different products have. This variation introduces differences in the cost of changing prices across products. For example, in the spice category, 25 items are linked together so implementing a single price change requires 25 new shelf tags in over 5,000 stores.[3]

Our empirical analysis focuses on more than 15,000 cost increases. These are events in which a vendor increases the normal wholesale cost of an item. In response, the retailer decides whether to change the regular retail price. We show that menu costs play a central role in this decision. When a product has only a single link (to itself) the probability of a price increase following a cost increase is 70.9%. But, if a product is linked to seven or more other product

[1] Several other prominent examples that built on this work are 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] Levy, Bergen, Dutta and Venable (1997) is a notable exception. This study calculates the cost of changing prices at five supermarket chains. They also report that the supermarkets are less likely to change prices when regulations mandate item-level pricing, requiring separate price tags on each unit of stock.

[3] The number of stores operated by this retailer is more than 5,000, but less than 10,000 (confidentiality prevents us from revealing the precise number of stores).

variants the probability of a price increase is just 58.8%. In the absence of these menu costs we estimate that cost increases would prompt 8% to 18% additional price increases.

Finally, we ask what explains why retailers offer more than one product variant. That is, why does a retailer need to offer both Parmesan and Cinnamon pita chips? Using a second dataset consisting of a large panel of historical consumer purchases, we show that this can be traced to heterogeneity in consumer preferences. As heterogeneity increases a retailer is more likely to adopt a broader set of product variants. We then use this measure to demonstrate a link between consumer heterogeneity and price stickiness.

A natural question is why an economist would care about these findings. Using the variation in the number of variants across products, we provide empirical evidence that menu costs do matter in the micro-data for pricing. We view this as an important finding, as it provides one of the first pieces of direct evidence for the menu-cost channel, which is central to all state-dependent models. Second, we show that products with the most variants are also the ones that sell the largest quantities and generate the most revenue. Thus, our results suggest that goods with bigger markets are "stickier." This implies that while we might empirically observe many goods changing prices, they could potentially account for a smaller fraction of the overall quantities and revenues sold. Finally, our work suggests a novel link between microeconomic behaviors (heterogeneity in consumer preferences) and macroeconomic outcomes (price stickiness).

The paper proceeds as follows. In Section 2 we describe the process involved in changing prices at this retailer. We then describe the data used in the study in Section 3. In Section 4 we review how often the firm increased prices in response to cost increases. In Section 4 we introduce the number of variants of each item as a measure of menu costs and investigate how this (and other factors) influenced the firm's response to a cost increase. In Section 5 we explore different factors that contribute to the number of variants an item has. Specifically, we show that the number of variants in part depends on the level of heterogeneity in customer preferences, and that this component of the variation is an important factor in determining whether the price increases after a cost increase. In Section 6 we focus on cost decreases, and identify differences in how the retailer responds to cost decreases versus cost increases. The paper concludes in Section 7 with a review of the findings.

2. Institutional Background

The analysis in this paper uses data provided by a major retailer. The retailer operates a large number of stores that sell items in grocery, health and beauty and general merchandise product categories.

Price changes at this retailer occur for three broad reasons. First, the firm makes changes to the "regular" price of items. These price changes are primarily prompted by changes in the wholesale prices charged by product manufacturers. Regular prices are also sometimes changed as a result of competitive price comparisons to more closely align prices with those charged by the firm's competitors. The second category of price changes is temporary price reductions. These typically occur as part of a negotiated schedule between the retailer and the manufacturer and are planned well in advance. Finally, price changes may occur as a result of a special event. For example, mergers with other retailers may lead to large-scale price changes in the newly acquired stores. The retailer also conducts occasional price tests, which are typically limited to a subset of the retailer's stores. Throughout this paper we will focus on the first category of price changes: changes to the regular price.[4]

In this section we describe the process involved in changing prices at this retailer. We review the management controls on this process together with the resource requirements and limitations that influence the overall frequency of price changes. We begin by describing how the firm organizes different variants of the same item.

Items With Multiple Variants

Many of the items sold by this retailer have multiple variants reflecting differences in colors, flavors, product types, or combinations of sizes. We provide four examples below.

PrimarySKU Description	Number of Variants	Examples of Individual SKU Descriptions
Cheerios	1	Cheerios 15oz
Stacy's Pita Chips	2	Parmesan 6oz Cinnamon 6oz
Johnson and Johnson Kids Bandages	4	Barbie 25 count assorted sizes Elmo 30 count assorted sizes Sponge Bob 20 count assorted sizes Dora 25 count assorted sizes
Gold Em Spices	25	Bacon Bits 2.82oz Ground Pepper 3.35oz Bay Leaves 0.25oz

[4] This distinction between regular price changes and temporary promotions is also made by other types of retailers. A former manager at a large chain of department stores revealed a similar policy. The manager cited an example that arose from the California electricity crisis. The operating costs for the retailer's California stores increased sharply during this period, to the extent that the firm expected to make no money in California that year. The firm investigated raising some apparel prices, but pricing decisions are made a year in advance with prices pre-printed on the price tags. While the firm could have theoretically changed the regular prices by replacing price tags on each garment, it did not have the labor capacity to implement this. In contrast, promotional discounts are easily implemented: the retailer simply puts a "Sale" sign on top of the merchandising fixtures highlighting the discounts.

This retailer assigns a common "Primary Stock Keeping Unit" (hereafter "PrimarySKU") to every variant, and then a stocking keeping unit number ("SKU number") to every individual variant. A key feature of the data is that prices and costs are the same for every SKU under the same PrimarySKU. However, each of the individual SKUs require a separate price sticker, and so changing the price of Gold EmSpices requires finding and changing 25 price stickers in each store, while changing the price of Cheerios only requires changing a single price sticker.

Because the retailer charges one price for all variants, any difference in the cost of changing prices due to the number of variants is an in-store implementation cost, not a management cost. The decision to charge a uniform price for all colors and variants of the same item is well-documented, and has been the subject of several studies. For example, Anderson and Simester (2008) investigate the tendency of women's apparel retailers to charge the same price across all sizes of an item. What makes this particularly surprising is that the retailers often pay higher wholesale prices for larger sizes. They present evidence that customers who need larger sizes perceive that it is unfair if retail prices are higher on larger sizes than on smaller sizes, and that this leads to lower overall sales. Other explanations for "uniform pricing" have focused on 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), demand uncertainty (Orbach and Einav 2007), simplifying the purchasing decision (Hauser and Wernerfelt 1990; Iyengar and Lepper 2000; and Draganska and Jain 2001), and avoiding an adverse quality signal for the lower-priced item (Anderson and Simester 2001; and Orbach and Einav 2007).

Changes to the Regular Price

The retailer asks manufacturers to provide sixty days notice of wholesale cost increases, which gives the retailer time to negotiate the cost increase and decide how to respond. The category manager responsible for this item decides whether or not to change the regular retail price and then contacts the Pricing Operations Team, which reviews the requested changes to ensure that it complies with the firm's pricing rules.[5] If approved, price changes are sent to the retail stores. To start the price change process an employee in the store prints a list of shelf stickers. This event triggers the database to update the point of sale (POS) system with the new regular price. The entire process generally takes less than 2 weeks but can take up to a month.

The Pricing Operations Team has an agreement with the Store Management Team that limits the number of regular price changes allowed in a day. It is currently set at five days per week (Tuesday through Saturday) and up to 100 price changes per day. These price changes are calculated at the SKU 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 syste 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. The annual bonuses received by the Pricing Operations Team depend upon these measures, with the team receiving smaller bonuses 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 just 91.8%, suggesting that the restriction on the frequency of price changes is not trivially satisfied.

Summary

This review of the process this retailer uses to change prices has identified several points that will be important for the analysis that follows. First, cost changes are a discrete event and the retailer makes a decision for each event whether to respond with a price change. In our analysis we investigate factors that influence the outcome of this decision. Second, in-store labor costs represent a substantial portion of this retailer's cost structure and the retailer very carefully monitors these costs. As a result, the frequency of price changes is closely observed and is a regular subject of negotiation and discussion between the category managers, the Pricing Operations Team and the Store Management Team. The focus on managing in-store labor costs has led to explicit policies limiting the frequency of price changes. These policies are implemented at the individual SKU level, so that price changes on items that have multiple variants are interpreted as multiple price changes. Below we investigate whether this results in a tendency for the retailer to avoid changing prices on items that have a larger number of variants.

3. Description of the Data

We obtained three sources of data that will be used throughout the paper. The first data source 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 and is carefully maintained. 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.[6] The price and cost

[6] As we discussed in the previous section, the price changes represent changes in the "regular" retail price. Temporary price changes, which typically last for just one week, are not included in these reports, and are managed through a separate process. The cost measures also ignore any "accrual" accounts that manufacturers provide to retailers to help fund promotions. We are confident that these accounts do not play an important role in deciding when to change the regular retail price since the retailer does not consider these funds in its internal systems when monitoring profit margins and regular price changes. In a previous study on a different topic the research team sought information about the value of these accrual accounts. The information was not readily accessible.

change reports also include the total unit volume for the item over the prior 12-months. In the Appendix we provide formal definitions and summary statistics for all of these variables.

The data includes 15,153 observations in which the cost increased and 5,867 observations in which it decreased. This is a complete census of every cost change during the 55-month period. For each of these cost changes we observe the initial and new cost, together with the initial and new retail price. In Table 1 we report the frequency with which these cost changes resulted in price changes, together with the average magnitude of the cost and price changes. Confidentiality concerns prevent us from reporting the actual profit margins or the magnitudes of the price and cost changes in percentage terms. However, in Section 4 we investigate how the firm's response to cost increases affected the profit margins on these items.

While cost increases often resulted in a price increase, cost decreases rarely led to price decreases. The average price change following a cost change is also asymmetric, with a ratio of 1.42 ($0.54 divided by $0.38) when cost increases, compared to just 0.56 ($0.40 divided by $0.71) when cost decreases.[7] These asymmetries have been recognized elsewhere in the literature (Peltzman 2000). 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 Section 6 we turn attention to cost decreases and highlight additional asymmetries in the retailer's response to cost increases versus cost decreases.

Table 1. Frequency and Magnitude of Cost and Price Changes

	Cost Increases	**Cost Decreases**
Frequency of Price Changes		
Price Increased	69.3%	5.7%
Price Decreased	1.1%	11.8%
No Price Change	29.6%	82.5%
Magnitude of Price and Cost Changes		
Average cost change	$0.38	-$0.71
Average price change	$0.54	-$0.40
Sample Size	15,153	5,867

The table reports the average frequency and magnitude of price and cost changes. The magnitudes are calculated using all of the observations (including cost changes in which the prices did not change).

[7] The average change in the price divided by the average change in the cost is not equal to the average change in the price/cost ratio. Our agreement with the retailer prevents us from directly reporting the price/cost ratio.

Our second source of data is a hierarchy mapping individual SKUs to PrimarySKUs. The hierarchy also assigns each PrimarySKU to a merchandising group and product category. There are seven merchandising groups, although over 99% of the PrimarySKUs fall into just five groups. In addition, there are 820 product categories, defined at a relatively fine level; for example, they distinguish between "whitening toothpaste" and "baking soda toothpaste". Throughout the paper (unless noted) we use a hierarchy dated July 2010, which results in some missing observations for items that were removed from the store. The intersection of the cost data and the product hierarchy yields a total 11,309 cost increases and 4,235 cost decreases for which we also observe the product hierarchy.[8]

The third data source is a large sample of individual transaction data, which we will use in Section 5 to calculate measures of heterogeneity in customer preferences. 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.

4. Measuring Menu Costs Using the Number of Variants

In this section we investigate how the number of variants of each item influenced the retailer's decision to change prices following a cost increase. We begin by describing the distribution of the number of variants across the sample of PrimarySKUs and then evaluate how the probability of a price change varies with this measure.

The Number of Variants

We restrict attention to the 9,310 PrimarySKUs for which there was either a cost change or retail price change in our 55-month data period. These 9,310 PrimarySKUs include a total of 14,633 individual SKUs, with an average of 1.60 SKUs per PrimarySKU and a maximum of 62 (which are different colors of a brand of nail polish). In Table 2 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

[8] We also have a product hierarchy from August 2006, but this introduces missing observations for items introduced after that date. Reassuringly, the pattern of findings remains robust, irrespective of which product hierarchy we use.

and units sold in the previous 12 months.[9] The frequency distribution reveals that, while PrimarySKUs with a single variant represent 80.3% of all PrimarySKUs, they only represent 59.3% of revenue and 50.2% of individual SKUs.

Table 2. Frequency Distribution of the Number of SKUs under Each PrimarySKU

NUMBER OF SKUS	PrimarySKU Frequency	SKU Frequency	Revenue Weighted	Units Weighted
1	80.3%	50.2%	59.3%	51.8%
2	9.1%	11.4%	13.4%	14.1%
3	4.0%	7.6%	8.2%	8.5%
4	2.3%	5.7%	5.3%	6.5%
5	1.1%	3.6%	3.4%	3.7%
6	0.9%	3.3%	2.6%	3.6%
7	0.5%	2.1%	1.6%	1.6%
8	0.4%	1.7%	0.8%	1.4%
9	0.3%	1.4%	0.6%	0.7%
10	0.1%	0.9%	0.3%	0.8%
11	0.1%	0.6%	0.3%	0.9%
12	0.1%	0.8%	0.4%	0.6%
13	0.1%	1.1%	0.7%	1.3%
14	0.1%	0.9%	0.3%	0.6%
15	0.1%	1.0%	0.3%	0.5%
15 or more	0.4%	7.8%	2.6%	3.5%

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 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.

To evaluate the importance of this effect 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 can address this issue by asking the following question: if the probability of a price increase

[9] We also investigated the distribution of *NUMBER OF SKUS* across each of the five major merchandising groups. All five of these merchandising groups include PrimarySKUs that have multiple variants, with the average *NUMBER OF SKUS* ranging from 1.3 for General Merchandise to 2.2 for Groceries.

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?

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

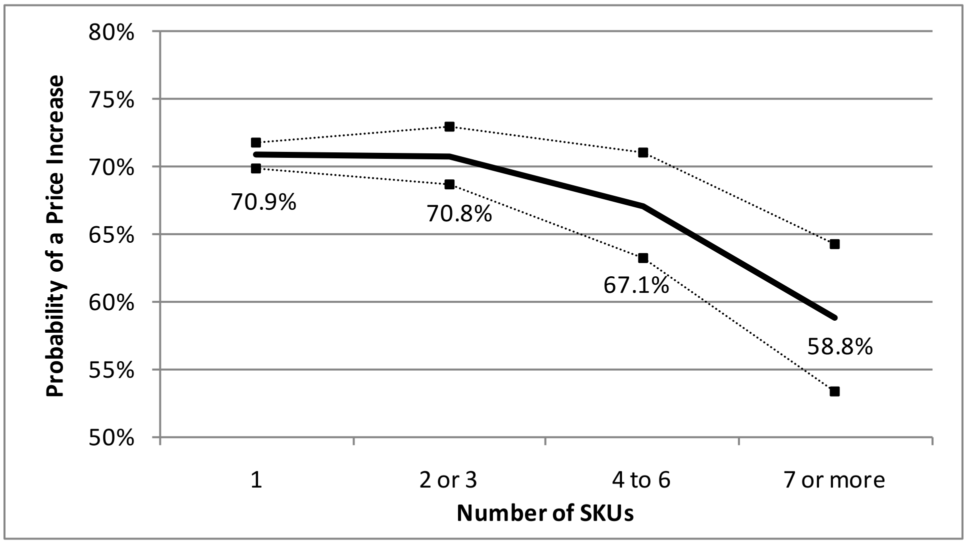

Figure 1b. Weighted by Prior Revenue

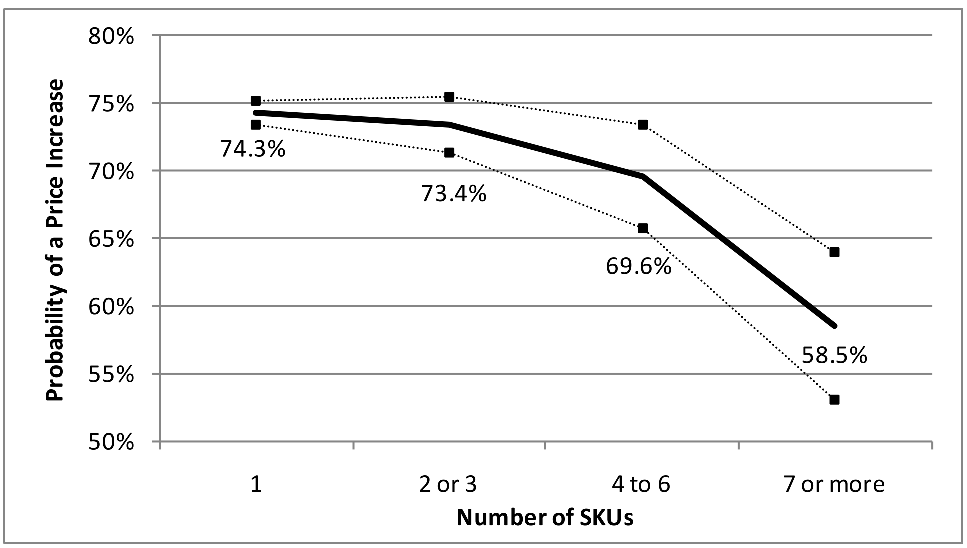

The figures report the probability of a price increase following a cost increase. The square markers indicate the 95% confidence interval. We report both weighted and unweighted averages, where the weighting uses total revenue over the prior 12-months. There are a small number of missing observations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

The Overall Size of the Effect

We begin by calculating the probability of a price increase following a cost increase for items that had only a single variant. We use the 6,844 PrimarySKUs that had at least one cost increase in our 55-month data period.

There were 8,596 cost increases on items with a single variant, and these cost increases resulted in 6,093 price increases. Therefore, the probability of a price increase following a cost increase on an item with a single variant is 70.9% (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 3.[10]

There were a total of 2,648 cost increases on items with multiple variants. If price increases occurred at the same rate on these items as on items with a single variant then we would have observed 1,877 price increases, which is 61 (3%) more than we actually observed. When weighting by revenue, there are 133 additional price increases (7% more).

Table 3. Overall Frequency of Price Increases on Items With Multiple Variants

	Unweighted		Weighted by Revenue	
	PrimarySKU Level	SKU Level	PrimarySKU Level	SKU Level
Cost Increases	2,648	10,744	2,648	10,744
Actual Price Increases	1,816	7,055	1,835	6,779
Projected Price Increases	1,877	7,617	1,968	7,986
Additional Price Increases	61	562	133	1,207

The table reports the actual number of cost and price increases on PrimarySKUs 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.

A limitation of this analysis is that it counts a price increase on an item with a single variant in the same way that it counts an item with 10 variants. This is not how the retailer itself counts price increases. Recall that the retailer monitors the frequency of price changes at the SKU level, so that price increases on two flavors of the same item count as two price increases. Moreover, we have presented evidence that the retailer is less likely to increase prices on items with more variants. The previous analysis does not account for this. To address both limitations, we also calculated the actual and projected frequency of price increases at the SKU level. Cost increases on items with multiple variants represented a total of 10,744 increases on

[10] When weighting by revenue the probability of a price increase on items with a single variant increases to 74.3%. We use this probability in the weighted analysis.

individual SKUs. If price increases occurred at the same rate as on items with a single variant then we would have observed 7,617 price increases, which is 562 (8%) more than we actually observed. When weighting by revenue, there would have been 1,207 additional price increases, or 18% more.

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 is a deterrent to price changes. 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. For example, if the size of the cost change is smaller on items that have more SKUs, then it is possible that we are observing a tendency to only increase prices when the cost increase is large. We address these concerns as follows. We first identify other observable factors that may have contributed to the decision to increase the price. Second, we develop a measure of consumer preference heterogeneity that can explain why a retailer offers multiple variants. We then develop an econometric framework to link consumer preferences with price stickiness.

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*).[11]

[11] 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

As preliminary analysis we grouped the cost increases into two sub-samples according to whether the firm increased the retail price and then calculated the average of the explanatory variables for each of the sub-samples. Comparing these averages provides a univariate measure of whether the explanatory variables are associated with the retailers' decision to increase the price following a cost increase. The findings are reported in Table 4, where to preserve confidentiality we index the profit margins.[12]

Table 4. Univariate Analysis

	No Retail Price Increase	Retail Price Increase	Difference	Standard Error of Difference
Size of Cost Change (%)	8.56%	9.23%	0.70%*	0.33%
Prior Profit Margin (indexed)	100.00	80.11	-19.89**	0.63
Prior 99-cent Price Ending (%)	57.14%	33.73%	-23.41%**	0.99%
Log Purchase Volume	10.24	10.15	-0.09	0.05
Sample Sizes	3,332	7,901		

The table reports the average of each variable, distinguishing between cost increases that led to a price increase and those that did not. The profit margins are indexed so that the average prior profit margin in the "no retail price increase sample" equals 100.

There are several findings of interest. First, as expected, the cost increases tend to be larger in the observations for which prices increased, indicating that 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.

Finally, the difference in the purchase volumes between the two samples is not significant. This last finding is surprising; we expected that the firm would consider the purchase volume when prioritizing price increases. Yet this univariate comparison indicates that the firm pays relatively little attention to the purchase volume when deciding whether to increase the retail price following a cost increase. To further investigate this result we also considered alternative specifications for the purchase volume variable but these alternative specifications did not

[12] The analysis is conducted using 11,233 observations, rather than the full sample of 11,309 cost increases. The difference reflects a small number of missing observations for some of the explanatory variables. In preliminary analysis we considered both dollar and percentage specifications for the *Prior Profit Margin* and *Cost Change* variables. The findings favored using the percentage specification. We also considered three additional factors: the prior *Retail Price*, the number of *Other Cost Increases* in that month, and the years since the *Last Cost Increase* for that PrimarySKU. However, the theoretical motivation for these variables is ambiguous and their inclusion has little impact on the coefficients of interest in the multivariate analysis that follows.

strengthen this relationship. However, we will see evidence that the firm is more likely to change prices on items with larger profit margins in our multivariate analysis.

Multivariate Analysis

In Table 5 we report the findings 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 findings for the different specifications are reported in Table 5 (to ease exposition we omit the year and month fixed effects from the table). In all of the models standard errors are clustered by the month of the observation.[13]

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 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 when there are 2 or 3 variants then the probability of a price increase following a cost increase is 2.1% lower than when the item has only a single variant. If there are 4, 5 or 6 variants the probability is 8.5% lower compared to a single variant, and if there are 7 or more variants the probability difference is 15.7%.

We also report a fourth model in Table 5, where we add category fixed effects to the linear probability model. With the inclusion of category fixed effects we control for all of the cross-category variation in the decision to increase the price, and so the coefficient of interest just captures the impact of within category variation in *NUMBER OF SKUS*. Comparing these results with Model 3 allows us to evaluate how much of the overall effect is contributed by variation in *NUMBER OF SKUS* within a category and how much reflects variation across categories. We illustrate the differences between Models 3 and 4 in Figure 2, where we describe the implied probability of a price increase. The findings confirm that the relationship between *NUMBER OF SKUS* and the probability of a price increase survives controlling for category fixed effects. When the *NUMBER OF SKUS* is large (4 or more) the slope is attenuated by approximately 50%. This indicates that approximately half of the effect is contributed by within-category variation in *NUMBER OF SKUS*, while the remainder of the effect reflects cross-category variation.

---

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

[14] These are the same groupings 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 OLS**	**Model 4 Fixed Effects**
NUMBER OF SKUS	-0.0555** (0.0177)			
Log NUMBER OF SKUS		-0.3390** (0.0772)		
NUMBER OF SKUS = 2 or 3			-0.0210 (0.0124)	-0.0351** (0.0098)
NUMBER OF SKUS = 4 to 6			-0.0852** (0.0256)	-0.0538* (0.0263)
NUMBER OF SKUS = 7 or more			-0.1567** (0.0416)	-0.0861** (0.0372)
Prior 99-cent Price Ending	-1.0188** (0.1110)	-1.0358** (0.1091)	-0.1965** (0.0209)	-0.1735** (0.0135)
Size of Cost Change	2.2971** (0.5400)	2.2943** (0.5402)	0.1461** (0.0472)	0.1597** (0.0454)
Prior Profit Margin	-4.4638** (0.3792)	-4.5498** (0.3789)	-0.7883** (0.0641)	-0.9418** (0.0818)
Purchase Volume (log)	0.0087 (0.0196)	0.0192 (0.0184)	0.0048 (0.0039)	0.0133** (0.0052)
Model	logistic	logistic	OLS	OLS
Log pseudolikelihood	-5909	-5894		
R^2 or pseudo R^2	0.1346	0.1369	0.1451	0.3572
Sample Sizes	11,233	11,233	11,233	11,233

The table reports coefficients from logistic and OLS models 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 (Model 4 also includes category fixed effects), but are omitted from this table. Standard errors are in parentheses. The standard errors are clustered by the month of the observation (month*year).

Magnitudes of the Cost and Price Changes

If the retailer is less likely to increase prices on items with more variants then profit margins will tend to decrease on these items, unless there is an offsetting difference in the magnitude of the cost and/or price changes. In Table 6 we report the relationship between the *NUMBER OF*

SKUS and: (a) the size of the cost increase; (b) the size of any price increase; and (c) the resulting change in the profit margins. We use the same set of explanatory variables as we used in Table 5.

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

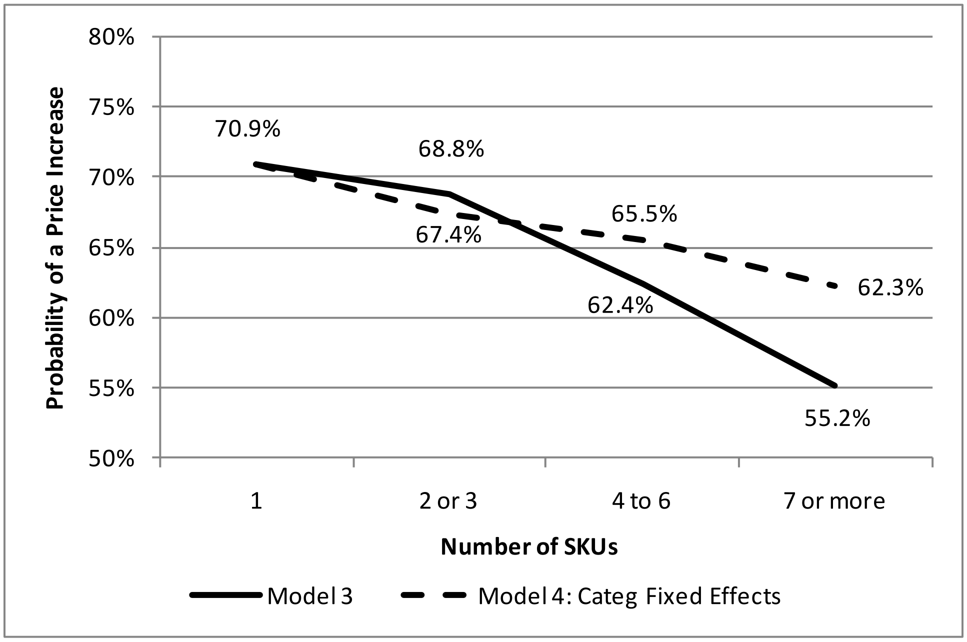

The figure interprets the coefficients from Table 5 by reporting the implied probability of a price increase following a cost increase. We index the findings by setting the probability of a price increase for an item with just one variant at 70.9% (the actual probability).

The findings confirm that there is no evidence that the number of variants contributed to systematic variation in either the size of the cost change or the size of the resulting price change. The coefficients for *NUMBER OF SKUS* are not significant in either model.[15] However, there is evidence that the cost increases resulted in lower profit margins on items with more variants. We conclude that the tendency to forgo price increases on items with more variants is not offset by reductions in either the magnitude of the cost increases or the magnitudes of the price increases (if any).

The evidence that the number of variants does not affect the size of the cost changes suggests that manufacturers do not consider the number of variants when negotiating the size of their cost increases. However, it is still possible that the number of variants affects the frequency of cost changes; if manufacturers realize that retailers are less likely to increase prices following a cost change then cost increases may be more frequent on these items. To investigate this possibility we calculated the number of cost increases for each PrimarySKU in the database. We then regressed the number of cost increases on a range of explanatory factors, including *NUMBER OF SKUS*. The model included all 4,950 PrimarySKUs that had at least one cost change

[15] We repeated this analysis using the log of the *NUMBER OF SKUS* and obtained the same pattern of results.

or price change in our 55-month data period and that were present in the store throughout this period. The findings are reported in the Appendix. They reveal no significant relationship between the number of cost increases and the *NUMBER OF SKUS*. We conclude that retailers do not appear to consider the lower probability of a price change on items with more SKUs when negotiating the frequency (or magnitude) of their cost increases.

Table 6. Magnitude of the Cost and Price Changes

	Cost Change	**Price Change**	**Change in Overall Margin**
NUMBER OF SKUS	-0.0007 (0.0005)	0.0001 (0.0003)	-0.0005** (0.0001)
Prior 99-cent Price Ending	0.0067 (0.0036)	0.0037 (0.0050)	-0.0077** (0.0015)
Size of Cost Change		0.4745* (0.2034)	-0.0816** (0.0206)
Prior Profit Margin	0.1377** (0.0241)	-0.0284 (0.0242)	-0.0438** (0.0048)
Purchase Volume (log)	-0.0031** (0.0011)	0.0005 (0.0008)	0.0010** (0.0003)
Adjusted R^2	0.0384	0.4482	0.1875
Sample Size	11,233	7,901	11,233

The table reports coefficients from an OLS model. Fixed year and month effects were included, but are omitted from this table. The dependent variables are percentage cost change (Model 1); percentage retail price change (Model 2); and percentage change in the profit margin (Model 3). In Model 2 we only include observations in which there was a price increase. Models 1 and 3 include all 11,233 observations for which there was a cost increase. Standard errors are in parentheses. The standard errors are clustered by the month of the observation (month*year).

In this section we 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 the next section 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 macroeconomic outcomes (price stickiness).

5. Sources of Variation in the Number of SKUs

We motivate the discussion in this section using a simple example to illustrate the factors that may contribute to differences across items in the number of variants. One factor that we identify is heterogeneity in customer preferences. 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 our 2-year sample of detailed transaction data. We then estimate how these and other factors describe variation in the *NUMBER OF SK*US and contribute to the retailer's decision to increase prices following a cost increase.

Motivating Example

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 if the retailer sells blue. 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.[16] 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.[17] The question of interest is: will the firm also want to sell the less popular blue color? If customers buy at most a single unit then the firm will sell the less popular blue variant iff $k < \alpha M(p\text{-}c)$.

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 $(1\text{-}\alpha\text{-}\beta)M$ 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: $(\alpha+\beta)M(p\text{-}c)$. 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:

---

[16] We previously cited studies that explain why retailers charge the same price for different variants.

[17] This implies that $k < (1\text{-}\alpha)M(p\text{-}c)$.

1. The cost of introducing an additional variant is lower (*k* is smaller).
2. There is more heterogeneity between customers in their preference (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:

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

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:

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

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:

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

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.[18] This yields an intersection of 934 PrimarySKUs. Summary statistics for these 934 PrimarySKUs are provided in the Appendix.

In Table 7 we report the pair-wise correlation between NUMBER OF SKUS and each of these measures. There are several findings of interest. First, and most importantly, there 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.

Second, there is a very strong relationship 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.

[18] 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.

Table 7. 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.

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 lines 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. We address this concern in Table 9, where we report the findings from a multivariate analysis using OLS. In Models 1 and 3 the dependent variable is NUMBER OF SKUS, while in Models 2 and 4 we use the log of this measure.

The multivariate analysis largely replicates the univariate findings. Higher purchase volumes and smaller physical sizes are both associated with an increase in the *NUMBER OF SKUS*. We also see strong evidence that if sales are more evenly distributed across the two most popular variants then the retailer tends to offer more variants. Both sources of sales for the second item appear to contribute to this result. Recall that our measure of *Heterogeneity* describes sales of the second variant to different customers than those who purchased the most popular variant, while *Variety Seeking* measures purchases by customers who purchased both variants. The coefficient for *Heterogeneity* is significant when using either dependent measure, while the coefficient for *Variety Seeking* is only significant when using the log of *NUMBER OF SKUS* (Model 4). We note that the Adjusted R^2 values are higher in Models 3 and 4, indicating that the model is better at explaining variation in the log transformation of *NUMBER OF SKUS*.

Table 8. *NUMBER OF SKUS*: Multivariate Analysis

	Model 1 NUMBER OF SKUS	**Model 2** log of NUMBER OF SKUS	**Model 3** NUMBER OF SKUS	**Model 4** log of NUMBER OF SKUS
Overall Sales Parity	4.2614** (0.7700)	0.7436** (0.1064)		
Heterogeneity			5.2079** (0.7868)	0.7868** (0.1148)
Variety Seeking			-0.0468 (1.6119)	0.5471* (0.2237)
Profit Margin	0.0230 (0.0767)	-0.0081 (0.0106)	-0.0099 (0.0772)	-0.0096 (0.0107)
Purchase Volume	0.1452** (0.0214)	0.0241** (0.0030)	0.1579** (0.0217)	0.0247** (0.0030)
Physical SKU Depth	-0.2305** (0.0619)	-0.0205* (0.0086)	-0.2116** (0.0620)	-0.0196* (0.0086)
Intercept	1.4674* (0.6631)	0.4753** (0.0958)	1. 4407* (0.6901)	0.4741** (0.0958)
Adjusted R^2	0.0965	0.1305	0.1045	0.1305
Sample size	934	934	934	934

The table reports OLS coefficients. The sample size for each coefficient is 934. Standard errors are in parentheses. The dependent variable in Models 1 and 3 is *NUMBER OF SKUS*, in Models 2 and 4 the dependent variable is the log of *NUMBER OF SKUS*.

Validating the *Heterogeneity* and *Variety Seeking* Measures

We caution that it is possible that our calculation of the *Heterogeneity* and *Variety Seeking* measures may mechanically introduce correlation with *NUMBER OF SKUS* by construction. In particular, it is possible that when the retailer adds an additional variant of an item, this may lead to different rates of substitution from the existing variants. If the additional variant results in a greater rate of substitution from the most popular variant than the second most popular variant, then an increase in *NUMBER OF SKUS* may lead to an increase in both *Heterogeneity* and *Variety Seeking*. We investigate this issue in the Appendix, where we divide our two years of detailed transaction into two samples: Year 1 and Year 2. This revealed several examples of PrimarySKUs for which the number of variants changed between Year 1 and Year 2. We then compared how our measures of *Heterogeneity* and *Variety Seeking* changed for these

PrimarySKUs between the two years. Our null hypothesis is that the *Heterogeneity* and *Variety Seeking* measures should be stable between the two years. The alternative hypothesis is that the increase in *NUMBER OF SKUS* led to uneven substitution from the most popular variants, resulting in a change in the *Heterogeneity* and *Variety Seeking* measures between the two years.

The findings indicate that the heterogeneity and variety seeking measures were not affected by increases or decreases in the *NUMBER OF SKUS*. This evidence is reassuring, and suggests that the positive coefficients reported in Table 8 (when regressing *NUMBER OF SKUS* on the measures) does not reflect a mere mechanical correlation introduced when constructing the measures.

We can also ask a related question. If higher (lower) levels of preference heterogeneity and variety seeking motivate retailers to introduce (remove) variants, then we might expect that variation in *Heterogeneity* and *Variety Seeking* could be used to predict changes in the *NUMBER OF SKUS*. In particular, do our measures in Year 1 help predict which PrimarySKUs will see an increase or decrease in the number of variants in Year 2? These findings are also reported in the Appendix. They confirm that our measures of preference heterogeneity and variety seeking are predictive of the change in *NUMBER OF SKUS*. In particular, there is strong evidence that the retailer discontinues the second variant when sales are low compared to the most popular variant.

In the next sub-section we will 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:

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

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

The unit of analysis is a cost increase event, and *Retail 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; $\mathbf{X}_i$ describes the matrix of other explanatory variables; and **B** is a vector of coefficients. 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 6 suffer from an omitted variables problem.[19] 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 2) 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). A 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).

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. The findings for each measure are reported in Table 9.

Table 9. Heterogeneity and Variety Seeking
Correlations Within and Between Product Categories

Basis Used to Select Item Pairings	Heterogeneity	Variety Seeking	Number of Pairs (Sample Size)
Within Product Categories	0.3015**	0.5145**	355
Across Entire Sample	-0.0204	0.0167	467

The table reports Pearson pair-wise correlation coefficients between randomly selected pairs of items. Missing observations arise when an odd number of items in a product category prevent matching of a final pair.

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 these findings as evidence that heterogeneity in preferences and variety seeking are similar across different items in the same category.

[19] We controlled for a range of observable factors that are likely to affect the retailers' decision in Table 6. However, there may be other unobservable factors that are correlated with both *NUMBER OF SKUS* and the retailer's decision.

Results

In Table 10 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 again 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 *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. It is this finding that supports our claim that there is a link between microeconomic behavior (preference heterogeneity and variety seeking) and macroeconomic outcomes (price stickiness).

Table 10. GMM Results

	Model 1 **NUMBER OF SKUS**	**Model 2** **log of** **NUMBER OF SKUS**
NUMBER OF SKUS	-0.0785** (0.0298)	
Log NUMBER OF SKUS		-0.3742** (0.0986)
Prior 99-cent Price Ending	-0.1681** (0.0244)	-0.1658** (0.0227)
Cost Change	0.1881** (0.0456)	0.1820** (0.0682)
Prior Profit Margin	-0.8352** (0.0912)	-0.9164** (0.0915)
Purchase Volume (log)	0.0225* (0.0102)	0.0345** (0.0097)
Wald Chi^2 (20 d.f.)	690.41	657.50
First Stage Adjusted R^2	0.0923	0.1694
Sample Size	7,204	7,204

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).

These findings also offer an additional source of reassurance about the analysis in the previous section. Recall that in Section 4 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 4 reflect the omission of unobservable factors that are jointly correlated with both variables. The findings in Table 10 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.

The construction of our *Heterogeneity* and *Variety Seeking* measures suggest two reasons that the findings reported in Table 10 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*. As we saw in Table 5 and Figure 2, within-category variation appears to contribute approximately half of the relationship between *NUMBER OF SKUS* and the decision to increase the price following a cost increase. The findings in Table 10 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*. More generally, if we identified alternative metrics that explain a larger portion of the variation in *NUMBER OF SKUS* it is possible that the findings reported in Table 10 would be strengthened.

Further Decomposing Variety Seeking

One approach to improving the instruments is to decompose the source of customers' variety seeking. The transaction data includes some occasions in which customers purchased the two most popular variants on the same shopping trip, together with other examples in which they purchased these variants on different shopping trips. It is possible that the retailer reacts to these two types of variety seeking in different ways. For example, customers' preference for variety seeking may be more apparent to the retailer if customers purchase both items on the same trip. By comparing how the findings change when we decompose variety seeking into two these types of behaviors we can evaluate the robustness of the findings and whether this decomposition strengths the findings. The findings from this decomposition are reported in the Appendix. When using these alternative instruments for *NUMBER OF SKUS* our findings are replicated, and slightly strengthened. This is reassuring and indicates that the findings are not overly sensitive to the construction of the instruments.

Our analysis has so far focused on cost increases and studying when the retailer responded by increasing the retail price. While we might expect a similar pattern of findings if we study the response to 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. We search for evidence of asymmetries in the next section, where we measure how often the retailer decreases prices in response to a cost decrease.

6. Cost Decreases

The price and cost change database includes 5,867 examples of cost decreases. Just 11.8% of these cost decreases resulted in price decreases (5.7% led to price increases). In comparison, recall that cost increases resulted in price increases 69.3% of the time. While cost increases often lead to a price increase, cost decreases rarely result in price decreases. This asymmetry has been recognized elsewhere in the literature. For example, Peltzman (2000) reports findings from a large sample of consumer and producer goods that exhibit similar patterns.[20]

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 in Table 11.

The results 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.

Using the linear OLS model (Model 3) we can also compare the magnitudes of the coefficients for the other explanatory variables with the findings previously reported for price increases (Table 5). The *Prior 99-cent Price Ending* coefficient was -0.1965 in Table 5, compared to 0.0188 in this model. 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).

---

[20] 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 a Chicago supermarket chain. He attributes this null finding to a distinction to individual firm decisions and market outcomes. The findings that we report 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 11. Factors that Contribute to the Decision to Lower the Price

	Model 1	Model 2	Model 3 OLS	Model 4 Fixed Effects
NUMBER OF SKUS	-0.0130 (0.0344)			
Log NUMBER OF SKUS		0.1152 (0.1482)		
NUMBER OF SKUS = 2 or 3			0.0244 (0.0181)	0.0211 (0.0158)
NUMBER OF SKUS = 4 to 6			0.0131 (0.0315)	-0.0029 (0.0365)
NUMBER OF SKUS = 7 or more			-0.0263 (0.0281)	-0.0147 (0.0435)
Prior 99-cent Price Ending	0.2768 (0.2283)	0.2768 (0.2284)	0.0188 (0.0154)	0.0013 (0.0098)
Absolute Size of Cost Change	5.3057** (0.5541)	5.3381**		

are now different; the firm appears to pay more attention to the size of the cost change when deciding whether to decrease prices. On the other hand, it pays less attention to the prior profit margin.

We conclude that the number of variants plays an important role in deciding when to increase prices following a cost increase. However, it does not appear to influence the decision to decrease prices after a cost decrease. This asymmetry is consistent with other evidence that the firm uses different criteria to evaluate price increases and price decreases. Most notably, 69.3% of cost increases resulted in a price increase, but just 11.8% of cost decreases led to a price decrease.

7. 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 documenting that menu costs have a direct effect 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 by estimating that, absent these menu costs, the number of price changes would increase by 8% to 18%. Our 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 cost of changing prices across products. We supplement this analysis by investigating why the number of variants differs across items. Using a large panel of historical consumer purchases, we show that variation in consumer preferences partially explains why firms offer multiple product variants, identifying a novel link between consumer heterogeneity and price stickiness. Taken together, these findings establish one of the first micro quantifications of the menu-cost channel.

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

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.

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.

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.

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.

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.

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.

Appendix

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	39.7%	0.4%	15,153
Size of the Cost Change	9.1%	0.2%	15,153
Prior Profit Margin	$3.66	$0.04	15,153
Log Purchase Volume	10.14	0.02	15,053

This table reports summary statistics for the 15,153 cost increases in the cost and price change database. All of these variables are reported in that database. Missing observations reflects missing data in the database.

***NUMBER OF SKUS* Analysis: Summary Statistics**

	Average	**Standard Deviation**
Profit Margin	$2.51	$2.26
Purchase Volume (100,000 units)	4.31	8.15
Physical SKU Size		
Width (inches)	57.59	110.04
Height (inches)	24.56	29.55
Depth (inches)	5.75	2.72
Preference Heterogeneity and Variety Seeking		
Overall Sales Parity	70.20%	21.86%
Heterogeneity	56.53%	20.37%
Variety Seeking	13.67%	10.86%

The table reports summary statistics for the 934 PrimarySKUs used in Section 5. All of these PrimarySKUs have at least two variants and appear at least once in the price and cost change reports.

Number of Cost Increases

	Number of Cost Increases
NUMBER OF SKUS	-0.0004 (0.007)
Initial Retail Price	0.004 (0.002)
Initial Profit Margin	-0.953** (0.111)
Log Purchase Volume	0.102** (0.010)
Intercept	0.828** (0.118)
Adjusted R^2	0.0302
Sample Size	4,950

The table reports coefficients from an OLS model in which the dependent variable is the number of cost increases in the 55-month period. The sample includes all 4,950 PrimarySKUs that had at least one cost change or price change in our 55-month data period and that were present in the store throughout this period.

Validating the *Heterogeneity* and *Variety Seeking* Measures

It is possible that when the retailer adds an additional variant of an item, this may lead to different rates of substitution from the existing variants. In particular, if the additional variant results in a greater rate of substitution from the most popular variant than the second most popular variant, then an increase in *NUMBER OF SKUS* may lead to an increase in both *Heterogeneity* and *Variety Seeking*. More formally, let:

Y_A = the original demand of the most popular SKU

Y_B = the original demand of the most popular SKU

Δ_A = the % of demand substituted from the most popular SKU

Δ_B = the % of demand substituted from the second popular SKU

If $\Delta_A > \Delta_B$ then $Y_B / Y_A < (1- \Delta_B)Y_B / (1- \Delta_A) Y_A$. For this reason, it is possible that the positive association between *NUMBER OF SKUS* and *Heterogeneity* and *Variety Seeking* does not reflect the retailer adjusting its product range in response to the variation in customer preferences. Instead, it is possible that we are simply measuring are an artifact of unobserved sources of variation in *NUMBER OF SKUS*.

To investigate this possibility we divided our two-years of detailed transaction into two samples: Year 1 and Year 2. This revealed several examples of PrimarySKUs for which the number of variants changed between Year 1 and Year 2. We then compared how our measures of *Heterogeneity* and *Variety Seeking* changed for these PrimarySKUs between the two years. Our null hypothesis is that the *Heterogeneity* and *Variety Seeking* measures should be stable between the two years. The alternative hypothesis is that the increase in *NUMBER OF SKUS* led to uneven substitution from the most popular variants, resulting in a change in the *Heterogeneity* and *Variety Seeking* measures between the two years.

The transaction data includes a total of 2,215 PrimarySKUs for which it was possible to calculate the *Heterogeneity* and *Variety Seeking* measures in each of the two years.[21] For 148 of these PrimarySKUs the *NUMBER OF SKUS* was smaller in Year 2 than in Year 1; for 26 of them the *NUMBER OF SKUS* was larger in Year 2; and for the remaining 2,041 PrimarySKUs there was no difference between the years in the number of variants. In the table below we report the values of the *Overall Sales Parity*, *Heterogeneity* and *Variety Seeking* measures in each year for these three groups of PrimarySKUs.

[21] These PrimarySKUs satisfied the following criteria: the two most popular variants were available in both years and total sales of these two variants exceeded 100 units across the two years.

Change in the Preference Heterogeneity and Variety Seeking Measures

	Difference in Measures (Year 2 – Year 1)			
Change in NUMBER OF SKUS	**Overall Sales Parity**	**Heterogeneity**	**Variety Seeking**	**Sample Size**
Decreased	-5.67% (6.27%)	-4.58% (6.20%)	-1.09% (0.57%)	148
Unchanged	-0.12% (0.65%)	-0.09% (0.64%)	-0.02% (0.12%)	2,041
Increased	-9.91% (12.61%)	-11.53% (11.55%)	1.62% (2.54%)	26
Overall	-0.60% (0.75%)	-0.53% (0.74%)	-0.08% (0.12%)	2,215

The table reports the difference in the *Overall Sales Parity*, *Heterogeneity* and *Variety Seeking* measures, calculated as: Year 2 minus Year 1.

Overall (across all 2,215 PrimarySKUs) the measures are very stable between the two years with no significant difference in any of the measures. Because sales in each year are measured separately this is reassuring and can be interpreted as a measure of reliability. For the sample of PrimarySKUs in which the *NUMBER OF SKUS* changed, the smaller sample sizes result in larger standard errors. However, none of the differences between the years are significantly different from zero. More importantly, there is no apparent trend: the change in the *NUMBER OF SKUS* is not associated with a systematic increase or decrease in any of the three measures. For completeness we also calculated the pair-wise correlation between the change in *NUMBER OF SKUS* and the change in each of the three measures. The three correlations range from 0.007 to 0.008.

We conclude that it does not appear that the heterogeneity and variety seeking measures were affected by increases or decreases in the *NUMBER OF SKUS*. This evidence suggests that the positive coefficients reported in Table 9 (when regressing *NUMBER OF SKUS* on the measures) do not reflect a mere mechanical correlation introduced when constructing the measures.

We can also ask a related question. If higher (lower) levels of preference heterogeneity and variety seeking motivate retailers to introduce (remove) variants, then we might expect that variation in *Heterogeneity* and *Variety Seeking* could be used to predict changes in the *NUMBER OF SKUS*. In particular, do our three measures in Year 1 help predict which PrimarySKUs will see an increase or decrease in the number of variants in

Year 2? We report the findings in the table below, where we focus on the 1,033 PrimarySKUs that had exactly two variants in Year 1.[22]

Predicting the Change in *NUMBER OF SKUS*
Using Year 1 Preference Heterogeneity and Variety Seeking Measures

	Decrease	Unchanged	Increase
Year 1 Overall Sales Parity	18.65% (5.05%)	67.46% (0.84%)	50.33% (14.44%)
Year 1 Heterogeneity	15.76% (4.44%)	56.85% (0.76%)	40.95% (13.21%)
Year 1 Variety Seeking	2.9% (0.89%)	10.62% (0.31%)	9.38% (2.68%)
Sample Size	34	991	8

The table reports the Year 1 heterogeneity and variety-seeking measures for PrimarySKUs with a decrease, increase and no change in *NUMBER OF SKUS* between Year 1 and Year 2. Standard errors are in parentheses.

For 34 of the 1,033 PrimarySKUs the retailer discontinued the less popular variant in Year 2. As we anticipated, the Year 1 *Overall Sales Parity*, *Heterogeneity* and *Variety Seeking* are all significantly lower for these 34 items than for the 999 items for which the retailer continued selling the less popular variant. These differences are very large and are consistent with the retailer re-arranging its product line to remove poorly selling variants.

There are only 8 items for which the retailer introduced additional variants in Year 2, which is too few to draw reliable conclusions. However, using all 1,033 items we can calculate the correlation between the change in *NUMBER OF SKUS* between the years, and the variation in the three Year 1 measures. These correlation coefficients are 0.234, 0.217 and 0.115 for the *Overall Sales Parity*, *Heterogeneity* and *Variety Seeking* measures (respectively). All three measures are significantly different from zero ($p < 0.001$). We conclude that our measures of preference heterogeneity and variety seeking are predictive of the change in *NUMBER OF SKUS*. In particular, there is strong evidence that the retailer discontinues variants when sales of the variant are relatively low compared to sales of the most popular variant.

[22] Two variants was the modal number of variants in our sample of 2,041 PrimarySKUs.

Sources of Variety Seeking

When introducing our *Heterogeneity* and *Variety Seeking* measures we used an example in which customers purchased 1,000 units of the most popular variant, and 800 units of the less popular variant. These 800 units included 200 purchases by customers who purchased both variants. We can extend this example by assuming that on 80 of those 200 purchases customers bought both variants on the same visit to the store. The remaining 120 purchases occurred on visits in which customers only purchased the less popular variant (their purchases of the more popular variant occurred on a different shopping trip). This suggests two new measures:

$$\text{Same Visit} = \frac{\text{Units of SKU B on visits that SKU A was also purchased}}{\text{Total units of SKU A}}$$

$$\text{Different Visit} = \frac{\text{Units of SKU B by customers who purchased SKU A on a different visit}}{\text{Total units of SKU A}}$$

Recall that SKU A is the more popular variant and so both measures are bounded by 0 and 1. Moreover, buying on the same visit and a different visit are mutually exclusive and collectively exhaustive, which implies: *Variety Seeking* = *Same Visit* + *Different Visit*. Intuitively, *Same Visit* represents the additional sales contributed by customers purchasing different variants of an item on the same shopping trip, while *Different Visit* represents variety seeking across different shopping trips. In our example, *Same Visit* has a value of 0.08 and *Different Visit* has a value of 0.12.

In the table below we report the GMM instrumental variables estimates when we decompose *Variety Seeking* into *Same Visit* and *Different Visit*. When using these alternative instruments for *NUMBER OF SKUS*, our findings are replicated, and slightly strengthened. This is reassuring and increases our confidence that the findings are not sensitive to the construction of the instruments.

GMM Instrumental Variable Results
Decomposing the Sources of Variety Seeking

	Model 1	**Model 2**
NUMBER OF SKUS	-0.0633** (0.0198)	
Log NUMBER OF SKUS		-0.3726** (0.0890)
Prior 99-cent Price Ending	-0.1733** (0.0227)	-0.1657** (0.0226)
Cost Change	0.1869** (0.0686)	0.1788** (0.0666)
Prior Profit Margin	-0.8224** (0.0885)	-0.9204** (0.0915)
Purchase Volume (log)	0.0176* (0.0072)	0.0345** (0.0089)
Wald Chi^2 (20 d.f.)	738.27	665.03
First Stage Adjusted R^2	0.1056	0.1767
Sample Size	7,204	7,204

The table reports GMM estimates from an instrumental variables model. Fixed year and month effects were included, but are omitted from this table. In Model 1 the endogenous variable is *NUMBER OF SKUS* and in Model 2 it is the log of *NUMBER OF* SKUS. Standard errors are in parentheses. The standard errors are clustered by the month of the observation (month*year).