Peter Boatwright & Joseph C. Nunes

Reducing Assortment: An Attribute-Based Approach

Most supermarket categories are cluttered with items, or stockkeeping units (SKUs), that differ very little at the attribute level. Previous research has found that reductions (up to 54%) in the number of low-selling SKUs need not affect perceptions of variety and therefore sales, significantly. In this research, the authors analyze data from a natural experiment conducted by an online grocer, in which 94% of the categories experienced dramatic cuts in the number of SKUs offered, particularly low-selling SKUs. Sales were indeed affected dramatically, increasing an average of 11% across the 42 categories examined. Sales rose in more than two-thirds of these categories, nearly half of which experienced an increase of 10% or more; 75% of households increased their overall expenditures after the cut in SKUs. In turn, the authors examine how different types of SKU reductions—defined by how the cuts affect the available attributes or features of a category (e.g., the number of brands)—affected purchase behavior differently. The results indicate that consumers experienced divergent reactions to the reduction in sizes, but they uniformly welcomed the elimination of clutter brought on by the reduction in redundant items. In addition, of households that were loyal to a single brand, size, or brand–size combination that was eliminated, nearly half continued purchasing within the category. Also, contrary to previous research on SKU reductions, the authors find that category sales depend on the total number of SKUs offered. The authors extend the previous research by showing that (1) category sales depend on the availability of key product and category attributes and (2) two particularly important attributes to consumers in an assortment are brand and flavor.

Regarding product assortment, the conventional wisdom among supermarket managers has been that "more is better." Consequently, the average number of stockkeeping units (SKUs) at a supermarket has grown from 6000 a generation ago to more than 30,000 items today (Drèze, Hoch, and Purk 1994; Food Marketing Institute 1993). Recognizing that a reduction in SKUs can clear away clutter and lower costs, grocery retailers have been under immense pressure in recent years to begin offering a more efficient assortment by simply eliminating the low- or non-selling items within a category. Yet retailers are generally reluctant to cut items for fear of losing consumers who will be unhappy with their offerings.

Most grocers realize that consumers often prefer stores that carry large assortments of products for several reasons (Arnold, Oum, and Tigert 1983). For one, the larger the selection, the more likely consumers are to find a product that matches their exact specifications (Baumol and Ide 1956). In addition, more products mean more flexibility, which is important if the consumer has uncertain preferences (Kahn and Lehmann 1991; Koopmans 1964; Kreps 1979; Reibstein, Youngblood, and Fromkin 1975) or is predisposed to variety seeking (Berlyne 1960; Helson 1964; Kahn 1995; McAlister and Pessemier 1982).

Peter Boatwright is Assistant Professor of Marketing, Graduate School of Industrial Administration, Carnegie Mellon University. Joseph C. Nunes is Assistant Professor of Marketing, Marshall School of Business, University of Southern California. Both authors contributed equally and are listed alphabetically. The authors thank a major online grocery shopping and delivery service. The authors thank Jeffrey Inman, Ajay Kalra, Xavier Drèze, and the three anonymous *JM* reviewers for their valuable comments. They also thank the participants of the Sheth Symposium (1999) at the Katz Graduate School of Business and the participants of the 1999 Southern California UCI/USC/UCLA Marketing Symposium for their suggestions.

Recent research, however, suggests that consumer choice is affected by the perception of variety among a selection, which depends on more than just the number of distinct products on the shelves. The consumer's perception of variety can be influenced by the space devoted to the category, the presence or absence of the consumer's favorite item (Broniarczyk, Hoyer, and McAlister 1998), the arrangement of an assortment and the repetition of items (Hoch, Bradlow, and Wansink 1999), and the number of acceptable alternatives (Kahn and Lehmann 1991). Therefore, many observers in industry and academia believe that, if they do it properly, grocers can make sizable reductions in the number of SKUs offered without negatively affecting sales. In a study by Drèze, Hoch, and Purk (1994), aggregate sales went up nearly 4% in eight test categories after experimenters deleted 10% of the less popular SKUs and dedicated more shelf space to high-selling items. This experiment lasted 16 weeks and tracked sales at 30 test stores and 30 control stores.[1]

Broniarczyk, Hoyer, and McAlister (1998) examine the link between the number of items offered, assortment from the consumer's perspective, and sales. They find that reductions (up to 54%) in the number of low-selling SKUs need

[1]In addition, an industry study by the Food Marketing Institute (1993) examined reductions in SKUs in six categories (toothpaste, salad dressing, toilet tissue, pet food, cereal, and spaghetti sauce) at eight stores of three participating supermarkets (24 stores in total). Items deemed duplicates were cut, yet shelf space remained constant and was allocated to high-volume products. The depths of cuts (the precise levels of which are unreported) were "limited," "recommended," or "extreme." The recommended level of cuts led to an increase in sales of 2%, but more extreme cuts reportedly led to a .76% reduction in sales.

Journal of Marketing
Vol. 65 (July 2001), 50–63

not affect perceptions of variety and therefore sales. In a field study, the researchers eliminated approximately half of the low-selling items in five categories (candy, beer, soft drinks, salty snacks, and cigarettes) in two test convenience stores while holding shelf space constant. Neither sales nor consumers' perceptions of variety differed significantly between two test stores and two control stores. As Broniarczyk, Hoyer, and McAlister (1998, p. 175) point out, however, their findings are limited by the extent to which their results would generalize to other categories and how the specific features of the category might make consumers more or less sensitive to SKU reductions (e.g., a category with a small number of brands). This research addresses both of these issues directly.

First and foremost, we examine how different types of SKU reductions—defined by how they affect the attributes available in a category (i.e., the number of brands, sizes, and flavors)—affect sales differently. Fader and Hardie (1996, p. 451) pose the question, "Is it sufficient to drop the lowest-selling items or is it wiser to eliminate all items sharing an ineffective SKU attribute level?" Few retailers would eliminate any item that is selling briskly, but category managers must wonder whether all low-selling SKUs can be cut with equal (or lack of equal) impact on sales and, if not, which SKUs should be the first to go. We believe that the answer can be found by tracking changes in attribute offerings. To date, no work we are aware of has examined the effect on sales resulting from different types of SKU reductions or, more specifically, cuts based on product characteristics other than being low- or nonselling.

Second, although the increases in sales were not significant at Broniarczyk, Hoyer, and McAlister's (1998) test stores, our results are similar to Drèze, Hoch, and Purk's (1994) in that sales can change significantly following a reduction in the number of SKUs offered (see Figure 1 and Table 1). Unlike Drèze, Hoch, and Purk (1994), who found that aggregate sales went up nearly 4% in 8 categories, we observed an average sales increase of 11% across the 42 categories we examined. In 35 categories, sales changed by more than 4%, and nearly half of all categories experienced an increase of 10% or more. Our research is the first to examine how these changes depend on the particular nature of the cuts (i.e., which attributes were affected). Our model finds that the resulting change in sales depends on how the reduction in items affects the features of the category (i.e., the availability of alternative attributes such as brands and flavors). Our model also allows for divergent consumer reactions to the assortment reduction, measuring not just the results for the average consumer but also the percentage of consumers who embrace (or reject) different aspects of the change in assortment.

Third, our results indicate that even while sales increased, purchase probabilities decreased slightly. Not all consumers welcomed a reduction in their selection, and these consumers stopped purchasing or purchased less frequently. But the majority of consumers bought larger quantities, and their increased purchases easily outweighed the loss in sales due to the minority that preferred the greater selection. Even so, our results raise questions about customer retention, and we therefore investigate why some households ceased purchasing in some categories. Although we document that customers who lose their favorite brand are likely to stop or reduce category purchases, we find that a large percentage (nearly 40%) of consumers who were

FIGURE 1
Effect of SKU Reduction on Sales

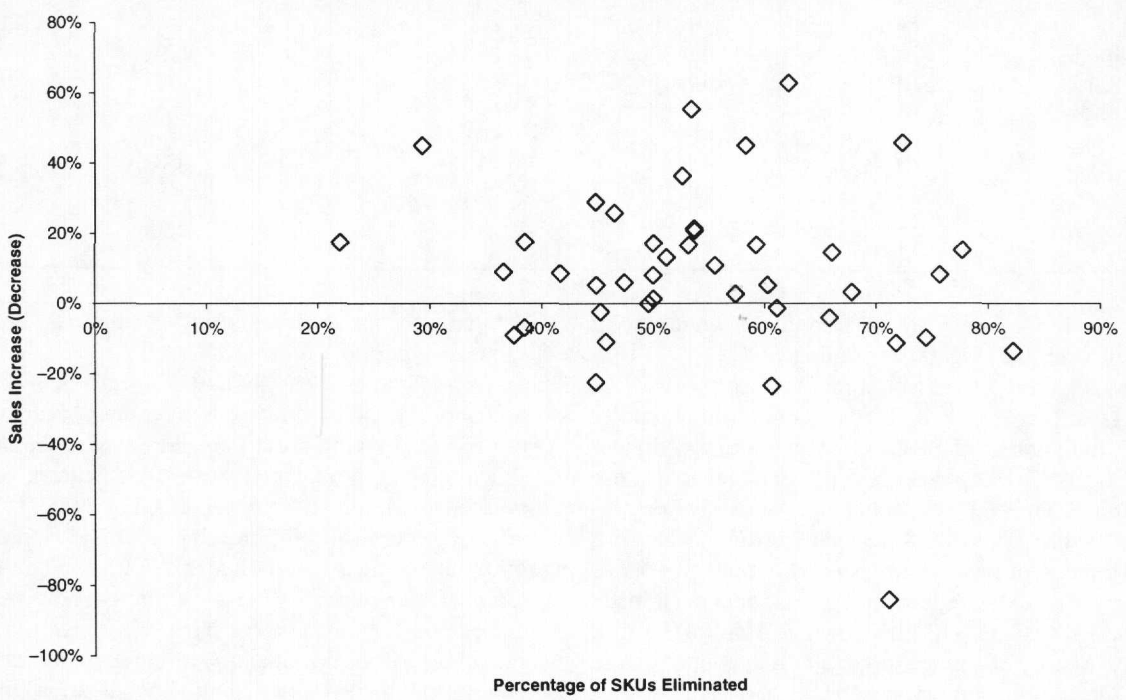

TABLE 1
Summary of SKU Reductions

	Percent Reductions of					
Category	**Items**	**Market Share**	**Brand**	**Size**	**Flavor**	**Brand–Size**
Cat food	22%	13%	15%	33%	27%	35%
Baby food	29	13	29	30	17	25
Facial tissue	37	20	17	25		31
Water	38	18	29	20		36
Tuna	38	29	0	36	0	38
Butter	38	3	25	0		20
Low-fat milk	42	3	50	20	0	44
Long-life milk	45	10	14	0		29
Laundry detergent	45	17	28	39		38
Soup	45	18	27	14	38	23
Milk	45	10	44	0	0	35
Flavored water	46	12	50	30	24	50
Paper towels	47	43	31	0		39
Bread	47	23	37	12	37	39
Cereal	50	24	55	16	31	41
Dish soap	50	37	0	40		47
Orange juice	50	10	50	33	0	56
Eggs	50	4	71	33	0	67
Waffles	51	40	40	43	23	47
Plastic bags	53	34	33	31		46
Frozen dinners	53	23	37	29		38
Refrigerated pasta	53	20	25	0	50	20
Soda/pop	54	25	29	54	40	46
Diet/sugar-free pop	54	21	27	50	37	47
Breakfast bars	56	14	33	33	35	41
Yogurt	57	29	27	33	54	29
Garbage bags	58	25	57	60		63
Coffee	59	38	38	26	67	39
Pasta	60	36	61	50	0	53
Cigarettes	61	42	53	0	0	53
Toilet paper	61	54	14	29		50
Cat litter	62	47	31	20		50
Juices	66	32	67	43	64	50
Crackers	66	40	56	41	58	53
Shredded cheese	68	52	40	33	25	50
Candy	71	27	45	58	45	67
Spaghetti sauce	72	47	67	68	50	74
Ice cream bars	72	42	45	53	58	66
Chunk cheese	74	55	40	42	53	50
Ice cream	76	42	40	0	66	61
Cookies	78	44	61	54	60	69
Diapers	82	62	50	61		74
Average	54	28	38	31	33	46

extremely loyal to a single brand continued to purchase in the category after that brand was eliminated.

Our investigation differs from previous work for several other important reasons. First, the data come from a retailer that (1) cut the number of SKUs offered in almost all categories (383 of 407 categories) and (2) implemented these cuts indefinitely. Second, the data come from an online grocer, which enables us to track the purchase behavior of a large, stable panel of regular shoppers who made purchases both before and after the reduction. These shoppers all made real purchase decisions with their own real dollars. Third, the purchase history of a much larger control panel of online customers, for whom the assortment did not change, provided a base for comparison. This control group's purchase data came from consumers in the same market during the same time period.

We should be clear about the objectives of this research up front. We do not attempt to measure consumers' perceptions of assortment, nor do we address issues related to shelf space allocation and product display. Instead, our primary goal is to explore how different types of SKU reductions affect category sales differently. Fader and Hardie (1996, p. 450) have argued that individual SKUs can be described in terms of a small set of discrete attributes and that "consumer choice is often made on the basis of these attributes." If the consumer views the category in terms of its attributes, we believe that the marketer should pay attention to changes in the availability of options among the attributes within a category. By monitoring how additions and deletions affect the presence or absence of attribute options, we believe that managers can better anticipate how category sales will respond to changes in SKU offerings and thus fine-tune their selection of SKUs in a more profitable manner. Our model tests how various changes in the availability of options among meaningful attributes, which are brought on by changes in product offerings, affect category sales.

We also believe that an understanding of which attributes are meaningful to consumers can help explain how different types of SKU reductions can affect category sales differently. Our secondary goal is to determine which attributes, if any, are meaningful across a wide variety of categories and are susceptible to changes that have a direct impact on assortment. Just as we believe that attributes should matter, we believe that not all attributes should matter equally. We recognize that a reduction in the availability of certain attributes could have a positive or negative effect and that the direction probably depends on the number of options originally offered within a category. It stands to reason, however, that changes in the availability of meaningless attributes will have no significant effect on category sales, whereas the magnitude of the effect of meaningful attributes will differ. For this work, we chose only attributes we believed would be meaningful in the vast majority of categories, yet we still expect that some will be more meaningful than others.

The rest of this article is organized as follows: In the next section, we briefly discuss how we selected the particular attributes we test. We then describe the data in more detail and explain why they are remarkably appropriate for this research. In the section that follows, we specify our measures and outline the development of a formal model that is used to examine how the particular nature of different SKU cuts affects sales differently. The results suggest that changes in category sales can be explained by changes in the availability of attributes. Finally, we discuss some secondary findings and the noteworthy managerial implications that emerge from these results. We conclude by reviewing some limitations of this work and offering suggestions for further research.

Meaningful Attributes

The first step in our analysis was to determine what would be treated as an SKU attribute. Marketing research firms typically use three important criteria (Fader and Hardie 1996). First, an SKU attribute must be consumer recognizable, that is, immediately observable to the consumer. Second, it must be objective, or precise. Third, it must be collectively exhaustive; in other words, it must apply to every SKU in the category.

At the onset, we attempted to be as inclusive as possible, yet the nature of the research, which examines issues that affect all categories, required that the attributes we chose transcend product categories. For this reason, we began by including brand, size, flavor, package, and form, all five of the key product characteristics articulated by the Food Marketing Institute in its 1993 report on variety and duplication. None of the categories examined included a significant number of different packages (e.g., Marlboro soft versus hard pack) or forms (e.g., Grape Nuts versus Grape Nut flakes), and therefore the categories could not experience significant changes along these attributes. Consequently, the three SKU attributes included in our final analysis were brand, size, and flavor. These attributes conveniently mimic the product descriptions onscreen at most online grocers (see Figure 2) as well as those Broniarczyk, Hoyer, and McAlister (1998) mention in their discussion of the cognitive aspects of assortment perceptions.

Not every product cut will lead to the elimination of a brand, size, or flavor. Merchandise managers might purposefully or inadvertently eliminate entire brands and sizes from a selection, or they might just trim several specific brand–size combinations. For example, if 12-packs of Sprite are eliminated but the consumer can still get 12-packs of alternative sodas (e.g., 7UP) and other sizes of Sprite (e.g., 2-liter bottles), then only a specific brand–size combination was cut; the brand and size themselves are still available. The total number of brand–size combinations cut could be viewed as indicative of a reduction in redundant items as long as no brands or sizes are eliminated. Consequently, we expect these types of cuts to have a positive effect on sales, because eliminating clutter makes it easier for consumers to find what they are looking for. If too much choice is truly demotivating (Iyengar and Lepper 1998), eliminating redundant items should help boost category sales (a simple positive relationship).

In addition, we included a variable that indicated the percentage of total items cut within each category. Unlike previous research that included either one level (10% in Drèze, Hoch, and Purk's [1994] work) or low, moderate, and high levels of cuts (25%, 50%, and 75% in Broniarczyk, Hoyer, and McAlister's [1998] work), the categories in our data experienced 27 different levels of reductions. Although this variable is identical in its operationalization to previous research, we recognize that it incorporates many of the differences among products that are specific to a particular category. For example, the Bounty brand of paper towels comes in white, medley, or fun prints. Strictly speaking, pattern is neither a package or form (e.g., extra strength, shorter sheets) difference, but if cut, a unique attribute disappears from the assortment. Consequently, we would expect a change in the number of items to be correlated with a change in sales (i.e., unexplained variance) because of these types of product distinctions disappearing. These attributes differ by category, occur in small numbers, and may or may not be valued by consumers. Therefore, a priori, we cannot predict the effect of item cuts on sales. We test the effect of varying degrees of cuts among brands, sizes, flavors, items, and brand–size combinations with data from a natural experiment that was conducted independently by an online grocer.

The Data

Electronic commerce has begun to change the nature and economics of food marketing radically as millions of computer-savvy consumers have begun shopping for groceries online. At the time of the experiment, 435,000 of the 53.5 million U.S. consumers online had purchased food products over the Internet, and a study by Anderson Consulting predicted that the online market for groceries and related products would reach $85 billion by 2007 (Thompson 1998). Most online grocers were reliant on the professional shop-

FIGURE 2
Typical Display at an Online Grocer: Sample from the Spaghetti Sauce Aisle

Spaghetti Sauce – 80 items		☐ Remember Sort		Alphabetical ▼
	Description	Size	Unit Price	Price
Buy	Barilla Pasta Sauce Marinara	26 OZ JAR	$0.11/OZ	$2.87
Buy	Barilla Pasta Sauce Sweet Pepper Garlic	26 OZ JAR	$0.11/OZ	$2.87
Buy	Barilla Pasta Sauce Tomato Basil	26 OZ JAR	$0.11/OZ	$2.87
Buy	Barilla Pasta Sauce Mushroom & Garlic	26 OZ JAR	$0.11/OZ	$2.87
Buy	Classico Di Capri Sundried Tomato Sauce	26 OZ JAR	$0.12/OZ	$3.17
Buy	Classico Di Firenze Spinach Cheese Sauce	26 OZ JAR	$0.12/OZ	$3.07
Buy	Classico Di Genoa Tomato Pesto Sauce	26 OZ JAR	$0.12/OZ	$3.07
Buy	Classico Di Liguria Tomato Alfredo Sauce	26 OZ JAR	$0.12/OZ	$3.07
Buy	Classico Di Napoli Tomato Basil Sauce	14 OZ JAR	$0.14/OZ	$1.97
Buy	Classico Di Napoli Tomato Basil Sauce	26 OZ JAR	$0.12/OZ	$3.17
Buy	Classico Di Parma 4 Cheese Pasta Sauce	26 OZ JAR	$0.12/OZ	$3.17
Buy	Classico Di Sorrento Roasted Garlic Sauce	26 OZ JAR	$0.12/OZ	$3.07
Buy	Classico Ripe Olive Mushroom Pasta Sauce	26 OZ JAR	$0.12/OZ	$3.17
Buy	Classico Spicy Red Pepper Sauce	26 OZ JAR	$0.12/OZ	$3.17
Buy	Classico Sweet Pepper Onion Pasta Sauce	26 OZ JAR	$0.12/OZ	$3.07
Buy	Contadina Pasta Ready Oil Garlic Spice	14.5 OZ CAN	$0.07/OZ	$0.95
Buy	Enrico's Pasta Sauce Garlic Lover	26 OZ JAR	$0.15/OZ	$3.79
Buy	Francisco Rinaldi Pasta Sauce Dolce 3 Cheese	26 OZ JAR	$0.08/OZ	$2.08
Buy	Five Brothers Alfredo Pasta Sauce	17 OZ JAR	$0.17/OZ	$2.87
Buy	Five Brothers Alfredo w/ Mushroom	17 OZ JAR	$0.17/OZ	$2.87
Buy	Five Brothers Marinara Burgundy Pasta Sauce	26 OZ JAR	$0.11/OZ	$2.87
Buy	Five Brothers Romano & Garlic Pasta Sauce	26 OZ JAR	$0.11/OZ	$2.87
Buy	Five Brothers Roast Garlic Onion Pasta Sauce	26 OZ JAR	$0.11/OZ	$2.87
Buy	Five Brothers Tomato Basil Pasta Sauce	26 OZ JAR	$0.11/OZ	$2.87

pers they hired to pick up items that were ordered for delivery at a local supermarket. This system contributed to above-average operating expenses, which hovered between 12 and 23% in 1998. Online grocers needed to cut costs desperately, and most planned to do so by moving to a system with central warehouses and bin shelving, where a severe reduction in the number of SKUs was seen as essential and inevitable. Industry analysts and executives agreed that the product offerings at online grocers should be cut roughly in half (Food Marketing Institute 1997).

In November 1996, a national online grocer began servicing a major market on the east coast. For eight months, all of its customers were offered the exact same selection online as was available at a local, affiliated grocery chain. In July 1997, this grocer undertook an enormous experiment, severely reducing the number of products available to a test panel or experimental group of consumers.[2] The company simultaneously monitored the purchases of everyone else in the market (i.e., the control group), for whom the assortment remained unaltered. The reduction did not appear to affect

[2]This test panel consists of 292 active consumers, where "active" means that (1) the consumers placed at least five orders in the eight months prior to the SKU reduction and (2) the consumers had not cancelled their service before the month for which we calculated category sales.

attrition; only 9 of the 292 (3%) active consumers in the experimental group quit the service. In the control group, 28 of the 455 (6%) customers quit.

In January 1998, the company provided the authors with all the recorded information for 1997 (e.g., SKU, quantity, date of purchase, unit price, shopping basket) for both panels. Accordingly, we compared purchases made by households in the two groups for six months before and five months after the SKU reduction.[3] To measure the effect of the SKU reduction on sales within a category, we modeled dollar sales to individuals in the experimental group compared with sales to a control group. The use of a control group increases the reliability of the statistical results and simplifies the analysis by establishing credible purchasing baselines and eliminating the need for a large number of covariates. Because we compare data from identical time periods for the two consumer panels, we do not need to account separately for any effect on sales caused by holidays (e.g., Thanksgiving), seasons (e.g., summer), or any other variables (e.g., price changes) that we can assume affected both panels equally. In addition, the simplified shopping environment online (no shelves, no lines, and so forth) offers a unique opportunity to explore assortment as a separate issue from these complexities.

To make the task more manageable, the analysis focused on 42 of the 47 top-selling categories, each of which had at least $5,000 in sales for the period in question.[4] On average, the households in the test panel shopped online 1.56 times per month and spent $72.32 dollars per shopping "trip" among just those 42 categories examined. Within these categories, the test panel placed 5924 separate orders for 135,979 items (an average of 23 items per order). The panel chose among 4181 SKUs for the period before the cut versus 1852 SKUs afterwards (a 56% reduction). As might be expected, the 2329 SKUs eliminated were mainly small-share items (see Figure 3). Individual categories lost between 20 and 80% of the products previously offered (see Table 1); none of the 42 categories was eliminated entirely. This variance helps provide a rich source of real-world data on consumers' reactions to a broad spectrum of different types of SKU reductions.

Measures and the Model

In this experiment, the assortment was simultaneously altered in a large number of categories. At the most basic level, our model compares sales before and after the manipulation, which is the format for the standard t-test for the difference in group means. Although the t-test offers a nice conceptualization of our model, the t-test itself is simplistic.

[3]Although we have data for the test panel for November and December of 1996, our control panel data exclude these two months. We therefore restrict our analysis to overlapping time periods (i.e., 1997). Also, purchases for July, the transition month, were excluded from the analysis.

[4]The top 47 categories each included more than $5,000 in sales during the period in question. The categories of produce, fish, deli, bakery, and meats were eliminated from the analysis because it was impossible to apply the requisite attributes (i.e., there were few, if any, standard brands, sizes, or flavor classifications in these categories).

FIGURE 3
Market Shares of Retained and Eliminated Items

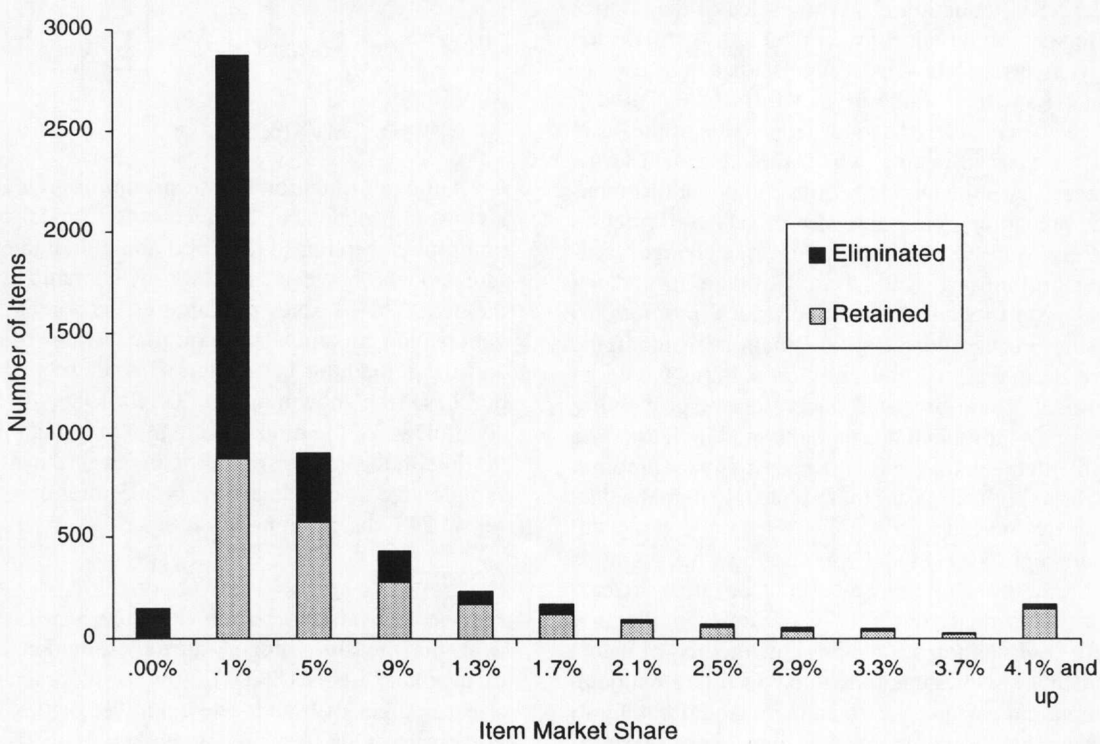

It requires the assumption that any difference in the sales in the two time periods could be attributed to the experiment, whereas we desire to measure effects of covariates on the change in sales. Some of the covariates are manipulations of the experiment, and others are outside the experiment. Because this experiment was not conducted in a lab but in a real purchasing environment, many factors are not accounted for in a simple t-test. Probably most important to our experiment are seasonal consumption patterns. For example, coffee sales are much higher in winter months than in summer months, and bottled water sales peak in late spring and early summer months. Because we have only one year of data, the periods of the experiment are potentially confounded with these seasonal consumption patterns. In addition, the prices of the individual products were not held constant but varied with the typical plethora of product promotions. Furthermore, these seasonal and promotional effects vary by category as well. In addition, the pre-cut period is six months, whereas the post-cut period is five months. Given the high variance in sales at the category level, the single month of additional observations can dramatically affect the variance of the estimate of mean sales, meaning that the mean sales estimates of the two periods do not have the same variance. Whatever statistical test is used to compare sales across the two time periods must account for differing variances.

Because of these complications, our basic model is in the following form:

$$(1) \qquad S_{jt} = \kappa\theta^{I_{REDUCT}} C_{jt}^{\gamma} \exp\left(\sum_{i=1}^{q} \phi_i x_{ijt}\right) \varepsilon_{jt},$$

where S_{jt} is sales of the experimental group in category j for month t; I_{REDUCT} is an indicator variable that takes the value of 1 after the SKU reduction; C_{jt} is the sales of the control group in category j in month t; $\{x_{ijt}\}$, i = 1, ..., q are covariates; and ε_{jt} is an error term. Linear regression can be used to estimate the parameters of the logged form of the model. This model can be considered a modified version of the basic t-test for differences in means, where the additional terms make the necessary adjustments for the complications mentioned in the preceding paragraph. Use of sales information from a control group (C_{jt}) adjusts S_{jt} for price changes, seasonality, and promotions. Both the control group and the experimental group received the same prices, promotions, and advertising. Furthermore, both populations reside in the same city, so their seasonal consumption patterns would be roughly identical. Therefore, as C_{jt} increases, the probability of purchase for an individual household would increase. The use of monthly data makes the homoskedasticity assumption on $\log(\varepsilon_{jt})$ (for hypothesis tests) more plausible than the data aggregated to two periods.

The basic regression model ignores the diverse tastes and needs of the consumers, however. Consumer heterogeneity is an additional source of variation, which is not accounted for by Equation 1. One possible method of incorporating consumer heterogeneity would be a mixture model, in which we would assume the consumers would belong to one of k segments of homogeneous consumers. Another possibility is to allow consumers all to differ from one another through a random effects model specification. A priori, there are no particular reasons to expect the consumer preferences to fall readily into one of a small number of segments. Furthermore, Allenby, Arora, and Ginter (1998) provide evidence that consumer heterogeneity is often much greater than that measured by a mixture model (latent class model). To allow for consumer heterogeneity, we model sales at the household level using a random effects model specification (variance components):

$$(2) \qquad S_{hjt} = \kappa_h\theta_h^{I_{REDUCT}} C_{jt}^{\gamma} \exp\left(\sum_{i=1}^{q} \phi_{ih} x_{ihjt}\right) \varepsilon_{hjt},$$

$$\Theta_h \sim MVN(\mu, \Sigma),$$

where S_{hjt} is sales to household h of category j in month t; $\Theta_h = [\kappa_h, \theta_h, \phi_{1h}, \ldots, \phi_{qh}]$, h = 1, ..., H; H is the total number of households in the data; and $\mu = [\bar{\kappa}, \bar{\theta}, \bar{\phi}_1, \ldots, \bar{\phi}_q]$ is a vector of the population means. Note that the household-level model includes an exponent γ for C_{jt}; because of the asymmetry of the log-normal distribution, individual household sales will not be exactly proportional to sales aggregated over a group of consumers. For simplicity, we specify Σ to be diagonal; that is, $\Sigma = \text{diag}(\sigma_\kappa^2, \sigma_\theta^2, \sigma_{\phi 1}^2, \ldots, \sigma_{\phi q}^2)$, the assumption that characterizes this model as variance components. Such models have been used extensively in the sciences to allow for variance at different levels in the data (Searle, Casella, and McCulloch 1992).

An additional complication arises with the model at the household level. Many households do not purchase from the online service in every category in every month, meaning that our data contain a large number of observations where $S_{hjt} = 0$. We therefore use a simple mixture model:

$$(3) \qquad S_{hjt} = \begin{cases} 0 \text{ w.p. } (1-p) \\ \kappa_h\theta_h^{I_{REDUCT}} C_{jt}^{\gamma} \exp\left(\sum_{i=1}^{q} \phi_{hi} x_{ijt}\right) \text{ w.p. } p \end{cases}$$

$$\Theta \sim MVN(\mu, \Sigma).$$

Although Equation 3 appears similar to a Tobit, it is not a censored model but an uncensored model mixed with a point mass. Similar to the Tobit and unlike latent class mixture models, however, mixture membership is observed in the data. The parameter p captures the percentage of observations that contain a purchase (the "when" of Gupta 1988), and the remaining parameters characterize the quantity of purchase (the "how much" of Gupta 1988). We estimate the coefficients of the model using maximum likelihood, where the likelihood for the model is given in the Appendix. The sample size of our data was 74,305 total observations, and for 43,219 observations $S_{hjt} > 0$.

Covariates

For the q covariates $\{x_i\}$ in Equation 3, we construct variables that measure category attributes: market share, number of available items (SKUs), number of brands, number of sizes, number of brand–sizes, number of flavors, and category price. Our measure of market share ($MKTSHR_{jt}$) is defined as the category dollar market share of the eliminated items in category j at time t. Product market shares, w_{nj}, are calculated relative to the category, or

$$(4)\qquad w_{nj} = \frac{\sum_{t=1}^{T_1} s_{njt}}{\sum_{m=1}^{N_j} \sum_{t=1}^{T_1} s_{mjt}},$$

where s_{njt} is the dollar sales of item n (n = 1, ..., N_j) in category j for week t (t = 1, ..., T_1) for the experimental group before the assortment cut,[5] and $\Sigma_{m=1}^{N_j} \Sigma_{t=1}^{T_1} s_{mjt}$ is the dollar sales of category j for the experimental group of consumers before the assortment cut. The share of the eliminated items in category j is then

$$(5)\qquad \text{MKTSHR}_{jt} = \frac{\sum_{l=1}^{L_{jt}} w_{lj}}{\sum_{n=1}^{N_j} w_{nj}},$$

where l = 1, ..., L_{jt} indexes the products eliminated from category j at time t. Because all items were available at first, MKTSHR_{jt} equals 0 before the SKU reduction.

The variable ITEMS_{jt} is the percentage of eliminated items:

$$(6)\qquad \text{ITEMS}_{jt} = 1 - \frac{\text{count(SKUs in category j) at time t}}{\substack{\text{count(SKUs in category j)} \\ \text{before assortment change}}}.$$

Similar to MKTSHR_{jt}, ITEMS_{jt} equals 0 before the SKU reduction, because all items were available at that time. Similarly, BRAND_{jt}, SIZE_{jt}, FLAVOR_{jt}, and BRNDSZ_{jt} are the percentages of brands eliminated, sizes eliminated, flavors eliminated, and brand–size combinations eliminated, respectively. Finally, PRCCHG_{jt} is the ratio of the category price at time t to that before the SKU reduction. Because we need the category price of offered products rather than purchased products, we calculated category price using the data for the universe of consumers rather than for our experimental or control group.[6] To calculate category price p_j, we used the share-weighted average of item prices,

$$(7)\qquad p_j = \sum_{n=1}^{N_j} w^u_{nj} \frac{1}{T_1} \sum_{t=1}^{T_1} p_{njt},$$

where n = 1, ..., N_j indexes the items within category j; w^u_{nj} is the market share of item n in category j calculated over the universe of consumers; t = 1, ..., T_1 indexes weeks before the assortment reduction; and p_{njt} is the price of item n in category j during week t. Our measure of category price change is then

$$(8)\qquad \text{PRCCHG}_{jt} = \frac{\sum_{m=1}^{M_{jt}} w^u_{mj} \frac{1}{T_1} \sum_{\tau=1}^{T_1} p_{mj\tau}}{p_j},$$

where products m = 1, ..., M_{jt} represent the products retained in category j. Before the SKU reduction, PRCCHG_{jt} takes the value of 1.

To allow for increasing/decreasing marginal effects of the independent variables, we also include the squares of the independent variables in our model. The loss of one brand versus zero brands of ten total may not affect sales, but the loss of seven brands versus six of ten total may affect sales greatly.

To facilitate interpretation of θ_h, we recode the variables so that they are mean-centered after the SKU reduction and equal to 0 before the SKU reduction. The parameter θ_h then can be interpreted as the average change in sales coinciding with the SKU reduction. Also, in 13 categories, flavors were either nonexistent (e.g., diapers) or ill-defined (e.g., frozen dinners). We therefore multiply the flavor covariates by an indicator variable I_F, which takes the value of 1 when flavors exist, 0 otherwise. In the hypothesis tests on flavor, we adjust the degrees of freedom to correct for the reduced number of observations used to estimate the flavor parameters.

Therefore, $[x_{1jt}, \ldots, x_{qjt}]$ = $[\text{BRAND}_{jt}, \text{BRAND}_{jt}^2, \text{SIZE}_{jt}, \text{SIZE}_{jt}^2, I_F \times \text{FLAVOR}_{jt}, I_F \times \text{FLAVOR}_{jt}^2, \text{ITEM}_{jt}, \text{ITEM}_{jt}^2, \text{BRNDSZ}_{jt}, \text{BRNDSZ}_{jt}^2, \text{MKTSHR}_{jt}, \text{MKTSHR}_{jt}^2, \text{PRCCHG}_{jt}]$. This matrix can also be summarized as [ATTRIBUTES, ITEMS, COVARIATES], where ATRRIBUTES is a matrix composed of the brand, size, and flavor variables; the ITEMS matrix is composed of the two items variables; and the remaining variables are included in the COVARIATES matrix.

Results

The results can be divided into two general classes: population mean parameters and heterogeneity (variance) parameters. The estimates of the population mean parameters $\mu = [\bar{\kappa}, \bar{\theta}, \bar{\phi}_1, \ldots, \bar{\phi}_q]$ and γ of the model are presented in the top portion of Table 2. Almost all of these parameters are statistically significant. At least one component (linear or quadratic) of market share, brand, and flavor is significant at the .05 level; however, size is not significant.[7]

These results indicate that (1) sales are affected by changes in the number of available SKUs; (2) the market share of items eliminated, the number of brands eliminated, and number of flavors eliminated all affected sales; (3) on average, category sales increased on the order of 11% in the experimental group; (4) a decline in sales due to brand reductions can be partially offset by reductions in brand–size combinations; and (5) reductions of assortment characteristics do not have a simple proportional relationship to sales.

[5]We used the experimental group here for our share calculation to ensure that we accounted for its preferences rather than those of the universe of consumers.

[6]The universe of consumers is the union of the experimental and control groups.

[7]The F-test for joint significance of the linear and quadratic factors for size also finds size to be insignificant.

TABLE 2
Parameter Estimates in Model of Sales (S_{hjt})

	Estimate	Standard Error	*p*-Value
$\bar{\kappa}$	1.592	.0197	<.0001
$\bar{\theta}$	.104	.0139	<.0001
γ	.324	.0083	<.0001
p	.582	.0018	
$\bar{\phi}$ *for*			
ITEMS	–8.601	.5447	<.0001
ITEMS2	9.805	.5614	<.0001
SIZE	–.208	.1536	.1766
SIZE2	.400	.2553	.1180
BRANDSIZE	2.419	.4480	<.0001
BRANDSIZE2	–2.629	.5194	<.0001
BRAND	1.192	.2107	<.0001
BRAND2	–2.620	.2740	<.0001
MKTSHR	1.063	.2631	<.0001
MKTSHR2	–3.540	.4539	<.0001
FLAVOR	.335	.1261	.0084
FLAVOR2	–1.040	.2423	<.0001
PRCCHG	2.783	.2958	<.0001
Heterogeneity Parameters			
σ^2_{κ}	.0953	.009	
σ^2_{θ}	.0241	.009	
σ^2_{ITEMS}	.2951	.174	
$\sigma^2_{ITEMS\wedge 2}$	.0000		
σ^2_{SIZE}	.6399	.099	
$\sigma^2_{SIZE\wedge 2}$	.0000		
$\sigma^2_{BRNDSIZE}$	.2136	.129	
$\sigma^2_{BRNDSIZE\wedge 2}$	.0000		
σ^2_{BRAND}	.2601	.126	
$\sigma^2_{BRAND\wedge 2}$	.0207	.196	
σ^2_{SHARE}	.1922	.137	
$\sigma^2_{SHARE\wedge 2}$	.2656	.365	
σ^2_{FLAVOR}	.1695	.039	
$\sigma^2_{FLAVOR\wedge 2}$	.0000		
$\sigma^2_{PRCCHNG}$	7.05	1.68	

Notes: The *p*-value on the linear and quadratic flavor terms of $\bar{\phi}$ have been adjusted for fewer degrees of freedom. In addition, all the signs and significances of the parameters are the same when the model is fit without categories with flavors.

The coefficient estimate for θ_h is .104, indicating that sales increased, on average, by $e^{.104}$, or approximately 11% after the assortment reduction, in which the average effect is across both consumers and categories. Overall, therefore, the experiment of reducing the assortment increased rather than decreased sales. For the effects of category attributes, most involve multiple parameter estimates. To clarify these effects, we have included graphs of some of these nonlinear relationships in Figure 4.

As shown in the graphs, initial reductions in share, the number of brands, and the number of flavors tend to increase sales, whereas further cuts engender sales declines. On average, the retailer in our data enjoyed maximal sales increases for reductions of approximately 25% of the brands, 18% of the flavors, and items accounting for 15% of the market share. Because the majority of the eliminated SKUs were small-share items (see Figure 3), a 15% reduction is already a large reduction of market share. Reductions in the clutter of brand–size combinations (when the number of brands and the number of sizes are held constant) increase sales of the category, though the marginal benefit of the reductions is decreasing. Finally, item reductions (not shown in Figure 4) decrease category sales, for which the negative effect is marginally decreasing. In all panels of Figure 4, the data histograms are included to show that the curves are estimated using data on both sides of the inflection points.

Brand reductions and brand–size reductions have opposite effects. Because of the close relationship of these two factors, along with their opposing impacts, it is theoretically plausible that brand reductions end up having a net effect close to zero. Although it is possible to eliminate brand–sizes without affecting the number of brands and/or sizes

FIGURE 4
Effects of Attribute Cuts on Sales

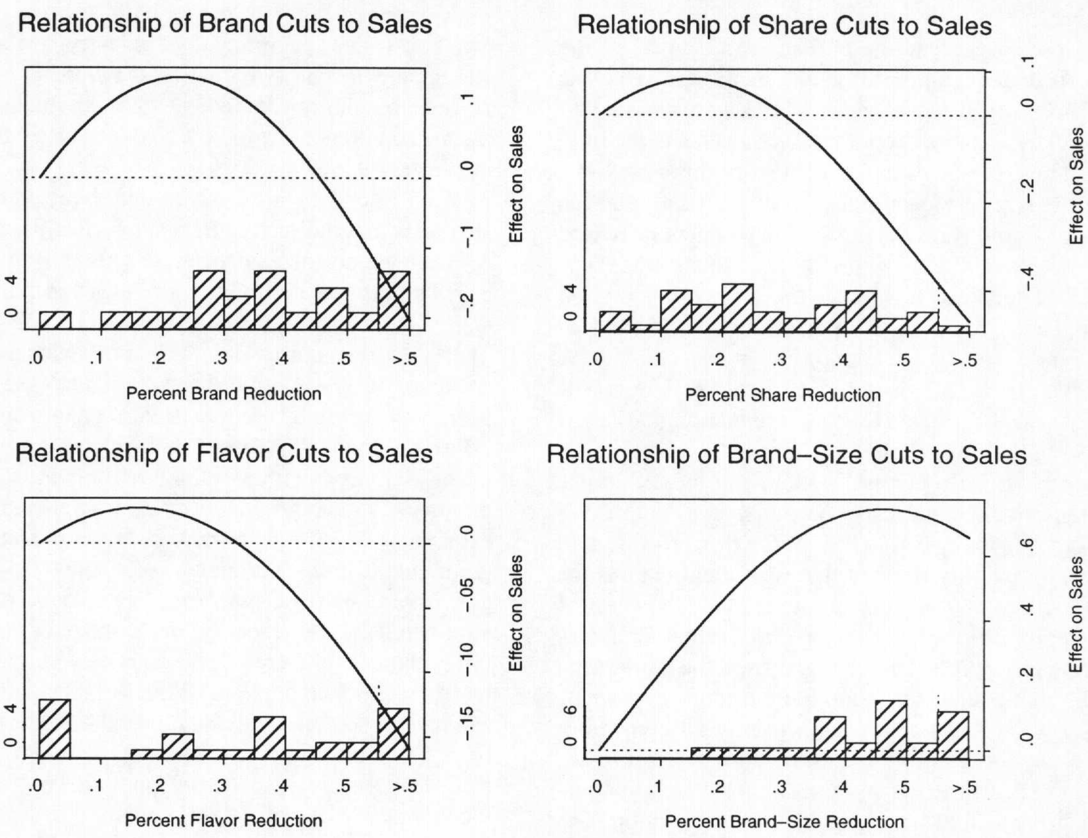

available, it is impossible to eliminate whole brands or sizes without affecting the number of brand–size combinations available. Typically, a grocery retailer does not carry all the brand–size combinations that are available from the manufacturer. In our data, a 10% reduction in brands resulted, on average, in approximately a 4% reduction in brand–size combinations (holding the number of available sizes constant). This relationship can be used to test the multivariate hypothesis $L\beta = 0$, where $L = [1, 1, .4, .4^2]$ and β is a vector of the parameter estimates of brand, quadratic–brand, brand–size combinations, and quadratic–brand–size effects, respectively. The F-statistic for this hypothesis was significant ($p < .0001$), meaning that, on average, a reduction in brands affects sales even after the positive effect of the reduction of brand–size combinations is taken into account.

The estimates of consumer heterogeneity are included in the lower portion of Table 2. Of particular importance is the heterogeneity on the change in sales due to the assortment reduction, θ_h, which is estimated as $\theta_h \sim N(.104, .0241)$. Of the mass of this normal density, 75% is greater than zero. Therefore, in addition to a significantly positive estimate for the population mean (i.e., average sales increased), the heterogeneity distribution estimates that 75% of consumers increased their category expenditures after the assortment cut.

For heterogeneity with respect to the elimination of product attributes, the results show that consumers are heterogeneous in their reactions to cuts of each attribute.[8] The largest heterogeneity parameter among the linear components of brand, size, and flavor is that for size, indicating that a reduction in the assortment of sizes results in a mixed reaction from consumers. The large degree of heterogeneity indicates that some consumers were upset with the reduction, but others welcomed the shorter list. Thus the heterogeneity parameter highlights an important result that is masked by the statistical insignificance of the population mean for size: The loss in sizes affects some members of the population quite strongly. Sales from some consumers decreased with the size deletions, but sales from other consumers balanced those losses such that the overall effect (the population mean) was close to null.

Although the heterogeneity parameter on the size reduction indicates a divergence of reactions, the small degree of heterogeneity indicated by the brand–size parameter suggests uniformity with respect to the elimination of clutter. When the estimated distribution of the brand–size linear term is N(2.419, .2136), a reduction in brand–size combina-

[8]Although some of the heterogeneity parameter estimates are zero for the quadratic effects, the results indicate that consumers are heterogeneous with respect to the linear portion of all the factors. The estimation algorithm (restricted maximum likelihood) is an iterative procedure wherein variance parameters are set to zero when they are, in essence, very close to zero.

tions elevates sales for the whole population, not just for a portion of it. In essence, no consumers are upset with the reduction of brand–sizes, provided that there is little change in the number of available brands and sizes.

We can also investigate the purchase probability of the population (i.e., the estimate of p of Equation 3, which is shown in Table 2). In some respects, this is a measure of category-level customer retention, in that it reflects the likelihood that a customer will make a category purchase, irrespective of the size of the shopping basket. (Recall that the distribution of θ_h indicates that 75% of consumers increased their category expenditures after the assortment cut.) The parameter p is purchase incidence, the probability that an observation (sales for a random household in a random month and random category) is not zero. In our data, 58.2% of the observations were nonzero. It is also possible to estimate p before and after the assortment reduction and test if the purchase probability changes because of the manipulation in assortment, for p is a proportion, and the difference in proportions is asymptotically normally distributed. We estimate p_1 to be .595 and p_2 to be .566, where p_1 refers to the purchase probability prior to the assortment cut. As for the hypothesis that $p_1 - p_2 = 0$, the *p*-value for the t-test of difference in proportions is less than .0001, which indicates that purchase probability has changed over the same time period as the change in assortment. For the control group, $\hat{p}_1 = .551$ and $\hat{p}_2 = .635$, which are significantly different from each other as well as from the experimental group estimates for p_1 and p_2. Therefore, during a season in which the purchase probability would be expected to increase, purchase probabilities for the experimental group significantly decreased. Again, the parameter p is a timing measure, a probability of purchase. This parameter decreases if some households buy less frequently in some categories. Our results indicate that some households did buy less frequently in some categories after the assortment reduction, though the loss in sales due to a smaller p was outweighed by an increase in purchase quantities, because overall sales increased.

In summary, we found that category sales increased, though the likelihood of making a purchase decreased, and that the attributes of a collection of products affected sales of the individual items within the collection. In addition, consumers' responses to changes in the availability of attributes differed across attributes (e.g., more mixed for size, uniform for brand–size combinations). In short, changes in the availability of attributes (i.e., attribute-based assortment) affect consumers differently and can help explain how category sales respond to changes in the SKUs offered.

Discussion and Implications

The primary objective of this research was to test the relationship between assortment and purchase behavior by examining how different types of changes in an assortment (i.e., SKU reductions) might affect sales differently. Unlike previous research on SKU reductions, our results indicate that a modest reduction, even one focusing on small-share items, can have a sizable impact on sales. Overall, we found that category sales tended to increase rather than decrease, on average, as a result of the SKU reduction. Broniarczyk, Hoyer, and McAlister (1998) and Drèze, Hoch, and Purk (1994) have suggested that SKU reductions can result in sales increases, and our model indicates that pragmatic cuts can boost sales significantly. Industry analysts have discussed how inventory reductions can lead to substantial cost savings, but our results indicate that pragmatic reductions in assortment do not just reduce costs but also can significantly increase sales.

Our secondary goal was to determine some meaningful attributes that indirectly affect sales, notably those that transcend a variety of categories and directly affect attribute-based assortment. Eliminating brands and flavors to a small degree helped sales, but deep cuts led to a decrease in sales. Therefore, the number of brands and flavors available within a category was meaningful and should be taken into account when category managers consider changes in their product offerings. Managers should be aware that consumers value the availability of their preferred brands and flavors, but pragmatic cuts in product offerings can increase overall category sales. The results may be summarized by a simple guideline: Eliminate redundant attributes while, on the margin, minimizing the number of brands, sizes, and flavors eliminated. More generally, the results suggest that simplifying choice can increase the sales of Web-based firms.

As mentioned previously, the formal attrition rate for the test panel was smaller than that for the control group, on the basis of membership cancellations. Yet there was a significant decrease in the category purchase probability despite the increase in overall sales, indicating that some consumers stopped purchasing in some categories even though the average consumer welcomed the assortment reduction by buying more. The category purchase probability is conditional on the customer returning to the store, but retailers may also be interested in the likelihood that a customer will not return. It is extremely difficult to obtain a reliable estimate of this measure of customer retention, as it must be estimated using interpurchase times, which in our data exhibit strong periodicity (people tend to purchase in weekly intervals). We are not aware of an applicable model that allows for this type of periodicity in the statistics or marketing literature. Rather, given our findings on the importance of category attributes, we investigated retention at the attribute level. We did this by examining the willingness of consumers to make purchases within a category in which their favorite brand, size, or brand–size combination was eliminated. Although the SKUs eliminated were typically small-share items, they frequently constituted the sole class of products along a particular attribute purchased within a particular category for a household. By focusing on only those purchase records in which a household bought a singular brand, size, or brand–size combination within a category before the SKU reduction, we are most likely to capture the cases in which a household lost its favorite attribute.

When a household lost the only brand it had purchased within a category before the cut, that household returned to make a purchase in the category 38.7% of the time (see Table 3). As a point of comparison, the proportion of households that bought a singular brand that was not cut and returned to buy in the category was 69.4%. This difference was highly

TABLE 3
Category Repurchase Rates for Households Loyal to a Single Attribute

	Households Loyal to One		
	Brand	**Size**	**Brand–Size Combination**
Within-category purchase records	3583	3422	2615
Purchased Item Eliminated			
Records in which sole item was cut	150 (4.2%)	150 (4.4%)	208 (8.0%)
Continued purchasing in category	58 (38.7%)	76 (50.7%)	111 (53.4%)
Purchased Item Retained			
Records in which sole item was not cut	3433 (95.8%)	3272 (95.6%)	2407 (92.0%)
Continued purchasing in category	2381 (69.4%)	2218 (67.8%)	1616 (67.1%)

significant (z-score = –7.57, $p < .0001$). Consumers were much less likely to buy in a category when their favorite brand was eliminated. Similarly, fewer consumers returned to shop in a category if their preferred size was eliminated (50.7%) than if it remained (67.8%), and this difference, though much smaller than the brand effect, was also highly significant (z-score = –4.11, $p < .0001$). In the case in which consumers purchased only one brand–size combination before the cuts and that combination was eliminated, 53.4% came back to buy in the category. Those whose sole brand–size combination was not cut were far more likely (67.1%) to buy in the category again during the post-cut period (z-score = –3.83, $p < .001$).

The loss in category purchasers seems large, yet retailers may be surprised to learn that nearly 40% of consumers who were extremely loyal to a singular brand and 50% who bought a single size or brand–size combination exclusively returned to purchase in the category after that attribute was eliminated. We should also note that eliminating brands had a more profound effect than either size or brand–size combination on leading these consumers to stop shopping in the category, which is not surprising given the results of our model regarding the significant impact of brand cuts on sales.

Overall, because of the heterogeneity of consumer preferences, elimination of any items will almost certainly cause some consumers to stop purchasing in the category. However, many of even the most product-loyal buyers will switch to an alternative product if their favorite item is eliminated. Furthermore, pragmatic product cuts can lead to sales increases, which greatly outweigh the loss in sales from the relatively few consumers who leave the category.

Limitations and Further Research

Although we develop a functional approximation of how sales are affected by changes in the availability of particular category attributes, our goal is exploratory and descriptive rather than explanatory. We use a random effects model that contains only a few category attributes and several covariates across a large number of categories to determine (1) if the availability of category attributes affects sales and (2) what some of those attributes are. Therefore, our research provides guidance to managers who aspire to alter their assortment. One common rule of thumb for an assortment reduction has been to eliminate small-share items. Although individually, small-share items contribute less to category sales, we find this rule of thumb to be overly simplistic. We find that judicious cuts of small-share items can increase sales rather than maintain current category sales levels. Cuts that eliminate redundancy boost sales, whereas cuts that eliminate key attributes reduce sales. For example, consider a simplistic category containing, among other items, 30-ounce regular Cascade, 30-ounce lemon-scented Cascade, and 75-ounce regular Cascade dishwasher detergent. Assuming that all these items have the same market share, the best item to eliminate would be the 30-ounce regular Cascade, because the consumer still would have access to lemon-scented and regular Cascade, and the consumer also would have the choice of 30 ounces or 75 ounces. Although our results do not offer optimal selections for individual categories, we find that the online grocery retailer would benefit by offering a smaller selection than the typical full-service bricks-and-mortar grocer offers.

Our results also raise several additional managerial issues worthy of further research. As would be expected, overly sparse assortments yield low sales. Although new products are typically viewed as either beneficial or neutral, our results imply that adding redundant new products may be harmful to category sales. We emphasize that our findings are based solely on item reductions. However, if these results hold for category additions, additional items (not necessarily new items) may depress overall category sales rather than increase them. We do not know of research that has investigated how the addition of category attributes (e.g., new brands) has affected category sales. Given the frequent introduction of new products into mature categories, we view this topic as one that is relevant to both managers and academics.

Future work could also investigate meaningful category-specific attributes, which might be tied to specific product types (e.g., food, nonfood). Although our work finds brand and flavor to be meaningful attributes, individual categories or clusters of similar categories should be considered if the goal is to develop a more accurate functional relationship between attributes and sales rather than the general reduced form function proposed here. For example, the relevant attributes in the toilet paper category (e.g., softness, color/

pattern, ply, brand, size) might be very different from those in the ice cream category (e.g., flavor, fat content, brand, size) but very similar to those in the paper towels category.

Although we find that brand, for whatever reason, is one of the most important category attributes, an analysis at the category level would be better suited to exploring specific brand associations. Consider the following example from Pan and Lehmann's (1993, p. 77) work: "Sony electronics are expensive (compared to other brands)." Or, a more appropriate example for this research: Hefty bags are "strong" plastic bags. Therefore, eliminating Hefty brand bags from an assortment might be synonymous with eliminating strong plastic bags. This research has shown that consumers respond strongly to the loss of brands, yet we recognize that brand may act as a surrogate for or may be confounded with other category attributes. In addition, the various associations with brand might differ across product categories. Exploring these associations is well beyond the scope of this article, yet in spite of this potentially high degree of variation in brand associations, brand appears as an important attribute across categories.

Future work in this area could also explore whether an optimal selection of items exists, which could depend on the availability of attributes and attribute options within a category. In other words, is there an absolute ideal for an assortment, or are consumers affected only by changes in items available (absolute or relative)?[9] Furthermore, do consumers tend to adapt to the selection offered by a retailer, as adaptation theory suggests (Bawa, Landwehr, and Krishna 1989; Helson 1964)? Another potential issue for further research pertains to the role of small grocery stores. Small grocery stores may end up serving a dual role—not only are they usually closer to consumers and decrease driving time, but they also may reap the demand-side benefits of reduced assortment (as well as the benefit of smaller inventory costs). Grocery chains may find it desirable to have multiple types of stores serving the same consumers—the megastore where everything can be found, as well as the streamlined outlet where the consumer can quickly and easily find general items.

More work could also be done using a model such as that developed by Rossi, Allenby, and Kim (2000) to determine how readily consumers substitute across attributes. Our research focuses on existing regular customers. Future work may more fully explore how assortment affects customer acquisition and retention. Finally, our results and those of other studies (Broniarczyk, Hoyer, and McAlister 1998; Drèze, Hoch, and Purk 1994; Food Marketing Institute 1993) show that sales are increased or at least unaffected by reductions in item counts. To understand fully how consumers respond to different assortments and changes in assortment, researchers must begin modeling the mental process consumers employ when confronted with an assortment of items in a category in which they are considering, or are intent on, making a purchase. We have already begun conducting a series of laboratory experiments and are making significant progress in determining the link among attribute availability, assortment perceptions, and purchase intention.

[9]This question was offered up by Debbie MacInnis.

Appendix

The parameters of Equation 3 can be estimated using conventional techniques, as the likelihood can be written as a binomial likelihood for p and a linear random effects model for the remaining parameters. The model is

$$\text{(A1)} \qquad S_{hjt} = \begin{cases} 0 \text{ w.p. } 1 - p \\ \kappa_h \theta_h^{I_{REDUCT}} C_{jt}^{\gamma} \exp\left(\sum_{i=1}^{q} \phi_{hi} x_{hijt}\right) \varepsilon_{hjt} \text{ w.p. } p \end{cases}$$

$$\Theta_h \sim MVN(\mu, \Sigma),$$

where $\Theta_h = [\kappa_h, \theta_h, \phi_{1h}, ..., \phi_{qh}]$, $\mu = [\bar{\kappa}, \bar{\theta}, \bar{\phi}_1, ..., \bar{\phi}_q]$, $\Sigma = \text{diag}(\sigma^2_\kappa, \sigma^2_\theta, \sigma^2_{\phi 1}, ..., \sigma^2_{\phi q})$, and $\log(\varepsilon_{hjt}) \sim$ i.i.d. $N(0, \sigma^2_e)$.

Given the sales S_{hjt} for household h = 1, ..., H of category j = 1, ..., J in month t = 1, ..., T along with covariates x_{hijt}, the likelihood is

$$\text{(A2)} \qquad L(\{\phi_{hi}, \kappa_h, \theta_h\}, p, \gamma | \{S_{hjt}, x_{ijt}\})$$

$$= \prod_{S_{hjt}=0} (1 - p) \prod_{S_{hjt} \neq 0} pg(\{\phi_{hi}, \kappa_h, \theta_h\}, \gamma | \{S_{hjt}, x_{ijt}\}).$$

Because g( ) is not a function of p, the distribution of p is proportional to the binomial distribution, meaning that the maximum likelihood estimate of p is

$$\text{(A3)} \qquad \hat{p} = \frac{n_1}{n_1 + n_2},$$

where n_1 is the number of observations where $S_{hjt} > 0$, and n_2 is the number of observations where $S_{hjt} = 0$.

The remaining parameters are all contained in the function g( ), the density of a random effects model. In general notation, the random effects model is as follows (Searle, Casella, and McCulloch, 1992):

$$\text{(A4)} \qquad y = X\beta + Zu + e,$$

where y is an $n_1 \times 1$ vector of observations, β represents fixed effects, u represents random effects, and Zu is partitioned as

$$\text{(A5)} \qquad Zu = [Z_1, \ ..., \ Z_r] \begin{bmatrix} u_1 \\ \vdots \\ u_r \end{bmatrix} = \sum_{i=1}^{r} Z_i u_i.$$

The length of u_i is q_i, where for our data $q_i = n_1$ for all i, the total number of observations of nonzero sales. In the customary random model, the random effects represented by u_i have the properties $E(u_i) = 0$ and $\text{var}(u_i) = \sigma^2_i I_{qi}$ for all i, and $\text{cov}(u_i, u'_h) = 0$ for $i \neq h$. In addition, $E(e) = 0$, $\text{var}(e) = \sigma^2_e I_{n_1}$, and $\text{cov}(u_i, e') = 0$ for all i. These assumptions lead to

$$\text{(A6)} \qquad E(y) = X\beta$$

and

$$\text{(A7)} \qquad V = \text{var}(y) = \sum_{i=1}^{r} Z_i Z'_i \sigma^2_i + \sigma^2_e I_{n_1}.$$

A notational convenience is to define $u_0 = e$ and $Z_0 = I_{n_i}$. This gives

$$\text{(A8)} \qquad y = X\beta + \sum_{i=0}^{r} Z_i u_i \quad \text{and} \quad V = \sum_{i=0}^{r} Z_i Z_i' \sigma_i^2.$$

The likelihood for all parameters but p is therefore

$$\text{(A9)} \quad L = L(\beta, V|y) = \frac{\exp\left[-\frac{1}{2}(y - X\beta)' V^{-1}(y - X\beta)\right]}{(2\pi)^{\frac{1}{2}n_i}|V|^{\frac{1}{2}}}.$$

We used Proc Mixed in SAS to obtain the maximum likelihood estimates of the parameters of this likelihood, and we report them in Table 2.

REFERENCES

Allenby, Greg M., Neeraj Arora, and James L. Ginter (1998), "On the Heterogeneity of Demand," *Journal of Marketing Research*, 35 (August), 384–89.

Arnold, Stephen J., Tae H. Oum, and Douglas J. Tigert (1983), "Determinant Attributes in Retail Patronage: Seasonal, Temporal, Regional, and International Comparisons," *Journal of Marketing Research*, 20 (May), 149–57.

Baumol, William J. and Edward A. Ide (1956), "Variety in Retailing," *Management Science*, 3 (1), 93–101.

Bawa, Kapil, Jane T. Landwehr, and Aradhna Krishna (1989), "Consumer Response to Retailers' Marketing Environments: An Analysis of Coffee Purchase Data," *Journal of Retailing*, 65 (4), 471–95.

Berlyne, D.E. (1960), *Conflict, Arousal, and Curiosity*. New York: McGraw-Hill.

Broniarczyk, Susan M., Wayne D. Hoyer, and Leigh McAlister (1998), "Consumers' Perceptions of the Assortment Offered in a Grocery Category: The Impact of Item Reduction," *Journal of Marketing Research*, 35 (May), 166–76.

Drèze, Xavier, Stephen J. Hoch, and Mary E. Purk (1994), "Shelf Management and Space Elasticity," *Journal of Retailing*, 70 (4), 301–26.

Fader, Peter S. and Bruce G.S. Hardie (1996), "Modeling Consumer Choice Among SKUs," *Journal of Marketing Research*, 33 (November), 442–52.

Food Marketing Institute (1993), "Variety or Duplication: A Process to Know Where You Stand," report prepared by Willard Bishop Consulting and Information Resources Inc. in cooperation with Frito-Lay. Washington, DC: Food Marketing Institute.

——— (1997), "Home Shopping Programs: A Strategic Guide," report prepared by Robert E. Linneman and Patrick J. Kirschling, Department of Food Marketing, Saint Joseph's University. Washington, DC: Food Marketing Institute.

Gupta, Sunil (1988), "Impact of Sales Promotions on When, What, and How Much to Buy," *Journal of Marketing Research*, 25 (November), 342–55.

Helson, Harry (1964), *Adaptation-Level Theory*. New York: Harper and Row.

Hoch, Steven J., Eric T. Bradlow, and Brian Wansink (1999), "The Variety of an Assortment," *Marketing Science*, 18 (4), 527–46.

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 (December) 995–1006.

Kahn, Barbara E. (1995), "Consumer Variety-Seeking Among Goods and Services: An Integrative Review," *Journal of Retailing and Consumer Services*, 2 (3), 139–48.

——— and Donald R. Lehmann (1991), "Modeling Choice Among Assortments," *Journal of Retailing*, 67 (3), 274–99.

Koopmans, Tjalliong C. (1964), "On the Flexibility of Future Preferences," in *Human Judgments and Optimality*, M.W. Shelly and G.L. Bryan, eds. New York: John Wiley & Sons, 243–56.

Kreps, David M. (1979). "A Representation Theorem for Preference for Flexibility," *Econometrica*, 47 (3), 565–77.

McAlister, Leigh and Edgar Pessemier (1982), "Variety Seeking Behavior: An Interdisciplinary Review," *Journal of Consumer Research*, 9 (3), 311–22.

Pan, Yigang and Donald R. Lehmann (1993), "The Influence of New Brand Entry on Subjective Brand Judgments," *Journal of Consumer Research*, 20 (June), 76–86.

Reibstein, David J., Stewart A. Youngblood, and Howard L. Fromkin (1975), "Number of Choices and Perceived Decision Freedom as a Determinant of Satisfaction and Consumer Behavior," *Journal of Applied Psychology*, 60 (4), 434–37.

Rossi, Peter E., Greg Allenby, and J. Kim (2000), "Modeling Consumer Demand for Variety," working paper, Graduate School of Business, University of Chicago.

Searle, Shayle R., George Casella, and Charles E. McCulloch (1992), *Variance Components*. New York: John Wiley & Sons.

Thompson, Maryann Jones (1998), "Grocery Shopping Online," *The Standard*, (September 7), (accessed April 20, 2001), [available at http://www.thestandard.com/research/metrics/display/0,2799,10005,00.html].