

MIT Open Access Articles

Advertising in a Competitive Market: The Role of Product Standards, Customer Learning, and Switching Costs

The MIT Faculty has made this article openly available. **Please share** how this access benefits you. Your story matters.

Citation: Anderson, Eric T., and Duncan Simester. "Advertising in a Competitive Market: The Role of Product Standards, Customer Learning, and Switching Costs." *Journal of Marketing Research* 50, no. 4 (August 2013): 489–504.

As Published: <http://dx.doi.org/10.1509/jmr.11.0538>

Publisher: American Marketing Association

Persistent URL: <http://hdl.handle.net/1721.1/90833>

Version: Original manuscript: author's manuscript prior to formal peer review

Terms of use: Creative Commons Attribution-Noncommercial-Share Alike



Advertising in a Competitive Market: The Role of Product Standards, Customer Learning and Switching Costs^{*}

May 2013

Eric T. Anderson
Northwestern University

Duncan Simester
MIT Sloan School of Management

Standard models of competition predict that firms will sell less when competitors target their customers with advertising. This is particularly true in mature markets with many competitors selling relatively undifferentiated products. We present findings from a large-scale randomized field experiment that contrast sharply with this prediction. The field experiment measures the impact of competitors' advertising on sales at a private label retailer of apparel, home furnishings and sporting goods. Surprisingly, for a substantial segment of customers the competitors' advertisements increased sales at this retailer. This is a robust effect, obtained through experimental manipulation and measuring actual purchases by large samples of randomly assigned customers. The effect is also large with customers ordering over 4% more items in some categories in the Treatment condition (compared to the control). We compare how these positive spillovers vary across product categories to illustrate the importance of product standards, customer learning, and switching costs. The findings have the potential to change our understanding of competition in mature markets.

Keywords: advertising, product standards, customer learning, switching costs

^{*} We thank our research assistant Hojin Jung for his contribution to this work. This paper has benefitted from seminar participants at the 2010 Yale Customer Insights Conference, and the 2011 Yale Marketing Industrial Organization Conference, together with audiences at the following universities: Alberta, Berkeley, Carnegie-Mellon, MIT, and Wisconsin.

1. Introduction

We report the findings from a large scale field experiment to investigate the impact of advertising in a competitive market. The field experiment was conducted with the cooperation of a large national retailer that sells private label apparel, home furnishings and sporting goods. Although we cannot reveal the identity of the retailer, for ease of exposition we will label it “Retail World”. In the field experiment approximately 370,000 of Retail World’s past customers were randomly assigned to two experimental conditions. The 185,000 customers in the Treatment condition received over 1.1 million targeted direct mail advertisements from Retail World’s close competitors. Most of these advertisements consisted of competing catalogs containing product information, prices and images for hundreds of products. Customers in the Control condition did not receive these advertisements.

The field experiment takes advantage of a common direct-mail practice: sharing of customer information between competing retailers. Sharing is often done on a reciprocal basis, so that firms exchange information for an equivalent number of customers. Firms use this information to identify where to send direct mail advertising, and this provides a unique opportunity to design a field experiment to measure the impact of competitors’ advertising on a firm’s sales. We measure this impact by comparing purchases at Retail World between the two experimental conditions using 20-months of transactions. This period encompasses the twelve month treatment period and eight subsequent months, and includes over one million orders and almost \$100 million in Retail World revenue.

Surprisingly, we find that for a substantial segment of customers the competitors’ advertisements increased sales at Retail World. This effect is large and it is obtained through comparison of actual purchases by large samples of randomly assigned customers. What makes the outcome particularly surprising is that the apparel, home furnishings and sporting goods categories are mature and comprised of many retailers selling relatively undifferentiated products.¹ The maturity of the markets ensures that all customers are aware of the product categories. Customers can easily search and purchase products at competing firms, customers do not have to learn how to use the products, and compatibility between products is generally not specific to a firm. These are product categories in which we might expect the threat from competition to be particularly strong. Yet, we find positive spillovers from the competitors’ advertising.

The size of the positive spillovers varied across customers. The largest increase in Retail World sales was from customers who had not made a recent Retail World purchase. The effect also varied across product categories. To explain these positive spillovers and the variation across both categories and customers

¹ The US retail apparel market contributes almost 3% of the US gross domestic product. In 2007 customers in the US purchased approximately 20.1 billion garments and 2.4 billion pairs of shoes, representing approximately \$371 billion in revenue (AAFA 2008). Sales of apparel and shoes are particularly strong in catalog and Internet channels. The Direct Marketing Association (DMA 2006) reports that the apparel category has a higher volume of Internet and catalog orders than any other product category. In 2006 approximately 32% of US adults ordered an item of apparel from a direct retailer (including 23% of males and 40% of females). Moreover, fully 24% of the 100 largest catalog retailers sell apparel. Not surprisingly, almost all major apparel retailers now use direct mail advertising and catalog and Internet channels to complement their traditional retail stores.

we will distinguish between the decision to make a purchase in a category, and the selection of which retailer to purchase from. Competitors' advertising may influence the first decision by priming customers to think about the category. However, even though the need to purchase may have been primed by competitors' advertising, some customers returned to Retail World to satisfy this need.

The focus of the paper is on the choice of which retailer to purchase from. In preliminary interviews for this research some customers described examples of products that they only buy from the Retail World because they have learned which sizes fit. One customer described how he had mailed an outline of his foot to Retail World, which the company used to determine his appropriate size. When he orders Retail World footwear he believes it will fit. However, he also knows that sizing differs across brands and so he avoids buying footwear elsewhere. A key feature of this example is the standardization of product sizes. While footwear sizes are standardized within Retail World's (private label) brand, they are different from sizes at competing brands.² If sizes were standardized across all brands, this customer could use Retail World's size information to also order from other retailers.

Standardization in sizes does not extend to all categories. For example, swimwear sizes are not standardized even within Retail World's brand. Customers may need a size 8 when purchasing one style of Retail World swimwear and a size 10 when purchasing another style. As a result the fit of past swimwear purchases may not help predict the fit of a new swimsuit. We will show that variation in how much customers learn about sizes from past purchases helps to explain where the impact from the competitors' advertising was largest.

While the customer interviews yielded useful initial insights, formally documenting these effects required that we construct measures of how much customers learn from past purchases. Fortunately this market provides an ideal measure of customer learning: changes in customer return rates. The main reason that customers return items is due to concerns about product fit. We compare the response to the competitors' advertising with the change in return rates between customers' first and last orders in the category. We find a clear relationship: categories with the larger decreases in return rates due to poor fit are the categories with the most positive treatment effects in the competitive advertising experiment. Across the 35 largest categories, the pair-wise correlation between the *Change in Return Rates* and the response to the competitors' advertising is -0.41 (the rank-order correlation is -0.50). We interpret this as evidence that learning about product sizes created switching costs in favor of Retail World. These switching costs led customers to purchase from Retail World, even though their need to purchase was primed by the competitors.

Related Literature

Spillover effects have often been studied using Feldman and Lynch's (1988) accessibility-diagnostics framework. This theory predicts that "an earlier response will be used as a basis for another subsequent response if the former is accessible and if it is perceived to be more diagnostic than other accessible inputs" (at page 421). In our application, products are associated through a network that includes both retail brands and categories. A retail advertisement may activate the retail brand, a category need and in-turn competing retail brands. This model has been used to explain both positive and negative spillovers among competing brands. For example, Janakiraman, Sismeiro and Dutta (2009) develop an

² Discrepancies in footwear and apparel sizes across brands are well-documented (see references on page 5).

empirical framework to investigate positive spillovers between competing brands in the pharmaceutical industry. They find that customers' quality beliefs about the first entrant can influence customers' prior beliefs about the quality of later entrants, but only if the products are similar. They do not find any evidence of positive spillovers through advertising. In particular, physicians do not appear to use detailing activities for one drug to form quality perceptions of competing drugs.

Examples of positive advertising spillovers have been established in the context of umbrella brands (among non-competing brands). For example, Balachander and Ghose (2003) show that these positive spillovers may not only project from the parent brand to the brand extension, but that reciprocal spillovers may also operate in the opposite direction (see also Morrin 1999).

Negative spillovers among competing brands have also been documented. Roehm and Tybout (2006, 2009) demonstrate how a product scandal may spillover from a brand like Vioxx, which was recalled due to safety concerns, to damage a competing brand like Celebrex. Analogously, Dark and Richie (2007) illustrate how deceptive advertising can have negative consequences for firms who are unrelated to the original deception. Our findings contribute to our understanding of spillovers by documenting both the existence of *positive advertising* spillovers between *competing* brands, and highlighting the moderating role of consumer switching costs.

Our findings also contribute to an extensive literature establishing the existence of state dependencies (consumer loyalty) in consumer choice (Dubé, Hitsch, Rossi and Vitorino 2008; Erdem 1996; Keane 1997; and Seetharaman, Ainslie and Chintagunta 1999). Switching costs represent an important potential source of state dependence. The earliest papers in the switching cost literature assumed that switching costs were exogenous and common across firms. Under these conditions switching costs tend to reduce competition and increase profits in equilibrium (Klemperer 1987). More recently it has been recognized that switching costs may be endogenous. For example, Bergemann and Välimäki (1996) endogenize switching costs in a multi-armed bandit problem in which buyers learn about product quality through experience. If firms can control the level of switching costs we might expect that they would prefer to increase them in order to dampen competition. However, the literature often predicts the opposite result. Anticipation that high switching costs leads will lead to higher second-period profits induces firms to compete aggressively in the first period. If the resulting price competition erodes profits, then firms may reduce switching costs to dampen first-period competition (see for example Caminal and Matutes 1990, Marinoso 2001, Bouckaert and Hans Degryse 2004, and Cabral and Villas-Boas 2005).

These distinct goals of "investing" to capture sticky customers and "harvesting" the surplus from them once captured makes it difficult to predict whether the presence of switching costs will increase or decrease the intensity of competition. This has led to recent empirical research evaluating the contrasting predictions. Dubé, Hitsch and Rossi (2009) provide a prominent recent example. They study the refrigerated orange juice and margarine categories and report that the presence of switching costs can act to amplify the intensity of competition in these grocery markets. Their study focuses on price competition among heterogeneous firms and they find that intermediate levels of switching costs may lead to lower prices. When switching costs are much larger, their numerical simulations show that prices can rise as competition is dampened.

Our findings complement these results by extending our understanding of how switching costs impact competitive intensity. The standard result in this context is that large switching costs create lock-in that insulates customers from competitors' actions. The surprise in our findings is that in the presence of

switching costs competitors' actions can generate positive spillovers that are large enough to outweigh competitive substitution. Moreover, the extant switching cost literature has largely focused on price competition in markets for low-involvement fast-moving consumer goods, while we focus on competitive advertising in markets for higher-priced consumer durables.

Dubé, Hitsch and Rossi (2009) attribute the state dependence in their study to a “psychological” switching cost, including cognitive dissonance in which customers change their preferences to rationalize previous choices. Elsewhere the literature has distinguished between “transaction costs” and “learning costs.” Examples of transaction costs include the cost of travelling to different gasoline stations or the time required to evaluate alternative suppliers and establish new account information. Learning costs describe the time and effort customers incur when learning how to use a new product (Erdem and Keane 1996, Akerberg 2003, Crawford and Shum 2005). Nilssen (1992) argues that this distinction is important. He predicts that because learning costs are only incurred once for each supplier while transaction costs are incurred every time a customer switches, transaction costs lead to a larger increase in equilibrium prices and contribute to a greater loss of social welfare.

The switching costs we identify in this paper do not fit neatly into either category. Because customers must learn about product fit when purchasing from a retailer they exhibit some of the characteristics of a learning cost. However, the cost to the customer also includes the risk of purchasing the wrong item, which is not the time and effort typically associated with learning. If customers purchase items that do not fit they may incur a transaction cost to return an item.

The findings provide support for recent work on the role that the Internet plays in lowering search costs. The fit of apparel products is a feature that is not easy for customers to search on over the Internet. To evaluate fit customers must actually try the products on, either by visiting a store or ordering the items. Lal and Sarvary (1999) study markets in which customers can search for some product features using the Internet, but other product features require a physical inspection.³ They show that the presence of non-digital product features can increase customer loyalty, as customers may not risk physically searching for products with better non-digital attributes, and instead, remain with the product that they are familiar with. This is essentially the same intuition that we use in this paper to explain why customers are more loyal to Retail World in categories in which they must learn about product sizes.

The results in this paper can also be compared with the literature on product standards, which has demonstrated that competition can result in the adoption or preservation of socially inefficient standards (see for example Farrell and Saloner 1985 and 1986; Katz and Shapiro 1985 and 1994; and Bessen and Farrell 1994). While the body of theoretical work is now very extensive, there has been recent recognition that there is need for more empirical work on this topic (Suarez 2005; Birke 2009). Much of the literature has focused on the role of network effects in technology markets (see for example Gupta, Jain and Sawhney 1999; Basu, Mazumdar and Raj 2003; Sun, Xie and Cao 2004; Dubé, Hitsch and Chintagunta 2010; and Wang and Xie 2011). Our findings can be interpreted as evidence that competition can lead to the adoption of socially inefficient standards even in non-technology markets and without the presence of strong network effects. There have been many large-scale studies

³ Related work has attributed higher price sensitivity on the Internet to the lower cost of search. Examples include evidence that airline demand is more sensitive to price changes in the Internet channel than in the traditional channel (Granados, Gupta and Kauffman 2012); and that apparel and home furnishing demand is more sensitive to charging sales tax on the Internet channel than in the catalog channel (Anderson, Fong, Simester and Tucker 2010).

documenting the variation in sizes across US apparel brands. For example, Kinley (2005) sent a team of research assistants into 20 retail stores in a Southwestern city to measure the sizes of 1,011 pairs of pants. They report large discrepancies across brands at every size level.⁴

The variation in apparel sizes survives despite numerous attempts to introduce standard sizes. As early as 1969 the International Standards Organization proposed a very flexible system of standard apparel sizes that promised to greatly reduce customer confusion. Other researchers have also proposed standard sizing systems (see for example Ashdown 1998; Caldwell 1996; Chun-Yoon and Jasper 1996; and Tamburrino 1992). However, none of these sizing systems have been widely adopted (Kinley 2005). Even though apparel retailers could in principle standardize the sizes of their products they choose not to do so. Industry commentators have attributed this reluctance to adopt a common standard to a desire among retailers to make it more difficult for customers to switch stores:⁵

“Retailers and clothing makers thrive on sizing confusion. Consumers who find a brand that fits are likely to stick with it, and a standard sizing system would encourage them to visit other stores.”

Barbaro (2006 at page 1)

A similar explanation is offered by Ellison and Ellison (2009) to explain why firms may want to obfuscate, in order to make it difficult for customers to compare prices. In their setting firms forgo standardization to obstruct price search by customers, resulting in less competition (see also Bergen, Dutta and Shugan 1996; and Gaudel and Sugden 2012). Our findings provide evidence that in apparel markets the absence of standardization yields a similar outcome by making it more difficult for customers to switch to competing retailers.

Structure of the Paper

The paper proceeds in Section 2 with a detailed description of the study design. We present aggregate results in Section 3 and then in Section 4 investigate the role of customer loyalty and product standardization. The paper concludes in Section 5 with a review of the findings.

2. Study Design

The study was conducted with the cooperation of a medium-sized retailer (“Retail World”) that sells apparel, home furnishings and sporting goods. With few exceptions, all of the products sold by this retailer are private label products carrying the firm’s own brand.⁶ Although many competitors sell

⁴ Delk and Cassill (1989), Workman (1991), Tamburrino (1992) and Workman and Lentz (2000) report similar findings.

⁵ The failure to adopt standard sizing systems is also sometimes attributed to “customer vanity”. Manufacturers of more expensive clothing reportedly tend to increase the physical dimensions to satisfy customers who want to believe that they fit a smaller size. There is evidence to support these claims; Kiley (2005) Sieben and Chen-Yu (1992) both report that more expensive products tend to have larger dimensions at a given size.

⁶ Private label apparel represents over 40% of all apparel sales (Cohen 2009) and up to 50% of sales in department stores (Trefis Group 2010). For many large retailers such as The Gap, J Crew, and Abercrombie and Fitch private label sales represent almost all of their sales. Moreover, direct mail advertising is one of the largest forms of advertising; in 2010 direct mail advertising spending increased 3.1% to \$45.2 billion (Levey 2011). This compares

similar products, Retail World's private label products are exclusively available from its own catalog, Internet and retail store channels. The products are moderately priced (\$48 on average) and past customers return to purchase relatively frequently (1.2 orders containing on average 2.4 items per year).

Competitive Advertising

Like many retailers in this market, Retail World regularly participates in competitive advertising programs that facilitate acquisition of prospective customers. Retailers who decide to join one of these programs agree to provide contact information (name and mailing address) and purchase histories of past customers. In return for providing data, a participating firm has access to customer information from other firms that join the program. A third-party acts as a clearinghouse and limits information flow between the firms. Participating firms do not generally have direct access to the names and purchase histories of other firms' customers. Instead, each firm makes a request for customers with a specific demographic profile and transaction history. Three examples of actual mailing requests from the competitive advertising event that we study are summarized in Table A1 in the appendix (the examples are disguised to protect confidentiality).

In the competitive advertising event that we study a total of 64 of the Retail World's competitors mailed at least one advertising piece to Retail World customers across a period of approximately 12-months (May 1, 2006 through May 14, 2007). Many of these competitors mailed multiple advertising pieces, and so a total of 217 different competitive advertising pieces were sent to the Retail World customers. Because the companies in some cases requested two different samples for the same advertising piece (e.g. 10,000 female buyers and 15,000 male buyers), the total number of requests exceeded the number of mailings. In particular, there were 301 requests spread across the 217 advertising pieces.

Design of Field Experiment

We use a field experiment to investigate how participation in this event affected purchases from Retail World. Prior to the start of the field experiment Retail World selected a sample of almost 370,000 customers and then randomly assigned them to Treatment and Control conditions. The randomization was done using the customers' account numbers, which ensured that the randomization was done at a customer level (rather than across geographic regions or some other aggregate level). The randomization resulted in a Treatment sample of 184,455 customers and a Control sample of 184,625 customers. All of these customers had made at least one purchase from Retail World in the 12-months before the start of the competitive advertising event. Customers in the Treatment condition were included in the cooperative mailing pool given to the third party. The customers in the Control condition were withheld from this pool of names.

To confirm that the allocation of customers to the two conditions was random we compared historical purchases from Retail World by the two samples. In particular, we compared the average units purchased and average revenue prior to the start of the competitive advertising event. If the assignment were truly random, we should not observe any systematic differences in historical sales

with \$56 billion spent on television advertising and just \$27.7 billion on digital (Internet) advertising. Despite these expenditure levels, direct mail has received much less attention in the literature than digital advertising.

between the two samples. Reassuringly, none of the differences are significant despite the large sample sizes (we report these historical comparisons in Table 1 in the next section).

Customers in the Treatment and Control conditions both continued to receive catalogs from Retail World. Retail World used the same mailing policy for all customers irrespective of the experimental condition they were in. It is also helpful to note that Retail World maintains a single price policy. Although the price of an item may vary over time, all customers purchasing that item at any point in time pay the same price. This policy ensures that there were no differences in the prices that customers paid across the two experimental conditions.

The opportunity to compare exogenously manipulated treatment and control groups is an important characteristic of the study that overcomes a problem that has plagued many previous studies of advertising. In general advertising decisions are endogenous, and are often at least partly determined by expectations about sales levels. As a result, simply measuring the level of association between advertising sales does not allow us to draw causal inferences about how advertising affects sales (Bagwell 2007). It is generally even more difficult to identify how sales are affected by competitors' advertising. In most cases it is not possible to distinguish which customers are targeted by competitors, or to account for the factors that may have contributed to the targeting decisions.

The Treatment

In total approximately 1.15 million competitive mailing pieces were sent to the 184,455 customers in our sample, or approximately 6.23 competitive mailing pieces per customer. We do not know which customers received which competitor's mailing. However, we do know the criteria that competitors used to request customers for each mailing. We summarize these criteria in Table A2, which is located in the appendix. The most common criteria were the recency of customers' last purchases and the amount spent on those purchases. Of the 301 requests, 87% included a recency criterion, while 80% included an "amount spent" criterion. Other common criteria included the type of catalog that customers had purchased from (50% of requests), the product category (43%), the customer's gender (37%) and the zip code they lived in (17%).⁷

We can also compare our eligibility estimates with measures describing customers' actual historical purchases from Retail World. These results are reported in Table A3 of the Appendix. They reveal a clear pattern: the competitors tended to request Retail World's most valuable customers. These include customers who had ordered more recently, more frequently and had placed larger orders with Retail World.

We caution that we cannot use the mailing requests to construct probabilities that an individual customer received a specific competitive mailing. Because many of the requests included a criteria that previously used customers were excluded (see the examples in Table A1), the actual mailing samples were not independent across the mailing efforts. In general, this "no previously used customers" criteria will reduce the heterogeneity in actual mailing frequencies across customers. Because the eligibility for the competitive mailings varied according to customers' purchasing histories, we also need

⁷ The frequency of customers' purchases was not an available selection criterion.

to be careful when evaluating how the results of the experiment varied according to customer characteristics. If we observe a larger effect among customers who were eligible to receive more mailings then we cannot easily distinguish whether it was the customer characteristics that amplified the effect or the increase in the number of competitive mailings that these customers received.⁸

3. Initial Results

Retail World did not receive data describing which of its customers made purchases from any of the competing firms. However, it does maintain complete records of customers' purchases of its own products. We compare purchases made by customers in the two experimental conditions. We compare purchases over a period of approximately 20-months (the "Measurement Period"). This period starts from the first mailing date in the competitive advertising event (May 1, 2006) through December 28, 2007 (almost eight months after the event finished on May 14, 2007). The average units purchased and average revenue in the two conditions are reported in Table 1.

Table 1 about here

Customers in the Treatment condition purchased 7.67 units, compared to 7.66 units in the Control condition. The difference in these averages (0.015) is not significant. There is also no significant difference in the average revenue per customer. To help interpret this null result it is helpful to ask how big the difference in these averages would need to be in order for the difference to be statistically significant. The standard error of the difference in the two sample means are \$1.5865 (revenue) and 0.0475 (units). Therefore, given the sample means for the control condition, the difference in the means would need to be equal \$3.11 (revenue) and 0.09 (units) to be significantly different.

From a practical perspective this result implies that on average there is no cost to the firm in allowing its competitors to advertise to its customers, even when these competitors are able to select the most valuable customers. We were surprised by this result. Given the number of competitors, maturity of the market, and relatively low levels of differentiation between products, we anticipated that allowing competitors to advertise to Retail World's customers would result in substitution of demand from Retail World to the competitor. The findings in Table 1 do not support this prediction. If anything, the differences in average sales are in the opposite direction.

One possible interpretation of this null overall result is that customers have self-selected to their favorite firms, and their demand is not sensitive to the actions of the competing firms. However, we will demonstrate that we should not conclude from this null overall result that the competitors' advertising did not have any effect on Retail World's sales. Instead we will show that the null result is caused by aggregating positive effects in some categories with negative effects in others. The remainder of the

⁸ We also note that because eligibility for the competitive mailings was primarily determined by historical purchasing history, eligibility changed as customers purchased. For example, consider a mailing that was only sent to customers who had purchased within the last 3-months. At the start of the event a customer may not have been eligible for this mailing. However, if this customer purchased from Retail World then the customers would have become eligible immediately after that purchase.

paper explores these interactions by comparing how the outcome varied across different customers and product categories. We begin by comparing the outcome for different segments of customers.

Time Since Last Purchase

One factor that may contribute to how customers respond to advertising is the extent to which their demand is already depleted. We might expect that customers who have just purchased may have satiated their demand, making it less likely they will react to an advertisement. To investigate this possibility we use the timing of customers' purchases prior to the start of the competitive mailing to measure whether the effect of competitive advertising is influenced by the time since a customer's prior purchase.

Recall that all of the customers in the study had made at least one purchase from Retail World in the twelve months before the start of the field experiment. We grouped customers into three mutually exclusive segments by dividing these twelve months into three time periods. The first segment includes customers who had made a purchase from Retail World within four-months of the start of the event. Customers in the second segment had made their most recent purchase between four and eight months prior to the start of the event. Customers in the third segment had made their most recent purchase between eight and twelve months of the start of the event. For each segment we then calculated the difference between the Treatment and Control conditions in average sales across the twenty-month Measurement Period. The findings are summarized in Table 2 (where we also report the historical sales comparisons for these three segments).

Table 2 about here

The competitive advertising did not have a significant effect on sales at Retail World among customers who had recently purchased from that firm. However, among customers whose prior purchases were the least recent (Segment 3), competitive advertising led to a significant *increase* in sales at Retail World. This is the opposite of what we would expect if the competitive advertising led solely to substitution to the competitors.

It is helpful to remember factors that cannot explain the difference in sales for Segment 3. We can rule out the possibility that the interaction between the response to the competitive mailing and the timing of customers' prior purchases reflects heterogeneity in the treatment itself. As we reported in Table 1, the competitors tended to request customers who had purchased more recently. In particular, customers in Segment 3 were eligible to receive an average of just 8.2 competitive mailings. This is much lower than the averages of 50.0 and 28.3 eligible requests in Segments 1 and 2 (respectively). It seems implausible that the positive effects in Segment 3 (higher orders, items, units and revenue) can be attributed to a smaller experimental treatment.

The difference in units sold in Segment 3 also cannot be explained by differences in the customers themselves. Our analysis compares customers in the two experimental conditions and random assignment ensures that there are no systematic differences between customers in the two conditions. Note that the findings in Table 2 confirm that for all three segments the historical sales to customers in the Treatment and Control conditions were not significantly different. Given the very large sample sizes this confirms that the random assignment of customers to the Treatment and Control conditions was performed correctly.

In Segment 3 we do observe slightly (although not significantly) higher sales in the Treatment condition than in the Control condition. To confirm that customer differences do not explain the differences during the Measurement Period we can use a multivariate approach to control for any historical differences between the customers. This also provides an opportunity to investigate other factors that may have contributed to the results. Our multivariate analysis is conducted in two parts. In the first part we measure purchases by each customer across all categories, while in the second part we measure purchases by each customer in each category.

Customer-Level Analysis

The dependent variable in this first part of the analysis measures the number of units that customer i purchased during the Measurement Period. Because $Units\ Purchased_i$ is a “count” variable, we estimate a negative binomial model. To estimate the impact of the timing of a customer’s prior purchase in the category we calculate $Recency_i$, which is defined as the number of days (in hundreds) since a customer’s last purchase (prior to the treatment). We then interact $Recency_i$ with a binary variable identifying customers in the treatment condition ($Treatment_i$). We also estimate a separate specification in which we group customers into Segments 1, 2 and 3 (as defined earlier) and estimate the impact of the treatment separately for these three segments.

As additional controls we considered what other factors might be correlated with the timing of customer’s prior purchases. There are two obvious candidates. First, we would expect the time since a customer’s last purchase in the category to be affected by how frequently the customer purchases. We measure this by simply counting the number of prior units purchased ($Prior\ Units_i$). Notice that this also explicitly controls for any differences in the historical units purchased between the Treatment and Control conditions. Second, we might expect that the timing of customers’ last purchase may be affected by the seasonality in their purchasing patterns. Recall that the treatment started on May 1 and so customers who only purchase at Christmas will tend to have a different $Recency_i$ measure than customers who purchase at other times of this year.⁹ To control for the impact of this seasonality we calculated the percentages of a customer’s (pre-treatment) purchases that are made in each calendar month ($Month\ 1, \dots, Month\ 12$).

To ensure that the relationship between the experimental treatment and the timing of customers’ prior category purchases ($Recency_i$) is not due to either the frequency with which customers purchase or the seasonality of their purchases, we include interactions between the treatment and these controls in our model specification. In particular, we estimate the following model:

$$\lambda_i = \beta_1 Treatment_i + \beta_2 Treatment_i * Recency_i + \beta_3 Treatment_i * Prior\ Category\ Units_i + \sum_{m=1}^{11} \varphi_m Treatment_i * Month\ m_i + \mathbf{BX}$$

⁹ We thank an anonymous reviewer for this suggestion.

The vector \mathbf{X} contains a complete set of control variables, including main effects for *Recency_i*, *Prior Units_i*, and the *Month m_i* variables. We also include main effects for the demographic variables.¹⁰ The coefficient of interest is β_2 . This measures how the effect of the treatment varies according to the timing of the customer's most recent prior purchase. The findings are reported in Table 3 (to demonstrate robustness we also include a base model without covariates). For ease of exposition we do not report the coefficients for our monthly control variables (seasonality) or their interaction with the treatment.

Table 3 about here

The results confirm that competitive advertising led to a large increase in purchases from Retail World among customers whose last purchases in the category were less recent. A 100-day increase in the time since a customer's last purchase is associated with a 1.61% increase in the impact of the competitive advertising. Moreover, we now see that the significant impact of the competitive advertising extends beyond Segment 3. For customers in Segment 1 (who had last purchased within 4-months of the start of the test) the competitive advertising is associated with a 2.53% increase in unit sales, which is not statistically significant. For Segments 2 and 3 the estimated increases in unit sales are 3.75% and 7.12% (respectively). The effects for these last two segments are both statistically significant ($p < 0.05$).

Customer x Category-Level Analysis

Notice that the timing of a customer's last purchase will generally vary across categories. While a customer may have purchased a pair of pants recently, they may not have purchased footwear. While customers may have depleted their demand for one category (pants), they may not have depleted demand for other product categories (footwear). To investigate these possibilities we re-analyzed sales at the category level.

A category is defined using two dimensions: product type (e.g. footwear, tops, swimwear, pants) and customer gender (for items where gender is relevant). We focus on the 35 largest categories, which represent 81% of all purchases (the remaining purchases are spread across approximately 200 other categories). We stack the data so that for each customer we obtain 35 observations, one for each category. The dependent variable measures the number of units that customer i purchased from category c during the test period. To account for category differences we estimate a quasi-maximum likelihood (QML) Poisson model with (conditional) category fixed effects (Wooldridge 1999). This estimator is consistent under very general conditions. For example, in contrast to the regular Poisson model the estimator is consistent even if there is over-dispersion or under-dispersion in the latent variable model. Moreover, the robust variance-covariance matrix described by Wooldridge (1999) allows for deviations from the Poisson distribution together with arbitrary category-level fixed effects.

¹⁰ We are missing demographic observations for a small number of customers. For these observations we set the demographic variables at zero and identify the missing data using binary flags (that we also include as controls).

To estimate the impact of the timing of a customer’s prior purchase in a category we calculate *Category Recency*_{*i,c*}, which is defined as the number of days (in hundreds) since a customer’s last purchase in the category (prior to the treatment). We then estimate the QML Poisson model including *Category Recency* both as a main effect and interacted with the binary *Treatment* flag. As additional controls we measure the number of prior units purchased at the category level (*Prior Category Units*_{*i,c*}). We also again control for the impact of seasonality and estimate the following system of independent variables:

$$\lambda_{i,c} = \beta_1 \text{Treatment}_i + \beta_2 \text{Treatment}_i * \text{Category Recency}_{i,c} + \beta_3 \text{Treatment}_i * \text{Prior Category Units}_{i,c} + \sum_{m=2}^{12} \varphi_m \text{Treatment}_i * \text{Month } m_i + \mathbf{BX}$$

The vector **X** again contains a complete set of control variables, including main effects for the demographic variables. Recall also that the model explicitly controls for arbitrary category-level differences using (conditional) category fixed effects. The coefficient of interest is β_2 . This measures how the effect of the treatment varies according to the timing of the customer’s most recent prior purchase in that category. In Table 4 we report findings when using *Category Recency* (Model 1) and when using a log transformation of *Category Recency* (Model 2). We also report a model in which we include both *Category Recency* and overall *Recency* both as main effects and interacted with the *Treatment* variable (Model 3). The standard errors are clustered by category to account for any correlation in the errors across customers within the category.

Recall that in our prior analysis every customer had made a prior purchase from the firm, and so *Recency_i* is well-defined when we calculate it at the aggregate level. However, not all customers have a prior purchase in all 35 categories. Because the behavior of these customers does not inform us about how the timing of prior category purchases influenced the outcome, when estimating the impact of *Category Recency* we restrict attention to observations in which customers have made a purchase in the category. In later analysis (Section 4) we use a different approach that relaxes this restriction.

Table 4 about here

The findings in Table 4 replicate the earlier results. The competitive advertising led to a larger increase in purchases from Retail World among customers whose last purchases in the category were less recent. In Model 1 we see that a 100 day increase in the time since a customer’s last purchase in the category is associated with a 0.33% increase in the response (in that category) to the competitors’ advertising.

Interestingly, in Model 3 where we include *Category Recency* and overall *Recency*, both interactions are significant. This indicates that demand depletion is not completely specific to a category. It appears that among customer who have purchased recently, it is more difficult to prompt an additional purchase, even in categories they have not recently purchased. For example, a recent purchase of pants appears to make it more difficult to prompt an additional purchase of footwear (for at least some customers). This could be explained by these customers having a budget constraint for apparel generally, rather than separate constraints for footwear and pants.

Our calculation of the time since a customer's last category purchase does not account for variation in the inter-purchase intervals between categories. However, there is considerable variation across categories in the typical interval between purchases, which may also influence whether a customer's demand for products in a category is depleted. For example, while pants have an average inter-purchase interval of approximately eighteen months, the average inter-purchase interval for underwear is closer to three years. Therefore, if the last purchase occurred two years ago, this may indicate that a purchase of pants is overdue, but customers are not yet ready to purchase underwear. To investigate this possibility we calculated the average inter-purchase interval for each category. We then re-analyzed the time since a customer's last purchase by re-scaling the time since a customer's last category purchase according to the number of inter-purchase intervals. The results are also reported in Table 4 (Models 4 and 5).

We report two specifications. First, we treat the number of (category) inter-purchase intervals since a customer's last purchase in the category as a continuous measure (*Intervals Since Last Purchase_{i,c}*) and interact this continuous measure with the treatment (Model 4). For every additional inter-purchase interval since a customer's last category purchase the impact of the competitive advertising increased by an average of 2.01%. Second, in the same way that we grouped customers into segments according to the timing of their last purchases, we can also group them according to the category inter-purchase intervals since their last purchase. We compare customers whose purchases were less recent with benchmark customers whose previous purchases were most recent.¹¹ In particular, our benchmark includes customers who made a prior purchase within 1 inter-purchase interval of the start of the treatment. If a customer's last category purchase was between 1 and 2 inter-purchase intervals prior to the start of the treatment, then the impact of the competitive advertising was 2.86% higher than the benchmark customers. If their last category purchase was over 2 inter-purchase intervals ago, the effect increased by 5.85% compared to the benchmark customers.

Summary

An initial comparison of purchases during the Measurement Period indicates that there is no significant difference in the outcomes in the Treatment and Control condition. However, further investigation reveals that we do observe a significant difference in the outcomes for customers who had not purchased recently. We illustrate this interaction using a variety of approaches, including both the time since a customer made any purchase from the firm, and the timing of purchases in specific categories. Our analysis includes controls for a variety of customer differences, together with arbitrary category-level effects.

We interpret the interaction between the timing of a customer's last purchase and the response to the competitive mailing as evidence of demand depletion. Satiation among customers who had recently purchased appears to have muted their reaction to the competitors' advertising, limiting any increase in primary demand. However, we do not have data describing purchases from competing firms and so we cannot directly measure the change in primary demand.

¹¹ Notice that the treatment effect for these benchmark customers can be calculated from the *Treatment* coefficient (-0.99%), but must also be conditioned on the number of prior category units that the customer has purchased.

Instead, we focus on understanding why customers responded to the competitors' advertising by purchasing from Retail World rather than the competitors. We turn to this issue in the next section where we explore how the impact of the competitive advertising varied across different product categories.

4. Switching Costs and Customer Learning

As we discussed in the Introduction, we interviewed a sample of Retail World's customers as background research for this study. An issue that arose frequently in these interviews was the role of product sizes. Some customers described examples of products that they only buy from Retail World because they have learned which sizes fit them. They expressed reluctance to buy from competitors because they are less certain about fit. A key feature of this argument is the specificity of product fit. For example, while footwear sizes are standardized within the Retail World (private label) brand, sizes typically vary across brands. If the sizes of footwear were standardized across brands customers could use Retail World's size information to also order products from other firms. Similarly, if there is no standardization in sizes, either because sizes are unimportant (e.g. backpacks), or because fit varies across products even within a firm (e.g. swimwear), then standardization will again not contribute to customer loyalty.

To evaluate the extent to which customers learn standardized sizes over time we constructed measures of the change in product return rates over time. In particular, comparing the return rate on a customer's first purchase in a category and that customer's most recent purchase may provide a measure of learning about product fit. The more learning that occurs the larger the expected reduction in return rates.

To obtain reliable estimates of the return rates we used the complete purchase histories for a sample of 3,634,695 Retail World customers.¹² Their purchase histories include a total of over 93 million purchases representing over \$3.6 billion in revenue. The historical data for these customers describes when a customer returns an item and the reason for the return. In particular, the sample includes a total of over 15.1 million returns. For approximately 78% of these returns we have a variable describing the reason that the item was returned.¹³ In Table 5 we list each of the return reasons and describe the frequency with which they appear. The most common reason that items are returned is that they are the wrong size, which contributes approximately half (46.6%) of all reasons for returns. The next most frequent reason is that the customer did not like the item when it arrived, generally because they did not like the color, material or styling. Defects contribute only a small portion of the returns.

¹² We obtain similar results if we calculate return rates using historical transactions by the 369,080 customers in the study. However, the additional sample size yields more precise estimates of the return rates.

¹³ The firm has used the same set of explanations for why items are returned throughout the transaction history. Approximately half of the returns with missing reasons occur before February 23, 1996, which was the first date that the firm started recording reasons for returns. After this date they occur with slightly diminishing frequency over time. To ensure that the *Return Trend* is not affected by missing return reasons, when calculating the *Return Trend* for each reason we omit all observations for a customer in a category if a return on the first or last category order is not accompanied by a reason.

Table 5 about here

There are some items for which we would expect more learning about product sizes than other items. These categories provide an opportunity to investigate our interpretation that the reduction in return rates provides a way to measure how much customers learn about product sizes. First, we might expect that customers will learn less about the size of “Childrens” items versus other items. Because childrens’ sizes change, experiences with the fit of past purchases may not provide as much information as the fit of future purchases in childrens’ product categories than in other product categories. Second, there are many items that do not have sizes. This includes non-apparel items (e.g. backpacks) and some apparel items for which “one size fits all” (e.g. scarves).

In Figures 1 and 2 we report the average difference between a customer’s first purchase in a category and their last purchase in a category according to whether: (a) the item is in the Childrens’ category or not and (b) whether the items in the category have sizes or not. The unit of observation is a customer in a category, and we use all of the observations for which the customer made at least two purchases in the category (so that we can calculate the change in the return rate). We report the difference in return rates for each of the four return reasons.

Figures 1 and 2 about here

There are several findings of interest. First, for categories without sizes there is essentially no change (0.06%) in the return rates due to size between a customer’s first and last category order. Further investigation confirms that the first and last return rates due to size are essentially zero in these categories. This is precisely what we would expect and offers reassurance that the return rate for sizes does not include returns for other reasons. Second, for those categories with sizes we see a large reduction (-1.45%) in the return rate due to sizes between customers first and last orders. This is consistent with our interpretation that customers learn about sizes as they place more orders in these categories. Third, the reduction in the return rate due to sizes is considerably smaller for Childrens’ categories than for other categories (-0.88% versus -1.50%). This is consistent with customers learning less about sizes from past purchases in Childrens’ categories than in other categories.

The change in return rates for the other three return reasons are also interesting. For these other return reasons the differences between categories with and without product sizes, and between Childrens’ and other categories are much smaller than we observe for returns due to size. However, the sample sizes are very large and so even these smaller differences are statistically significant. It appears that when a customer returns an item because of poor fit this is generally categorized accurately as a size-related return. However, the larger reduction in returns on categories with sizes for these other three return reasons could indicate that some of the size returns spillover to these other return reasons. These attributions are not necessarily errors. Notice that a customer who finds that an item does not fit could quite reasonably indicate that they do not like the item, or interpret the poor fit as a defect. Any measurement errors that result from these spillovers will hinder (not help) validation of our prediction that the response to the competitive advertising event is associated with the change in the size return rate.

Return rates due to defects and because customers do not like the items exhibit a slight increase in categories without sizes. One possible explanation is that Retail World has a generous returns policy

and the increase in return rates in these two categories may reflect customers learning about this policy. If this explanation is correct then, in the absence of learning about sizes, we might expect return rates in every product category to be higher for the last purchase than for a customer's first purchase.

The Change in Return Rates and the Response to the Competitors' Advertising

Our analysis again focuses on the 35 largest product categories. However, in this analysis we add an additional refinement to our definition of a category. Preliminary investigation revealed that for some apparel items Retail World uses different "size types" within a category. For example, women's pants come in 3 different size types:

- {S, M, L XL, XXL ...}
- {1, 2, 3, 4, 6, 8, ...}
- 23", 24", 25", 26", 28", 30", 32", 34"

We would expect that learning would only occur within one of these narrow product categories. For example, learning about {S, M, L XL, XXL ...} sizes may not help customers learn about {1, 2, 4, 6, 8, ...} sizes. For this reason we use our separate sample of 3.6 million customers to identify the most common size type within each category and in our *Size Learning* analysis restrict attention to purchases and returns within this size type.¹⁴

To construct a category-level measure of the amount of learning about sizes in each category we calculated the difference in the return rate between a customer's first and last purchase in each category and then averaged across customers. Because there is heterogeneity in the number of orders that customers place in different categories, we used two different approaches. In the first approach we averaged all customers who placed at least two orders in the category, while in the second approach we restricted attention to customers who had made exactly three orders in the category. The second approach was suggested by a reviewer and controls for heterogeneity in the frequency of orders across categories. As we would expect the two measures are highly correlated; the pair-wise correlation between them across the 35 categories is 0.97.

In Table 6 we report the pair-wise correlation between these two measures and the response to the competitors' advertising event measured by the number of units purchased during the twenty month Measurement Period. We report the pair-wise correlations separately for all four return reasons. The key finding is that the faster the reduction in the return rate due to sizes between a customer's first and last orders, the larger the response to the competitive advertising. This is consistent with the explanation that customers returned to purchase from Retail World because they had learned about product sizes from previous purchases. The rank-order correlation between the change in returns due to size and the response to the competitive advertising is even larger (0.50), indicating that the relationship cannot be attributed to mere outliers.

Table 6 and Figure 3 about here

¹⁴ The demand depletion interpretation that motivated our analysis of the recency of customer's prior category purchases is not sensitive to these size types. For example, purchasing a size L pair of pants may deplete demand for size 10 pants as well as size L pants.

In Figure 3 we provide a scatter plot of the response to the competitors' advertising (in each category) and the change in return rates due to size. The relationship is clearly observable in this figure. The figure highlights why we observe a null overall effect; the impact of the competitive advertising is positive in some categories and negative in others. The figure also illustrates the magnitude of the interaction effect. In particular, if we group the 35 categories into quintiles (of 7 categories each) according to the change in returns due to size, the mean response to the competitors' advertising in each quintile is monotonically decreasing. The first quintile (the leftmost portion of Figure 3) has the largest decrease in returns due to size and a mean response of 0.87%. In contrast, this effect declines to -1.8% and -3.1% in quintiles 4 and 5 (the rightmost portions) which have the smallest decrease in returns due to size. Importantly, this segmentation of the data clearly highlights the competing positive and negative effects of competitive advertising. We do observe negative advertising effects, but these are limited to categories with low consumer loyalty (the reduction in returns is smallest). Positive spillovers exist, but are limited to categories with high consumer loyalty (the reduction in returns is largest).

While learning about sizes appears to play an important role in explaining the variation in the response to the competitive advertising, learning about the other return reasons is apparently less important. We do see a negative correlation between the other three return reasons and the response to the competitive advertising. However, none of these correlations are statistically significant.

Number of Prior Purchases

Our interpretation of the change in returns measure suggests that customers who have made more prior purchases in the category should have learned more about products sizes. However, we have not directly tested how the response to the competitive mailings varies according to the number of prior purchases a customer has made in a category. Doing so requires that we overcome an important confound. A customer who has made more prior purchases from Retail World in a category may be more loyal to Retail World for reasons other than their knowledge of sizes (to aid exposition we will label this "generic loyalty"). An increase in purchases from Retail World as a result of receiving competitors' advertising may just reflect this generic loyalty.

To demonstrate that the relationship between prior purchases and the response to the competitive advertising (if any) is due to learning about sizes rather than just generic loyalty, we will again rely on a comparison across categories. In particular, the generic loyalty explanation operates for every category, while we only expect prior purchases to contribute to learning about sizes in some categories. If we can show that the relationship between prior purchases and the response to the competitor's advertising is stronger in categories in which there is more learning about product sizes, we can claim that the relationship is attributable to customer learning rather than other sources of loyalty.

We begin by constructing a measure of the rate at which customers learn about product sizes in each category. In particular, for each category we estimate how the change in returns (between the first and last order) varies according to the total number of orders a customer has made in the category. The slope of this relationship provides a measure of how quickly customers learn from previous orders. In categories in which customers learn more about the firm's products we would expect this slope to be more negative, so that greater learning leads to a faster reduction in returns. For ease of exposition we

label this slope the “*Return Trend*”.¹⁵ A large positive *Return Trend* indicates that customers quickly increase their rate of returns (across orders), while a large negative trend indicates that the rate of returns quickly decreases.

We then construct a measure of how much we expect each customer to have learned about sizes in each category by multiplying this measure of the rate of learning with the number of orders that customers have placed in each category:

$$Size\ Learning_{i,c} = Return\ Trend_c * Prior\ Category\ Orders_{i,c}$$

We add this variable to the QML Poisson customer x category-level model used to estimate the recency interactions in the previous section. In particular, we include both *Size Learning* as a main effect, and an interaction between *Size Learning* and the *Treatment* variable. It is the coefficient for this *Treatment* interaction that is the coefficient of interest. We caution that the interpretation requires care because this is a 3-way interaction between the *Return Trend*, the number of *Prior Category Orders*, and the *Treatment* variable. The coefficient measures how an increase in the number of prior customer orders affects the response to competitors’ advertising in categories in which there is a lot of learning about product sizes compared to categories without a lot of learning about product sizes.

Recall from our earlier discussion that *Category Recency* is not well-defined if a customer does not have a prior purchase in a category. It was for this reason that when we investigated the *Treatment * Category Recency* interaction in the previous section we omitted these observations. In order to retain *Category Recency* as a control in our *Size Learning* model we use two different approaches. First, as in Section 3 we estimate the model when omitting observations for customers with no prior purchases in a category (Model 1). In the second approach we include these observations and constructed a binary flag identifying customers with no prior purchases in a category.¹⁶ We include this binary flag as both a main effect and interacted with the *Treatment* variable (Model 2). This second approach recognizes that while these observations do not contribute to our understanding of the *Category Recency* interaction, they may provide insight in this model. In particular, customers with no prior category purchases cannot have learned any size information from past purchases. Our interpretation that learning about sizes contributes to customer loyalty suggests that the competitors’ advertising should have a smaller effect on purchases from Retail World for these observations.

The findings are reported in Table 7 (we again cluster the standard errors by category). In both models the coefficient for the interaction between *Treatment* and *Size Learning* is negative and highly significant. This indicates that the response to the competitive advertising was higher among customers who had more prior experience in categories where the *Return Trend* is more negative. It is consistent

¹⁵ In each category we use OLS to estimate the *Return Trend* using the separate sample of 3.6 million customers. We use all customers who placed at least two orders in the category. The average sample size for these estimates is approximately 143,000, with a minimum of 10,000 observations in any category.

¹⁶ Notice that in Model 1 the sample size is smaller than the sample size for the models reported in Table 4. This reflects the narrower definitions of a category to focus on a single size type.

with the interpretation that prior purchases lead to more loyalty in categories in which customers learn more quickly about product sizes.

Table 7 about here

We also observe a negative interaction between the *Treatment* variable and *No Prior Category Purchase*. This is consistent with our prediction that these customers would have a less favorable response to the competitive advertising. Notice that for these customers both *Category Recency* and *Prior Category Units* equal zero, and so the estimated net impact of the *Treatment* on category sales for these customers is -6.62% (adjusting the 3.28% main *Treatment* coefficient with just the -9.90% *No Prior Category Purchase* interaction coefficient). This result highlights the mix of positive and negative treatment effects. We observe a positive main effect from advertising but a negative effect among customers who have no history in a category.

5. Conclusions

Standard models of competition predict that firms will sell less when competitors target their customers with advertising. We have presented findings that contrast with this prediction and demonstrated that advertising can lead to positive spillovers for other firms even in a competitive market. We investigate how these positive spillovers vary across product categories. The findings reveal that the product categories with the largest spillovers are the categories in which customers appear to learn about product attributes from prior purchases. Detailed data on the reasons that customers return items allowed us to investigate what customers are learning. We show that the rate that customers learn about size information is closely related to the response to the competitive advertising.

We also investigated how the response to the competitive advertising varied according to how many times customers had previously purchased in a category. We find that customers who have made more prior purchases were more likely to respond to the competitors' advertising by purchasing at Retail World. However, this only holds in product categories in which there is evidence that customers learn about product sizes from prior purchases. This distinction is important; it suggests that the relationship between past purchasing and the positive spillovers can be attributed to learning about product sizes, rather than other generic sources of customer loyalty.

Our findings provide unique measures of the impact of advertising in a competitive market, and we believe that the results have the potential to change our understanding of competition in mature markets. First, we show that advertising by a competitor may increase another firm's sales. This result can be compared with previous studies of competitive advertising, which have generally focused on negative spillovers. For example extensive research in the tobacco industry has studied whether tobacco advertising leads to brand switching or an increase in primary demand: does Marlboro advertising lift Marlboro sales by increasing smoking, or by switching smokers from Rothmans to Marlboro? Few empirical studies anticipate the possibility of positive spillovers from competitors' advertising, such as a positive impact of Marlboro advertising on sales of Rothmans.

Second, the findings highlight the role of product standards. Our evidence that customers can use past experience to learn about product sizes does not extend to all categories. Instead it is limited to categories in which product sizes are standardized within a brand but not across brands. This may help to explain why firms do not adopt common standards in industries where doing so would offer obvious customer benefits. The customer benefits are particularly true in the apparel market. It is estimated that 70% to 80% of apparel products may not meet customers' size expectations, and this contributes to 50% of women and 62% of men not being able to find products that fit (Abend 1993 and DesMarteau 2000). The outcome is not just lost sales but a large volume of returns. At the apparel retailer that provided data for this study an average of 10.6% of purchases are returned. This increases to over 20% for swimwear, imposing a monetary cost in managing logistics and depreciation of the returned items, and possibly also an intangible cost due to customer dissatisfaction.

Our findings may also extend our understanding of the role of switching costs. Learning about product fit is relatively easy, returning requires only nominal effort, and this effort is only incurred if the product does not fit. Yet in a mature market with many close competitors we provide evidence that these switching costs are sufficient to greatly reduce the threat from competition.

Finally, the findings also offer guidance to retailers. When selecting which customers to make available in rented mailing lists or cooperative advertising events retailers should focus on (1) customers who have not purchased recently, and (2) customers who have made repeat purchases in categories where learning about sizes is important (such as footwear). We caution that this does not imply that these are also the best types of customers to request when mailing to a competitors' customers. This answer obviously depends upon sales at the competing firms (which we do not observe). Because the competing firms do observe these outcomes, the evidence that they tend to choose customers who have purchased from Retail World most recently and most frequently, suggests that it is the more recent and more frequent purchasers that a competitor should request.

Our results also call for future research on the role of advertising with consumer switching costs. We argue that positive spillovers may arise when competitors' advertising reminds consumers about a primary need for a category that is later fulfilled at a preferred retailer. More research is required to understand this mechanism and the implications for advertising competition.

References

- AAFA (2008), *Trends: An Annual Statistical Analysis of the U.S. Apparel and Footwear Industries*, American Apparel and Footwear Association, Arlington VA.
- Abend, Jules (1993), "Our Fits Over Fit," *Bobbin*, 34, 78-79.
- Ackerberg, Daniel A. (2003), "Advertising, Learning, and Consumer Choice in Experience Good Markets: A Structural Empirical Examination," *International Economic Review*, 44, 1007–1040.
- Anderson, Eric T., Nathan M. Fong, Duncan I. Simester, and Catherine E. Tucker (2010), "How Does An Obligation To Collect Sales Tax Affect Consumer And Firm Behavior?" *Journal of Marketing Research*, 47(2), 229-239.

Bagwell, Kyle (2007), "The Economic Analysis of Advertising," *Handbook of Industrial Organization*, Vol. 3, edited by M. Armstrong and R. Porter, Elsevier B.V., 1703-1843.

Balachander, Subramanian and Sanjoy Ghose (2003), "Reciprocal Spillover Effects: A Strategic Benefit of Brand Extensions," *Journal of Marketing*, 67(1), 4-13.

Barbaro, Michael (2006), "Standard Sizing Gives Clothing retailers Fits," *The New York Times*, March 31, <http://www.nytimes.com/2006/03/31/business/worldbusiness/31iht-size.html>.

Bergemann, Dirk and Juuso Välimäki (1996), "Learning with Strategic Pricing," *Econometrica*, 64, 1125-1150.

Bergen, Mark, Shantanu Dutta and Steven M. Shugan (1996), Brand variants: A Retail Perspective," *Journal of Marketing Research*, Vol. XXXIII, 9-19.

Bessen, Stanley M. and Joseph Farrell (1994), "Choosing How to Compete: Strategy and Tactics in Standardization," *Journal of Economic Perspectives*, 8, 117-131.

Birke, Daniel (2009), "The Economics of Networks: A Survey of the Empirical Literature," *Journal of Economic Surveys*, 23(4), 762-793.

Borenstein, Severin (1991), "Selling Costs and Switching Costs: Explaining Retail Gasoline Margins," *RAND Journal of Economics*, Vol. 22, 354-369.

Bouckaert, Jan and Hans Degryse (2004), "Softening Competition by Inducing Switching in Credit Markets," *The Journal of Industrial Economics*, Vol. 52, 27-52.

Braithwaite, Dorothea (1928), "The Economic Effects of Advertisement," *Economic Journal*, 38, 16-37.

Cabral, Luis M.B. and Miguel Villas-Boas (2005), "Bertrand Supertraps", *Management Science*, (April), vol. 51 no. 4 599-613.

Caminal, Ramon, and Carmen Matutes. "Endogenous switching costs in a duopoly model." *International Journal of Industrial Organization* 8.3 (1990): 353-373.

Chouard, Hayley H. and Jeffrey M. Perloff (2007), "Gasoline Price Differences: Taxes, Pollution Regulations, Mergers, Market Power, and Market Conditions," *The B. E. Journal of Economic Analysis & Policy*, Vol. 7, Article 8.

Chun-Yoon, Jongsuk, and Cynthia R. Jasper. "Key Dimensions of Women's Ready-to-Wear Apparel: Developing a Consumer Size-Labeling System." *Clothing and Textiles Research Journal* 14.1 (1996): 89-95.

Cohen, Marshal (2009), "The Challenges of Private Label Apparel," <http://www.licensemag.com/licensemag/Chartbuster/The-Challenges-of-Private-Label-Apparel/ArticleStandard/Article/detail/618844>.

Crawford Gregory S. and Matthew Shum (2005), "Uncertainty and Learning in Pharmaceutical Demand," *Econometrica*, 73(4), 1137-1173.

Darke, Peter R., and Robin JB Ritchie. "The defensive consumer: Advertising deception, defensive processing, and distrust." *Journal of Marketing Research* (2007): 114-127.

Delk, Ann E., and N. L. Cassill. "Jeans sizing: problems and recommendations." *Apparel Manufacturer* 1.2 (1989): 18-23.

Deltas, George (2008), "Retail Gasoline Price Dynamics and local Market Power," *The Journal of Industrial Economics*, Vol. 56, 613-628.

DesMarteau, Kathleen (2000), "Let the Fit Revolution begin," *Bobbin*, 42(2), 42-56.

DMA (2006), *Statistical Fact Book*, Direct Marketing Association, New York NY.

Ellison, Glenn and Sara Fisher Ellison (2009), "Search, Obfuscation, and Price Elasticities on the Internet," *Econometrica*, 77(2), 427-452.

Erdem, Tulin and Michael Keane (1996), "Decision-Making Under Uncertainty: Capturing Dynamic Brand Choice Processes in Turbulent Consumer Goods Markets," *Marketing Science*, 15, 1–20.

Farrell, Joseph and Garth Saloner (1985), "Standardization, Compatibility and Innovation," *Rand Journal of Economics*, 16, 70-83.

Farrell, Joseph and Garth Saloner (1986), "Installed Base and Compatibility – Innovation, Product Preannouncements, and Predation," *American Economic Review*, 76(5), 940-955.

Feldman, Jack M. and John G. Lynch Jr. (1988), "Self-Generated Validity and Other Effects of Measurement on Belief, Attitude, Intention, and Behavior," *Journal of Applied Psychology*, 73(3), 421-435.

Gaudel, Alexia and Robert Sugden (2012), "Spurious Complexity and Common Standards in Markets for Consumer Goods," *Economica*, Vol 79, Issue 314, 209-225.

Granados, Nelson, Alok Gupta and Robert J. Kauffman (2012), "Online and Offline Demand and Price Elasticities: Evidence from the Travel Industry," *Information Systems Research*, Vol. 23(1), 164-181.

Janakiraman, Ramkumar, Catarina Sismeiro, and Shantanu Dutta (2009), "Perception Spillovers Across Competing Brands: A Disaggregate Model of How and When," *Journal of Marketing Research*, XLVI (August), 467-481.

Katz, Michael L. and Carl Shapiro (1985), "Network Externalities, Competition and Compatibility." *American Economic Review*, 424-440.

Katz, Michael L. and Carl Shapiro (1985), "Systems Competition and Network Effects." *Journal of Economic Perspectives*, 8(2), 93-115.

- Kim, Moshe, Doron Kliger and Bent Vale (2003), "Estimating Switching Costs: the Case of Banking," *Journal of Financial Intermediation*, 12, 25-56.
- Kinley, Tammy R. (2005), "Size Variation in Women's Pants," *Clothing and textiles Research Journal*, Vol. 21(1), 19-31.
- Klemperer, Paul (1987), "Markets with Consumer Switching Costs," *Quarterly Journal of Economics*, Vol. 102, 375-394.
- Klemperer, Paul (1995), "Competition When Consumers Have Switching Costs: An Overview with Applications to Industrial Organization, Macroeconomics, and International Trade," *Review of Economic Studies*, 62, 515-539.
- Knittel, Christopher R. (1997), "Interstate Long Distance Rates: Search Costs, Switching Costs, and Market Power," *Review of Industrial Organization*, Vol. 12, 519-536.
- Lal, Rajiv and Miklos Sarvary (1999), "When and How is the Internet Likely to Decrease Price Competition," *Marketing Science*, 18(4), 485-503.
- Lee Jongsu, Yeonbae Kim, Jeong-Dong Lee and Yuri Park (2006), "Estimating the Extent of Potential Competition in the Korean Mobile Telecommunications Market: Switching Costs and Number Portability," *International Journal of Industrial Organization*, Vol. 24, 107-124.
- Levey, Richard H. (2011), "Acquisition Spending, Direct Mail on the Rise: Winterberry Report," January 26 2011, <http://chiefmarketer.com/direct/news/winterberry-biegel-direct-mail-digital-rise-0126/>.
- Marinosa, Begona Garcia (2001), "Technological Incompatibility, Endogenous Switching Costs and Lock-In," *The Journal of industrial Economics*, Vol. 49, 281-298.
- Morrin, Maureen (1999), "The impact of Brand Extensions on Parent Brand Memory Structures and Retrieval Processes," *Journal of Marketing Research*, 36, November, 517-525.
- Nilssen, Tore (1992), "Two Kinds of Consumer Switching Costs," *RAND Journal of Economics*, Vol. 23, 579-589.
- Ozga, S. Andrew. "Imperfect markets through lack of knowledge." *The Quarterly Journal of Economics* (1960): 29-52.
- Robinson, Joan, et al. *The economics of imperfect competition*. London: Macmillan, 1933.
- Roehm, Michelle and Alice Tybout (2006), "When Will a Brand Scandal Spill Over, and How Should Competitors Respond?" *Journal of Marketing Research*, 43 (August), 366-373.
- Sharpe Stephen A. (1997), "The Effect of Consumer Switching Costs on Prices: A Theory and Application to the Bank Deposit Market," *Review of Industrial Organization*, Vol. 12, 79-94.
- Shy, Oz. "A quick-and-easy method for estimating switching costs." *International Journal of Industrial Organization* 20.1 (2002): 71-87.

Sieben, Wanda A., and Hsiu-I. Jessie Chen-Yu. "The Accuracy of Size Information on Men's Prewashed Jeans." *Clothing and Textiles Research Journal* 11.1 (1992): 74-82.

Stigler, George J. (1961), "The Economics of Information," *Journal of Political Economy*, 69, 213–225.

Stigler, George J. and Gary S. Becker (1977), "De Gustibus Non Est Disputandum," *American Economic Review*, 67, 76–90.

Suarez, Fernando F. (2005), "Network Effects Revisited: The Role of Strong Ties in Technology Selection," *Academy of Management Journal*, 48(4), 710-720.

Tamburrino, N. (1992), "Sized to Sell," *Bobbin*, 33, 69-74.

Trefis Group (2010), "Private Label Surge Threatens Polo Ralph Lauren," <https://www.trefis.com/company?article=18432#>.

Tybout, Alice and Michelle Roehm (2009), "Let the Response Fit the Scandal" *Harvard Business Review*, (December), 12, 82-88.

Viard, V. Brian. "Do switching costs make markets more or less competitive? The case of 800 -number portability." *The Rand Journal of Economics* 38.1 (2007): 146-163.

Workman, Jane E. (1991), "Body Measurement Specifications for Fit Models as a Factor in Clothing Size Variation," *Clothing and textiles research Journal*, 10(1), 31-36.

Workman, Jane E. and Elizabeth S. Lentz (2000), "Measurement Specifications for Manufacturers' Prototype Bodies," *Clothing and textiles Research Journal*, 18, 251-259.

Table 1. Average Sales at Retail World: All Customers

	Treatment Condition	Control Condition	Difference
Measurement Period			
Units Purchased	7.67 (0.03)	7.66 (0.03)	0.01 (0.05)
Total Revenue	\$260.61 (\$1.10)	\$260.59 (\$1.15)	\$0.02 (\$1.59)
Historical Sales			
Units Purchased	33.56 (0.14)	33.54 (0.14)	0.02 (0.19)
Total Revenue	\$1,276.17 (\$5.04)	\$1,277.02 (\$5.40)	-\$0.85 (\$7.39)
Sample Size	184,455	184,625	

The table reports the averages for each Retail World sales measure in the Treatment and Control Conditions. The outcomes are calculated using the entire twenty-month Measurement Period. The Historical Sales measures are calculated using the customers' entire transaction histories prior to the start of the treatment. Standard errors are in parentheses. *Significantly different from zero, $p < 0.05$, **significantly different from zero, $p < 0.01$.

**Table 2. Difference in Sales Between Experimental Conditions
by Timing of Customers' Prior Purchases**

	Segment 1	Segment 2	Segment 3
	0 to 4 months	4 to 8 months	8 to 12 months
Measurement Period			
Units Purchased	-0.40% (0.87%)	0.21% (0.83%)	5.30%** (1.94%)
Total Revenue	-0.53% (0.86%)	-0.04% (0.83%)	5.88%** (1.93%)
Historical Sales			
Units Purchased	-0.10% (0.84%)	0.04% (0.76%)	1.73% (1.51%)
Total Revenue	-0.10% (0.85%)	-0.20% (0.75%)	1.19% (1.57%)
Sample Size	135,480	187,538	46,062

The table describes the difference in the average sales measures between the Treatment and Control conditions (as a percentage of sales in the Control condition). Positive (negative) values indicate customers in the Treatment condition placed more (less) orders on average. The outcomes are calculated using the entire twenty-month Measurement Period. The Historical Sales measures are calculated using the customers' entire transaction histories prior to the start of the treatment. Standard errors are in parentheses. *Significantly different from zero, $p < 0.05$, **significantly different from zero, $p < 0.01$.

Table 3. The Recency Interaction Estimated at the Customer Level

	Model 1	Model 2	Model 3
Treatment	-1.72% (1.00%)	1.32% (1.91%)	
Treatment x Recency (100s days)	1.73%** (0.61%)	1.61%* (0.72%)	
Recency (100s days)	-55.12%** (0.43%)	-31.59%** (0.51%)	
Treatment x Segment 1			2.53% (1.93%)
Treatment x Segment 2			3.75%* (1.50%)
Treatment x Segment 3			7.12%** (2.38%)
Segment 2			-39.62%** (0.91%)
Segment 3			-72.96%** (1.39%)
Treatment x Prior Total Units (100s)		-3.97% (3.84%)	-4.42% (3.82%)
Prior Total Units (100s)		364.20%** (2.74)	371.46%** (2.72)
Age (years)		-0.24%** (0.02%)	-0.25%** (0.02%)
Estimated Household Income (\$10,000s)		0.57%** (0.04%)	0.57%** (0.04%)
Female Head of Household		7.61%** (0.56%)	7.67%** (0.56%)
Married Head of Household		4.95%** (0.56%)	5.03%** (0.56%)
Number of Kids		-0.10% (0.31%)	-0.11% (0.32%)
Log Likelihood	-1,041,263	-989,125	-989,496
Sample Size	369,080	369,080	369,080

The table reports coefficient from a negative binomial model estimated at the customer level, where the dependent variable measures the number of units purchased in the Measurement Period by customer i . Estimated coefficients that are omitted from Models 2 and 3 include the controls for seasonality (and their interactions with the *Treatment*), and the binary flags identifying missing demographic data. Standard errors are in parentheses.

Table 4. Recency Interaction Estimated at the Customer x Category-Level

	Model 1	Model 2	Model 3	Model 4	Model 5
Treatment	-1.64% (1.90%)	-2.53% (1.91%)	-6.26%** (2.30%)	-1.55% (1.91%)	-0.99% (0.63%)
Treatment * Category Recency (100s days)	0.33%** (0.07%)		0.30%** (0.06%)		
Category Recency (100s days)	-6.37%** (0.49%)		-5.86%** (0.41%)		
Treatment * Log Category Recency		2.72%** (0.43%)			
Log Category Recency		-34.28%** (0.98%)			
Treatment * Intervals Since Last Purchase				2.01%** (0.43%)	
Intervals Since Last Purchase				-41.24%** (1.77%)	
Treatment * 1-2 Intervals					2.86%** (1.04%)
Treatment * 2 Intervals or more					5.85%** (0.98%)
1-2 Intervals					-63.54%** (3.71%)
2 Intervals or more					-127.13%** (6.40%)
Treatment * Category Recency (100s days)			4.84%** (0.55%)		
Recency (100s days)			-39.63%** (2.02%)		
Treatment * Prior Category Units	0.62%** (0.05%)	0.57%** (0.06%)	0.59%** (0.05%)	0.62%** (0.05%)	0.63%** (0.05%)
Prior Category Units	0.46%** (0.04%)	0.44%** (0.04%)	0.44%** (0.04%)	0.46%** (0.04%)	0.46%** (0.04%)
Age (years)	0.20%** (0.05%)	0.16%** (0.05%)	0.18%** (0.05%)	0.20%** (0.05%)	0.18%** (0.05%)
Estimated Household Income (\$10,000s)	0.42%** (0.06%)	0.39%** (0.06%)	0.39%** (0.07%)	0.43%** (0.06%)	0.42%** (0.06%)
Female Head of Household	2.72% (10.60%)	2.76% (10.04%)	1.77% (11.22%)	2.37% (10.41%)	2.54% (10.88%)
Married Head of Household	3.43%* (1.68%)	3.34% (1.87%)	1.93% (1.82%)	3.48%* (1.67%)	3.55%* (1.68%)
Number of Kids	-2.88% (2.31%)	-2.40% (2.36%)	-2.85% (2.41%)	-2.85% (2.29%)	-2.76% (2.34%)
Log Likelihood	-3,089,978	-3,079,526	-3,031,000	-3,080,093	-3,103,969
Sample Size	2,214,658	2,214,658	2,214,658	2,214,658	2,214,658

The table reports coefficients from quasi-maximum likelihood Poisson models with (conditional) fixed product category effects. The dependent variable measures the number of units of category *c* purchased in the Measurement Period by customer *i*. Estimated coefficients that are omitted from Models 2 through 5 include the controls for seasonality (and their interactions with the *Treatment*), and the binary flags identifying missing demographic data. The standard errors are clustered by category and reported in parentheses.

Table 5. Return Reasons: Summary Statistics

Return Reason	%	Return Reason	%
Did Not Like The Item		Defective	
Did Not Like Color	6.6%	Defective Construction	2.7%
Did Not Like Material	4.4%	Care Instructions Failed	0.6%
Did Not Like Styling	12.3%	Shipping Damage	0.2%
Did Not Perform As Expected	2.4%	Miscellaneous	
Finish Unacceptable	0.2%	Miscellaneous	19.2%
Not As Described	0.5%	Difficult To Assemble	0.1%
Not As Pictured	1.1%	Price Too High	0.5%
Quality Unsatisfactory	1.5%	Wrong Item Shipped	1.5%
Item was the Wrong Size		Total	100%
Ordered Wrong Size	2.8%		
Too Large	22.1%		
Too Small	21.7%		

The table describes the frequency with which each return reason appears in the transaction data for the separate sample of 3.6 million customers.

Table 6. Change in the Return Rate and the Response to the Competitive Advertising

	Wrong Size	Did Not Like the Item	Defect	Other
Pearson Correlation Coefficients				
3 category orders	-0.38*	-0.01	-0.24	-0.20
All customers with at least 2 category orders	-0.41*	-0.01	-0.22	-0.09

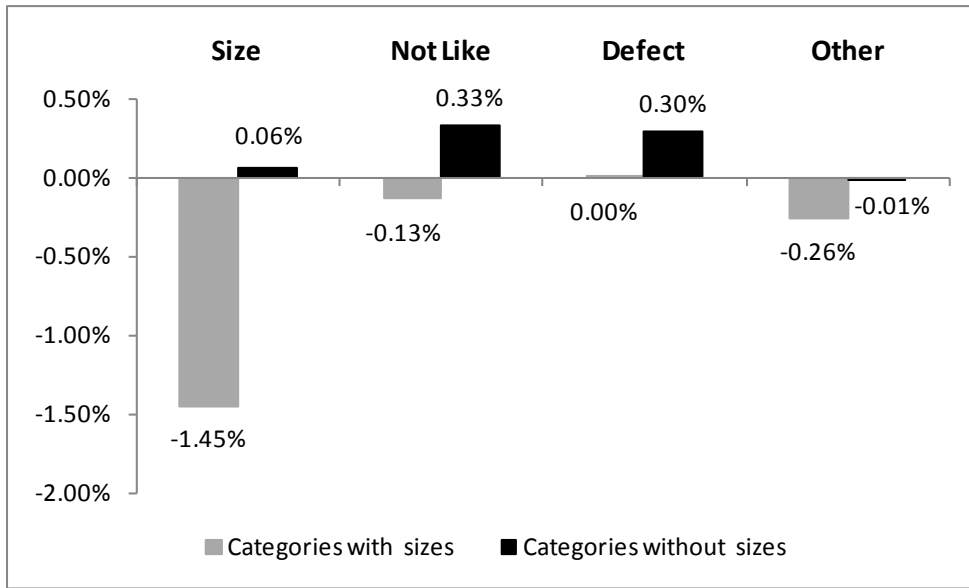
The table reports the pair-wise correlation between the response to the competitive advertising and both the *Change in Return Rate*. The *Change in Return Rate* is estimated using the separate sample of 3.6 million customers. The response to the competitive advertising measures the % difference in units purchased in each category by customers in the Treatment and Control conditions. The sample size for each correlation is 35 (categories). *Significantly different from zero, $p < 0.05$, **significantly different from zero, $p < 0.01$.

**Table 7. Size Learning Interactions
Estimated at the Customer x Category-Level**

	Model 1	Model 2
Treatment	-4.54%* (1.91%)	3.28%** (1.10%)
Treatment * Size Learning	-5.60%** (1.53%)	-5.29%** (1.38%)
Size Learning	-5.14%** (0.67%)	-5.10%** (0.68%)
Treatment * No Prior Category Purchase		-9.90%** (1.25%)
No Prior Category Purchase		243.70%** (11.15%)
Treatment * Category Recency (100s days)	0.39%** (0.05%)	0.38%** (0.06%)
Category Recency (100s days)	-6.08%** (0.49%)	-5.95%** (0.49%)
Treatment * Prior Category Units	0.39%** (0.05%)	0.41%** (0.05%)
Prior Category Units	0.38%** (0.04%)	0.38%** (0.04%)
Age (00s years)	0.18%** (0.05%)	0.06% (0.09%)
Estimated Household Income (\$10,000s)	0.45%** (0.05%)	0.56%** (0.05%)
Female Head of Household	3.28% (10.67%)	7.00% (10.13%)
Married Head of Household	3.46% (1.78%)	5.94%** (1.87%)
Number of Kids	-2.68% (2.52%)	-0.34% (2.89%)
Log Likelihood	-2,798,281	-5,175,369
Sample Size	2,127,691	12,917,800

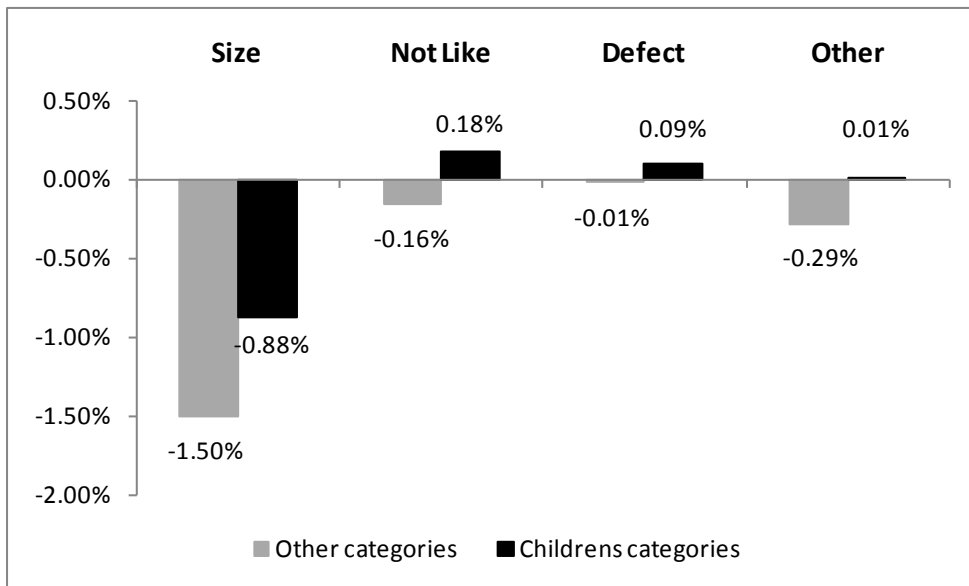
The table reports coefficients from a quasi-maximum likelihood Poisson model with (conditional) fixed product category effects. The dependent variable measures the number of units of category c purchased in the Measurement Period by customer i . Estimated coefficients that are omitted from the table include the controls for seasonality (and their interactions with the *Treatment*), and the binary flags identifying missing demographic data. The standard errors are clustered by category and reported in parentheses.

Figure 1. Change in Return Rates in Categories With / Without Item Sizes



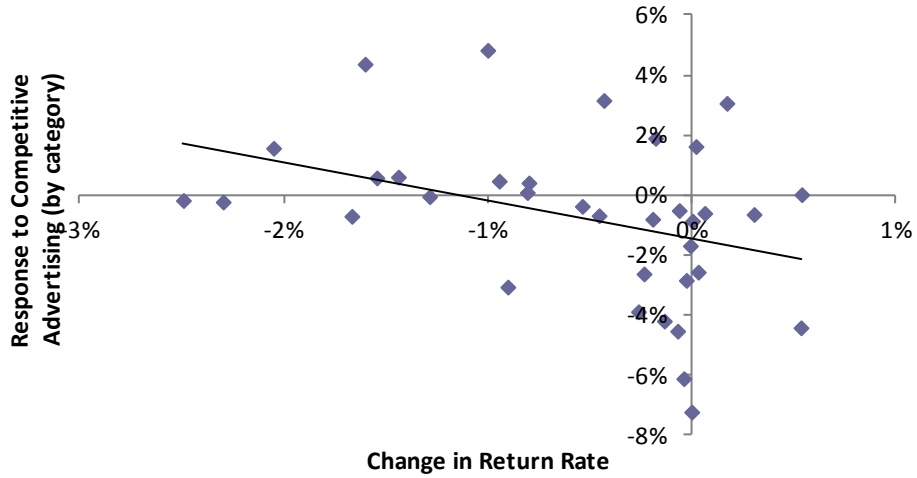
The figure reports the average change in the return rates in categories with / without item sizes. The change in the return rates is measured as the difference in the return rates on a customer’s first and last order in the category. The unit of analysis is a category x customer and the sample includes all customers who placed at each two orders in a category. The sample sizes are 5,160,972 for items with sizes and 1,245,168 for items without sizes. The differences in the two averages are statistically significant ($p < 0.01$) for all four return reasons.

Figure 2. Change in Return Rates in Childrens’ vs. Other Categories



The figure reports the average change in the return rates in Childrens’ categories vs. Other categories. The change in the return rates is measured as the difference in the return rates on a customer’s first and last order in the category. The unit of analysis is a category x customer and the sample includes all customers who placed at each two orders in a category. The sample sizes are 429,128 for the Childrens’ categories and 4,731,844 for the other categories. The differences in the two averages are statistically significant ($p < 0.01$) for all four return reasons.

Figure 3. The Response to the Competitive Advertising by Category and the Change in the Rate of Size Returns



The y-axis measures the percentage difference in the average number of units purchased in each category by customers in the Treatment and Control conditions. Positive (negative) values indicate customers in the Treatment condition ordered more (less) units on average. The x-axis measures the change in return rates between customers first and last orders in the category. The measure for each category is averaged across customers who had placed at least 2 orders in the category using the separate sample of 3.6 million customers. The figure includes a linear trend line.

Appendix Table A1. Examples of Mailing Requests

	Request 20	Request 21	Request 22
Company	K	P	AJ
Request Date	2/20/06	2/9/06	3/13/06
Mail Date	3/25/06	3/28/06	4/2/06
Mailing Piece	Spring 2006 catalog	Spring preview catalog	March catalog
Selection Criteria	Customers who spent \$100 or more dollars in the last 6 months from the childrens' catalogs. Exclude previously used customers.	Female customers who spent \$75 or more since November 2005 on women's' apparel products	Female customers who live in mid-West states and purchased products other than apparel in the last three months.
Number Ordered	28,867	20,000	44,600
Number Mailed	25,905	20,000	42,011

These examples are disguised to protect confidentiality.

In the first example (Request Number 20), on February 2, 2006 Company K asked for its Spring 2006 catalog to be sent to 28,867 customers who had spent \$100 or more in the last six-months on items from Retail World's childrens' catalogs. Company K also specified that any customers included in previous requests were to be excluded. The third-party is responsible for both identifying the qualifying customers and fulfilling the mailing request. In this case the third-party firm found 25,905 of Retail World's customers that qualified under these criteria and organized for Company K's Spring 2006 catalog to be mailed to these customers on March 25, 2006. Company K never receives the actual names of the customers who were mailed. Because the 25,905 customers believe that they were mailed a catalog directly by Company K, any orders they place are made directly with this company.

Appendix Table A2. Competitors' Mailing Request Criteria

	Frequency	Percentage of Requests
Recency: Time Since Last Purchase		
Less than 3 months	96	32%
Less than 6 months	54	18%
Within 7 to 12 months	27	9%
Specific months	68	23%
Any recency criterion	262	87%
Amount Spent		
Over \$50	36	12%
Over \$75	46	15%
Over \$100	83	28%
Over \$150	62	21%
Any amount spent criterion	240	80%
Other Criteria		
Purchasers from a specific catalog-type	150	50%
Purchases from a specific category	128	43%
Customer gender	112	37%
Specific zip codes	50	17%
Sample Size		
Total number of mailing requests	301	

This table summarizes the criteria used in the 301 mailing requests from competitors.

**Appendix Table A3. Eligible Mailing Requests
by Customers' Historical Purchases From the Retail World**

	Average Number of Eligible Requests	Sample Size
Recency: Time Since Last Order		
Less than 4 months	50.0	135,480
4 to 8 months	28.3	187,538
8 to 12 months	8.2	46,062
Frequency: Number of Prior Orders		
1 to 3 prior purchases	20.4	137,040
4 to 6 prior purchases	28.5	55,877
7 to 12 prior purchases	34.5	66,554
13 or more prior purchases	52.7	109,609
Average Order Size		
0 to \$50	16.4	86,602
\$50 to \$75	21.9	44,923
\$75 to \$100	25.7	41,155
Over \$100	45.8	196,400

This table summarizes the estimated number of competitors' mailing requests (out of a total of 301) that each customer was eligible for. Customers are grouped according to their historical purchases from Retail World. We restricted attention to transactions that occurred prior to the start of the competitive advertising event.