HARVARD | BUSINESS | SCHOOL

The Impact of Increasing Search Frictions on Online Shopping Behavior: Evidence from a Field Experiment

Donald Ngwe
Thales Teixeira

Working Paper 18-114

The Impact of Increasing Search Frictions on Online Shopping Behavior: Evidence from a Field Experiment

Donald Ngwe
Harvard Business School

Thales Teixeira
Harvard Business School

Working Paper 18-114

The impact of increasing search frictions on online shopping behavior: Evidence from a field experiment

Donald Ngwe
Assistant Professor
Harvard Business School
Morgan Hall 179
Boston, MA 02163
(617) 496-6585
dngwe@hbs.edu

Thales Teixeira
Associate Professor
Harvard Business School
Morgan Hall 165
Boston, MA 02163
(617) 495-6125
tteixeira@hbs.edu

The impact of increasing search frictions on online shopping behavior: Evidence from a field experiment

Abstract

Online retail accounts for a rapidly growing proportion of revenues in many industries. While selling online broadens firms' access to consumers, operating margins are often lower in online stores than in physical stores. There are well-recognized reasons for this discrepancy: prices are easy to compare online, discount coupons and codes have high uptake, and sellers often bear the cost of shipping products to buyers. In addition to these factors, online selling precludes many methods of price discrimination exercised in offline environments. Many online stores present few barriers to accessing discounted products. We propose that deliberately increasing search frictions by placing obstacles to locating discounted items can improve online retailers' margins and increase conversion. We demonstrate using a simple theoretical framework that inducing consumers to inspect higher-priced items first can simultaneously increase the average selling price and the overall purchase rate. We test these predictions in a series of field experiments conducted with an online fashion and apparel retailer. Using information from historical transaction data about each existing consumer, we demonstrate that price-sensitive shoppers are more likely to incur search costs in order to locate discounted items. Our results show that adding search frictions can be used as a self-selecting price discrimination tool to match high discounts with price-sensitive consumers and full priced offerings with price-insensitive consumers.

Keywords: e-commerce, online retailing, friction, effort, search costs, price discrimination

Online retail accounts for a rapidly growing proportion of revenues in many industries. As of the fourth quarter of 2017, e-commerce accounted for 8% of total U.S. retail sales, from less than 4% in 2008 (U.S. Census Bureau 2017). While online retail broadens firms' access to consumers through an addition channel, operating margins are often lower in online stores than in physical stores. Amazon, the largest online retailer in the U.S., averaged 1.3% in operating margins from 2011 to 2013 while its brick-and-mortar counterparts typically experienced 6% to 10% (Rigby 2014). Reasons for this difference are well recognized: prices are easy to compare online, discount coupons and codes have high uptake, and sellers often bear the cost of shipping products to buyers.

These factors are in part a consequence of many online sellers' efforts to make online shopping as convenient as possible. Pure online sellers such as Amazon, Wayfair.com, and Overstock.com are known for minimizing the search, transaction, and delivery costs for shoppers in an attempt to lure them from offline channels. In industries ranging from consumer electronics to flight and hotel booking, third-party sites reduce search costs even further by enabling cross-website comparisons.

This trend is in stark contrast with the practice of brick-and-mortar retailers, who have long embraced the deliberate use of search frictions to improve in-store margin performance. In making lower-priced and lower-margin items harder to locate, by placing the sale section in the back of the store, using bargain bins and discount racks, or having a separate outlet store, physical stores induce a self-selection among consumers who are heterogeneous in their willingness to pay (e.g. Coughlan & Soberman 2005; Ngwe 2017).

In this paper, we seek to challenge the prevailing assumption that minimizing search frictions, i.e., facilitating consumer search across a retailer's entire assortment, is always the optimal strategy for online retailers (Bakos 1997; Brynjolfsson & Smith 2000). We contend that just as in physical selling contexts, careful incorporation of search frictions can facilitate price discrimination in online retail.

The existing literature has typically conceived of search costs as the time, effort, and money required to physically identify and consider additional options prior to making a purchase decision (Bell, Ho & Tang 1998). In offline retail contexts, these often take the form of travel costs or time spent shopping in measurable intervals. Given the ease and immediacy of online shopping, it is perhaps unsurprising that equivalent search costs have not been studied as tools a firm would use to implement price discrimination. We identify and explore the potency of search frictions—deliberately built-in search costs—in online settings: the effort behind clicking an additional link, viewing an additional page, scrolling through a catalogue of items, or mentally calculating the percentage discount on a sale item.

Our hypothesis is that under certain conditions, an online seller can improve its average selling price and overall purchase rate by adding search frictions associated with finding and purchasing discounted items on its website. The first condition is that increasing search frictions by eliminating retail website elements such as navigation tools changes the order in which consumers consider items. The second condition is that upon encountering these added search frictions, price-insensitive consumers would substitute from discounted items to full priced items. The third condition is that price-sensitive consumers would exert the extra effort required to find discounted items on the website, in a similar manner as has been observed in physical settings (e.g. Seiler & Pinna 2017). We offer a simple model of search that generates these predictions under a basic set of assumptions.

In order to test our hypothesis, we run a series of field experiments with an online fashion and apparel retailer. This category is particularly appropriate for our purposes because it features moderate frequency and value of purchase. Consumers are broadly aware of price points for apparel but are not completely certain of product assortment on any given purchase occasion. Item prices are material but not the sole consideration for most shoppers. Lastly, it is common practice for apparel retailers to frequently offer sizeable discounts in order to acquire and retain customers.

In our first experiment, we randomly assign new visitors to the online store over the course of nine days to either a control group or one of three treatment groups. Product availability and prices are held constant across all conditions. Each treatment involves increasing search frictions in some way: (i) removing the direct link to the outlet section, which is a catalog for highly discounted products; (ii) removing the option to sort product listings by discount size; and (iii) removing item-specific visual markers that indicate discount percent. We find that the average discount of purchased items is significantly lower in the treatment groups than in the control group while the conversion rate is higher. These results demonstrate that our experimental manipulations have significant effects on purchase behavior and provide preliminary evidence for our hypothesis.

In the succeeding analyses, we aim to establish the relationship between consumers' price sensitivity and their responses to added search frictions. We use historical transaction data for existing customers to pre-classify them according to their price sensitivity. We do so by regressing the average discount of their most recent transaction on demographic and past purchase variables, then using the predicted values as a proxy for price sensitivity. We validate this classification by showing that consumers we identify as price-sensitive are more likely to click on randomly assigned discount-oriented versus full price-oriented email newsletters.

In our second experiment, we randomly assign nearly 350,000 visitors—new and existing customers—to the online store over the course of two weeks to either a control group or one of four treatment groups. Again, product availability and prices are held constant across all conditions. We carry over the treatments from our first experiment and include the replacement of discount banners with non-discount banners as an additional condition. We identify the presence of self-selection among the firm's entire customer base, utilizing past purchase behavior to measure heterogeneous treatment effects.

We find that as in the first experiment, the addition of search costs lowers the average discount of purchased items and enhances the conversion rate. Moreover, gains from increased selling prices are mostly attributed to price-insensitive consumers buying disproportionately more full priced items in the treatment groups. These results imply that placing obstacles to locating discounted items cause price-insensitive consumers to switch to full priced items. They also show that the main effects we capture are stable across varying demand conditions, as our second experiment included new as well as existing customers and was run more than a year after our first experiment during the sale season.

RELATED LITERATURE

Our work relates to the literature on search costs in online settings, much of which explores the effect of diminishing information frictions on consumer choice, welfare, and market structure. Early papers associated decreases in search costs with an increase in price competition (e.g. Bakos 1997). Subsequent work has found that online search may actually be more costly than originally understood, possibly due to online shoppers having higher search costs than offline shoppers, and that there may be substantial heterogeneity in search costs among online shoppers (Brynjolfsson & Smith 2000).

Lynch and Ariely (2000) find conditions under which increased quality transparency online can mitigate price competition, providing support for their claim that "in a competitive environment, the strategy of keeping some search costs high is arguably doomed to fail." Yet others have shown that some firms have found it profitable to artificially increase search costs in an attempt to make price comparison more effortful. Ellison and Ellison (2009) show that firms participating in a marketplace have an incentive to obfuscate the actual price and quality information of goods sold online so as to reduce search-sensitive consumers' inclination to exhaustively compare prices across firms.

As examples of obfuscation, Ellison and Ellison (2005) show how online retailers profitably offer complicated menus of prices, products that seem to be bundled but are not, hidden prices, complex product descriptions, and other tactics designed "to make the process of examining an offer sufficiently time-consuming." They claim that many advances in search engine technology that presumably were intended to facilitate consumer information gathering have been subsequently accompanied by firms' investments in hampering search.

In our approach, we introduce search frictions in order to manipulate the order in which consumers consider items in an online store. Our findings thus relate to research on how position, with respect to search order, influences consumer choice (e.g. Weitzman 1979; de los Santos & Koulayev 2014; Narayanan & Kalyanam 2015; Armstrong 2017). Recent papers have used field experiments to cleanly identify the impact of search order and rankings on purchase outcomes in single-category settings (e.g. Choi & Mela 2016; Ursu 2017). We build on this stream of literature by demonstrating the relevance of website navigation links on search, particularly where a multi-category seller offers a large number of choices that cannot be ordered in a single list.

Our work relates to research that links search and price discrimination. Varian (1980) formulates a model of price discrimination that allows a firm to extract surplus from uninformed consumers via high prices and, concurrently, sell to the informed consumers via low prices using a mixed-strategy pricing equilibrium. More recent models featuring heterogeneous search also find equilibria in which sellers adopt a mixed strategy (Ratchford 2009). For example, Stahl (1996) finds a mixed strategy equilibrium that, when implemented by two retailers, creates a separation in the market such that fully informed consumers always buy from the lower priced firm while uninformed consumers stop short of comprehensive search and pay higher prices. Our findings show that in addition to strategic incentives for obfuscation, there are benefits to obfuscation that accrue to monopolistically competitive firms.

We focus specifically on cases where the monopolistic seller designs an optimal menu of products such that consumers self-select according to their willingness to pay (Mussa & Rosen 1978). Previous empirical research has demonstrated the occurrence of this practice in such settings as Broadway theater (Leslie 2004) and coffee shops (McManus 2007). We build on this research by demonstrating through online field experiments how second-degree price discrimination can be exercised in e-commerce settings by altering website design.

Related research has explored the use of effort to allocate discounts to price-sensitive consumers, particularly in the context of discount coupons (e.g. Narasimhan 1984) and waiting in line (e.g. Nichols et al 1971). Crucially, our implementation of increasing search frictions differs from earlier work in that part of the added effort induced on consumers involves inspecting additional items sold by a multiproduct firm, which we show has salutary effects on purchase likelihood.

THEORETICAL FRAMEWORK

We sketch a simple model of search to illustrate our hypotheses and inform the design and analysis of our field experiments. The basic setup follows recent work in the search literature that examines obfuscation in settings with monopolistic competition (e.g. Petrikaite 2014; Gamp 2016; Armstrong 2016). We introduce features to the basic setup that reflect search in an online retail setting and correspond to the treatment conditions in our field experiments. In the following exposition, we focus on describing consumer responses to different seller decisions rather than characterizing equilibrium outcomes.

Suppose there is a mass of consumers that is normalized to one. A consumer has unit demand and wants to buy one of two horizontally differentiated products. The net utility of consumer i who buys product j is denoted by u_{ij}, which equals the difference between her match value ϵ_{ij} and the price p_j of the product, given the consumer's price sensitivity α_i:

$$u_{ij} = \epsilon_{ij} - \alpha_i p_j$$

The match value indicates the valuation of variety j by consumer i; it is consumer- and product-specific. Match values are distributed independently and identically among consumers and products according to a uniform distribution over the interval $[0,1]$. Both products are sold by a single online seller at prices such that $p_1 < p_2$. Note that products are ex-ante identical apart from their prices, so henceforth we shall be referring to each product as p_1 and p_2.

Consumers arrive at the website and consider products sequentially. Product prices are common knowledge but match values are only realized when products are inspected (e.g. Gu &

Liu 2013). Consumers can inspect the first product for free but pay a search cost s to inspect the second product. By manipulating the website, the seller can change the order in which consumers consider products. Suppose that the seller has chosen to present the lower-priced product p_1 first. The search process follows this sequence:

1. The consumer arrives at the website and views p_1
2. She realizes match value ϵ_{i1} and calculates her product utility u_{i1}
3. If the consumer's expected payoff from inspecting the second product is greater than her payoff from buying the first product or not purchasing anything, i.e. $\max(u_{i1}, 0) \leq E(u_{i2}) - s$, then she inspects the second product
4. The consumer purchases the product, among those inspected, that offers the highest utility
5. If all inspected products have utility less than zero, then the consumer does not make a purchase

Note that the consumer inspects the second product if

$$\max(u_{i1}, 0) \leq E(u_{i2}) - s$$

$$\max(\epsilon_{i1} - \alpha_i p_1, 0) \leq E(\epsilon_{i2} - \alpha_i p_2) - s$$

$$\max(\epsilon_{i1} - \alpha_i p_1, 0) \leq 0.5 - \alpha_i p_2 - s$$

If $s > 0.5 - \alpha_i p_2$, the expected utility from the second product is lower than the search cost and no consumers inspect the second product. For $s \leq 0.5 - \alpha_i p_2$, consumers search further if their match value from the first product is low enough, i.e. $\epsilon_{i1} \leq 0.5 + \alpha_i p_1 - \alpha_i p_2 - s$.

We can now characterize demand. Demand for the first product equals the sum of two probabilities: the probability that a consumer buys the first product without inspecting the second product ($\epsilon_{i1} \geq 0.5 + \alpha_i p_1 - \alpha_i p_2 - s$) and the probability that a consumer inspects both products and buys the first product because its utility is higher and non-negative ($\max\{\epsilon_{i2} - \alpha_i p_2, 0\} \leq \epsilon_{i1} - \alpha_i p_1 < 0.5 - \alpha_i p_2 - s$). As a result, demand for the first product equals

$$d_1 = \begin{cases} 1 - (0.5 + \alpha_i p_1 - \alpha_i p_2 - s) + \int_{\alpha_i p_1}^{0.5 - s - \alpha_i p_2 + \alpha_i p_1} (\epsilon - \alpha_i p_1 + \alpha_i p_2) d\epsilon, & \text{if } s \leq 0.5 - \alpha_i p_2 \\ 1 - \alpha_i p_1, & \text{if } s > 0.5 - \alpha_i p_2 \end{cases}$$

The consumer buys the second product if she pays the search cost (which implies that $\epsilon_{i1} < 0.5 - \alpha_i p_2 - s + \alpha_i p_1$) and finds that $\max(\epsilon_1 - \alpha_i p_1, 0) \leq \epsilon_2 - \alpha_i p_2$. The probability of this event constitutes the demand for the second product:

$$d_2 = \begin{cases} (0.5 - s - \alpha_i p_2 + \alpha_i p_1)(1 - s - \alpha_i p_2) + \int_{\alpha_i p_2}^{0.5 - s} (\epsilon - \alpha_i p_2 + \alpha_i p_1) d\epsilon, & s \leq 0.5 - \alpha_i p_2 \\ 0, & s > 0.5 - \alpha_i p_2 \end{cases}$$

The mass of consumers who choose not to purchase is $d_0 = 1 - (d_1 + d_2)$. We illustrate the relationships between the seller's choice of search order, the magnitude of search costs, and seller performance through a numerical example. Setting prices $p_1 = 0.1$ and $p_2 = 0.2$, we compute purchase outcomes for consumers with price sensitivity $\alpha_i = 1$ and present the results in Figure 1. The three subplots show revenues, conversion rates, and accrued search costs for the two possible search orders as chosen by the seller over varying levels of s.

[Insert Figure 1 about here.]

Several interesting relationships arise from our simple model. The first subplot presents revenues $p_1 d_1 + p_2 d_2$ for the two different search orders. When p_1 is presented first (labeled as p1 -> p2), revenues decline as s increases because fewer consumers choose to inspect the second, more expensive item. When p_2 is presented first, revenues initially increase with s, but drop at the point where no consumers consider the second item. Presenting the more expensive item first always results in higher revenues. This finding echoes similar results as in Jing (2016).

The second subplot shows conversion $1 - d_0$, or the portion of consumers who choose to buy one of the two products rather than choosing the outside option. At low levels of s, conversion is very high because all consumers who have negative utility from the first option choose to inspect the second option. For $0.3 < s \leq 0.4$, presenting the less expensive option first results in lower conversion than presenting the more expensive option first. In the first ordering, consumers have a low expected utility from inspecting the second product because it is more expensive, whereas in the second ordering, consumers are willing to inspect the second product because of its low price. When search costs are high enough such that consumers inspect only the first product in either case, conversion simply reflects the demand for either good.

The third subplot shows ex-post search costs incurred by consumers in either case. The inverted U shape reflects the impact of the magnitude of s on search behavior. Importantly for our paper, the graph shows that switching the order of product presentation such that more expensive options are considered first increases the ex-post search costs experienced by consumers.

Next, we examine how the effect of the order of product presentation varies with the consumer's price sensitivity α_i. Figure 2 maps out the purchase likelihood of the higher-priced item for a range of price sensitivity, and given the two orderings and fixing search cost to $s = 0.1$. The purchase likelihood of the higher-priced item is weakly higher when it is presented first, regardless of price sensitivity. Consumers of mid-range price sensitivity ($2 < a_i < 5$) have strictly positive likelihood of purchasing the higher-priced item only when it is presented first, as they do not opt to inspect the product when it is presented second. Overall conversion is weakly higher in the second ordering for consumers with low price sensitivity ($\alpha_i < 4$), but is lower for the remainder of consumers.

[Insert Figure 2 about here.]

Our theoretical framework shows that under a basic set of assumptions, a monopolistic seller may be better off presenting high priced products first when consumers search sequentially. The potential benefit is twofold. First, this ordering increases the average selling price by increasing demand for the high priced product relative to the low priced product. Second, this ordering induces more consumers to inspect both products, causing consumers to absorb higher search costs but increasing the likelihood that certain consumers will purchase either product instead of opting for the outside option. In the remainder of the paper, we describe how we test these predictions empirically.

OVERVIEW OF EMPIRICAL ANALYSIS

We employ both field experiments and analyses of historical purchase data in our empirical analysis. All data are provided by an online fashion and apparel retailer in the Philippines. The online retailer sells branded merchandise as well as items under its private label. The firm offers the widest online selection of men's and women's fashion items in the country. While statistics on industry concentration are not available, according to its managers the firm accounts for about

40% of overall online fashion retail in the country, with the next largest competitor having less than 5% market share. The market structure is therefore one in which there is dominant firm, with hundreds of smaller firms comprising a competitive fringe.

Items are listed on the website under three catalogs: main, sale, and outlet. The main catalog contains all full priced offerings as well as some lightly discounted items. The sale catalog contains moderately discounted items, while the outlet catalog contains heavily discounted items, where the precise cutoffs between light, moderate, and heavy discounting varies over time. All products offered by the firm are first listed in the main catalog, then gradually discounted and listed in the other catalogs as newer products are introduced, a common industry practice.

We run two field experiments through the online store. The first is an exploratory study in which new visitors to the website are exposed to added search frictions. A key objective in running this pilot experiment is to assess which, if any, of our basic set of treatments represent nonzero search costs that induce changes in purchase activity.

In our second field experiment, we expose both new and returning customers to the online store to added search frictions and measure how the treatment effects differentially impact purchase rates, discounts, and margins, depending on a shopper's estimated price sensitivity. This experiment complements our initial field experiment to achieve multiple objectives:

- It increases the sample size to include the firm's existing customers
- Because it is run during a sale season, the second experiment more closely aligns with the assumption of price being common knowledge in our theoretical framework
- It allows us to disentangle and explore additional means of implementing search frictions in an online store

- It allows us a replication of results on conversion and average selling price
- It allows us to assess whether varying levels of consumer information, particularly on prices, has impacts on the findings

The available data from the firm consists chiefly of transaction records containing product and customer attributes. Limited and high-level browsing data such as average session length and average number sessions in each condition are available from a third-party web analytics provider; however, we do not have access to granular clickstream data and server logs. In order to measure heterogeneous treatment effects from our second field experiment, we use a regression model to classify existing consumers according to their price sensitivity using historical transaction data. We validate our classification by means of measuring response rates to randomly assigned email newsletters. Details for these classification steps are found in the Appendix.

We turn to a discussion of the design and results of each field experiment.

FIELD EXPERIMENT I

In our first experiment, we seek preliminary evidence that specific changes to website design can have significant effects on shopper behavior and purchase outcomes. We vary the presence of website features that potentially facilitate shopper inspection of discounted items. We include only new visitors to the desktop version of the online store in order to mitigate potentially negative effects on the firm's performance and to control for prior information among consumers. In evaluating the outcomes, we are particularly interested in the treatment effects on the discount levels of completed transactions and the overall conversion rate. In line with our theoretical framework, our prediction is that deliberately increasing search costs decreases the likelihood that price-insensitive customers will incur the search cost required to inspect discounted products, and that these customers will substitute to full priced items instead.

We run the experiment on the retailer's website over a period of nine days.[1] During this period, all new visitors to the website were randomly assigned to the control group or one of three treatment groups with equal probability. New visitors are defined as customers who do not have the website's cookies on their computers and sign up for a new account before making any purchase. Only visitors who were using a desktop, laptop, or tablet computer were included in the study. In total, 104,605 shoppers were included in the experiment. Posterior randomization checks confirm that the firm correctly implemented the randomization of new visitors to treatment and control groups. (See Table W3 in the Web Appendix for randomization checks.)

Additionally, only consumers who arrived at the site via the main landing page are included; we exclude consumers who visit the site via an email coupon, newsletter, or link from a third-party website. During the experiment no other changes were made to the website. Descriptions of the control and treatment conditions follow. In each of the treatment conditions, neither the available product assortment nor any product prices were different from the control condition.

Control

The control condition is simply the website as is at the time of the study. The website features elements designed to facilitate consumer search for discounted items. Customers have three ways to find discounted items: clicking on a prominent link from the landing page to the outlet catalog, sorting products according to discount level in each catalog through a drop-down

[1] Field Experiment I ran on June 17-June 25, 2015.

option, and viewing markers that highlight discounts above 40% (see Figure W3 in the Web Appendix). In each of the treatment conditions, we eliminate each of these elements with the objective of reordering consumers' search paths and increasing the required effort to locate discounted items.

Treatment 1: No link to outlet catalog from main landing page

In this condition, we eliminate the most direct path to discounted items: the outlet link from the landing page. The remaining links to the outlet catalog are within the website's sale section, requiring at least one additional click from a shopper to arrive at the outlet catalog relative to those in the control group. In line with our theoretical framework, this would require consumers to view higher-priced items before viewing items with the largest discounts.

Treatment 2: No discount filter and no discount markers

Here we remove the ability of consumers to order product listings according to discount percentages. We also remove the accompanying discount markers, which provide visual cues for identifying high discounts. These elements are widely used together by online retailers to facilitate shopper search and navigation. Similarly to Treatment 1, this would cause consumers to view higher-priced items first.

Treatment 3: No outlet link, no discount filter, and no discount markers

In our third treatment we simultaneously remove all website elements taken out piecemeal in the first two treatments. Our objective is to induce variation across our treatments in the magnitude of the search cost associated with inspecting discounted items. As with the first two treatments, visitors can still find discounted items with the requisite effort in typical locations throughout the site. Eliminating links, filters, and markers merely adds to the number of clicks, browsing time, and page views required to locate discounted items.

In evaluating the effects of each treatment we consider several outcome variables[2]:

1. Average discount: the average ratio of selling prices to original prices[3] over items bought in each treatment group. Given that each treatment makes locating discounts more difficult, we expect percent discounts to be lower in treatment conditions relative to the control on average. We use this variable in place of selling prices as a means of enabling a consolidated presentation of results given the multi-category setting.
2. Percent full priced purchases: the proportion of purchased items sold without discounting. Historically, about 50% of purchases on the website are made at full price. Similarly to the average discount, an increase in this variable in our treatment groups would be supportive of our hypothesis.
3. Conversion rate: the percent of consumers who opt to make a purchase on the website within the testing period. Our theoretical framework suggests that by inducing consumers to inspect more products, our manipulations may result in a higher conversion rate. We track the conversion rate to assess if this is true given consumer preferences in our empirical setting.

[Insert Table 1 about here.]

Table 1 contains the results of the experiment. Customers in all three treatment groups purchased items at significantly lower discounts on average (11.5 to 12.3% off versus 15.7% off), and purchased more full priced items (64.2% to 67.5% versus 59.8%).[4] The average selling prices of items purchased in the three treatments were significantly higher than in the control

[2] We note that two of our outcome variables (average discount and percent full priced purchases) are defined conditional on purchase; hence, they are computed using selected samples by construction. While not ideal for evaluating experimental results, this is a direct consequence of the behavior being studied and is in line with our theoretical framework (e.g. Sahni, Wheeler & Chintagunta 2018).
[3] The firm does not inflate original prices.
[4] The average discount across the three treatment groups was 11.8%, which is significantly different from the control at p=0.000.

condition, indicating that consumers do not substitute to lower-priced items and possibly lower margin items in response to the treatments.

We show in our theoretical framework that overall conversion may either increase or decrease depending on the sensitivity of consumers to price and search costs. A natural concern is that if search frictions for finding discounted items are increased too much, then the expected result of reducing discounted purchases could also be accompanied by lower conversion rates. This is of particular concern for first time shoppers, who may be unaware of the availability of lower-priced items on the website. Yet, we found no significant decrease in conversion rates, as measured by the number of transactions completed in any of the treatment conditions versus the control group. As shown in Table 2, conversion rates were slightly higher in treatments 1 and 2.

As a robustness check of our main finding, we performed a comparison across treatments and control groups at the basket level (versus item level) to compare differences in shopping behavior with respect to total purchase value and composition decisions. Confirming the main item-level results, Table 2 shows that the average discount of purchase baskets in two out of the three treatments is significantly lower than for the control group (11.6 to 12.2% versus 14.5%). For treatment 3, it is marginally significantly lower. Average basket sizes in any of the treatment groups are not significantly smaller than the control group, although in treatment 3 there were fewer items in each basket on average.

[Insert Table 2 about here]

These results provide supporting evidence for our hypothesis and demonstrate that our manipulations have a measurable impact on consumer choice. Given that discounts in fashion and apparel retail are directly tied to gross margins, our results additionally show that online retailers can increase their margins without sacrificing conversion by slightly increasing search frictions associated with their discounted offers. (See Table W4 in the Web Appendix for measurements of the impact on profitability of the treatment conditions.)

In a setting without search frictions, we contend that price-insensitive consumers inspect discounted options "for free." By adding search frictions, online retailers can direct these consumers toward full priced options while still making discounted options available for price-sensitive shoppers willing to incur the extra search cost to find them. In the following sections, we verify that such heterogeneous treatment effects underlie our main results.

FIELD EXPERIMENT II

In our second experiment we investigate the pattern in which shopper responses to added search frictions vary according to price sensitivity. Analogous to Field Experiment I, we expose consumers to different versions of the online store, each with a manipulation designed to increase search frictions. We compare the impact of the treatments versus control on retailer performance measures and use consumer purchase histories to characterize the heterogeneity in shopper responses. In addition, this experiment also serves as a replication of our first experiment, run during a sale season approximately one year afterward with a wider set of treatments and a broader set of subjects. [11]

Experimental Design

We ran this experiment for two weeks on the desktop and tablet versions of the online store. All consumers were randomly assigned to either the control group or to one of four treatment groups with equal probability. Whereas in Field Experiment I we included only new

[11] Field Experiment II was run on June 1-15, 2016.

visitors entering through the main landing pages, here we include new as well as returning consumers regardless of which page they view first. This is a strong test of our model-based predictions as it seeks to find the conjectured shopper behavior in a population with presumably high variation in information. A total of 348,110 visitors are included in this experiment (see Table W6 in the Web Appendix for randomization checks.) The treatment conditions are as follows:

Treatment 1: Removal of links from main pages to outlet and sale sections of the website

Treatment 2: Removal of discount markers

Treatment 3: Removal of discount sorting option

Treatment 4: Replacement of discount-oriented banners with non-discount oriented banners

In contrast to Field Experiment I, we assign the removal of discount markers and sorting options into two different treatments to separately measure the effects of each intervention. We also add a fourth treatment, the use of non-discount oriented banners throughout the site (see Figure W6 in the Web Appendix for examples of discount banners). In practice, discount-oriented banners serve a dual purpose: to communicate the existence of marked down items as well as a navigation tool to access the relevant product listings.

[Insert Table 3 about here.]

Results

Before assessing the impact of price sensitivity on shoppers' propensities to find and buy discounted items, we conduct the same analysis as in Table 1, which pools all types of consumers and present the results in Table 3. Further validating the main findings in Field

Experiment I, this time including existing in addition to new customers, we find that removing discount markers, sorting by discount and discount-oriented banners (Treatments 2 to 4) decreases both the average discount of items purchased as well as the incidence of purchasing items on discount.[12] As before, this is achieved without a decrease in conversion rates. An exception to this, and counter to the findings in Field Experiment I, is the null effect of the removal of outlet and sales links from the home page (Treatment 1). A possible explanation is that current customers were not as deterred as new shoppers from finding the high discounts in the outlet portion of the website, in the sense of differences in perceived search costs.[13] With the exception of this treatment, the addition of search costs impacts new and current customers in a qualitatively very similar manner.

[Insert Table 4 about here.]

A test of the remainder of our predictions entails measuring the interaction between shoppers' price sensitivity and their willingness to incur search costs to find discounted items. Using repeat customers, while changing the navigability of the website is, in our view, a stronger test of this prediction. In line with our theoretical assumptions, existing customers are more likely to recall that heavily discounted items exist on the platform. We classify customers according to their price sensitivity using historical transaction data. (Details are in Appendix A.) While using observational data as we do to infer price sensitivity has limitations, we note that any deficiency in predictive accuracy would bias our results toward a false negative in detecting heterogeneous treatment effects.

[12] The average discount across the three treatment groups was 17.0%, which is significantly different from the average discount in the control condition of 18.2% at p=0.0042.
[13] See Web Appendix C for a description of the difference in outcomes for new versus existing customers.

In Table 4, we group consumers into three quantiles according to their price sensitivity as indicated by predicted values from the full model (column 4 of Table A3 in the appendix). In an additional validation of our classification model, we find that low price sensitivity consumers are directionally more likely to purchase full price items across the board (first row). In three of four treatment conditions, we observe statistically significant increases in the proportion of full-priced items bought by customers with low price sensitivity. Equally notable, this is not the case for medium or high price sensitivity consumers, who willingly incur search costs to avail of discounts. This result provides additional evidence, by including current users and adding other forms of search costs to the website, that online retailers can improve their margins and, thus, profitability, by deliberately adding small frictions to the shopping process. We present actual profitability measures in Table W7 in the Web Appendix using item-level marginal cost data provided by the firm.

Our theoretical framework posits that higher conversion is a consequence of more products inspected in regions of moderate search costs and price sensitivity. Unfortunately, we lack the granular data required to precisely measure the number of products inspected by consumers. We do, however, have access to aggregate data on browsing behavior available through the firm's web analytics provider. Table 5 presents key measures of browsing behavior across each treatment group. On average, visitors falling within each treatment group visited the website more times during the testing period, spent more time during each visit, and viewed more pages. Longer visit times, more pages viewed, and more sessions per user are consistent with our proposed mechanism for employing search frictions. Specifically among price-sensitive individuals, locating discounted items would necessarily involve more time spent on the website in the treatment conditions. We conjecture that these changes in browsing behavior may have concurrently resulted in shoppers considering a broader set of products, thus leading to higher conversion rates.

[Insert Table 5 about here.]

Caveats and Limitations

Our two field experiments provide convergent evidence that increasing search frictions may be profitable for a monopolistic online seller in the short run. While we replicate our results under different demand conditions and consumer profiles, natural questions remain pertaining to the generalizability of our results to other online retail contexts and the long-term viability of increasing search frictions. In this subsection we provide suggestive evidence that points to probable answers to some of these questions, and to potential directions for future research.

First, we consider the conjecture that increasing search frictions may cause consumers to switch to alternatives in the presence of competing firms. While our data provider has no similar-sized competitors online, they offers items under a private label in addition to items from national and global brands. The firm faces less competition for its private label products, which are available only on their online store, than for the rest of its products, which are available from offline sellers. We examine the outcomes within each of these categories in an effort to find evidence of heterogeneous treatment effects, which would be indicative of the role of competition. If consumers are prone to switching to competitors upon facing added search frictions, then they should be less likely to buy national and global branded products in the treatment conditions. Table 6 presents the portion of items sold in each treatment group that are store brand products.

[Insert Table 6 about here.]

We find that in none of the treatment conditions is the ratio of store brand to national/global brand items sold significantly higher than in the control, which suggests that the presence of competitors for branded products does not play a significant role in mitigating our results within this setting.

We turn our attention to the potential long run effects of increasing search frictions in an online store. Transaction and web browsing data are available for several periods beyond the end of our experimentation, and may point to specific long run effects. We consider cohorts of consumers falling into each of the randomly assigned groups from Field Experiment II and track their revisit and conversion rates over time. Table W8 in the Web Appendix shows the number of consumers who visited the website during our sample period of two weeks in 2016. For each half-month period thereafter, until the end of the year, we list the number of users from the original sample who visit the website and their corresponding conversion rate.

We find no evidence that consumers in the treatment conditions had lower revisit rates than that in the control. In fact, for every single month and treatment group, the revisit rate is higher among the treatment cohorts than in the control cohorts at $p < 0.01$. Qualitatively similar measurements are obtained from groups in Field Experiment I. This is perhaps unsurprising given that the conversion rates were higher in the treatment groups rather than in the control condition.

CONCLUSION

Online retail represents a rapidly growing proportion of overall retail sales. However, margins in online retail can often be smaller than in offline retail. One conjecture for this discrepancy is that online sellers are less able to price discriminate compared to their offline counterparts. In this paper, we explore how imposing additional search costs on online shoppers can improve gross margins by increasing the number of items inspected and serving as a sorting mechanism among customers.

We find encouraging evidence that minor changes to the design of an online store can substantially improve its margins and profitability. By increasing search frictions in simple ways—removing selected links, narrowing down product sorting options, and limiting visual markers—online sellers can achieve more full priced sales from price-insensitive shoppers who face higher search costs. While substantial enough to affect price-insensitive shoppers' purchasing behavior, these search frictions are minor enough such that price-sensitive shoppers exert the extra effort to locate discounted items. As a result, the average selling price increases due to a higher proportion of full priced items sold.

Our theoretical model shows how increasing search frictions may either increase or decrease conversion rates depending on consumer preferences. Through our field experiments, we find that conversion rates increase upon the addition of search frictions in a typical online retail setting. Inspecting browsing behavior in the treatment groups suggests that as visitors spend more time on the website given higher search frictions, they may also be considering a larger set of products.

Our results have direct implications for online sellers. Without changes to product prices or assortment, online sellers can improve their margin performance by implementing subtle changes to website design. We note with particular excitement that this is an essentially no-cost manipulation with low data requirements, and one that can deliver large gains in margin. Our findings imply that by indiscriminately prioritizing ease of search and purchase, online sellers may be giving up gross margins by unwittingly giving away discounts to price-insensitive consumers and curtailing consumer exploration of the product assortment.

Our research also suggests that some online browsing behavior can be effortful or time-consuming enough for shoppers to prefer paying higher prices. We consider it a fruitful area for future research to determine which specific properties of online interaction consumers find most effortful. This can provide helpful guidance for a wide array of applications, from online store design to digital advertising.

REFERENCES

Armstrong, M. (2017). Ordered consumer search. *Journal of the European Economic Association*, *15*(5), 989-1024.

Bakos, J. Y. (1997). Reducing buyer search costs: Implications for electronic marketplaces. *Management Science*, *43*(12), 1676-1692.

Bell, D. R., Ho, T. H., & Tang, C. S. (1998). Determining where to shop: Fixed and variable costs of shopping. *Journal of Marketing Research*, 352-369.

Brynjolfsson, E., & Smith, M. D. (2000). Frictionless commerce? A comparison of Internet and conventional retailers. *Management Science*, *46*(4), 563-585.

Choi, H., & Mela, C. F. (2016). Online marketplace advertising.

Coughlan, A. T., & Soberman, D. A. (2005). Strategic segmentation using outlet malls. *International Journal of Research in Marketing*, *22*(1), 61-86.

De los Santos, B., & Koulayev, S. (2017). Optimizing Click-Through in Online Rankings with Endogenous Search Refinement. *Marketing Science*, *36*(4), 542-564.

Ellison, G., & Fisher Ellison, S. (2005). Lessons about Markets from the Internet. *Journal of Economic Perspectives*, *19*(2), 139-158.

Ellison, G., & Ellison, S. F. (2009). Search, obfuscation, and price elasticities on the internet. *Econometrica*, *77*(2), 427-452.

Ellison, G., & Wolitzky, A. (2012). A search cost model of obfuscation. *The RAND Journal of Economics*, *43*(3), 417-441.

Gu, Z., & Liu, Y. (2013). Consumer fit search, retailer shelf layout, and channel interaction. *Marketing Science*, *32*(4), 652-668.

Jing, B. (2015). Lowering Customer Evaluation Costs, Product Differentiation, and Price Competition. *Marketing Science*, *35*(1), 113-127.

Kim, B. D., Srinivasan, K., & Wilcox, R. T. (1999). Identifying price sensitive consumers: the relative merits of demographic vs. purchase pattern information. *Journal of Retailing*, *75*(2), 173-193.

Kopalle, P., Biswas, D., Chintagunta, P. K., Fan, J., Pauwels, K., Ratchford, B. T., & Sills, J. A. (2009). Retailer pricing and competitive effects. *Journal of Retailing*, *85*(1), 56-70.

Lal, R., & Bell, D. E. (2003). The impact of frequent shopper programs in grocery retailing. *Quantitative Marketing and Economics*, *1*(2), 179-202.

Lambrecht, A., & Skiera, B. (2006). Paying too much and being happy about it: Existence, causes, and consequences of tariff-choice biases. *Journal of Marketing Research*, *43*(2), 212-223.

Leslie, P. (2004). Price discrimination in Broadway theater. *RAND Journal of Economics*, 520-541.

Lynch Jr, J. G., & Ariely, D. (2000). Wine online: Search costs affect competition on price, quality, and distribution. *Marketing Science*, *19*(1), 83-103.

McManus, B. (2007). Nonlinear pricing in an oligopoly market: The case of specialty coffee. *The RAND Journal of Economics*, *38*(2), 512-532.

Monga, A., & Saini, R. (2009). Currency of search: How spending time on search is not the same as spending money. *Journal of Retailing*, *85*(3), 245-257.

Mussa, M., & Rosen, S. (1978). Monopoly and product quality. *Journal of Economic Theory*, *18*(2), 301-317.

Narasimhan, C. (1984). A price discrimination theory of coupons. *Marketing Science*, *3*(2), 128-147.

Narayanan, S., & Kalyanam, K. (2015). Position effects in search advertising and their moderators: A regression discontinuity approach. *Marketing Science*, *34*(3), 388-407.

Ngwe, D. (2017). Why outlet stores exist: Averting cannibalization in product line extensions. *Marketing Science*, *36*(4), 523-541.

Nichols, D., Smolensky, E., & Tideman, T. N. (1971). Discrimination by waiting time in merit goods. *The American Economic Review*, *61*(3), 312-323.

Ratchford, B. T. (2009). Online pricing: review and directions for research. *Journal of Interactive Marketing*, *23*(1), 82-90.

Rigby, D. (2011). The future of shopping. *Harvard Business Review*, *89*(12), 65-76.

Sahni, N. S., Wheeler, S. C., & Chintagunta, P. (2018). Personalization in Email Marketing: The Role of Noninformative Advertising Content. *Marketing Science*.

Seiler, S., & Pinna, F. (2017). Estimating search benefits from path-tracking data: measurement and determinants. *Marketing Science*, *36*(4), 565-589.

Stahl, D. O. (1996). Oligopolistic pricing with heterogeneous consumer search. *International Journal of Industrial Organization*, *14*(2), 243-268.

Ursu, R. (2017). The power of rankings: Quantifying the effect of rankings on online consumer search and purchase decisions. *Marketing Science*, forthcoming.

U.S. Census Bureau News. 2018. https://www2.census.gov/retail/releases/historical/ecomm/17q4.pdf. Accessed March 19, 2018.

Van Heerde, H. J., Gijsbrechts, E., & Pauwels, K. (2008). Winners and losers in a major price war. *Journal of Marketing Research*, *45*(5), 499-518.

Varian, H. R. (1980). A model of sales. *The American Economic Review*, *70*(4), 651-659.

Weitzman, M. L. (1979). Optimal search for the best alternative. *Econometrica: Journal of the Econometric Society*, 641-654.

TABLES

Table 1: Results of Field Experiment I

Group	Sample size	Average discount of purchased items	Percent full-priced purchases	Average selling price (Philippine Pesos)
Control	26,014	15.74% (0.81)	59.77%	751.75 (27.20)
Treatment 1	26,199	11.54%*** (0.71)	67.48%***	907.05*** (42.08)
Treatment 2	26,343	12.32%*** (0.69)	64.19%*	1,137.07*** (64.44)
Treatment 3	26,049	11.49%*** (0.75)	66.04%**	1,035.87*** (61.45)

H_0: value is equal to that in the control condition. *** p<0.01, ** p<0.05, * p<0.1

Table 2: Basket level results from Field Experiment I

Group	Average discount	Average basket size	Transactions completed	Number of items
Control	14.53% (1.01)	1,609.87 (74.15)	318	2.14 (0.14)
Treatment 1	11.56%** (0.88)	1,775.78 (115.74)	355*	1.96 (0.08)
Treatment 2	12.23%** (0.86)	2,280.55*** (203.33)	355*	2.01 (0.08)
Treatment 3	11.96%* (0.88)	1,826.73 (158.13)	334	1.76*** (0.06)

H_0: value is equal to that in the control condition. *** p<0.01, ** p<0.05, * p<0.1

Table 3: Results of Field Experiment II

Treatment Group	Number of visitors	Average discount of sold items	Percent of items bought at full price	Number of transactions (Conversion rate)
Control	68,343	18.25% (0.42)	48.83%	1,351 (1.98%)
Treatment 1 No outlet and sales links	70,058	17.32%* (0.37)	50.15%	1,599 (2.28%)***
Treatment 2 No discount markers	70,025	16.69%*** (0.36)	51.78%**	1,605 (2.29%)***
Treatment 3 No discount sorting	69,859	17.09%** (0.37)	51.47%**	1,605 (2.30%)***
Treatment 4 No discount banners	69,825	16.80%*** (0.35)	52.58%***	1,713 (2.45%)***

H_0: value is equal to that in the control condition. *** p<0.01, ** p<0.05, * p<0.1

Table 4: Proportion of items bought at full price

Price sensitivity	Control	Treatment 1 No outlet and sale links	Treatment 2 No discount markers	Treatment 3 No discount sorting	Treatment 4 No discount banners
Low	58.7%	67.8%**	66.6%**	63.9%	67.5%***
Medium	54.0%	52.1%	57.0%	53.1%	57.0%
High	36.3%	40.8%*	38.4%	32.6%	33.2%

H_0: value is equal to that in the control condition. *** p<0.01, ** p<0.05, * p<0.1

Table 5: Browsing behavior

Treatment group	Avg. sessions per user	Avg. session duration (seconds)	Avg. pages viewed per session
Control	2.13	438.99	7.37
Treatment 1: N**o outlet and sale links**	2.21	457.44	7.59
Treatment 2: N**o discount markers**	2.27	456.65	7.59
Treatment 3: N**o discount sorting**	2.26	458.63	7.65
Treatment 4: N**o discount banners**	2.27	459.25	7.63

Disaggregated data and standard deviations are unfortunately not available, precluding hypothesis testing. A session is defined as a set of interactions, no two of which occur more than 30 minutes apart.

Table 6: Store branded products as percent of items sold

	No. of store brand items sold as percent of total number of items sold	p-value
Control	18.97	
Treatment 1: No outlet and sale links	18.88	0.9242
Treatment 2: No discount markers	16.71	0.0224
Treatment 3: No discount sorting	17.22	0.0809
Treatment 4: No discount banners	19.14	0.8639

H_0: value is equal to that in the control condition.

Figure 1: Simulated results given variation in search cost

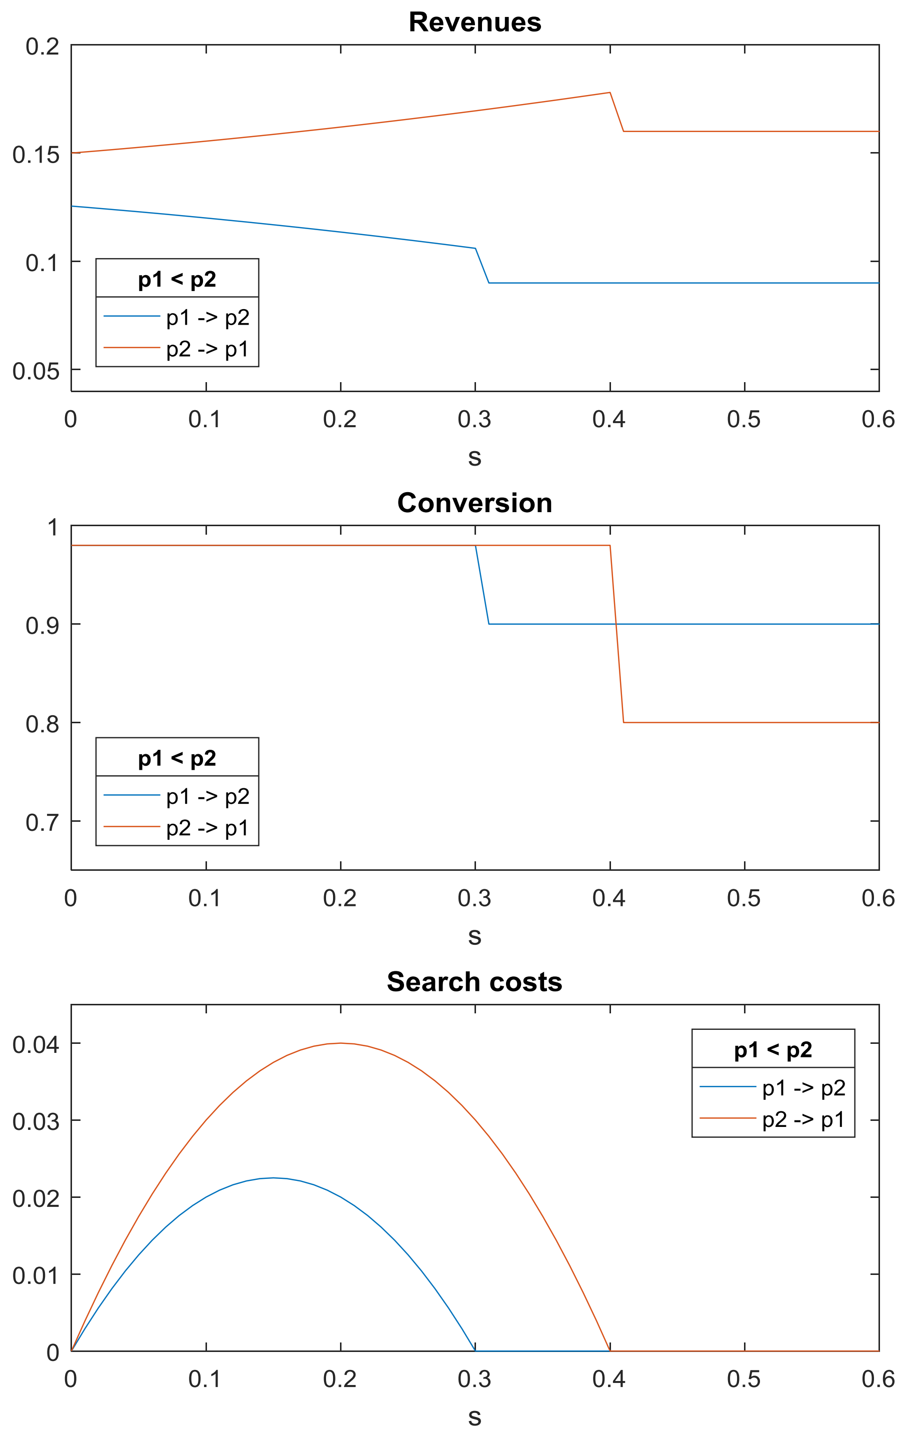

Figure 2: Simulated results given variation in price sensitivity

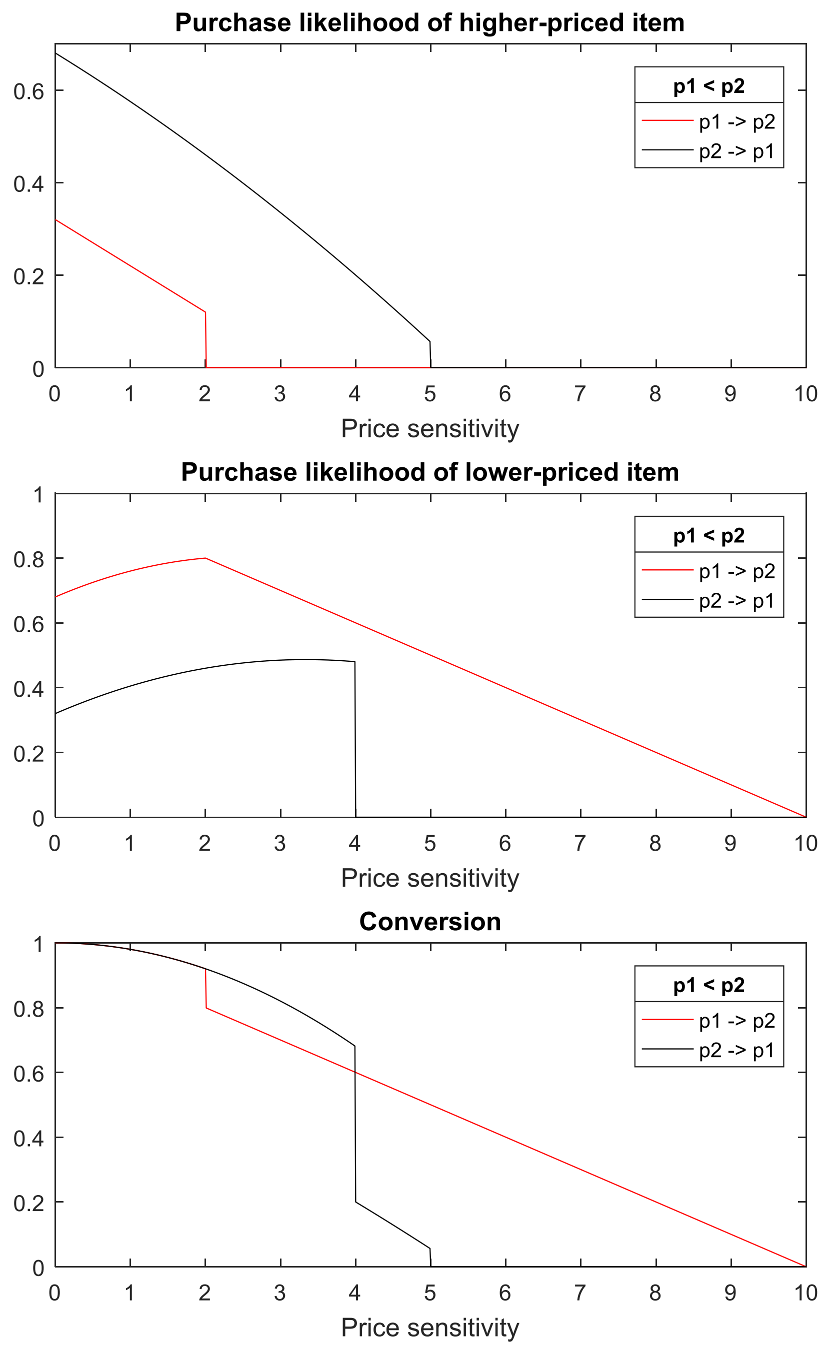

APPENDIX A: Approximating Price Sensitivity

Our theoretical framework predicts differential effects of search frictions on price-insensitive versus price-sensitive shoppers. The predicted impact on purchase behavior should disproportionally affect the former. We develop a parsimonious empirical model of price sensitivity for shoppers in our setting. Its purpose it to provide us with a means of classifying consumers according to their baseline appetite for discounts. We use the predicted values of this model as a proxy for price sensitivity. After estimating the model, we assess its predictive accuracy by comparing the behavior of pre-classified groups of shoppers in a validation field experiment.

Data

The data for this analysis consists of the retailer's historical transaction-level sales records from its inception in September 2012 until September 2015. Over 2.5 million individual items were sold whereby 418,039 consumers made over one million transactions during that period. Each record (an item sold) contains shopper attributes, product attributes, and transaction attributes. Tables A1 and A2 provide a description of the available data and basic transaction-level summary statistics.

Table A1: Classification data set summary statistics

Start date	3 September 2012
End date	30 September 2015
Records (items sold)	2,609,421
Transactions	1,099,683
Unique customers	418,039
Unique items	547,574
Unique brands	2,380

Table A2: Transaction summary statistics

	Mean	**S.D.**
Item selling price (in Ph. Pesos)	651.69	700.70
Item original price (in Ph. Pesos)	903.23	1,017.00
Discount percent	15.94	22.44
Items per transaction	2.37	2.50
Basket size (in Ph. Pesos)	1,546.38	1,940.27

Model

We estimate a simple model of price sensitivity in order to determine through a succeeding field experiment whether price sensitive shoppers are more willing to bear search costs to locate discounted items online. Since the primary objective of estimation is not to identify primitives of consumer utility, but merely to discriminate between price-insensitive and price-sensitive shoppers, we adopt a parsimonious model that aims to explain the basket-level discount of completed transactions. The underlying assumption is all else equal, that highly price-sensitive shoppers are more likely to purchase discounted items than price-insensitive shoppers. Categories of explanatory variables for price sensitivity include demographic characteristics, prior transaction behavior, and shopping conditions known to be associated with discount-seeking behavior.

These variables were chosen based on availability and management's expectation of their relationship to discount purchasing. We run a series of Tobit regressions of most recent average basket discount on these covariates and present estimates in Table A3. As per prior literature, we use a Tobit model given that discounts are a left-censored (at zero) proxy for price-sensitivity, our conceptual variable of interest (Lambrecht & Skiera 2006; Van Heerde, Gijsbrechts & Pauwels 2008). In order to evaluate the relative importance of demographics, prior transaction behavior, and current shopping conditions, we estimate separate regressions for each subcategory of explanatory variables in addition to the full model.

Results

In general we find that relationships between consumer attributes, prior shopping behavior, current shopping conditions, and current shopping behavior are strong and robust to the usage of different choices of covariates. Each category of explanatory variables (corresponding to columns in Table A3) improves the ability of the model to predict preference for discounts. Observed discounts are lower for men and older customers. They are higher for customers who have previously bought at higher discounts, used coupons, and bought more store branded items. Meanwhile, discounts are lower for consumers who redeem coupons in the current purchase instance, use store credit, and have more previously completed transactions.

We use the empirical model estimated in this section to pre-classify shoppers according to their levels of price sensitivity in order to articulate the mechanism behind our main result in

Field Experiment 1. In effect we use all of the available information on consumers to achieve this classification, assigning weights to each variable according to its estimated coefficient. We consider this to be an improvement over an ad hoc classification, say, by grouping shoppers according to the average discount in their purchase histories. However, we also recognize the shortcomings of this approach owing to the aggregation of information, the lack of information on visits that result in no purchase, and the changing assortment over time. In order to increase our confidence in the resulting classification, we seek to establish its external validity. In the next section, we describe how we validate our classification model by measuring responses to email newsletters in a field experiment.

Table A3: Tobit regression for price sensitivity

	Variables	1	2	3	4
Model	Male	-0.254***			-0.197***
		(0.009)			(0.011)
	Age	-0.126***			-0.111***
		(0.011)			(0.013)
	Prev average discount		0.396***		0.392***
			(0.002)		(0.002)
	Prev coupon usage[14]		0.325***		0.458***
			(0.009)		(0.009)
	No. of previous transactions		-0.645***		-0.500***
			(0.049)		(0.049)
	Prev store brand ratio		0.170***		0.125***
			(0.013)		(0.012)
	Time since first purchase		0.174***		0.149***
			(0.021)		(0.020)
	Store credit dummy			0.150***	-0.323**
				(0.014)	(0.147)
	Coupon dummy			-0.336***	-0.544***
				(0.008)	(0.010)
	Constant	0.141***	-0.166***	0.704***	0.278***
		(0.007)	(0.008)	(0.013)	(0.100)
Sigma	Constant	0.331***	0.310***	0.329***	0.307***
		(0.000)	(0.000)	(0.000)	(0.000)
	Billing region FE	yes			yes
	Month FE			yes	yes
	Observations	1,112,297	698,456	1,112,298	698,456
	Pseudo R2	0.00324	0.0576	0.0105	0.0729

Note: *** p<0.01, ** p<0.05, * p<0.1

[14] The dependent variable is the discount before any coupons; hence, this coefficient implies substitution by consumers between product-specific discounts and coupon-based discounts.

APPENDIX B: Validation Experiment

We test the external validity of our empirical model of price sensitivity. We send discount and non-discount oriented email newsletters to randomly assigned consumers and test whether the classification determined by the Tobit model of the previous section is indeed associated with a higher response rate for discount (versus non-discount) emails for price-sensitive (versus price-insensitive) shoppers.

Experimental Design

We include the firm's entire mailing list of 246,688 consumers in this experiment. A consumer gets on the mailing list by providing his or her email address to the firm through registering an account, signing up for updates, or requesting a coupon. Consumers were randomly assigned to two groups, henceforth Group 1 and Group 2. Each group received a schedule of both discount- and non-discount oriented newsletters as presented in Table B1 in order to counteract day-of-week effects.[15]

Table B1: Schedule of newsletter treatments

	Group 1 (50%)	**Group 2 (50%)**
Sunday	Control	Control
Monday	Discount	Full price
Tuesday	Discount	Full price
Wednesday	Discount	Full price
Thursday	Full price	Discount
Friday	Full price	Discount

[15] This experiment ran from October 4, 2015 through October 9, 2015.

The control newsletters sent out on Sunday were non-discount oriented and identical between groups whereas, within each day, only the discount versus non-discount messaging was different between groups. For each newsletter sent, we observe whether the email was opened and which link within the email, if any, was clicked by the recipient. We also observe all transactions on the website, which we can link to consumers in the experiment via their email addresses.

Product categories featured on the email newsletters varied between days, but were kept constant between control and treatment groups within days. Care was also taken to keep all creative elements on the newsletters constant, such that only discount messaging (e.g. “up to 40% off”) and any price information varied in the execution. This variation in messaging was also reflected in the subject line. For an example of a discount and full price email used, see Figure W5 in the Web Appendix.

Customers on the mailing list vary in the frequency with which they opt to receive newsletters. The breakdown is that 62.34% of subscribers receive them every day, 3.32% receive them three times a week, and 34.35% receive them once a week. The newsletter schedule shown in Table B1 was designed to gain maximum variation among consumers regardless of their frequency as well as minimize any day-of-week effects.[16]

In order to establish the validity of the classification of the Tobit model, we generate predicted values from the model given each consumer’s purchase history prior to the newsletter experiment in this paper.[17] We argue that our model is indicative of price sensitivity if consumers

[16] The original plan included a pair of newsletters for Saturday that followed Friday’s pattern; however due to a server failure the emails were never sent out.
[17] Only consumer-specific and purchase history variables (male dummy, customer age, previous discount, previous coupon dummy, number of previous transactions, store brand purchase ratio, time since first purchase) will be

we predict to be more price-sensitive have a higher propensity to open and click on discount-oriented versus non-discount-oriented emails. Since the final experiment will compare average shopping behavior across groups of consumers, price-sensitive and insensitive, the classification model only needs to be accurate at the group level rather than at an individual level, thus the choice of a parsimonious model specification.

Results

We regress newsletter outcome variables on our variables of interest. Each record in the following regressions is an email-customer pair. The dependent variables are binary, where success is either an opened or a clicked email. The independent variables are: a discount email dummy, predicted price sensitivity from the empirical model, the interaction between discount email and price sensitivity, and day of the week.[18]

Table B2: Regressions on newsletter response variables

VARIABLES	**(1)** **open**	**(2)** **click**	**(3)** **click | open**
Constant	0.207***	0.0328***	0.152***
	(0.00630)	(0.00324)	(0.0141)
Discount dummy*Price_sensitivity	0.125***	0.0505**	0.0728
	(0.0420)	(0.0216)	(0.0933)
Discount dummy	-0.0149*	-0.00103	0.0188
	(0.00811)	(0.00417)	(0.0181)
Price_sensitivity	0.0402	0.0248	0.0868
	(0.0297)	(0.0153)	(0.0671)
Day dummies	yes	yes	yes
Observations	106,534	106,534	22,043
R-squared	0.002	0.001	0.003

Standard errors in parentheses. *** p<0.01, ** p<0.05, * p<0.1

included in the prediction. The predicted values themselves will have no direct interpretation, but will be treated as sufficient statistics for price sensitivity.

[18] We include only consumers with at least one past purchase in this analysis. Hence, consumers in this regression are a subset of consumers in the classification analysis in Appendix A—namely those that have not opted out of receiving email newsletters.

We find that our measure of price sensitivity is significantly positively correlated with the probability of responding (e.g., click or open) to a discount-oriented email relative to a non-discount oriented email. Table B2 considers the relationship between price sensitivity and three response variables: (1) whether the email was opened, (2) whether a link in the email was clicked, and (3) whether a link was clicked conditional on the email being opened. We find that price sensitivity is positively associated with the first two measures.[19]

Figure B1 provides evidence of the relationship between our constructed price sensitivity measure and the likelihood of responding to each type of newsletter. There is a clear positive relationship between price sensitivity, as measured in decile membership, and the likelihood of clicking a discount newsletter relative to the probability of clicking a non-discount newsletter.

Figure B1: Newsletter click counts by price sensitivity decile

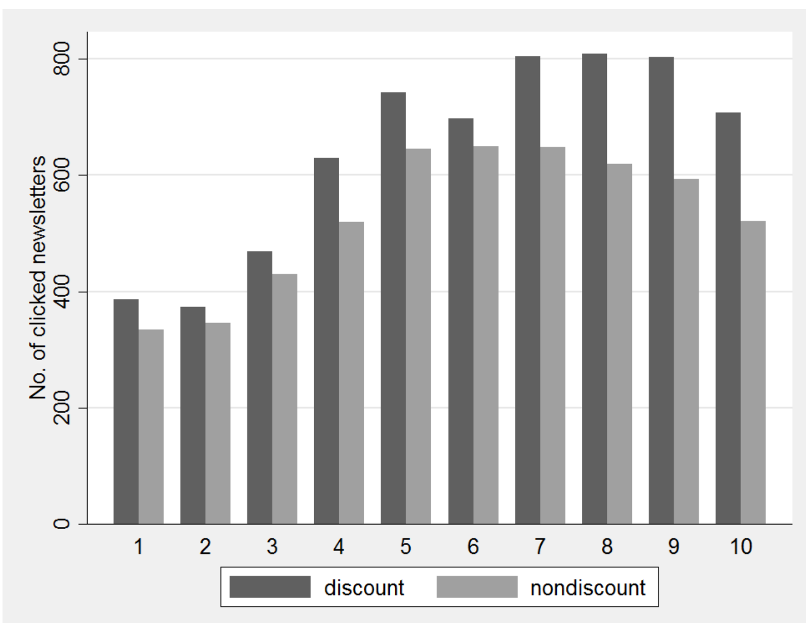

Note: Higher deciles correspond to higher price sensitivity.

[19] We include the third measure, click conditional on open, for completeness; however because random assignment is not preserved for this measure the associated parameters are not relevant for our purposes.

Web Appendix of

"The Impact of Increasing Search Frictions on Online Shopping Behavior: Evidence from a Field Experiment"

WEB APPENDIX A: Variation in non-price attributes

Our theoretical framework assumes the same distribution of non-price attributes for different price and discount levels. While not crucial to our main results, variation in non-price attributes may influence the results from our field experiments. For instance, if by re-ordering search paths we expose consumers to products that are of higher quality in addition to being higher priced, this may cause a further increase in purchase probabilities overall and of full priced products in particular. Our empirical setting of fashion and apparel retail limits our ability to generate non-price variables that can reliably be thought of as corresponding to product quality. However, the data include brand variables as well as indicators for whether an item was returned. We consider these variables as proxies for product quality.

We begin exploring this possibility by measuring the correlation between discounts and non-price attributes in the data. For this analysis, we use the entire data set of transactions available to us, excluding periods covered by our field experiments. To avoid conflating category-specific factors we focus our analysis on shoes, the most commonly purchased category.

Table W1: Average discount percentage (top entries) and transaction counts (bottom entries) by brand and return status

	Store brand	Not store brand	Total
Returned	24% 16,603	22% 59,048	22% 75,651
Not returned	26% 199,534	20% 852,845	21% 1,052,379
Total	25% 216,137	20% 911,893	21% 1,128,030

Inspecting average discounts shows that store brand shoes have higher discounts on average (statistically significant with $p<0.01$), by about a five percentage point spread. Return status seems to have a less straightforward relationship with discounts on average. Store brands are widely considered to be less desirable than branded items, and such is the case for our partner firm. Therefore, there does seem to be a relationship between quality and discounts in at least one dimension.

In order to gauge the impact of this variation on our findings, we split our purchase outcome measures by store-brand and non-store brand purchases and present counts in Table W2. Most relevantly for our research objectives, we find that there are more items sold for both store brand and non-store brand items in each of our treatment conditions relative to the control. This implies that the dominant impact is the increase in purchase likelihood as a results of more items being inspected by the average consumer, rather than a substitution from store brand to non-store brand items.

Table W2: Store brand and non-store brand products sold in Field Experiment II

Treatment Group	Number of items sold	
	Store brand	**Non-store brand**
Control	514	2,195
Treatment 1 No outlet and sales links	546	2,722
Treatment 2 No discount markers	622	2,673
Treatment 3 No discount sorting	546	2,625
Treatment 4 No discount banners	694	2,931

WEB APPENDIX B: Distribution of product discounts

We present histograms that show the distribution of product discounts from both of our field experiments. In each subplot of Figure W1, we pool observations from the treatment conditions and overlay the histogram with that of the control condition. In both field experiments, substitution between control and treatment occurs primarily between full priced purchases and discounted purchases more generally, with little substitution from higher discounts to lower discounts otherwise.

Figure W1: Histograms of product discounts

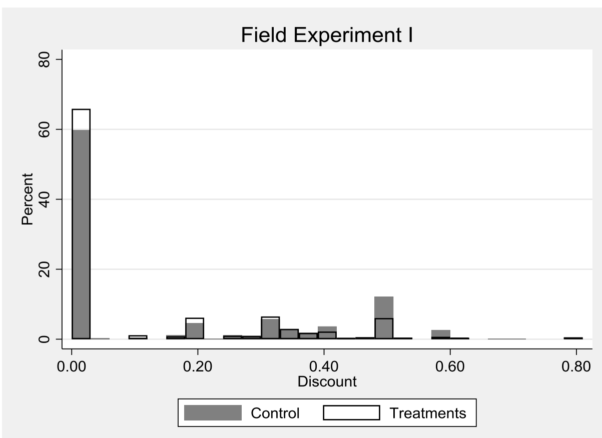

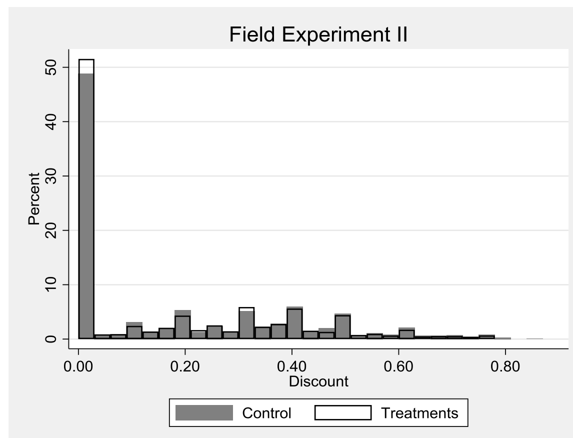

WEB APPENDIX C: New and Existing Customers in Field Experiment II

Including both new and existing customers in Field Experiment II allows us to explore possible heterogeneous treatment effects. While our theoretical framework does not provide sharp predictions in this regard, there is a reasonable argument that our assumptions on consumers' information sets—that consumers are aware of the distribution of match values and product prices—are more appropriate for existing than for new customers. One may also conjecture that search costs are intrinsically lower for existing customers, who are more familiar with the website and are selected based on having made a prior purchase.

Figure W2 graphs the share of full-priced purchases by old and new customers in each treatment group. We observe that both old and new customers buy more full-priced items in the treatment groups than in the control group, which is directionally consistent with our model's predictions. We also observe that there exists no consistent pattern in the relative effects between groups within each treatment condition. The differences we observe are potentially artifacts of the specific experimental manipulations in each condition; however, given that we lack the browsing data required to fully explore these differences and that this is not central to our research objectives, we leave an investigation of these differences for future research.

Figure W2: Full priced purchases by old and new customers

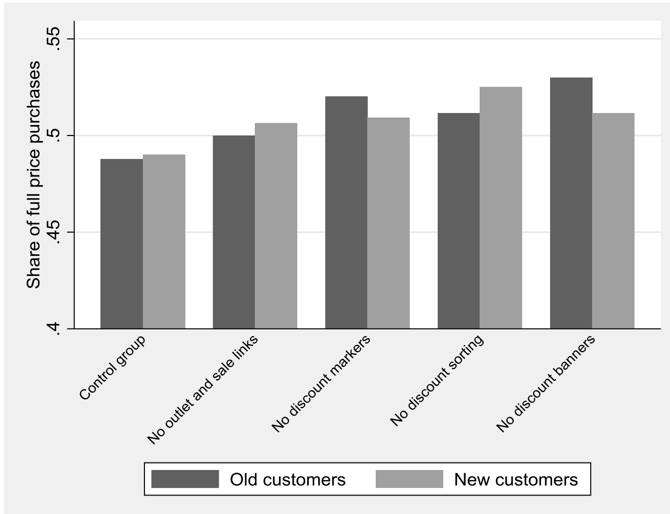

WEB APPENDIX D: Additional Tables and Figures

Table W3: Randomization checks for Field Experiment I

	Sample size	Ages 18-24	Chrome browser	Windows system	Visited on Thursday
Control	26014	4234	19538	23931	6319
Treatment 1	26199	4268	19616	23985	6452
Treatment 2	26343	4284	19732	24116	6348
Treatment 3	26049	4380	19622	23893	6420
p-value		0.2469	0.6042	0.2079	0.3845

H_0: Proportions are equal across all conditions.

Table W4: Margin measurements in Field Experiment I

	No. of items sold	Average product margin	p-value	Gross margin
Control	681	141.68 (9.14)		96,483.48
Treatment 1	695	174.20 (13.23)	0.0441	121,070.70
Treatment 2	712	287.20 (28.09)	0.0000	204,488.30
Treatment 3	589	211.00 (28.49)	0.0143	124,280.30

H_0: value is equal to that in the control condition. Gross margin is the sum of product margin over items sold in each condition.

Table W5: Randomization checks for validation experiment

	Sample size	No. of Female	Average Age	CLV	No. of transactions	Newsletter frequency
Group 1	123,346	86,032	31.22	6,169.42	3.92	1.72
Group 2	123,342	86,166	31.13	6,015.13	3.82	1.72
p-value		0.5486	0.0698	0.3875	0.2235	0.9968

H_0: value is equal to that in the control condition. CLV is computed as the sum of all previous basket sizes.

Table W6: Randomization checks for Field Experiment II

	Sample[20]	No. of new users	No. of ages 18-24	Used Chrome browser	Used Windows system	Visited on Thursday
Control	68,343	42,692	13,396	49,893	53,996	16,388
Treatment 1 No outlet and sale links	70,058	43,679	13,629	51,025	55,675	16,424
Treatment 2 No discount markers	70,025	43,549	13,645	50,673	55,423	16,465
Treatment 3 No discount sorting	69,859	43,734	13,696	50,722	55,315	16,435
Treatment 4 No discount banners	69,825	43,740	13,624	50,736	55,199	16,482
p-value		0.38800	0.9342	0.0874	0.2381	0.1425

H_0: Proportions are equal across five groups.

[20] The null that the number of subjects in the control condition is not equal to 20% of the total number of subjects, or that the number of subjects in the five groups are not equal, cannot be rejected at p<0.01. We have investigated possible sources for sampling discrepancies with the firm and its third-party testing platform but have found no satisfying explanations. Tests of balance find no significant differences between groups in available variables; hence, we argue that any systematic departures from assignment with equal probabilities has not resulted in observably selected samples.

Table W7: Profitability for Experiment #3

	No. of items sold	Average product margin	p-value	Average percent margin	p-value	Gross margin
Control	2,709	225.64 (6.33)		0.23		611,246.31
Treatment 1 No outlet and sales links	3,268	215.51 (5.01)	0.2043	0.25	0.0159	704,302.19
Treatment 2 No discount markers	3,295	218.23 (5.70)	0.3840	0.2472	0.0883	719,059.63
Treatment 3 No discount sorting	3,171	221.49 (5.65)	0.6245	0.2464	0.0959	702,355.94
Treatment 4 No discount banners	3,625	220.78 (5.46)	0.5616	0.2527	0.0119	800,344.38

H_0: Value is equal to that in the control condition.

Table W8: Repeat visits by consumers in the treatment conditions

	Control		*No outlet and sale links*		*No discount markers*		*No discount sorting*		*No discount banners*	
Dates	Visits	Conv rate	Visits	Conv rate	Visits	Conv rate	Visits	Conv rate	Visits	Conv rate
Jun 1 - Jun 15	68343	1.98%	70058	2.28%	70025	2.29%	69859	2.30%	69825	2.45%
Jun 16 - Jun 30	14157	5.88%	15205	6.15%	15337	5.73%	15167	5.70%	15302	7.07%
Jul 1 - Jul 15	10461	6.10%	11649	6.03%	11393	6.42%	11281	6.60%	11497	7.44%
Jul 16 - Jul 31	9058	9.03%	9870	8.11%	9956	7.63%	10037	7.91%	9962	9.13%
Aug 1 - Aug 15	7560	7.34%	8290	7.25%	8308	6.91%	8295	7.15%	8501	7.15%
Aug 16 - Aug 31	6367	7.81%	6815	7.18%	7025	6.28%	6918	6.79%	7091	7.63%
Sep 1 - Sep 15	5288	6.92%	5635	6.87%	5928	6.55%	5921	7.52%	5905	7.50%
Sep 16 - Sep 30	4922	6.85%	5306	6.20%	5553	6.86%	5569	6.12%	5520	8.13%
Oct 1 - Oct 15	4549	6.70%	4945	7.46%	5073	7.71%	5045	7.49%	5101	8.10%
Oct 16 - Oct 31	4364	7.03%	4658	7.11%	4784	7.48%	4771	7.76%	4755	7.38%
Nov 1 - Nov 15	3824	6.88%	4065	7.23%	4186	7.12%	4275	5.89%	4253	7.05%
Nov 16 - Nov 30	2931	6.28%	3308	7.04%	3414	5.27%	3342	5.69%	3294	7.13%
Dec 1 - Dec 15	3654	8.95%	4019	8.43%	4002	8.67%	4191	7.92%	3992	8.37%
Dec 16 - Dec 31	1764	6.58%	1889	5.93%	1985	6.25%	1980	6.92%	1970	5.43%

For all subsequent time periods after the duration of the experiment (June 1-15, 2016), the proportion of consumers revisiting is significantly higher in each treatment than in the control with $p < 0.01$. Conversion rates are the number of transactions divided by the number of visits. Note that the website reverts to the control condition after June 15, 2016 for all consumers.

Figure W3: Focal navigation elements in Field Experiment I

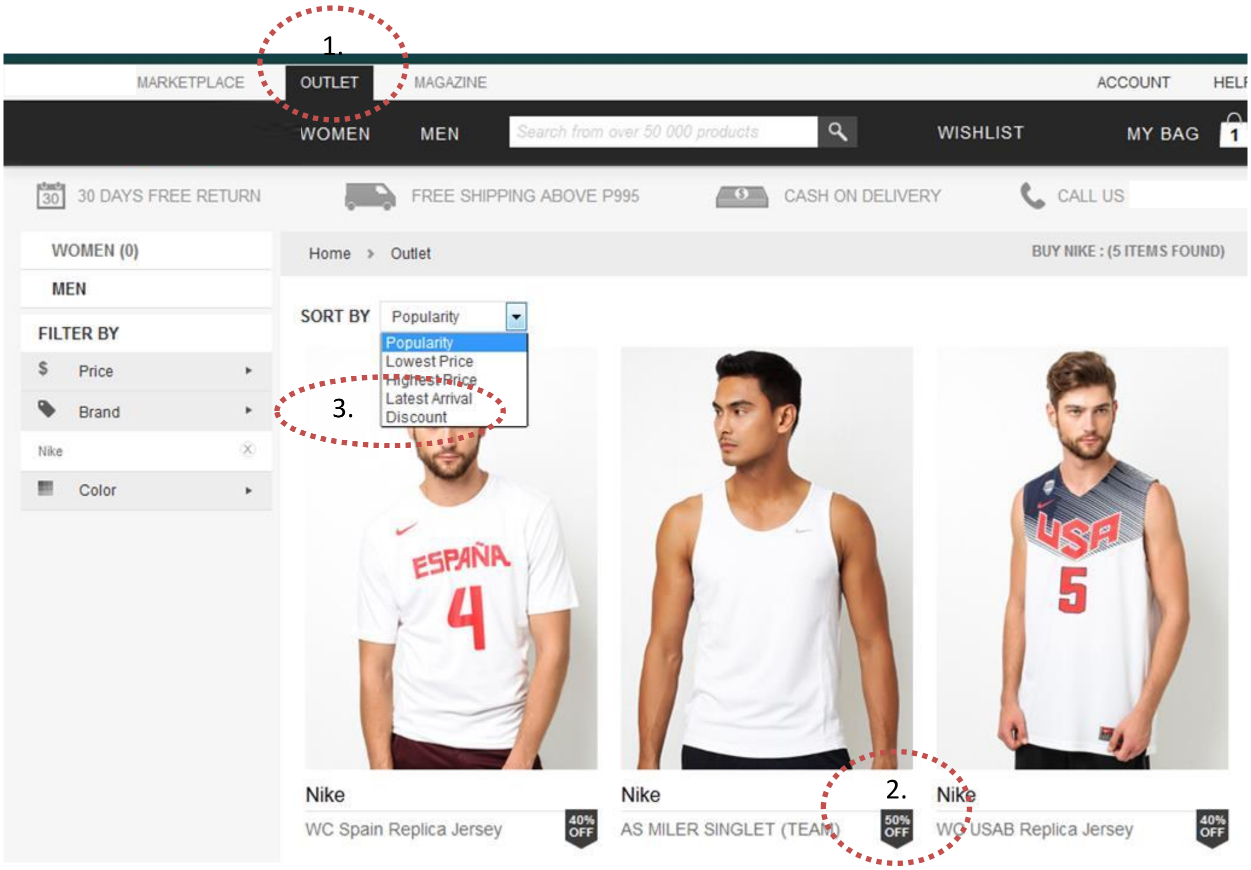

Figure W4: Regular newsletter schedule

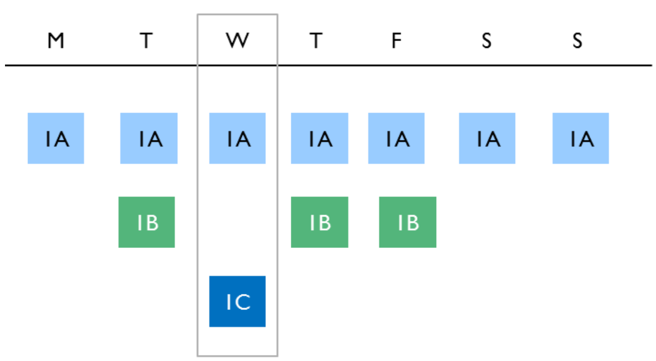

This diagram is taken from internal company documents. Customers who subscribe to the email newsletter opt to receive them daily, three times a week, or once a week.

Figure W5: Examples of discount and non-discount email newsletters

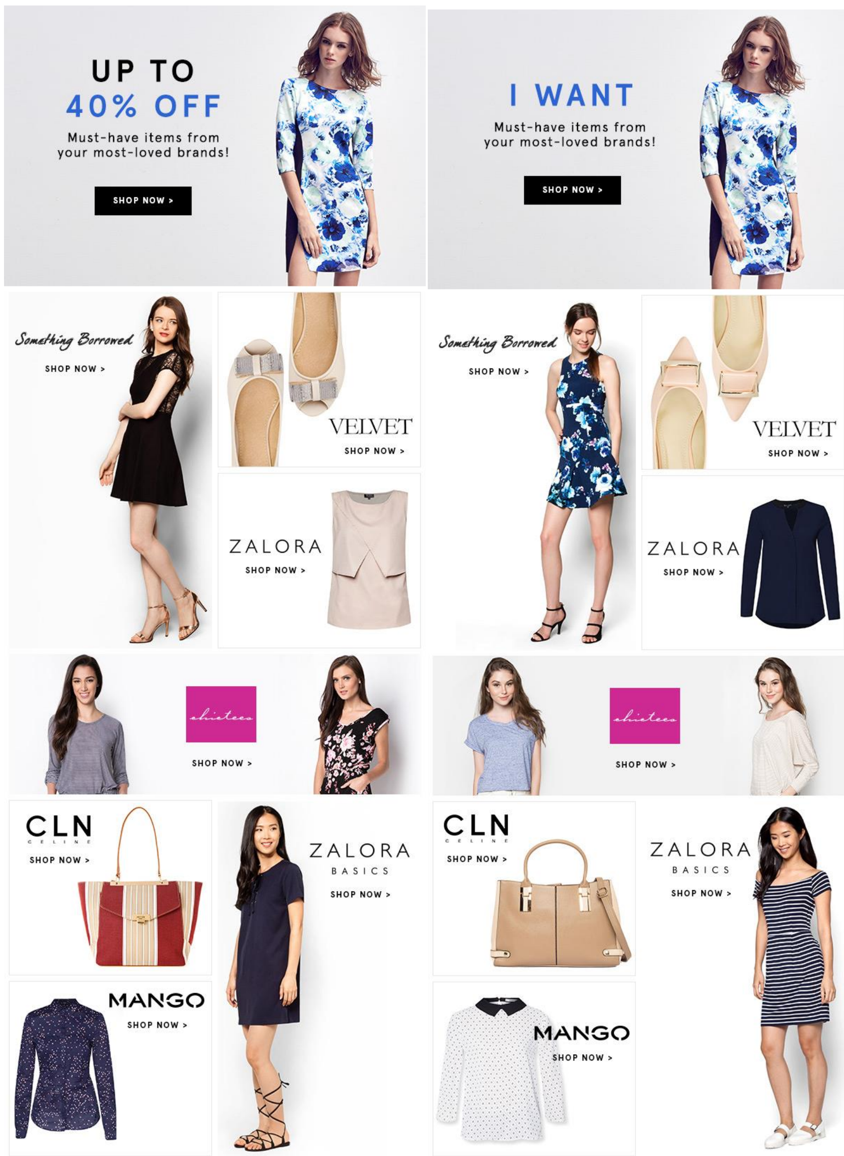

This pair of newsletters is representative of the array of discount- and non-discount newsletters sent out in the validation experiment. Product categories differed between days, as did newsletters sent to male and female recipients. These are two out of a total of 22 unique newsletters sent out during the experiment.

Figure W6: Examples of discount banners