Marketing Science Institute Working Paper Series 2019
Report No. 19-106

The Impact of Coupons on the Search-to-Purchase Funnel: Theory and Empirical Evidence

Arun Gopalakrishnan and Young-Hoon Park

The Impact of Coupons on the Search-to-Purchase Funnel: Theory and Empirical Evidence

Arun Gopalakrishnan and Young-Hoon Park*

June 2018

* Arun Gopalakrishnan is Assistant Professor of Marketing at the Olin School of Business, Washington University, St. Louis, MO 63130; email: agopala@wustl.edu. Young-Hoon Park is Sung-Whan Suh Professor of Management and Professor of Marketing at the Samuel Curtis Johnson Graduate School of Management, Cornell University, Ithaca, NY 14853; email: yp34@cornell.edu. We thank several managers at the participating company for their efforts in making this research possible. We thank seminar participants at the 2016 Missouri and Washington St. Louis Marketing Camp, 2017 Marketing Dynamics Conference, 2018 Marketing Science Conference, and University of Virginia for their comments.

**The Impact of Coupons on the Search-to-Purchase Funnel:
Theory and Empirical Evidence**

Abstract

Firms often use coupons to stimulate sales. While couponing is popular in practice, limited research has examined how coupons work, specifically in the entire search-to-purchase funnel of customers. This paper quantifies the effects of heterogeneous coupons for heterogeneous customers along the search-to-purchase funnel to explain the extent of their sales lift. We develop a model of customer search leading to purchase in the absence and presence of couponing among heterogeneous customers. In the absence of couponing, high-value customers are expected to dominate low-value customers all along the search-to-purchase funnel. Our theory suggests that couponing either has an upstream-only effect where customers are induced to search or a comprehensive effect all along the funnel. We test these propositions using data from a field experiment with an online retailer where customers were divided into two heterogeneous customer segments with two types of coupons (base value and better value). We find that couponing is effective in increasing customer purchases, primarily due to a promotion-as-advertising effect rather than coupon redemption. Customer purchases are moderated by customer heterogeneity. Our results show that while the base coupon has only an upstream effect in the search-to-purchase funnel for both customer segments, the better coupon has a comprehensive effect in the funnel for high-value customers, implying that coupon value acts as a moderator of how coupons affect customers. Finally, our analysis suggests that a bandwidth of $\pm$ 0.5 standard deviations is reasonable for effective regression discontinuity analysis when recovering causal estimates of promotion effects from non-random observational data.

Keywords: Promotion, Couponing, Consumer search, Search-to-purchase funnel, Customer heterogeneity, Field experiment, Regression discontinuity,

1. Introduction

Firms across many industries (e.g., casinos such as Harrah's, retailers such as Macy's and LL Bean) use coupons to stimulate sales. While couponing is a long-standing practice in brick-and-mortar retailing, its use in the digital economy has continued to grow at a steady clip, even as evidence for how these coupons affect consumer behavior remains limited in the literature and in practice. The traditional couponing literature has leveraged price discrimination (e.g., Narasimhan 1984) as the primary driver of sales lift – in other words, coupons are meant to generate sales from those who redeem them (e.g., Neslin 1990), albeit at a lower margin. This focus on coupon redemption (à la Neslin and Shoemaker 1983) has been challenged by recent research (e.g., Venkatesan and Farris 2012, Sahni et al. 2016) that suggests that coupons may serve as a form of advertising in generating sales lift from customers who do not redeem them, but who would not have bought in the absence of the coupon.

Even if we can decompose sales lift from coupons into redemption and advertising effects, a number of questions remain unclear. First, how does the advertising effect work – does it increase the likelihood of consumer search? Does it increase the depth of consumer search – either in terms of time spent shopping or number of products examined? Does it increase the chance of purchase conversion from search? Or some combination thereof? That very limited insights are available on these fundamental questions is one key motivator for this paper.

A second category of questions, beyond how a coupon affects consumer behavior, also emerges: how are the effects of couponing moderated by customer heterogeneity and coupon value? Should coupons be targeted in light of differential effects on customers varying in terms of RFM (recency, frequency, monetary value) score variables? Or should coupons be untargeted? Further, how does the face value of the coupon play into these considerations?

In this paper, we examine the above questions, which have important implications for firms investing a sizeable portion of their marketing budget in couponing, and particularly in targeted couponing based on past customer purchase histories (e.g., Rossi et al. 1996, Thomas et al. 2004, Hartmann et al. 2011). If such firms fail to understand the breakdown of redemption and advertising effects of coupons, whether or not they should pursue behavior-based targeting, and mechanism(s) by which coupons achieve their sales results, significant revenues may be left on the table due to suboptimal couponing strategies. Using advances in information technology that have allowed firms to integrate consumer data at the individual level and to tailor marketing offers to customers with considerable precision, our work looks to shed light on the inner workings of coupons (targeted or otherwise) on a heterogeneous customer base.

Central to our analysis is the observation of the entire "search-to-purchase funnel" of customers. The funnel at the top (or *upstream*) begins with reach – what proportion of customers are induced towards search incidence (beginning the search). At the middle of the funnel (or *midstream*), customers are engaging in search to some depth of product SKUs. At the bottom of the funnel (or *downstream*), searching customers decide whether to convert their search into purchase (what we refer to as purchase conversion). While conceptually this funnel could be applied equally to offline and online settings, we focus on an online setting in this paper as it allows for the tracking of data on customer search within a website as well as purchase information (e.g., Fong 2017) at a level of specificity not typically available in offline retailing (an exception being Seiler and Yao 2017).

Our research questions in a nutshell, then, are to quantify the upstream, midstream, and downstream effects of coupons to explain the extent of their sales lift (as measured by purchase incidence, purchase amount, and average revenue). We are fundamentally interested in heterogeneous treatment effects of coupons for customers with different past purchase histories and also how coupon effects may be moderated by coupon value. Put together, our work provides evidence to address the question: *to target or not to target*?

Our methodology is to develop a theoretical analysis to understand potential coupon effects, use this theory to guide the construction of a field experiment involving a large online retailer, collect data using the field experiment, present the overall sales lift from couponing using the typical outcome measures found in past literature, and, finally, to explain the sales lift figures using clickstream data that sheds light on the search-to-purchase funnel of customers. We leverage our theoretical work in interpreting the effects couponing on this funnel.

Our parsimonious theory is based on two heterogeneous customer segments (high type and low type) in which we assume that the high type has a greater mean utility from purchasing a product from a firm than the low type. In the absence of couponing, the high-type customers are expected to dominate the low-type customers all along the search-to-purchase funnel (i.e., they should have higher search incidence and depth along with higher conversion from search). Our theory suggests that couponing either has an *upstream-only* effect or a comprehensive effect all along the funnel (upstream, midstream and downstream).

We conducted a field test in collaboration with a large online retailer in Asia that sells personal care products to consumers. The experiment divides customers into high-value and low-value segments (matching our theory) with two types of treatments: a base-level coupon and a better-level coupon, offering a higher value than the base coupon. Customers are expected to meet a condition for the purposes of redemption (e.g., Lee and Ariely 2006). To understand the potential benefits (or drawbacks) of targeted couponing on customer value, one treatment group involved making a base coupon offer to both high-value and low-value segments. In addition, we have a treatment group with targeted couponing in which the high-value segment receives a better coupon offer. Finally, to provide a baseline for comparing treatment effects, we also have a control group that does not receive any coupons. Importantly, our design allows for the analysis at the segment level not just the aggregate level, which increases statistical power (Simester et al. 2017).

Our empirical work establishes several key findings relating to the effects of coupons. First, with regard to how couponing affects customer purchase behavior, we find that couponing is effective in influencing customer purchase behavior, despite redemption only accounting for less than 20% of purchases in all treatment conditions. In other words, the majority of the lift arises from an indirect advertising effect, consistent with the recent literature on couponing (e.g., Venkatesan and Farris 2012, Sahni et al. 2016). In addition, the base and better coupons lift purchase incidence relative to the control group, both at the group level and at the segment level. Purchase incidence is moderated by customer heterogeneity as the high-value segment experiences a greater lift in purchase incidence from the base coupon than the low-value segment. Further, purchase incidence is also moderated by coupon value as the better coupon further generates lift in customer purchase compared to the base coupon for the high-value segment, again mainly due to an advertising effect. In other words, our results suggest that the price discrimination effects of couponing only play a limited role in their overall effectiveness.

Second, in regard to how couponing affects customer search behavior, when examining the effect of the base coupon on customer search activity, we find a remarkably similar lift (about 23 percentage points) in search incidence for both high-value and low-value segments. However, the base coupon does not affect beyond the search itself, i.e., search depth and purchase conversion. Further, the virtually identical effect of base coupon on search incidence across the segments means that observed difference in purchase incidence across these segments is due to baseline heterogeneity and not couponing. In other words, the base coupon increases the reach at the top of the search-to-purchase funnel and has no further effect on subsequent search and purchase behaviors. On the other hand, the better coupon targeted at high-value customers has a more comprehensive effect than the base coupon. Its lift in search incidence is no different than the base coupon (suggesting a limit in tempting customers to start search based on coupon value alone) but its effect on search depth and purchase conversion is positive as well, leading to a much higher lift in purchase incidence. This suggests that higher coupon value generates higher engagement in the downstream search process such that coupon value moderates the parts of the search-to-purchase funnel at which a given coupon is effective.

The third key finding relating to the impact of coupons on customer search and purchase behaviors has to do with robustness. We estimate local average treatment effects using regression discontinuity (RD) analysis, as our approach to design a field test uses a threshold based on the monetary value in the pretest period that induces a discontinuity. This analysis serves both as a robustness check and a validation of RD methods since we are able to compare its estimates with the true effects obtained from randomization. Our analysis suggests that a bandwidth of $\pm$ 0.5 standard deviations is reasonable for effective RD analysis in recovering the true effects. We believe that this can serve as a baseline for practitioners and researchers who may non-randomly use targeted coupons to recover causal estimates.

Overall, our research makes several novel contributions to the couponing literature. To the best of our knowledge, it is the first paper to document the split between advertising and redemption effects of heterogeneous coupons to heterogeneous customers using a randomized field experiment. We also contribute to a newly emerging stream of literature that examines how firms can influence consumer search (e.g., Seiler 2013, Fong 2017) through promotions. Our findings on coupon effects throughout the search-to-purchase funnel (upstream, midstream, and downstream) and how these are moderated by customer heterogeneity and coupon value are also novel. Finally, we also demonstrate the robustness of RD analysis with the benefit of having the gold standard of randomized controlled trials for comparison, which bolsters the case to use RD methods in observational data when possible (e.g., Gordon et al. 2018).

The remainder of this paper is organized as follows. In Section 2, we develop our theoretical model on how heterogeneous customers search for products in the presence and absence of a coupon. In Section 3, we describe the experimental design. In Section 4, we present the results of the experiment. We conclude in Section 5.

2. A Model of Consumer Search and Purchase

In this section, we propose a model of consumer search leading to purchase. Extant literature has limited predictions that relate to our empirical setting. Dholakia (2006) and Anderson and Simester (2001) argue that high-value customers should not receive attractive coupons, because such customers do not require much persuasion to buy (e.g., Lal and Bell 2003, Musalem and Joshi 2008). Anderson and Simester (2004) and Shin and Sudhir (2010) argue that coupons should target new and low-share customers. Making a different argument, Acquisti and Varian (2005) show that a monopolist firm can do as well without using past purchase history to set prices in the presence of forward-looking consumers. Homburg et al. (2008), on the other hand, suggest targeting high-value customers. While these are interesting findings, they reflect a lack of convergence on a prescription for targeting. They also do not make predictions about how couponing affects the search-to-purchase funnel, which is our primary focus.

While the reader may skip Section 2 and still be able to understand the empirical results in Section 4, we believe that our theory helps examine the mechanisms through which coupons influence the search-to-purchase funnel, and that the theory helps support generalization that our findings may replicate in other settings. From the model we derive several testable predictions of the impact of coupons on the funnel. These predictions help design a field experiment and interpret the empirical evidence from the experiment. All proofs are provided in Appendix A.

2.1 Model Overview

The model represents consumer search as a means to find a fit between a product and a customer's preferences. For example, if a customer becomes interested in the products that a firm has to offer, she can engage in a search to find a suitable product that fits her preferences and make a purchase if she finds one that has a satisfactory fit. We assume that fit is not fully observed prior to consumer search. Otherwise, the consumer can simply purchase the product she has in mind without any search. This assumption is well suited to online retail settings and fast-moving consumer goods markets in which a large number of products are offered to customers and new products are introduced on a regular basis.

The assumption that fit is not fully known prior to search is also applicable to contexts in which customers receive coupons that apply to only a subset of a firm's brands and products. When receiving a coupon offer, customers may search more online if they are expected to meet redemption policies (e.g., $X off for select brands or products). The assumption may be less applicable to contexts in which coupons can be applied to any basket of purchases or campaigns that only feature a very small selection of products for promotion.

We assume that there are two types of customers: high-type and low-type consumers. Each type of customer draws her utility for a product from a utility distribution with a type-dependent mean and a variance that is common between the two types. The high type is assumed to have, on average, a higher utility for the product by the retailer, as compared to the low type. That is, $m_H > m_L$, where m_H (m_L) is the mean of the utility for the high (low) type.

In the model, we consider a two-stage process of consumer search prior to purchase at a firm. In stage 1, customers decide whether or not to search (search incidence) and then in stage 2, they decide how many products to search (search depth) in order to find a product that matches her preferences. We assume that the customer engages in simultaneous search (e.g., Stigler 1961, Honka and Chintagunta 2017). By virtue of choosing the optimal number of products to search, the customer is assumed to know the ex-ante purchase likelihood, implying that purchase conversion (i.e., proportion of customers who make a purchase after starting search) is endogenous. To develop the model, we begin with the customer's decision of search depth (stage 2) conditional on search incidence at an online retailer, and then move backwards to the decision of search incidence itself (stage 1). The reason we present the model in this order is because stage 1 presumes knowledge of optimal strategies in stage 2 via backward induction. We derive baseline predictions for high-type and low-type customers without a coupon in Section 2.2 and present predictions with a coupon in Section 2.3.

2.2 Model without a Coupon

Stage 2. Conditional on a customer searching at an online retailer, we assume that the customer makes a decision about the number of products to search. Searching through more products increases the chance of finding a product that matches her preferences. However, consumer search is costly in terms of time and cognitive resources required to process the information. This decision can be written as the following optimization problem:

$$\underset{K}{\operatorname{argmax}}\, E[y_{iK} \cdot I(y_{iK} > U^*)] - K \cdot c, \tag{1}$$

where K is the number of products to search, y_{iK} is the maximum utility customer type i obtains from searching over K products, U^* is the reservation utility at which the customer would actually make a purchase, and c is the search cost per product.

It is worth noting that equation (1) is parsimonious in that high-type and low-type customers have a common cost c and reservation utility U^*. The customer types only differ in terms of their mean utility over the products offered by the firm. Suppose each type's utility is normally distributed with type-dependent mean m_i and common standard deviation s. If a customer were to take a random draw of a product to search, her utility is assumed to be drawn from the type-dependent distribution. The customer is only interested in purchasing the product if her utility exceeds U^*.

Customers may search for more than one product. The maximum utility across all K products searched is not distributed as the original normal distribution. Instead, the maximum utility's distribution for a set of K i.i.d. draws is $y_{iK} \sim F(m_i, s)^K$, where F is the cumulative distribution function of a normal distribution with mean m_i and standard deviation s. While it is helpful to know that the distribution of maximum utility over K products is the product of the i.i.d. distributions, there is no analytical expression that we can obtain from this formula that would help explain how different types of customers may optimize the number of products to search. We therefore invoke the Fisher-Tippett-Gnedenko extreme value theorem (Gnedenko 1998), which shows that asymptotically the maximum order statistic is distributed as Extreme Value Type I when F has an exponential tail (as does the normal distribution). In practice, the asymptotic distribution provides a very close estimate of the actual distribution even for $K > 7$ in the case of the normal distribution. We define K_i^* as the optimal number of products searched by customer type i. In Proposition 1, we focus on the relationship between search depth and customer type.

Proposition 1: $K_H^* > K_L^*$.

Proposition 1 suggests that search depth is greater for the high type compared to the low type. In Appendix A, we show that $E[y_{HK} \cdot I(y_{HK} > U^*)] > E[y_{LK} \cdot I(y_{LK} > U^*)], \forall K$, which results in a higher K for the high type than the low type. It implies that search is more rewarding for the high type which means they are more willing to incur a higher search cost. For the low type, on the other hand, search is rewarding but saturates at a lower number of products.

Proposition 2 states the relationship between purchase conversion and customer type. It is a by-product of Proposition 1. In our search-to-purchase model, the likelihood of purchase conversion arises endogenously from the search process. For the same number of products searched, the high type has a higher probability for purchase conversion than the low type since it is easier for the high type to find a product that exceeds the reservation utility. Further, since the high type searches more than the low type (Proposition 1), the disparity in conversion likelihood will be even higher.

Proposition 2: $\gamma_{H,K_H^*} > \gamma_{L,K_L^*}$, where $\gamma_{i,K} = \Pr(y_{iK} > U^*)$.

It is important to note that both Propositions 1 and 2 follow from one key element of the model: the mean utility for the products offered by the firm differs between the two customer types. We assume all other aspects of preferences are the same across customer types because the difference in mean utilities alone suffices to generate both predictions. We also note that while the normal distribution was used for the distribution of utility for each type, the use of the asymptotic Extreme Value Type I distribution means that both propositions are robust to any distribution that has an exponential tail.

Stage 1. We next describe the customer's decision of search incidence in the first stage in which customers decide whether or not to search. Recall that, as in stage 2, this search is in the absence of a coupon and is intended to develop a baseline of how heterogeneous customers behave. We assume that type i is aware of the optimal number of products to search (K_i^*) and the likelihood of purchase conversion (γ_i). While the customer will be aware of the expected utility from optimal search, we also allow for a random shock with a mean of zero to affect each customer's perceived expected utility. Without this random shock in the first stage, all consumers would either search or not, which is unlikely to be true empirically. Since some proportion of high-type or low-type customers would choose not to search at all, our model assumes that these customers are likely to have received a negative shock compared to others in their type. Suppose that the firm has N customers who are either high type or low type. The utility from searching for customer n who is of type i is modeled as:

$$V_{i(n)}\left(K_{i(n)}^*\right) - K_{i(n)}^* \cdot c + \epsilon_{i(n),n}, \tag{2}$$

where $V_{i(n)}(\cdot)$ is the expected utility from the search with a given number of products, $K_{i(n)}^*$ is the optimal number of products searched in the second stage, and $\epsilon_{i(n),n}$ is the random shock for customer n who is of type i. It is important to note that two customers that are both of the same type will receive different random shocks.

We assume the utility from not searching is 0. The probability of search incidence is:

$$\Pr\left(V_{i(n)}\left(K_{i(n)}^*\right) - K_{i(n)}^* \cdot c + \epsilon_{i(n),n} > 0\right) = 1 - F_\epsilon\left(K_{i(n)}^* \cdot c - V_{i(n)}\left(K_{i(n)}^*\right)\right). \tag{3}$$

Since the second-stage search model is already solved, we define $\Delta_i = V_{i(n)}\left(K_{i(n)}^*\right) - K_{i(n)}^* \cdot c$ as the surplus for type i. A higher surplus Δ_i will lead type i to search with higher probability since $F_\epsilon(\cdot)$ is decreasing in Δ_i. The probability of searching for type i is $1 - \Pr(K_i^* > 0)$. In Proposition 3, we state the relationship between search incidence and customer type.

Proposition 3: $\Pr(K_H^* > 0) > \Pr(K_L^* > 0)$.

Proposition 3 implies that the high type has a higher probability of search incidence. The intuition is that the high type must have a higher surplus at K_L^* than the low type, and its surplus only increases since the high type searches more than K_L^*.

In summary, our model suggests that in the absence of a coupon, high-type customers have higher search incidence and depth than low-type customers. This in turn leads to higher purchases for the high type as compared to the low type. Necessarily, this leads to higher revenues for the high type compared to the low type. We test these predictions, derived from the model without a coupon, using data from a field experiment that includes a control group who does not receive any coupons.

2.3 Model with a Coupon

Couponing could affect consumer search and purchase behaviors in two different ways. First, the coupon could increase the mean of the utility distribution for a given customer type. From Propositions 1 and 2, this increases search depth and purchase conversion. From Proposition 3, it also increases search incidence. Second, it could affect search incidence in the first stage of the model without affecting the second stage. Similar to the model without a coupon, we present the impact of the coupon on the second stage and then on the first stage of the model.

Stage 2. Conditional on a customer searching at an online retailer, the coupon can shift the utility distribution for type i to the right, i.e., increase its mean. Hence, it will lead customers to increase search depth, which increases purchase conversion. To reflect the model parameters with and without a coupon, we define $m_{i,1}$ ($m_{i,0}$) as the mean of utility, $K_{i,1}^*$ ($K_{i,0}^*$) as the optimal number of products searched, and $\gamma_{i,1}$ ($\gamma_{i,0}$) as the purchase conversion by type i with (without) a coupon. Proposition 4 (5) states the relationship between search depth (purchase conversion) and customer type in the presence of a coupon, as compared to the case without a coupon.

Proposition 4: $m_{i,1} > m_{i,0} \leftrightarrow K_{i,1}^{*} > K_{i,0}^{*}$.

Proposition 5: $m_{i,1} > m_{i,0} \leftrightarrow \gamma_{i,1} > \gamma_{i,0}$.

Both propositions follow naturally from the implications of Propositions 1 and 2. A coupon that increases the mean of utility for type i increases search depth, which also increases purchase conversion. Further, an increase in search depth and conversion implies an increase in the mean utility for a given customer type. It is important to note that our model does not predict a partial effect of only the increase in search depth or purchase conversion. To the extent that partial effects are observed, this would be a critique of our model and may suggest that other mechanisms are at work. Should a coupon not affect these metrics at all, it would suggest that the coupon has no effect on the mean utility distribution of a given customer type. It should be noted that a third possibility also exists (albeit ex ante less likely than the ones above): a coupon decreases the mean utility such that both search depth and conversion decrease.

Different types of coupons can have different effects on consumer search and purchase behaviors. While our model does not specifically describe a certain type of coupon, one can posit two general variations. First, the coupon may not have any restrictions for redemption (e.g., $X off to any purchase made) or features a few products for promotion only. It would then naturally shift the utility distribution to the right for both types of customers, since the same products can be purchased at effectively lower prices. Second, the coupon may have “strings attached” that require the customer to satisfy a condition for the purposes of redemption. For example, a coupon may require her to purchase from a subset of brands or products to qualify for the redemption so that managers can focus on increasing demand for those brands or products without sacrificing margin across the board. In this case, the shift in mean utility may be reduced because consumers only benefit if they buy from the subset of products.

If a shift does occur in the mean utilities, its magnitude could be type-specific. For example, it is possible that low-type customers experience a higher mean shift than high-type customers such that the rank ordering on which type searches more products may be reversed. On the other hand, the shift could also be of similar magnitude for both customer types, or even higher for the high type, resulting in a further divergence in search behaviors. Our model allows for all of the above possibilities and empirical evidence is needed to understand if coupon effects are in line with our model and if so, the direction of those effects for each type.

Stage 1. We next describe the impact of a coupon on the first stage of consumer search. There are two possibilities for the coupon to affect the proportion of customers who search. First, the coupon may shift the utility distribution for each type to the right, i.e., increasing the mean of the distribution, potentially with type-dependent magnitudes. As discussed earlier, this will lead to higher search depth and conversion. It will also lead to a higher surplus for both high-type and low-type customers. As a result, the proportion of customers who search is also predicted to increase for both customer types, though the relative magnitude depends on the relative mean shift for each type. We label this as a "comprehensive" coupon effect as it affects all three metrics leading to purchase (search incidence, depth, and conversion).

Second, the coupon may add a systematic effect to the utility equation as follows, where δ_i is a type-specific increment in the utility of search incidence:

$$V_{i(n)}\left(K^*_{i(n)}\right) - K^*_{i(n)} \cdot c + \delta_i + \epsilon_{i(n),n}. \tag{4}$$

It should be noted that δ_i does not affect subsequent search behavior as it only appears in the equation for search incidence. Under this possibility for a coupon's effect, there may be a lift in search incidence that is type-dependent but no further change in search depth or purchase conversion. We label this as a "limited" coupon effect as it affects only the top of the search-to-purchase funnel (search incidence), while the remaining processes are unaffected by the coupon.

Which of the two scenarios is at work in the presence of a coupon can be tested (at least partially) with empirical evidence. The comprehensive effect requires a coupon to affect all three metrics, i.e., search incidence, depth, and conversion. The limited effect only affects search incidence. The limited or comprehensive effect could differ in magnitude for high-type and low-type customers. Further, whether the effect is limited or comprehensive may depend on coupon value. Thus, we empirically test these two competing possibilities using data from a field test that has different customer types with different coupons. Note that if both possibilities are at work, it may be difficult to tease them apart in the case where both increase search incidence.

In summary, we have developed a parsimonious theory that predicts high types to have higher baseline search incidence, depth, and conversion than low types in the absence of a coupon. There are two possible coupon effects that can be heterogeneous across types: a comprehensive effect on all measures and a limited effect on only search incidence. Neither prediction has been made in the extant literature. Our theory predictions help inform the design of the field experiment and examine the underlying mechanisms that can explain our findings.

3. Field Experiment

3.1 Experimental Setting

In this section we describe a field experiment that was conducted with the cooperation of an online retailer that sells consumer goods in the personal care category. The consumer goods market is an appropriate setting for the experiment because most retailers in this market sell customers a wide assortment of brands and products, ranging from low-end to high-end goods at varying prices. In addition, firms in this industry constantly introduce new products and spend a significant portion of their marketing budgets on promotions and personalized marketing offers. As a result, it is non-trivial for customers to find a product that could match her preferences by meeting a condition in the promotion.

The cooperating firm is a large online retailer in Asia that prefers to remain anonymous. A field experiment allows us to exploit variation generated from random assignment as opposed to collecting observational data on a firm's customer base based on the firm's own promotional policies, which requires a much more elaborate analysis to address endogeneity concerns. As discussed earlier, we see a field setting as critical as our goal is to make causal inferences on behaviors of heterogeneous customer segments as well as their underlying search processes.

The customers of the firm vary considerably in their purchase patterns. For the purpose of designing the experiment with different customer types, we used data pertaining to the 12-month purchase histories of the customers who had made at least one transaction over a period of 12 months prior to the experiment. Preliminary analysis reveals a skewed distribution of purchase amounts spent by each customer, with the top 20% of the customers accounting for approximately 75% of the firm's revenue. As the firm has a large number of customers, we divided the customers into 20 equal-sized groups, i.e., ventile, instead of using quintile or decile analysis, based on the purchase amounts spent by each customer in the past 12 months prior to the experiment. Figure 1 presents the purchase amount, monetary value, across 20 groups in the pretest period. After consulting with the managers at the firm, we categorized the top 4 groups (20% of the customers) as the high-value segment and the rest of the 16 groups (80% of the customers) as the low-value segment because customer behavior between the two segments differs qualitatively. For instance, we observed wide variation in the average purchase amount in the high-value segment, ranging from approximately $800 in ventile 1 to about $100 in ventile 4.[1] We also observed large variation in the average purchase amount in the low-value segment, ranging from about $80 in ventile 5 to only a few dollars in ventile 20. Figure 1 also presents the number of purchases (a measure of frequency) across 20 ventiles in the pretest period. As frequency and monetary value are highly correlated across the groups, a segmentation based on the monetary value is appropriate to examine the impact of coupons on the two segments.

As data are collected in the course of commercial operations, we looked at how to maximize the value of the field experiment while complying with the retailer's requirements. For instance, before randomly sampling participants from each segment for the experiment, we excluded ventile 1 (top 5%) in the high-value segment because the retailer noted that these customers differ considerably from the rest of the high-value segment. In addition, we excluded ventiles 15 to 20 (bottom 30%) in the low-value segment because the retailer restricted our attention to customers who spent at least $20 over the 12-month pre-experiment period. From the 20 equal-sized ventiles, therefore, we have the customers in ventiles 2 to 4 for the high-value type, and those in ventiles 5 to 14 for the low-value type. To proceed for the random assignment, we randomly sampled a total of 6,785 customers from the low-value segment and a total of 6,714 customers from the high-value segment. We oversampled high-value customers in a stratified design to ensure adequate power for estimating segment-level treatment effects (Simester et al. 2017). A total of 12,959 customers between the two segments, along with a set of targeted promotions to be described, were randomly assigned across the experiment and control groups.

[1] All transactions were recorded in the currency of the country in which the headquarters of the company was located. We converted purchase amounts to U.S. dollars using the average exchange rate over the data period.

Customers in the experiment were offered discounts if they made a purchase that met a condition for redemption (e.g., Lee and Ariely 2006). Examples of such types of promotions typically used by firms in the industry, including our corporate partner, are "$X off for purchases of $Y or more," "$X off for select brands or products," etc. In collaboration with the retailer, we chose "$X off for select products" for the experiment, in which a customer can redeem the coupon if she purchases a product priced at $20 or higher. We note that our condition is easier for both low-type and high-type customers to meet compared to a basket-spend requirement which may be out of reach for the low type. Using managerial guidance from the firm, we have two types of coupon with different discount amounts, a base-level offer with $7 discount and a better-level offer with $10 discount, if the customer purchases a product priced at $20 or higher. Thus, the customer can redeem the coupon by searching at the retailer's online store to find a product that meets the coupon condition while also matching her preferences.

3.2 Experiment Design

As our main objective was to compare the performance of targeting heterogeneous customers with coupons on their search and purchase behaviors, an experiment was conducted as follows. With the two segments of heterogeneous customers (the high-value and the low-value segments derived from the retailer's customer database), we conducted a randomized intervention in the form of a coupon that was sent to both customer segments. We then created two test groups and a control group. The first test group (T1) involved making a non-targeted coupon with base offer ($7 discount) to both high-value and low-value segments. The second test group (T2) involved the high-value segment receiving a coupon with better offer ($10 discount).[2] Finally, the control group (CG) comprised both high-value and low-value customers who did not receive a coupon. This group serves a baseline for assessing treatment effects.

[2] In T2, the firm was unwilling to offer the better coupon to the low-value segment.

We note that our control group features the absence of any communication from the firm rather than a generic communication that does not involve any coupons for two reasons. First, the firm preferred to retain the CG without any communication over a generic communication, because it helps measure the treatment effects of the coupon relative to no coupons. Second, recent studies (e.g., Sahni et al. 2018) suggest that seemingly trivial modifications to a communication can change its efficacy, thus opening up a large number of degrees of freedom in how to craft the message for the CG. We suggest that the cleaner standard is that of no communication to the control group and offer a few suggestions in the discussion section on future work that can explore these topics further.

Importantly, our design allows for the analysis at the segment level not just the aggregate level, which increases statistical power. In order to enhance power for the high-value segment, which is smaller in size by construction, we used a disproportionate stratified random sampling scheme in which the high-value segment is oversampled relative to its relative size in the customer database, i.e., population. Figure 2 shows the design of the study and the number of high-value and low-value customers and the coupon type across the conditions.

All customers in the experiment conditions were contacted by the firm on the same day at the same time via their mobile phones in October 2016 for a campaign spanning a one-week period. The campaign offered a price-discount digital coupon that recipients could redeem within the campaign period while purchasing any product priced at $20 or more. This coupon is made available to the consumer when they shop, thus simplifying the process for redeeming them compared to paper coupons (Andrews et al. 2016). The campaign communication (e.g., content, creative) was exactly the same between the test groups, except the discount depth. We controlled for other marketing efforts during the campaign period as all customers included in the study were not exposed to other promotions.

4. Results

In our field test, we divided customers into two segments based on a past 12-month monetary spend. Those customers spending less than $87.50 in this period are classified as the low-value segment and those spending $87.50 or higher are classified as the high-value segment. From each segment, we randomly sample a set of customers for each condition, yielding a total of 12,959 customers across all conditions. We begin by presenting randomization checks for the low-value and high-value segments in Tables 1a and 1b. The low-value segment is assigned to either a base coupon offer (T1) or a control group (CG). The high-value segment is assigned to a base coupon offer (T1), a better coupon offer (T2), or a control group (CG). These tables demonstrate the face validity of our randomization, as the mean purchase amount in the pretest period is not statistically different for a given segment across the conditions. Further, Tables 1a and 1b show that the order statistics also closely match such that there is no evidence of selection bias that may affect our results. In Appendix B, we also present randomization checks for recency and frequency metrics, and find that customer behavior in the pretest period did not vary across the conditions.

The rest of the section is organized as follows. In Section 4.1, we describe the purchase (outcome) behaviors that result from our experimental conditions. Our goal is to establish how coupons perform in our field test, both in terms of similarities to the existing literature and new findings. In Section 4.2, we examine the search (process) behaviors that lead to the outcomes and connect empirical findings with our predictions derived in Section 2. In Section 4.3, we conduct robustness checks that demonstrate alternative ways to analyze our data using RD analysis.

4.1 The Effect of Coupons on Customer Purchase

In examining the effects of coupons on customer purchase behavior, we observe multiple dependent measures: first, the number of purchases made, which is the main outcome measure of interest for the company. We also observe the instances in which the customers in the test groups redeem the coupon with a purchase. Note that these are two different measures: in the former, purchasers may or may not redeem the coupon but purchases were made; in the latter, purchases were made by redeeming the coupon. In addition, conditional on purchase, we observe purchase amount net of any discount. Based on this measure, we compute the average revenue per customer (ARPC) for each condition, which is defined as total revenue (net of coupon discounts) divided by the number of customers in the group.[3] This measure is used in the literature (e.g., Venkatesan and Farris 2012, Sahni et al. 2016) because it accounts for the effect of margin loss due to coupon redemption. Based these outcome measures, we decompose whether coupons work for the customer base as a whole and/or for each segment, and study potential heterogeneity in treatment effects by customer type and coupon value.

4.1.1 Group-level Coupon Effect

Table 2 shows the results from the experiment. The first three columns under "Study Sample" in the table correspond to the outcomes in each group of the study sample across both low-value and high-value segments. The next three columns under "Population" in the table present the adjusted estimates at the population level, which account for different probabilities of selecting a

[3] ARPC is similar to the average revenue per user (ARPU), which is defined as the total revenue divided by the number of subscribers in industries such as consumer communications, digital media, and network companies.

customer from each segment in our disproportionate stratified sampling design.[4] As the findings are similar to each other, we focus on the results from the study sample.

Base Coupon. The base coupon increases purchase incidence by 1.97% compared to the control group (p-value < 0.001). This finding leads to the following question: is this lift driven by coupon redemption or by the coupon serving as a form of advertising for the customers? The redemption rate in T1 of 0.58% shows that coupon redemption accounts for less than a third of the overall lift in purchase incidence due to the base coupon, as the redemption rate conditional on purchase is only about 16%.[5] Thus, this is in line with the recent literature that the advertising effect of the coupon can be much larger than its direct effect through redemption (e.g., Venkatesan and Farris 2012, Sahni et al. 2016). We note that this redemption rate is still higher than what is expected in offline settings (e.g., Dong and Kaiser 2005), probably because of the increased convenience of coupons in an online setting relative to traditional coupons.

Average spend conditional on purchase is not statistically different between the groups: \$34.04 in T1 (net of coupon redemption) and \$31.26 in CG (p-value = 0.507). Further, the net average spend among customers who redeemed the coupon is indistinguishable from that among purchasers who did not redeem (\$31.75 vs. \$34.47, p-value = 0.554). This result is in line with past purchase patterns of customers shown in Figure 1: purchase frequency and monetary spend are closely related, suggesting that the difference across the customers in the pretest period is not

[4] To compute the population-level estimates, we need to have (1) the actual proportion of high-value (low-value) customers in the sampling frame and (2) the probability a customer is selected for the high-value (low-value) segment. The population-level estimates are computed as follows: $\frac{\sum_{i=L}^{H} ATE_i \cdot N_i \cdot w_i}{\sum_{i=L}^{H} N_i \cdot w_i}$, where ATE_i is the average treatment effect, N_i is the sample size, and w_i is the weight of customer type i in a given condition, which is the inverse of the probability a customer is selected. Because of the nondisclosure agreement with the partnering company, we cannot disclose the weights in order to preserve anonymity of the firm's size of customer base.

[5] The redemption rate conditional on purchase is computed as the number of redeemers divided by the number of purchasers. For example, in the case of the base coupon for the study sample, it is computed as 0.58/3.65 = 15.9%.

expenditure per purchase, but the number of purchases made per year. It is therefore not surprising that the coupon did not significantly increase purchase amount.

The overall measure of group performance is the ARPC measure, which combines purchase incidence and amount. Managers (and researchers) focus on this metric as a simple way of gauging how much revenue to expect from the customer base on average. We observe that an increase of \$0.72 is observed in ARPC due to the base coupon, which is statistically significant (p-value < 0.001). As a counterfactual, suppose the advertising effect of the base coupon is completely removed and its only impact is through coupon redemption. In such a scenario, we would expect purchase rate to be 2.26%, i.e., 1.68% (purchase rate in CG) plus 0.58% (redemption rate in T1). If we assume that net purchase amount remains unchanged, we can compute the ARPC under this counterfactual as \$0.77 (= 2.26% × \$34.04). Compared to CG, this is only an increase of \$0.25 rather than the increase of \$0.72 that is actually observed in the data. This suggests that managers should not only expect an increase in ARPC due to coupon redemption but also due to the advertising effect.

Better Coupon. We next report the results with the better coupon from the study. The better coupon clearly outperformed the base coupon, i.e., an increase of 5.87% in purchase incidence (p-value < 0.001). However, this comparison is problematic because only high-value customers are included in the better coupon condition (T2). Nevertheless, we still find that the advertising effect plays an important role in assessing the impact of the coupons because only about 24% (1.89/7.84) lift in purchase conversion in T2 is due to coupon redemption, and the majority of the lift is due to a promotion-as-advertising effect. The better coupon still does not increase net spend conditional on purchase, but the large increase in ARPC to \$3.13 accrues from the large increase in purchase incidence.

Clearly, segment-level heterogeneity is ignored in the findings described above. The key benefit of our experimental design is that we can uncover how low-value and high-value customers differ in terms of outcome variables, which we describe next.

4.1.2 Segment-level Coupon Effect

We describe the findings with the base coupon and then the results with the better coupon. In Tables 3a and 3b, we present the segment-level results between the test and control groups from the experiment.

Base Coupon. Table 3a shows that there are statistically significant gains in purchase incidence and ARPC for the low-value segment, while the advertising effect continues to dominate the redemption effect: In T1, only 0.48% of customers redeem out of a total of 2.47% of customers who purchase. Table 3b shows that the gains with the base coupon are in fact larger for the high-value segment than the low-value segment, with both the purchase incidence and ARPC lifts being more than double what they are for the low-value segment. Interestingly, we find that the advertising effect is stronger for the high-value segment (2.39% lift) as compared to the low-value segment (0.98% lift), also reflecting a lower redemption rate (12.6% versus 19.4%) of the base coupon for the high-value segment.[6] The purchase amount in the high-value (low-value) segment rises slightly in T1 compared to CG, with an increase of \$3.20 (\$3.50), but these differences are not statistically significant. Hence, the ARPC gains in both segments are not due to changes in the basket spend, but driven by increases in purchase incidence.

These results are suggestive of heterogeneous treatment effects of the base coupon on the two customer segments, which could not be provided in the group-level analysis described

[6] The lift due to advertising effect is computed by taking the total lift and subtracting the percentage of customers who redeem the coupon. For example, in the case of the base coupon for the high-value segment, it is computed as 3.22 – 0.83 = 2.39%.

earlier. Though the high-value segment has a higher baseline level of purchase incidence than the low-value segment (as measured in CG), its lift from the base coupon is higher, and mostly arising from the advertising effect. The net result for the firm seems doubly positive: by marketing a base coupon to the high-value customers, a stronger lift in purchase incidence is obtained and most customers do not even redeem the coupon, which does not impact its margin.

Better Coupon. Next, we estimate the effect of the better coupon on customer purchase. In our experiment, the high-value segment also participated in the better coupon condition (T2). Given the results in T1 versus CG, it is interesting to examine if offering an even better deal to the high-value segment results in further improvement in the outcome measures. On the one hand, it may be reasonable to expect higher redemption when the coupon is worth more. On the other hand, if the increase in coupon redemption is mitigated by the advertising effect, this targeted coupon with a higher discount to the high-value segment may be justified. Table 3b contains the metrics for T2 and its comparison with both T1 and CG. Note, of course, that the measures for T2 are identical between Table 2 and Table 3b.

We note that the lift in purchase incidence from T1 to T2 is similar in magnitude (2.93%) to the one from CG to T1 (3.22%) for the high-value segment. Further, while this is partly due to a significant increase in coupon redemption (a lift of more than 1%), the advertising effect continues to dominate, accounting for about 64.5% (1.89/2.93) of the lift from the base to the better coupon. The basket spend continues to be no different from any of the other conditions. Thus, the ARPC increases to $3.13 due to the additional lift in purchase incidence.

Heterogeneous Effects of Coupons. While the redemption rate conditional on purchase is 19.9% for the better coupon targeted at the high-value segment (which is very similar for the base coupon for the low-value segment, 19.4%), the additional lift due to the advertising effect is worth further discussion. As discussed in Section 3, our control group featured no communication rather than a generic communication without any coupons. Might it therefore be that the lifts from the coupons in this experiment may act as a reminder for customers who have not made a purchase for a while? Our data can allow us to test this possibility in two different ways. First, under the reminder hypothesis and our finding of generally low redemption rates conditional on purchase (< 20%), one may expect that the lifts in purchase incidence in the base and better coupons for the high-value segments ought not be different. But the evidence runs counter to this hypothesis. As we presented, the better coupon generates a substantially larger lift in purchase incidence than the base coupon, which is mostly not accounted for by increased coupon redemption. Thus, it appears that the coupon value itself affects customer purchase. We explore this notion further in Section 4.2 when examining customer search across the conditions.

Second, as couponing may result in an additional purchase that may not have occurred in the absence of couponing or simply accelerate a future purchase (Neslin et al. 1985), we test how customer behavior in the pretest period may affect customer purchase and redemption. As customer loyalty can be a moderator of various marketing activities (e.g., Kumar and Shah 2004, Blattberg et al. 2008), we use the recency measure in the pretest period as a measure of customer loyalty.[7] We also have the monetary value in the pretest period because it reflects the randomization nature of couponing to customers in our data.[8] We report our results using a linear probability model and we also show robustness to a logit formulation in Appendix C. Table 4 shows our main results, building up to the full specification in column (3). Column (1) shows the general patterns in the data and confirms that the coupons, whether they have base or better

[7] For the recency measure, we checked the number of days since the last purchase with all purchases included and with purchases of a product priced at $20 or higher. We found the results from both metrics to be similar. We reported the results with the latter as the recency measure.
[8] We also used a binary variable for the monetary value, which indicates whether or not a customer belongs to the high-value segment, and found qualitatively similar results to those reported in the manuscript.

values for customers, increase purchase incidence and that higher monetary value in the pretest period results in greater lift in purchase incidence. The results in column (2) repeat the analysis with the control for the recency measure. We find that more recent customers have a higher purchase incidence than more distant customers. The results in columns (1) and (2) show that the pretest descriptors have a significant impact on the propensity of the customer to purchase within the campaign period. Note that after controlling for recency and monetary measures, the main result holds. This is reassuring and suggests that our result is not an artifact of a failure of randomization. In column (3), we test for heterogeneity in the effects of the base and better coupons on purchase incidence by interacting the coupon effects with the recency measure. The results indicate that the effect of both the base and better coupons improves with customer loyalty, with a stronger effect for the better coupon.

Our findings run counter to the reminder hypothesis, which would predict that more distant customers might be converted to purchasers by any form of communication by the firm. In contrast, customers are more likely to purchase under the coupon if they have purchased more recently. The result is indicative of how engaged customers may benefit from couponing.

In Table 5, we study heterogeneous effects of coupons on coupon redemption. For this analysis, we can only include groups that received a coupon since a customer in CG cannot by definition redeem the coupon. We find that more recent customers are also more likely to redeem a coupon, but the effect sizes, while statistically significant, are small.

4.1.3 Summary of Coupon Effect on Customer Purchase

We find that couponing is effective. The base and better coupons lift purchase incidence and ARPC relative to the control group, both at the group level and at the segment level. However, purchase amount is not statistically different due to couponing. Customer purchases are moderated by customer heterogeneity as the high-value segment experiences a greater lift in purchase incidence from the base coupon than the low-value segment. Further, customer purchases are moderated by coupon value as the better coupon further generates lift in purchase incidence and ARPC compared to the base coupon for the high-value segment.

4.2 The Effect of Coupons on Customer Search

Having established the effect of coupons on customer purchase behavior, we turn our attention to explore underlying mechanism(s) that can explain the pattern of our findings. Specifically, we estimate the effect of coupons on the search-to-purchase funnel. Customer behaviors of interest include the proportion of customers who initiate the search online, i.e., search incidence. We also observe two measures of search depth associated with the extent of customer search: the number of product pages searched and the (time) duration of the search at the customer level. Based on these measures, we link the empirical evidence with the predictions derived in Section 2. The tables report unconditional measures (i.e., across all customers in the condition)[9] and we also draw attention to related measures such as time spent per product page, number of pages searched, and purchase conversion, all conditional on search incidence.

4.2.1 Customer Search without a Coupon

Tables 6a and 6b present the segment-level results between the test and control groups from the experiment. Our findings indicate that the average time spent per product page is quite similar

[9] For careful comparison of customer search across the conditions, we focus statistical testing on unconditional measures that are obtained by averaging across all customers in a segment. We report conditional measures (e.g., purchase conversion conditional on search) for directional evidence but do not perform statistical tests on these measures since the combination of customers after conditioning on customer search in each group is no longer necessarily exempt from selection effects.

between the low-value (23.23 seconds) and high-value (25.93 seconds) segments in CG. This is in line with the model assumption that the search cost per product (c) is constant across segments.

We next focus on CG to compare the segments in terms of baseline search behavior. In line with our predictions, high-value customers are more likely to initiate the search (diff. = 9.06%, p-value < 0.001), search more products (diff. = 5.08, p-value < 0.001), and spend more time (diff. = 142.08 seconds, p-value = 0.001) as compared to the low-value customers. These findings are in line with other measures that suggest high-value customers are more engaged with the firm; for example, customers in the high-value segment are better in terms of recency and frequency measures in the pretest period than those in the low-value segment. Put together, our results are consistent with our theoretical explanation that the high-value segment has a higher expected utility for the firm's products.

Conditional on search, the high-value segment searches more products (44.08) than the low-value segment (34.43). On average, customers who initiate the search, spend a total of about 13 to 19 minutes at the online retailer. An increase in terms of search depth leads to a higher purchase conversion, 13.58% for the high-value segment, and 8.52% for the low-value segment. These metrics are consistent with the nature of the coupons used in this experiment: customers can engage in a search to find a suitable product that fits her preferences and meets the condition for the redemption.[10]

The baseline results described above have implications for the subsequent analysis of customer search under couponing. Specifically, if a coupon lifts search incidence for the high-

[10] There is considerable variation in terms of the number of products sold at the partnering company, because it is a large retailer selling a variety of products in the personal care category where new products are constantly introduced. Our partner firm sells more than a few thousand products online and approximately a third are priced at $20 or higher.

value segment, it benefits from inherently higher search depth and purchase conversion as compared to the low-value segment, even if the coupon has no effect downstream in the search-to-purchase funnel. Therefore, it is important to disentangle the upstream and downstream effects of a coupon on the search-to-purchase funnel and any heterogeneity in these effects across customer segments.

4.2.2 Customer Search with a Coupon

Base Coupon. The effect of the base coupon on search incidence for both segments is a lift of about 23% (see the column of T1-CG in Tables 6a and 6b). Thus, the base coupon attracts a large proportion of customers who would otherwise not have searched. The average time spent per page view remains at approximately 25 seconds for both segments; the coupon has no effect on search time per product. We find that page views and search time are significantly higher for both segments under the base coupon; however, these unconditional effects are driven mostly by higher search incidence.

When examining page views conditional on search for the low-value segment, there is a slight decrease from 34.4 in CG to 26.2 in T1. The low-value segment in T1 also displays a lower purchase conversion (a decline of 2.11%) conditional on search as compared to CG. The higher search incidence combined with lower page views and purchase conversion conditional on search is explained by our model as follows: there is both a positive upstream effect (on search incidence) and a negative downstream effect (on search depth and conversion) for the low-value segment. While our objective is not to explain why this occurs, a plausible explanation is that the low-value segment is motivated by the base coupon to initiate the search, but finds the search less fruitful due to the need to find a product that meets the coupon condition while also matching her preferences. We stress, however, that the lift in search incidence is large and dominates any negative effect downstream in the search process, leading to the higher ARPC result reported in Section 4.1.

When examining page views conditional on search for the high-value segment, there is no change between CG (44.1 page views) and T1 (43.3 page views). The high-value segment also displays no change in purchase conversion conditional on search in T1 (13.62%) as compared to CG (13.58%). The higher search incidence combined with unchanged page views and purchase conversion conditional on search is explained by our model as follows: there is a positive upstream effect (on search incidence) and no downstream effect (on search depth and conversion) for the high-value segment. It appears that the coupon motivates a greater proportion of high-value customers to begin search but does not otherwise affect downstream in the search process. The lift in ARPC reported in Section 4.1 for the high-value segment under the base coupon can therefore be attributed mostly to a lift in search incidence.[11]

When we examine how the base coupon affects customer search activity, we find a remarkably similar lift in search incidence for both segments, albeit from different baseline levels. Given that most of the purchasers who enter the search-to-purchase funnel do not redeem the coupon, we argue that the advertising effect of the coupon is roughly equal across the two segments. The apparent greater lift in purchase incidence (and ARPC) for the high-value segment is therefore driven by the baseline heterogeneity in search depth and purchase conversion from search rather than due to the coupon. However, it can be argued that the advertising effect of the base coupon is primarily to "get customers through the door" and we therefore find little evidence to support heterogeneous treatment effects for the base coupon based on search data.

[11] We found no statistical difference in the measures of search depth between redeemers and non-redeemers across customer segments.

Better Coupon. We now examine the effect of the better coupon on customer search activity. As described in Section 4.1, the better coupon increases the ARPC relative to the base coupon for the high-value segment, but also increases the redemption rate. Table 6b shows that as compared to CG, the lift in search incidence for the high-value segment is about the same for the better coupon (22.67%) as the base coupon (23.22%). However, both search depth and purchase conversion increase for the better coupon relative to the other two conditions. The overall evidence points towards a shift in the mean of the utility distribution (see Section 2) that lifts search incidence, search depth, and purchase conversion. Hence, the better coupon does about as well as the base coupon upstream in the search process, but also improves the downstream search behavior. Hence, the improved ARPC from the better coupon as compared to the base coupon for the high-value segment is not due to improvements in search incidence but in downstream search behavior.

Our findings suggest that the mechanism by which a coupon affects customer search is moderated by coupon value. A more attractive coupon ($10 off compared to $7 off) elicited more effort by high-value customers in the search process. Yet, the modest increase in coupon redemption (still less than 20%) for the better coupon suggests that this firm may improve customer revenue by offering better deals to the high-value segment expecting that redeemers will form a small proportion of buyers, thereby resulting in limited loss in margin due to coupon redemption.

Heterogeneous Effects of Coupons. Given the extensive amount of search undertaken by customers, we next examine how customer behavior in the pretest period may affect customer search behavior. For example, if the coupon acts as a reminder to make a purchase, might the customer check a familiar product and purchase it with minimal search? In line with our analysis for customer purchase behavior reported in Section 4.1, we conduct a regression analysis and report our results on search incidence and depth. Table 7 shows our main results, building on the full specification in Tables 4 and 5.[12]

Our findings indicate the opposite effect to the question described above. Namely, high-value customers who have bought more recently are more likely to initiate the search, search more products, and spend more time when the better coupon was offered. For the base coupon, however, recency has no effect on search incidence. But more recent purchasers search more products and spend more time online. This implies that the condition for redemption motivates search for those who are engaged with the firm's offerings and does not simply result in habituated purchases. This distinction can matter because habituated purchases could be arguably accelerated from a future purchase, such that overall sales from the customer do not change due to the coupon, just the timing is accelerated. While we do not have data to study the longer-term effect of coupons in this study, our finding suggests that customers are engaging in extensive searches in response to the coupons, which they would not otherwise have done.

4.3 Robustness Checks

In our field experiment, treatment (base or better coupon) was assigned based on a cutoff of the monetary value in the pretest period. The discontinuity in the treatment induces a discontinuity in the outcomes for individual customers at the threshold. We present additional analyses in this section that exploit the regression discontinuity (RD) at the threshold used to divide customers into segments.

[12] We also used a binary variable for the monetary value, which indicates whether or not a customer belongs to the high-value segment, and found qualitatively similar results to those reported in the manuscript.

Low-value and high-value customers can be thought of as relatively similar around a bandwidth of monetary values close to the threshold. We can exploit the threshold, which is unknown to customers, to compute a local average treatment effect at the threshold for the following set of comparisons: (1) CG for low-value segment versus T1 for high-value segment, (2) T1 for low-value segment versus CG for high-value segment, and (3) T1 for low-value segment versus T2 for high-value segment. All three comparisons feature an abrupt change in treatment at the threshold, which lends itself to a sharp RD analysis to estimate if the marginal customer at the threshold is affected by the change in treatment.

Unlike Hartmann et al. (2011), where the firm made targeted offers to casino customers in a non-experimental setting, our design involves an experiment such that we know the treatment effects on low-value and high-value segments. However, it is a useful counterfactual to construct the above three comparisons in which we can estimate effects even in the absence of randomization using RD analysis. This is useful as we can compare the RD effects with actual treatment effects to gain additional insights, and also to provide a guideline for practitioners and researchers that may non-randomly use targeted coupons to recover causal estimates. We also demonstrate how the results vary with the selected bandwidth in Appendix D.

Table 8 presents estimates of the RD estimator with a bandwidth of $40 (about 0.5 standard deviations of the monetary value in the pretest period) around the threshold ($87.50) dividing low-value and high-value segments. The outcome variables are purchase incidence and search incidence. We can estimate from these regressions the estimated treatment effect at the threshold for (1) the high-value segment receiving a base coupon (using the low-value control group to the left of the threshold as the counterfactual), (2) the low-value segment receiving a base coupon (using the high-value control group to the right of the threshold as the counterfactual), and (3) the high-value segment receiving a better coupon (using the low-value base coupon group to the left of the threshold as the counterfactual). The estimates simply involve adding or subtracting the appropriate terms in the regression to construct the difference between two groups.

Referring to the first column of Table 8 with purchase incidence, we see that the RD analysis is effective because the main and interaction effects of the high-value segment are insignificant (this should be the case since we are focused on data points around the discontinuity threshold). We find that the effect of the base coupon for the low-value segment at the threshold is estimated as 1.62% (p-value < 0.001), which is comparable in magnitude to the 1.47% overall effect size for the low-value segment under the base coupon (Table 3a).[13] The effect of the base coupon for the high-value segment at the threshold is estimated as 2.64% (p-value < 0.001), which is comparable in magnitude to the 3.22% overall effect size for the high-value segment under the base coupon (Table 3b). The effect for the better coupon for the high-value segment at the threshold is estimated as 3.24% (p-value < 0.001) relative to the base coupon, which is comparable in magnitude to the 2.93% overall lift for the high-value segment going from the base to better coupon (Table 3b).

We see similar patterns with search incidence, which is presented in the second column of Table 8. We find that the effect of the base coupon for the low-value segment at the threshold is estimated as 22.56% (p-value < 0.001), which is comparable in magnitude to the 23.73% overall effect size for the low-value segment under the base coupon (Table 3a). The effect of the base coupon for the high-value segment at the threshold is estimated as 21.14% (p-value <

[13] The estimate 1.62% can be obtained from Table 8 as follows: (1) compute the estimate for the base coupon to the low-value segment and the estimate for the high-value segment in the control group and (2) calculate the difference between the two estimates, which leads to 1.62%.

0.001), which is comparable in magnitude to the 23.22% overall effect size for the high-value segment under the base coupon (Table 3b). The effect size for the better coupon for the high-value segment at the threshold is estimated as -0.37% (not significant) relative to the base coupon, which is comparable in magnitude to the -0.55% overall change for the high-value segment going from the base to better coupon (Table 3b).

The RD approach with customer purchase and search behaviors revealed that RD estimates valid treatment effects for purchase incidence and search incidence. In this regard, a real-world implementation of targeted coupons can exploit the data around the discontinuity to reduce the sample size requirement for any future experiment. Alternatively, our experimental design can allow for testing of RD estimators relative to the gold standard of randomization. The robustness of the results at the threshold suggests that our treatment effects are found across the spectrum of customers in each segment and not only at particular monetary amounts. Finally, as we show in Appendix D, using smaller bandwidths still detects most of the true effects except in one case: the treatment effect of the base coupon on purchase incidence is smaller than the other effects and therefore is not statistically significant for the bandwidths lower than $40. Since we have the benefit of knowing that there is an effect of the base coupon, our analysis suggests that a bandwidth of $40 (± 0.5 standard deviation) is reasonable for effective RD analysis. Our analysis is complementary to the theoretical work on optimal bandwidth selection (Imbens and Kalyanaraman 2012).

5. Discussion

In this paper, we develop theory on how heterogeneous customers search and purchase and how this is affected by couponing. We then design and conduct a field experiment to generate a rich set of data to both estimate the effects of couponing at the aggregate and segment levels and examine the evidence in light of our theory. In doing so, we see our primary contributions as (1) documenting the effects of couponing on a heterogeneous customer base using data from a field test, and (2) examining the mechanisms driving the effects of coupons from the top to the bottom of the search-to-purchase funnel. Our study also provides researchers and practitioners a guideline for bandwidth selection in regression discontinuity analysis when recovering causal estimates of marketing effort from non-random observational data.

Our findings both corroborate and expand on existing literature. The literature documents a promotion-as-advertising effect in which coupons stimulate purchases that mostly do not involve redemption. Venkatesan and Farris (2012) study an offline setting (without randomized treatment and control groups) while in Sahni et al. (2016) the coupon is so specific there is practically no redemption. Our online commerce setting provides a much easier opportunity for consumers to redeem coupons if they can meet the condition for redemption. The fact that less than only about 20% of customers purchasing under the coupon conditions (for either high-value or low-value segment) redeem suggests that the advertising effect is pervasive and should be considered carefully by managers beyond the usual price discrimination rationale for couponing (e.g., Narasimhan 1984). In other words, couponing isn't just for consumers who redeem them, but for others who are stimulated into action as well (without eventual redemption). One direction for future research is to examine how a demand model estimated without consideration of the advertising effect may be biased.

Our paper, however, goes beyond documenting the advertising effect of coupons on outcome measures such as purchase incidence and ARPC. We also examine the effect of coupons on search behaviors using clickstream data from the customers involved in the experiment. We find that the coupons in our study (both base and better coupons) result in increasing search incidence across both low-value and high-value segments. As such, if the role of the coupon is advertising, its effect on both segments is similar in terms of the increase in adding more customers to the top of the purchase funnel, despite heterogeneity in treatment effects when examining average revenue alone. This apparent contradiction is explained by two factors: (1) the high-value segment performs better than the low-value segment at the downstream of the search-to-purchase funnel, with a larger number of products searched and higher purchase conversion even in the absence of a coupon; therefore, baseline heterogeneity means that improvement at the top of the funnel filters down to better lift in average revenue for the high-value segment, and (2) the low-value segment in fact experiences a negative effect later in the purchase funnel in the presence of a coupon with a lower number of products searched and lower purchase conversions, which further accentuates the effect on high-value customers. To the best of our knowledge, we see this paper as the first attempt to understand how marketing activity affects different stages of the search-to-purchase funnel among heterogeneous customers.

Augmenting these findings, we identified a key moderator of these customer search and purchase behaviors: coupon value. The base coupon for the high-value segment did not affect customer search beyond the upstream effect. However, the better coupon targeted at the high-value segment worked differently. Its upstream effect was similar to the base coupon, but it further increased the number of products searched and purchase conversion. Thus, coupon value moderates the explanation provided for the differences by our model: the better coupon shifts the utility distribution for the high-value segment such that positive effects are seen all through the purchase funnel, but the base coupon has the upstream effect only.

As discussed previously, several papers argue against rewarding high-value customers with better coupons than low-value customers (e.g., Anderson and Simester 2001, Lal and Bell 2003, Anderson and Simester 2004, Dholakia 2006, Musalem and Joshi 2008) with the argument that the firm would lose margin without much upside in terms of behavioral change from an already loyal segment. Our results with a conditional promotion suggest a different perspective. Because high-value customers are already heterogeneous in their search-to-purchase behaviors (in a direction beneficial to the firm) compared to low-value customers, any lift *upstream* in the funnel gets magnified in terms of revenue lift. Further, we find that increasing coupon value to the high-value segment results in further improvements *downstream* in the funnel without sacrificing much due to increased coupon redemption. Thus, the advertising effect of coupons changes the decision calculus of prescriptions for targeting.

We see our work as extending the analysis of consumer search models from identification of search model parameters (e.g., Honka and Chintagunta 2017) to studying how firms can influence consumer search (e.g., Fong 2017), and where in the search-to-purchase funnel the effects of marketing activity occur. Seiler and Yao (2017) find that offline feature advertising does not generate a top-of-the-funnel effect in grocery stores. In contrast, we find a top-of-the-funnel effect in an online setting. This suggests that effects may well be moderated by the differences in channels and/or in product categories, which is a rich area for further study to better understand the heterogeneity in treatment effects in various marketing contexts.

We note that our findings do not seem to be driven by a reminder effect of coupons getting people to purchase if they have not done so for a while. In contrast, customers who bought more recently are more likely to buy under the coupon condition. One caveat of our experiment is that we do not observe if the revenue lift under couponing reflect acceleration of purchases such that the long-term effect of the coupon may be lower (i.e., due to intertemporal substitution of purchases by customers). We suggest the exploration of temporal couponing effects as a fruitful area for future field experimentation.

We hope our work encourages further study on the topic of targeted coupons and other marketing activities and how they affect search and purchase behaviors of heterogeneous customers. We used a specific type of conditional coupon that required the customer to meet a given condition for redemption. The effects of promotions may be moderated by what conditions (if any) are included, which we suggest as an opportunity for future research to explore. For instance, it is quite likely that a coupon that can be applied to any basket will have a much higher redemption rate (if the purchase amount required is not too high). Coupons may also be targeted based on consumer preferences rather than their overall value to the firm (e.g., Fong 2017), which can give rise to different effects. Understanding the variety of effects across promotion types and targeting practices can lead to a richer understanding of how firms can influence consumer search and purchase behavior in a profitable manner.

References

Acquisti, Alessandro and Hal R. Varian (2005), "Conditioning Prices on Purchase History," *Marketing Science*, 24 (3), 367-381.

Anderson, Eric T. and Duncan I. Simester (2001), "Price Discrimination as an Adverse Signal: Why an Offer to Spread Payments May Hurt Demand," *Marketing Science*, 20 (3), 315-327.

Anderson, Eric T. and Duncan I. Simester (2004), "Long-Run Effects of Promotion Depth on New Versus Established Customers: Three Field Studies," *Marketing Science*, 23 (1), 4-20.

Andrews, Michelle, Jody Goehring, Sam Hui, Joseph Pancras, and Lance Thornswood (2016), "Mobile Promotions: A Framework and Research Priorities," *Journal of Interactive Marketing,* 34, 15-24.

Blattberg, Robert C., Byung-Do Kim, and Scott A. Neslin (2008), *Database Marketing: Analyzing and Managing Customer*, Springer, New York, NY.

Dholakia, Utpal M. (2006), "How Customer Self-Determination Influences Relational Marketing Outcomes: Evidence from Longitudinal Field Studies," *Journal of Marketing Research*, 43 (1), 109-120.

Dong, Diansheng and Harry M. Kaiser (2005), "Coupon Redemption and Its Effect on Household Cheese Purchases," *American Journal of Agricultural Economics*, 87 (3), 689-702.

Fong, Nathan. M. (2017), "How Targeting Affects Customer Search: A Field Experiment," *Management Science*, 63 (7), 2353-2364.

Gnedenko, Boris V. (1998), *The Theory of Probability*, 6th edition, CRC Press, Boca Raton, FL.

Gordon, Brett. R., Florian Zettelmeyer, Neha Bhargava, and Dan Chapsky (2018), "A Comparison of Approaches to Advertising Measurement: Evidence from Big Field Experiments at Facebook," Working Paper, Northwestern University.

Hartmann, Wesley, Harikesh S. Nair, and Sridhar Narayanan (2011), "Identifying Causal Marketing Mix Effects Using a Regression Discontinuity Design," *Marketing Science*, 30 (6), 1079-1097.

Homburg, Christian, Mathias Droll, and Dirk Totzek (2008), "Customer Prioritization: Does It Pay Off, and How Should It Be Implemented?," *Journal of Marketing*, 72 (5), 110-130.

Honka, Elisabeth and Pradeep Chintagunta (2017), "Simultaneous or Sequential? Search Strategies in the US Auto Insurance Industry," *Marketing Science*, 36 (1), 21-42.

Imbens, Guido and Karthik Kalyanaraman (2012), "Optimal Bandwidth Choice for the Regression Discontinuity Estimator," *Review of Economic Studies* 79 (3), 933-959.

Kumar, V. and Denish Shah (2004), “Building and Sustaining Profitable Customer Loyalty for the 21st Century,” *Journal of Retailing*, 80 (4), 317-330.

Lal, Rajiv and David E. Bell (2003), “The Impact of Frequent Shopper Programs in Grocery Retailing,” *Quantitative Marketing and Economics* 1 (2), 179-202.

Lee, Leonard and Dan Ariely (2006), “Shopping Goals, Goal Concreteness, and Conditional Promotions,” *Journal of Consumer Research* 33 (1), 60-70.

Musalem, Andres and Yogesh V. Joshi (2009), “How Much Should You Invest in Each Customer Relationship? A Competitive Strategic Approach,” *Marketing Science*, 28 (3), 555-565.

Narasimhan, Chakravarthi (1984), “A Price Discrimination Theory of Coupons,” *Marketing Science*, 3 (2), 128-147.

Neslin, Scott. A. and Robert W. Shoemaker (1983), “A Model for Evaluating the Profitability of Coupon Promotions,” *Marketing Science*, 2 (4), 361-388.

Neslin, Scott. A, Caroline Henderson, and John Quelch (1985), “Consumer Promotions and the Acceleration of Product Purchases,” *Marketing Science*, 4 (2), 147-165.

Neslin, Scott. A. (1990), “A Market Response Model for Coupon Promotions,” *Marketing Science*, 9 (2), 125-145.

Rossi, Peter E., Robert E. McCulloch, and Greg M. Allenby (1996), “The Value of Purchase History Data in Target Marketing,” *Marketing Science*, 15 (4), 321-340.

Sahni, Navdeep S., Dan Zou, and Pradeep K. Chintagunta (2016), “Do Targeted Discount Offers Serve as Advertising? Evidence from 70 Field Experiments,” *Management Science*, 63 (8), 2688-2705.

Sahni, Navdeep S., S. Christian Wheeler, and Pradeep K. Chintagunta (2018), “Personalization in Email Marketing: The Role of Noninformative Advertising Content,” *Marketing Science,* 37 (2), 236-258.

Seiler, Stephen (2013), “The Impact of Search Costs on Consumer Behavior: A Dynamic Approach,” *Quantitative Marketing and Economics*, 11 (2), 155-203.

Seiler, Stephen and Song Yao (2017), “The Impact of Advertising along the Conversion Funnel,” *Quantitative Marketing and Economics*, 15 (3), 241-278.

Shin, Jiwoong and K. Sudhir (2010), “A Customer Management Dilemma: When Is It Profitable to Reward One's Own Customers?,” *Marketing Science*, 29 (4), 671-689.

Simester, Duncan, Artem Timoshenko, and Spyros I. Zoumpoulis (2017), "Efficiently Evaluating Targeting Policies Using Field Experiments," Working Paper, Massachusetts Institute of Technology.

Stigler, George J. (1961), "The Economics of Information," *Journal of Political Economy*, 69 (3), 213-225.

Thomas, Jacquelyn S., Werner Reinartz, and V. Kumar (2004), "Getting the Most out of All Your Customers," *Harvard Business Review*, 82 (7-8), 116-123.

Venkatesan, Rajkumar and Paul W. Farris (2012), "Measuring and Managing Returns from Retailer-Customized Coupon Campaigns," *Journal of Marketing*, 76 (1), 76-94.

Table 1a: Pretest Descriptive Statistics of Purchase Amount ($): Low-Value Segment

	T1	CG	Difference T1-CG
Mean	40.30	39.88	0.42
			(0.48)
Min	17.60	18.60	-1.00
p10	20.96	20.83	0.13
p25	25.35	25.60	-0.25
p50	34.60	35.00	-0.40
p75	51.90	51.15	0.75
p90	69.80	67.70	2.10
Max	87.40	87.20	0.20
Observations	4787	1998	

Notes: *** $p < 0.001$; ** $p < 0.01$; * $p < 0.05$. Standard errors appear in parentheses.

Table 1b: Pretest Descriptive Statistics of Purchase Amount ($): High-Value Segment

				Difference		
	T1	T2	CG	T2-T1	T2-CG	T1- CG
Mean	160.49	160.75	162.26	0.26	-1.51	-1.77
				(1.79)	(2.47)	(2.65)
Min	87.50	87.50	87.50	0.00	0.00	0.00
p10	96.61	95.00	97.60	-1.61	-2.60	-0.99
p25	107.91	109.07	110.80	1.16	-1.73	-2.89
p50	139.75	141.13	140.36	1.38	0.77	-0.61
p75	201.36	200.05	204.10	-1.31	-4.05	-2.74
p90	263.45	261.62	264.80	-1.83	-3.18	-1.35
Max	325.08	333.00	381.32	7.92	-48.32	-56.24
Observations	1927	3446	801			

Notes: *** $p < 0.001$; ** $p < 0.01$; * $p < 0.05$. Standard errors appear in parentheses.

Table 2: Group-Level Customer Purchase

	Study Sample			Difference in Study Sample			Population		
	T1	T2	CG	T2-T1	T2-CG	T1-CG	T1	T2	CG
Purchase (%)	3.65	9.52	1.68	5.87***	7.84***	1.97***	3.42	—	1.55
				(0.48)	(0.60)	(0.39)			
Redeem (%)	0.58	1.89	—	1.31***	—	—	0.56	—	—
				(0.21)					
Purchase ($)	34.04	32.92	31.26	-1.12	1.66	2.78	33.04	—	29.68
				(2.89)	(4.09)	(6.32)			
ARPC ($)	1.24	3.13	0.52	1.89***	2.61***	0.72***	1.13	—	0.46
				(0.23)	(0.25)	(0.20)			
Observations	6714	3446	2799						

Notes: *** $p < 0.001$; ** $p < 0.01$; * $p < 0.05$. Standard errors appear in parentheses. For T2, we have only one stratum, only high-value customers. Therefore, we could not compute a population-level estimate that includes the low-value customers when there are none in T2.

Table 3a: Customer Purchase: Low-Value Segment

	T1	CG	Difference T1-CG
Purchase (%)	2.47	1.00	1.47***
			(0.38)
Redeem (%)	0.48	—	—
ARPC ($)	0.71	0.25	0.46***
			(0.14)
Observations	4787	1998	

Notes: *** $p < 0.001$; ** $p < 0.01$; * $p < 0.05$. Standard errors appear in parentheses.

Table 3b: Customer Purchase: High-Value Segment

	T1	T2	CG	Difference		
				T2-T1	T2-CG	T1- CG
Purchase (%)	6.59	9.52	3.37	2.93***	6.15***	3.22***
				(0.79)	(1.08)	(0.97)
Redeem (%)	0.83	1.89	—	1.06***	—	—
				(0.35)		
ARPC ($)	2.56	3.13	1.20	0.57	1.93***	1.36**
				(0.41)	(0.47)	(0.62)
Observations	1927	3446	801			

Notes: *** $p < 0.001$; ** $p < 0.01$; * $p < 0.05$. Standard errors appear in parentheses.

Table 4: Heterogeneous Effects of Coupons on Customer Purchase

	(1)	(2)	(3)
Base coupon	0.020***	0.019***	0.021***
	(0.0033)	(0.0033)	(0.0036)
Better coupon	0.048***	0.047***	0.039***
	(0.0063)	(0.0062)	(0.0059)
Monetary value	0.035***	0.026***	0.027***
	(0.0036)	(0.0037)	(0.0037)
Recency		-0.0002***	-0.00004
		(0.00002)	(0.00002)
Base coupon × Recency			-0.00009***
			(0.00003)
Better coupon × Recency			-0.0004***
			(0.00005)
Intercept	0.021***	0.023***	0.021***
	(0.0026)	(0.0026)	(0.0027)
Observations	12,959	12,959	12,959
R-squared	0.030	0.037	0.042

Notes: *** $p < 0.001$; ** $p < 0.01$; * $p < 0.05$. Results are from a linear model with purchase incidence as dependent variable. Robust standard errors appear in parentheses. The variable monetary value (00s) is the past purchase amount and has been centered with a threshold value ($87.50) for the customer segment. The recency variable is the number of days since the last purchase of a product priced at $20 or higher and has been mean-centered.

Table 5: Heterogeneous Effects of Coupons on Coupon Redemption

	(1)	(2)	(3)
Better coupon	0.010***	0.010***	0.007**
	(0.0026)	(0.0026)	(0.0025)
Monetary value	0.004*	0.003	0.003
	(0.0016)	(0.0017)	(0.0017)
Recency		-0.00002**	-0.000006
		(0.000006)	(0.000005)
Better coupon × Recency			-0.00006**
			(0.00002)
Intercept	0.006***	0.006***	0.006***
	(0.0010)	(0.0010)	(0.0010)
Observations	10,160	10,160	10,160
R-squared	0.004	0.005	0.006

Notes: *** $p < 0.001$; ** $p < 0.01$; * $p < 0.05$. Results are from a linear model with coupon redemption as dependent variable. Robust standard errors appear in parentheses. The variable monetary value (00s) is the past purchase amount and has been centered with a threshold value ($87.50) for the customer segment. The recency variable is the number of days since the last purchase of a product priced at $20 or higher and has been mean-centered.

Table 6a: Customer Search: Low-Value Segment

	T1	CG	Difference T1-CG
Search (%)	34.89	11.16	23.73^{***}
			(1.16)
Page views	9.15	3.84	5.31^{***}
			(0.70)
Search time	229.13	89.20	139.93^{***}
(seconds)			(20.19)
Observations	4787	1998	

Notes: *** $p < 0.001$; ** $p < 0.01$; * $p < 0.05$. Standard errors appear in parentheses.

Table 6b: Customer Search: High-Value Segment

				Difference		
	T1	T2	CG	T2-T1	T2-CG	T1- CG
Search (%)	43.44	42.89	20.22	-0.55	22.67^{***}	23.22^{***}
				(1.41)	(1.88)	(1.98)
Page views	18.81	20.89	8.92	2.08	11.97^{***}	9.89^{***}
				(1.56)	(2.14)	(1.84)
Search time	478.56	503.71	231.28	25.15	272.43^{***}	247.28^{***}
(seconds)				(44.98)	(61.26)	(59.04)
Observations	1927	3446	801			

Notes: *** $p < 0.001$; ** $p < 0.01$; * $p < 0.05$. Standard errors appear in parentheses.

Table 7: Heterogeneous Effects of Coupons on Customer Search

	Search Incidence	ln(Page Views+1)	ln(Search Time+1)
Base coupon	0.237***	0.598***	1.226***
	(0.0089)	(0.0287)	(0.0526)
Better coupon	0.191***	0.472***	0.973***
	(0.0125)	(0.0412)	(0.0747)
Monetary value	0.048***	0.233***	0.375***
	(0.0067)	(0.0235)	(0.0415)
Recency	-0.0002***	-0.0006***	-0.001***
	(0.00005)	(0.0002)	(0.0003)
Base coupon × Recency	-0.0001	-0.0005**	-0.001**
	(0.00006)	(0.0002)	(0.0004)
Better coupon × Recency	-0.0007***	-0.003***	-0.005***
	(0.00009)	(0.0003)	(0.0006)
Intercept	0.149***	0.454***	0.854***
	(0.0068)	(0.0224)	(0.0405)
Observations	12,959	12,959	12,959
R-squared	0.081	0.082	0.082

Notes: *** $p < 0.001$; ** $p < 0.01$; * $p < 0.05$. The variable monetary value (00s) is the past purchase amount and has been centered with a threshold value ($87.50) for the customer segment. The recency variable is the number of days since the last purchase of a product priced at $20 or higher and has been mean-centered.

Table 8: RD Estimates of Coupons on Customer Purchase and Search

	Purchase Incidence	Search Incidence
Base coupon	0.018*	0.218***
	(0.0093)	(0.0224)
Better coupon	0.049***	0.208***
	(0.0119)	(0.0286)
High-value	0.002	0.007
	(0.0133)	(0.0320)
Base coupon × High-value	0.006	0.001
	(0.0158)	(0.0379)
Intercept	0.013	0.162
	(0.0079)	(0.0189)
Observations	4,634	4,634
R-squared	0.009	0.034

Notes: *** $p < 0.001$; ** $p < 0.01$; * $p < 0.05$. Results are from a linear model with purchase incidence and search incidence as dependent variable. Robust standard errors appear in parentheses.

Figure 1: Pretest Descriptive Statistics

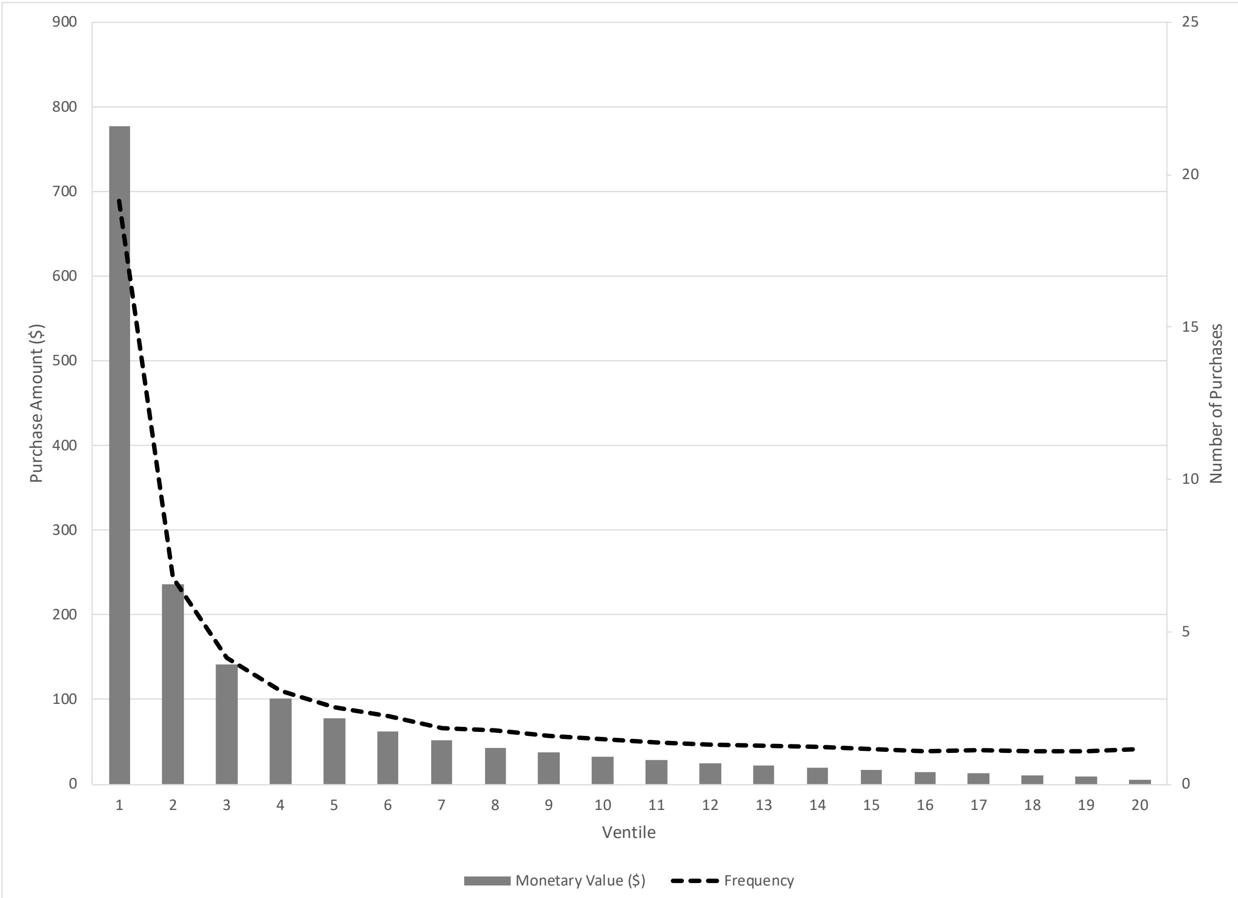

Figure 2: Design of the Study

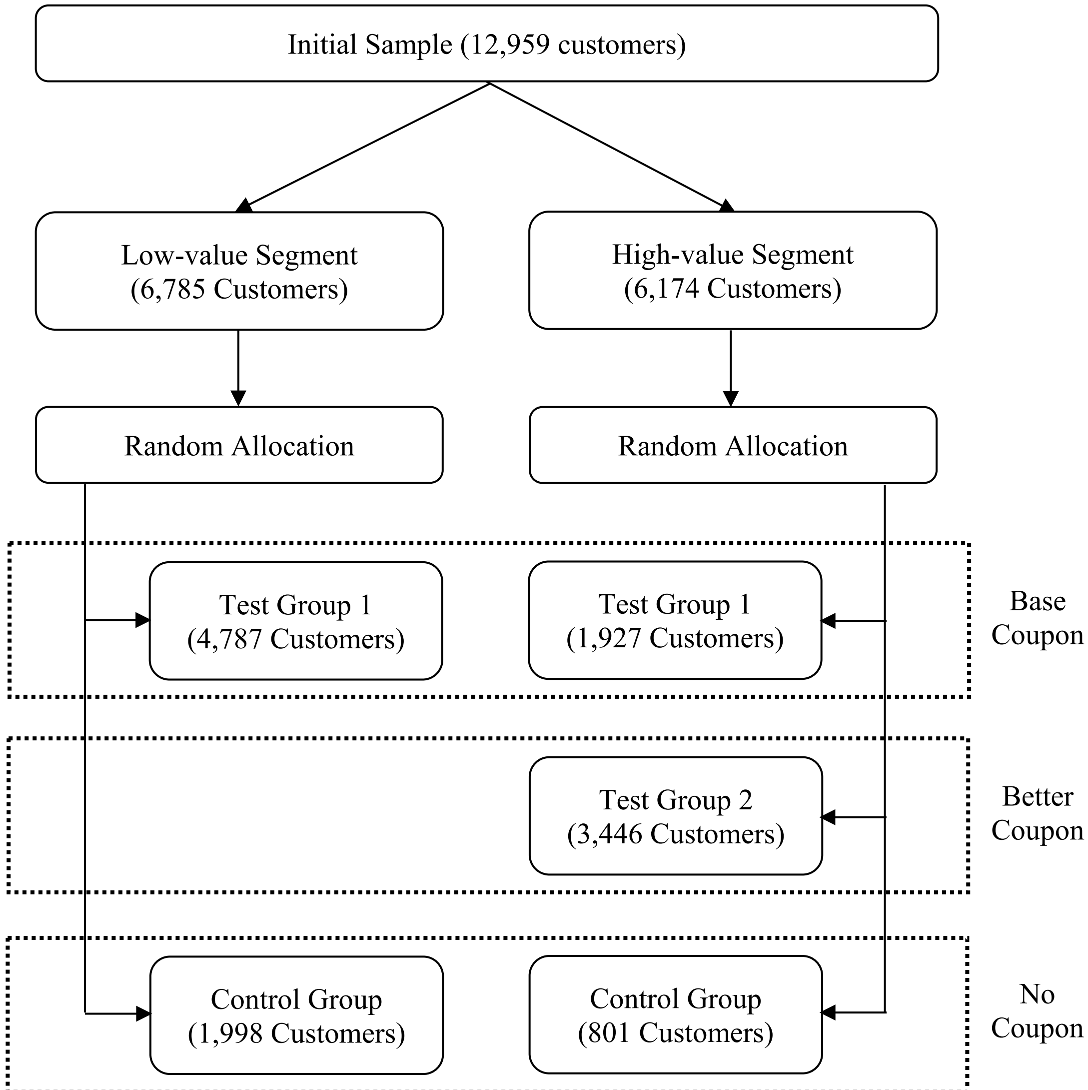

Appendix A: Proofs

Proof of Proposition 1: The primary difference between the high type and the low type is that we assume the high type has a higher mean of the utility distribution. For the purpose of illustration, suppose the distributions for both types follow normal distributions with type-dependent mean m_i and common variance s^2. The objective function for type i who has decided to search is:

$$\underset{K}{\operatorname{argmax}}\, E[y_{iK} \cdot I(y_{iK} > U^*)] - K \cdot c.$$

where y_{iK} is the maximum of K i.i.d. random variables $\{x_{i1}, x_{i2}, \dots . x_{iK}\}$ drawn from a normal distribution with mean m_i and variance s^2. U^* is the reservation utility at which the customer would actually make a purchase, otherwise utility is set to 0. While it is possible to write down the cdf for y_{iK} as the product of K normal distributions, this does not yield a tractable form for evaluation as it is difficult to compute expectations analytically or numerically.

We therefore use the Fisher-Tippett-Gnedenko extreme value theorem (Gnedenko 1998) that states that the asymptotic distribution for the maximum order statistic will be distributed Extreme Value Type I when the initial cdf has an exponential tail (as does the normal distribution). In practice, the asymptotic distribution provides a very close estimate of the actual distribution even for $K > 7$ in the case of the normal distribution.

Define a standardized variable $x'_{ij} = \frac{x_{ij} - m_i}{s}$. By the Fisher-Tippet-Gnedenko extreme value theorem, the maximum of K such variables, which we denote as v_{iK} can be approximated as: ExtremeValueDist(μ_K, σ_K), where $\mu_K = \Phi^{-1}\left(1 - \frac{1}{K}\right)$ and $\sigma_K = \Phi^{-1}\left(1 - \frac{e^{-1}}{K}\right) - \mu_K$.

For $K > 7$, note that by the property of the extreme value Type I distribution (also known as Gumbel), $E[v_{iK}] = \mu_K - \psi(1) \cdot \sigma_K$, where ψ is the digamma function and $-\psi(1)$ is the Euler-Mascheroni constant.

$E[y_{iK}] = s \cdot E[v_{iK}] + m_i$. This helps translate the standardized variables back to the original utility scale. Note that it is not necessary that the initial utility distribution is normally distributed since the extreme value Type I distribution is the limiting distribution for the maximum order statistic of any distribution that has an exponential right tail.

However, what we are after is $E[y_{iK} \cdot I(y_{iK} > U^*)]$, which includes the effect of censoring when $y_{iK} \le U^*$.

$$\begin{aligned}
E[y_{iK} \cdot I(y_{iK} > U^*)] &= E\left[s \cdot (v_{iK} + m_i) \cdot I\left(v_{iK} > \frac{U^* - m_i}{s}\right)\right] \\
&= s \cdot E\left[v_{iK} \cdot I\left(v_{iK} > \frac{U^* - m_i}{s}\right)\right] + s \cdot m_i \cdot E\left[I\left(v_{iK} > \frac{U^* - m_i}{s}\right)\right] \\
&= s \cdot E[v_{iK} \cdot I(v_{iK} > U_i^*)] + s \cdot m_i \cdot \Pr(v_{iK} > U_i^*)
\end{aligned}$$

where $U_i^* = \frac{U^* - m_i}{s}$ is decreasing in m_i.

$\Pr(v_{iK} > U_i^*)$ is decreasing in U_i^* and therefore increasing in m_i.

We now need to show that $E[v_{iK} \cdot I(v_{iK} > U_i^*)]$ is decreasing in U_i^* for $K \geq \underline{K}$ (where $\underline{K}$ is some lower bound), which would also imply that it is increasing in m_i. Analytically, this expression is not tractable. We use numerical integration over a large range of K's (from K = 8 to K = 200) to show that $E[v_{iK} \cdot I(v_{iK} > U_i^*)]$ has an inverse S-shaped curve for all these values of $K \geq 8$. Example curves are shown for various values of K below.

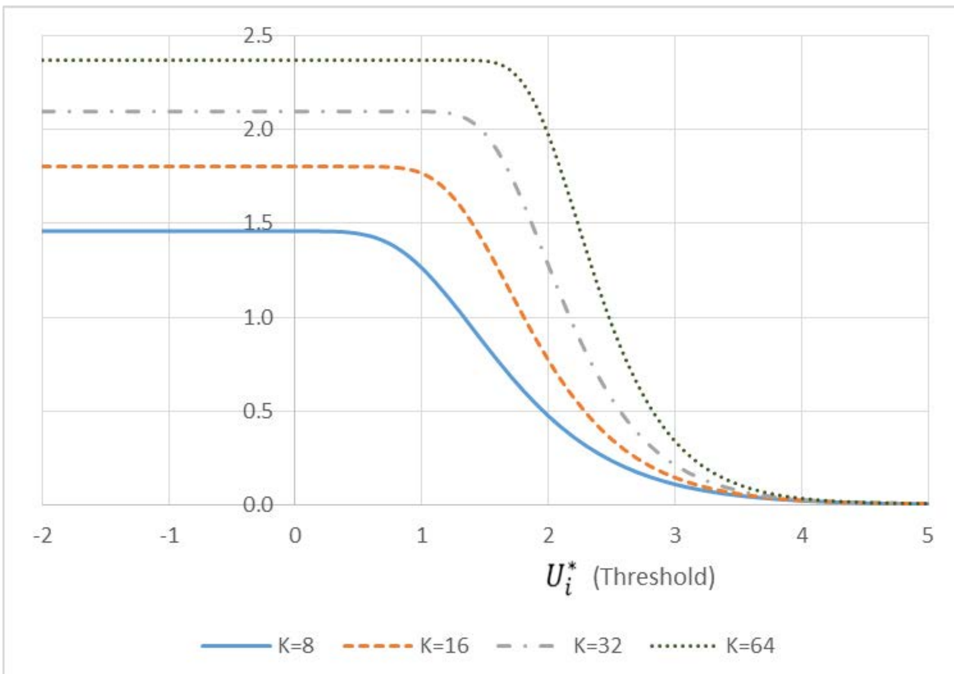

Figure A1: $E[v_{iK} \cdot I(v_{iK} > U_i^*)]$ as a Function of Threshold (U_i^*)

Thus, for the range of number of products searched that is practical, we observe that $E[v_{iK} \cdot I(v_{iK} > U_i^*)]$ is decreasing in U_i^* and therefore increasing in m_i. Therefore, $E[y_{iK} \cdot I(y_{iK} > U^*)]$ is increasing in m_i and $E[y_{LK} \cdot I(y_{LK} > U^*)] < E[y_{HK} \cdot I(y_{HK} > U^*)]$ $\forall\, K \geq 8$.

Hence, since the search cost is linear in K, this line will cross $E[y_{HK} \cdot I(y_{HK} > U^*)]$ at a higher value of K than $E[y_{LK} \cdot I(y_{LK} > U^*)]$ if $E[v_{iK} \cdot I(v_{iK} > U_i^*)]$ is concave in K. We again use numerical integration to examine the shape of $E[v_{iK} \cdot I(v_{iK} > U_i^*)]$ as a function of K (we start at K=8 per the analysis above) for various fixed values of U_i^*, as shown below. The concavity holds for $U_i^* < 1$ as well. As $U_i^* > 4$, the function's value becomes close to zero (as it is almost impossible to find a satisfactory product even with extremely large K) and concavity is violated. However, this is unlikely to occur as $U_i^* = 4$ represents a threshold that is four standard deviations above type i's mean utility.

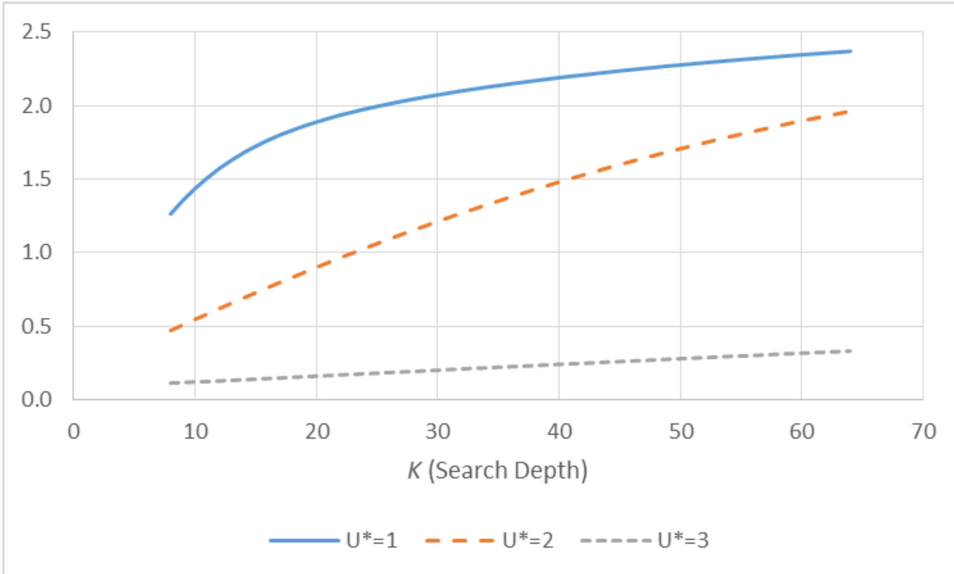

Figure A2: $E[v_{iK} \cdot I(v_{iK} > U_i^*)]$ as a Function of Search Depth (K)

Hence $K_H^* > K_L^*$ for $U_i^* \leq 4$.

Proof of Proposition 2: $\gamma_{i,K} = 1 - F_{EV1}(\frac{U^* - m_i}{s}; \mu_K, \sigma_K)$. $F_{EV1}(\cdot)$ is decreasing in m_i and therefore γ_{iK} is increasing in m_i. Further $\gamma_{i,K}$ is increasing in K. Since $m_H > m_L$ by assumption and $K_H^* > K_L^*$ from Proposition 1, $\gamma_{H,K_H^*} > \gamma_{L,K_L^*}$.

Proof of Proposition 3: From Proposition 1, $K_H^* > K_L^*$ and in the proof of Proposition 1 we have shown that $E[y_{LK} \cdot I(y_{LK} > U^*)] < E[y_{HK} \cdot I(y_{HK} > U^*)]$. Therefore, $E\left[y_{LK_L^*} \cdot I\left(y_{LK_L^*} > U^*\right)\right] < E\left[y_{HK_L^*} \cdot I\left(y_{HK_L^*} > U^*\right)\right]$ which means the surplus up to K_L^* is higher for the high type. The high type derives additional non-negative surplus from $(K_H^* - K_L^*)$ which ensures that the high type has a higher surplus than the low type. Consequently, the probability of search incidence is higher for the high type than the low type.

Proof of Propositions 4 and 5: The only difference between the expressions for $E[v_{iK} \cdot I(v_{iK} > U_i^*)]$ and γ_{iK} are that the types have different expected utilities m_i. We have shown in Propositions 1 and 2 that the high type searches over more products and has a higher purchase incidence than the low type. Given that only one variable drives these differences, if across two groups, one groups dominates the other both in terms of number of products searched and purchase incidence, it must be, according to our model, because the dominant group had a higher expected utility from their distribution of utilities for the firm's products.

Appendix B: Additional Pretest Descriptive Statistics

Appendix B presents additional pretest descriptives of the study which we discussed in Section 4. Table B1a (B1b) summarizes the individual-level data in the randomized field experiment for the low-value (high-value) segment between the conditions. As shown in both tables, the evidence provides further support that there is no systematic variation in the test.

Table B1a: Additional Pretest Descriptive Statistics: Low-Value Segment

			Difference
	T1	CG	T1-CG
Recency (days)			
All purchases	188.10	192.27	-4.17
			(2.96)
Purchases of a product	230.19	231.24	-1.05
priced at $20 or higher			(4.37)
Frequency			
All purchases	1.66	1.62	0.04
			(0.03)
Purchases of a product	1.25	1.24	0.01
priced at $20 or higher			(0.02)
Observations	4787	1998	

Notes: *** $p < 0.001$; ** $p < 0.01$; * $p < 0.05$. Standard errors appear in parentheses.

Table B1b: Additional Pretest Descriptive Statistics: High-Value Segment

				Difference		
	T1	T2	CG	T2-T1	T2-CG	T1- CG
Recency (days)						
All purchases	116.29	117.74	121.25	1.45	-3.51	-4.96
				(2.78)	(3.80)	(4.13)
Purchases of a product	123.77	126.74	132.23	2.97	-5.50	-8.46
priced at $20 or higher				(2.88)	(3.97)	(4.40)
Frequency						
All purchases	4.05	4.12	4.11	0.07	0.01	-0.06
				(0.08)	(0.11)	(0.12)
Purchases of a product	3.36	3.41	3.42	0.05	-0.01	-0.06
priced at $20 or higher				(0.06)	(0.08)	(0.08)
Observations	1927	3446	801			

Notes: *** $p < 0.001$; ** $p < 0.01$; * $p < 0.05$. Standard errors appear in parentheses.

Appendix C: Heterogeneous Effects of Coupons on Customer Purchase

Appendix C presents the results of the logit regression which we discussed in Section 4.1.

	(1)	(2)	(3)
Base coupon	0.808***	0.792***	0.718***
	(0.1615)	(0.1618)	(0.1643)
Better coupon	1.305***	1.261***	0.982***
	(0.1631)	(0.1631)	(0.1791)
Monetary value	0.646***	0.473***	0.480***
	(0.0532)	(0.0566)	(0.0571)
Recency		-0.004***	-0.001
		(0.0005)	(0.0011)
Base coupon × Recency			-0.003*
			(0.0012)
Better coupon × Recency			-0.005***
			(0.0013)
Intercept	-4.100***	-4.171***	-4.070***
	(0.1482)	(0.1498)	(0.1480)
Observations	12,959	12,959	12,959
Log likelihood	-2,307.2	-2,268.8	-2,261.8

Notes: *** $p < 0.001$; ** $p < 0.01$; * $p < 0.05$. Results are from a logit model with purchase incidence as dependent variable. Robust standard errors appear in parentheses. The variable monetary value (00s) is the past purchase amount and has been centered with a threshold value for the customer segment. The recency variable is the days since the last purchase of a product priced at $20 or higher and has been mean-centered.

Appendix D: RD Estimates of Coupons on Customer Purchase and Search

Appendix D presents the results of the RD estimates with varying bandwidths which we discussed in Section 4.3.

	Bandwidth: $30		Bandwidth: $20		Bandwidth: $10	
	Purchase	Search	Purchase	Search	Purchase	Search
Base coupon	0.020	0.200^{***}	0.009	0.177^{***}	-0.015	0.182^{***}
	(0.0123)	(0.0292)	(0.0161)	(0.0385)	(0.0257)	(0.0577)
Better coupon	0.045^{***}	0.232^{***}	0.060^{***}	0.220^{***}	0.079^{**}	0.225^{***}
	(0.0136)	(0.0324)	(0.0163)	(0.0390)	(0.0254)	(0.0570)
High-value	0.002	-0.034	-0.014	-0.045	-0.044	-0.038
	(0.0161)	(0.0384)	(0.0202)	(0.0483)	(0.0320)	(0.0719)
Base coupon × High-value	-0.001	0.033	0.017	0.051	0.029	0.058
	(0.0189)	(0.0451)	(0.0236)	(0.0565)	(0.0375)	(0.0842)
Intercept	0.014	0.177	0.020	0.194	0.044	0.178
	(0.0104)	(0.0248)	(0.0138)	(0.0330)	(0.0219)	(0.0492)
Observations	3285	3285	2214	2214	1066	1066
R-squared	0.007	0.031	0.011	0.025	0.020	0.026

Notes: *** $p < 0.001$; ** $p < 0.01$; * $p < 0.05$. Results are from a linear model with purchase incidence and search incidence as dependent variable. Robust standard errors appear in parentheses.