

# When Does Retargeting Work? Information Specificity in Online Advertising\*

Anja Lambrecht<sup>†</sup> and Catherine Tucker<sup>‡</sup>

May 6, 2013

---

\*We thank Havas Digital and particularly Katrin Ribant for access to data from Artemis and Marco Bertini for facilitating the contact to Havas Digital. We gratefully acknowledge financial support from the London Business School Centre for Marketing. We thank Kristin Diehl, Anindya Ghose, Avi Goldfarb, Brett Gordon, Duncan Simester, Catalina Stefanescu, Florian von Wangenheim and Ken Wilbur for their comments, as well as participants at the 2011 SICS conference and seminar participants at Cass Business School, ESMT, ESSEC, London Business School and the National University of Singapore.

<sup>†</sup>London Business School, London, UK; [alambrecht@london.edu](mailto:alambrecht@london.edu).

<sup>‡</sup>MIT Sloan School of Management, MIT, Cambridge, MA; [cetucker@mit.edu](mailto:cetucker@mit.edu). and NBER

# When Does Retargeting Work? Information Specificity in Online Advertising

## **Abstract**

Firms can now serve personalized recommendations to consumers who return to their website, based on their earlier browsing history on that website. At the same time, online advertising has greatly advanced in its use of external browsing data across the web to target internet ads. ‘Dynamic Retargeting’ integrates these two advances, by using information from earlier browsing on the firm’s website to improve internet advertising content on external websites. Consumers who previously visited the firm’s website when surfing the wider web, are shown ads that contain images of products they have looked at before on the firm’s own website. To examine whether this is more effective than simply showing generic brand ads, we use data from a field experiment conducted by an online travel firm. We find, surprisingly, that dynamic retargeted ads are on average less effective than their generic equivalent. However, when consumers exhibit browsing behavior such as visiting review websites that suggests their product preferences have evolved, dynamic retargeted ads no longer underperform.

# 1 Introduction

Innovations in the parsing and processing of individual-level browsing data enable firms to serve product recommendations in real time to consumers who return to their website. These personalized ‘recommendation systems’ often highlight the specific products that the consumer was browsing before they left the website, and may increase sales (Linden et al., 2003; Dias et al., 2008). However, consumers who browse products online often leave the website without buying and do not return. To reach out to such consumers, dynamic retargeted ads feature pictures of precisely the product consumers previously browsed.

At first blush, this makes sense: The marketing literature has emphasized that greater specificity of a firm’s interactions with consumers should increase relevance and consumer response (Hoffman and Novak, 1996; Komiak and Benbasat, 2006; Dias et al., 2008). Firms that offer retargeting services point to strong increases in advertising effectiveness. For example, Criteo (2010) reports that personalized retargeted ads are six times more effective than standard banner ads, and four times more effective than retargeting that uses generic ads. As a result, dynamic retargeting has attracted much enthusiasm among online advertising practitioners (Hunter, 2010; Hunter et al., 2010; Hargrave, 2011). For example, a single firm that sells retargeting solutions, the ‘Next Performance’ ad network, reports that it has served 30 billion retargeted impressions, analyzed 1 billion products for possible inclusion in a dynamic retargeted ad, and served dynamic retargeted ads to 500 million unique visitors.<sup>1</sup>

However, there is little empirical evidence that a personalized product recommendation is as effective when displayed on external websites, as it is when it is displayed internally on the firm’s own website. Personalized recommendation systems were designed to sell to consumers who are engaged enough to return to a firm’s website. Dynamic retargeting, on the other hand, tries to engage people who have not yet returned to the firm’s website.

---

<sup>1</sup>See <http://www.nextperformance.com>.

Despite much enthusiasm about dynamic retargeting, advertisers currently do not know whether this technique is indeed effective. Advertisers also do not know what information they can use to determine when to show such ads that feature content that is highly specific to an individual consumer. This research seeks to fill these gaps.

We empirically explore these questions using data from an online field experiment by a travel company. After consumers visited the travel company’s website and looked at hotels, an ad network showed banner ads on behalf of the travel company to these consumers when they subsequently browsed other websites. On these other websites, consumers randomly either saw an ad that contained an image of the specific hotel the consumer had previously browsed plus three similar hotels (dynamic retargeting) or an ad that showed a generic brand ad for the travel firm (generic retargeting). We find that, surprisingly, on average dynamic retargeting is *not* effective.

The crucial question for advertisers and ad networks, however, is when dynamic retargeting is effective in converting consumers to purchase. We suggest that the effectiveness of a retargeted ad depends on whether the concreteness of its message matches how narrowly consumers construe their preferences (Trope et al., 2007; Lee et al., 2010).

Consumers may initially have only a broad idea what they like. Their preferences are construed at a high level and they focus on higher-level goals. For example, they might want a ‘relaxing vacation’. Over time, consumers shift their focus to specific product attributes and they develop narrowly construed preferences. For example, they may look for a hotel with a large swimming pool near the beach. In using the term ‘narrowly construed preferences,’ we therefore refer to consumers having a detailed viewpoint on what kind of products they wish to purchase. We propose that consumers who focus on higher level goals respond better to advertising messages that address such higher level goals than to messages that display specific products. Only consumers with narrowly construed preferences are therefore likely to respond positively to the content of dynamically retargeted ads.

We empirically explore whether the effectiveness of a dynamically retargeted ad changes in parallel with consumers' browsing behavior which reflects such a shift in goals. We isolate browsing behavior which may indicate that a consumer has shifted to having narrowly construed preferences and may be more receptive to such highly-specific advertising. We use as a proxy whether a consumer has visited a travel review site. When searching product-specific information on a review site, such as TripAdvisor, a consumer is comparing and contrasting product features and confronting the trade-offs inherent in a product choice. A visit also provides evidence that consumers are prepared to evaluate products on a detailed level and indicates that the consumer is thinking deeply about specific products. Therefore we take a visit to a review site as a potential proxy measure that a consumer now has narrowly construed preferences.

We find that generic ads are most effective before consumers seek out product quality information at a review site. Dynamic retargeting becomes relatively more effective only after consumers have visited a product review site. We find that the greater effectiveness of retargeting further increases for consumers who are at that time also browsing category-level content. This is consistent with literature which suggests that the quality of the advertising message is mediated by consumers' involvement (Petty et al., 1983).

We acknowledge that visiting a travel review website, as well as being a proxy for having narrowly construed preferences, could also be a proxy for many other things. In the lab, we rule out alternative explanations such as privacy, reactance, social validation and sample selection that could also potentially explain why after visiting a travel review site consumers react more positively to a specific than to a generic ad.

Therefore, our findings suggest that dynamic retargeting is effective at encouraging consumers to purchase when consumers have visited a review site and are actively browsing other websites in the category. In all other settings, generic retargeting is more effective. Our findings about the optimal content of retargeted ads are important given the growth

of retargeting as an advertising tool. For example, large web platforms such as Facebook have now introduced retargeted ads into its members' newsfeeds as a central plank in their advertising strategy (Rusli, 2013). Additionally, our results provide operational insights for managers who are thinking about embracing this new technology. It suggests that dynamic retargeting is best employed when managers also have access to external browsing data that would help them identify if preferences have evolved and so *when* dynamic retargeting will be effective. Otherwise, more generic campaigns may work better. We next turn to discuss how this result, as well as being managerially useful, adds to the existing literature.

## 2 Relationship to Prior Literature

[Table 1 about here.]

Table 1 summarizes how our research relates to previous work on personalized recommendations, tailored communications and targeting.

Research on personalized recommendations on a firm's website has focused on both documenting their effectiveness (Dias et al., 2008) and on suggesting ways of improving their effectiveness (Linden et al., 2003). However, by their very nature these personalized recommendations are only shown to consumers who already decided to return to the firm's website. They do not reach consumers who do not return to the site.

Similarly, the literature on tailoring communications consistently finds that tailoring improves the performance of communications. Consumer characteristics can be used to identify appropriate segments to customize for, like segmenting on consumers' cognitive style (Hauser et al., 2009), celebrity affinity (Tucker, 2011), past browsing behavior such as previous ads clicked (Agarwal et al., 2009) or past purchases (Malthouse and Elsner, 2006). However, the focus of this literature is identifying which segment the individual belongs to and then showing the appropriate ad for that segment, rather than individualizing each ad.

A growing body of online targeting literature has attempted to define what kinds of data a

web-content publisher should use when deciding which ad to display to which consumer. This literature finds that data on consumer browsing behavior (Chen et al., 2009) or demographics (Joshi et al., 2011) can improve targeting. However, it does not look at whether individual advertisers might benefit from incorporating into the content of their ads information that is highly specific to individual consumers, such as their prior product interests.<sup>2</sup>

Our study builds on this literature in four different ways. First, we study personalized recommendations outside the firm’s website. Second, we focus on ad content personalized to individual browsing history, not segments. Third, we study the tailoring of ad content, not simply selecting who sees ads based on prior browsing behavior. Last, we analyze whether and how the effectiveness of such retargeted ads changes depending on whether browsing behavior suggests that their preferences have evolved. Our results show that online gathering of data on what consumers do outside the firm’s boundaries can be used not only to target but also to *time* when ads are shown.<sup>3</sup>

### 3 Empirical Setting and Data

#### 3.1 Retargeting

Retargeting is typically organized by an advertising network on behalf of a focal firm. Ad networks aggregate advertising space across multiple publishers of web content and then sell this space to advertisers. This means that advertisers do not have to manage relationships with a large number of web publishers, which significantly increases the efficiency in the market for online ad space.

Since retargeting is a new technology, we describe in more detail how it is typically implemented:

---

<sup>2</sup>Gatarski (2002) suggests an algorithm to optimize the content of banner ads within a given design format, but does not discuss tailoring the content to individual consumers or consumer segments.

<sup>3</sup>This tactic builds on work such as Lambrecht et al. (2011), that shows that detailed online data can be used to understand how different stages of a consumer’s purchase process interconnect.

1. *Product exposure*: The consumer visits the focal firm's website and views products. For each product page the consumer views, a pixel tag, that is a  $1 \times 1$  image, is downloaded automatically, recording the fact that that consumer was looking at a specific product. This information becomes part of the individual user profile held by the ad network on behalf of the focal firm, and tracked typically by cookies.
2. *Targeting Consumers*: The consumer browses the Internet. At some point, the consumer visits a website whose ads are provided by an advertising network that offers retargeted advertising. The advertising network uses the cookie to identify that the consumer has previously visited the website of the focal firm.
3. *Ad Design*: In the case of generic retargeting, the ad network uses the individual cookie profile simply to identify people who have visited the focal firm's website before and show them generic ads for the focal firm rather than showing them ads for another firm. Such generic ads typically display a picture broadly relating to the category. An airline, for example, might display their logo and a picture of a smiling air attendant, or a travel company might display a picture of a beach alongside its logo.

In the case of dynamic retargeting, the ad network additionally designs the ad to display the exact product the consumer had looked at before and sometimes other similar products the focal firm sells. In Figure 1(a) we display an example of dynamic retargeting. A consumer who browsed a pair of children's shoes at an online retailer (left panel), was later served an ad that displayed precisely this product alongside three other similar products (right panel). Dynamic retargeted ads use standardized designs where a predefined space is subdivided into multiple areas for images of specific products. This standardization reflects the need to incorporate a vast array of possible images and text in an ad using a sophisticated algorithm in real time. The standardized design means that as well as being personalized, dynamic retargeted ads are also more

complex in design than most banner ads.

4. *Purchase*: The consumer purchases from the focal firm’s website. The ad network records this purchase in its individual-level profile and links to any ad exposures. After a purchase, a consumer will typically not be retargeted unless they visit the website again.<sup>4</sup> The ad network is usually not given information about the precise product that was purchased.

### 3.2 Data

We use data on a travel website that sold hotel stays and hotel vacation packages to consumers. It advertised its services on external websites using several ad networks. The firm engaged in four types of targeted online advertising. These are summarized in Table 2.

[Table 2 about here.]

The firm conducted a field experiment, in cooperation with a major ad network. In this field experiment, the consumer was randomly exposed to a generic or a dynamic retargeted ad when they subsequently visited an external website where the ad network showed ads on behalf of the firm.<sup>5</sup>

The travel firm ran the field test for 21 days for the hotel category, which is its major product focus. All consumers who during the 21-day time period had viewed a specific hotel on the travel firm’s website were eligible for the field experiment. Reflecting the beach destination focus of the travel firm, the generic ad displayed an image that evoked a beach vacation alongside its brand logo. The dynamic retargeted ad displayed one hotel the consumer had browsed on the focal firm’s website, alongside three others that were similar

---

<sup>4</sup>Instances of post-purchase retargeting are mostly due to the ad network not properly recording the purchase or the firm aiming to convince consumers to repurchase.

<sup>5</sup>This particular retargeting network did not engage in real-time bidding for the pricing of its ads but instead used a previously-agreed rate. This reduces the potential for distortion that would result if the allocation of advertising were decided based on an auction network.

in location and star rating. We do not have information on which hotel was displayed.<sup>6</sup> In Figure 1(b), due to confidentiality agreements, we can show only an approximation of the design of the dynamic retargeting ads the firm used, though the real ads displayed were more expertly and attractively designed. This shows that the design of a dynamic retargeted ad is more complex than that of a generic banner ad. The field experiment should therefore be interpreted as a comparison of dynamic retargeting as commonly practiced relative to generic retargeting as commonly practiced.

[Figure 1 about here.]

If on any day the consumer visited multiple websites that were part of the ad network that implemented the field experiment, they would see multiple retargeted ads. However, the randomized trial was designed so that on any one day a consumer would see either only generic or dynamic retargeted ads. This means that the same consumer can be in different treatment groups on different days. This is one of our motivations for including a stock of previous ads the consumer is exposed to in our later regression analysis.

Table 3(a) summarizes the individual-level data. It covers ad exposure data collected by the firm for the 77,937 individual profiles who were part of the field experiment because they had visited *both* a part of the firm’s website devoted to a specific hotel and subsequently websites that were part of the ad network that implemented retargeting.<sup>7</sup> The firm did not store data on what products consumers initially browsed. Though only one ad network implemented retargeting, our data includes all ad exposures across all ad networks the firm collaborated with.

We also do not observe whether the consumer visits the firm’s website again. If a consumer revisits the firm’s website, then the firm may understand their preferences better and

---

<sup>6</sup>Usually, the dynamic retargeting algorithm focuses on the most recent product browsed on the website, but we do not have data to confirm that this is the case in this instance.

<sup>7</sup>As in prior research (Rutz and Bucklin, 2011), we do not have data to address measurement error introduced by multiple cookies on multiple computers leading to inaccurate matching of an individual consumer with an individual profile.

so may be better able to match products to a consumer’s preferences in a dynamic targeted ad relative to a consumer who does not revisit the firm’s website. This is an additional factor that our field data precludes us from measuring that could contribute to greater effectiveness of dynamic targeted ads over time.

*Purchase* reflects whether a consumer made a purchase online from the travel firm’s website within the time-frame of the study. In our data, 10% of consumers who had browsed a specific product and were later retargeted purchased within the 21-day observation period. Very few consumers clicked on ads. Indeed, clickthrough rates were on average less than 0.01%. However, as discussed by Dalessandro et al. (2012), clicks are not very relevant for understanding the effectiveness of the type of display advertising we study.

We do not know the type or quantity of product that the consumer purchased, but given the firm’s strong focus on selling hotels either individually or with flights, it is highly likely that it included a hotel room. We do not observe consumers’ purchases after the end of the campaign or purchases made offline..

*VisitedReviewSite* indicates that 40% of consumers were exposed to one of the firm’s ads when browsing a travel review site. As the firm was a major advertiser on the main travel review sites, it believes that most visitors to a travel review site would have been exposed to its advertising. There is a positive correlation between ever visiting a travel review site and the likelihood of purchase. 8.6% of consumers who do not visit a travel review site purchase the product. 14.6% of consumers who do visit a travel review site ultimately purchase. This suggests that a visit to a review site may not be random, something we address directly in our empirical specification.

Table 3(b) describes the data at a daily level over the 21 days, including the types and number of ads each consumer was exposed to. *RetargetedAd* summarizes that across the 21 days of the field experiment, a consumer had an 8.9% likelihood of seeing at least one retargeted ad per day. *RetargetedAd*  $\times$  *DynamicAdContent* reflects that roughly half of

these ads were dynamic retargeted ads. In total, over 149,000 generic retargeted ads and over 161,000 dynamically retargeted ads were displayed as part of the field test.

We checked the face validity of the randomization between generic and dynamic retargeted ads. There was no statistically significant relationship between whether a consumer was shown a generic or a dynamic retargeted ad ( $p=0.56$ ) on successive days. Also, consumers who had viewed a specific type of ad content on a particular day were not more likely to receive either a generic or a dynamic ad on that day (viewed travel website  $p=0.19$ , viewed news website  $p=0.21$ ). Importantly, how many ads they had previously seen also did not affect what type of retargeted ad they were shown on their next visit ( $p=0.46$ ). This evidence further supports that generic or dynamic retargeted ads were shown randomly and that there is no systematic variation in websites which showed the different types of ads that could explain our result.

*AnyAd* captures that on average, a consumer had a 21.4% probability of being exposed to at least one ad by the travel firm. *ContextualAd* captures that on 4.2% of days a consumer was exposed to a contextual targeted ad, and similarly, *OtherBehavioralAd* captures that on 12.2% of days a consumer was exposed to behavioral targeted ad. Neither of these latter categories displayed dynamic content. Similarly, we summarize the cumulative number of ads in each category that a profile viewed prior to that particular date across the 21 days of the field experiment.

Table 4(a) reports the same data as Table 3(a) but for all 2,818,661 consumers who were served any type of ad by the firm during the 21 days of the field experiment, not just those who were part of the field test. This means it includes consumers who had not viewed a specific product on the firm's website. The indicator variable *Eligible for Dynamic Retargeting Test* reflects whether or not the consumer was eligible to receive the retargeting campaign, and shows that only a small proportion of consumers were included in the field test, simply because relatively few consumers visited the firm's website and browsed its products.

Consumers who were eligible for the field test have a higher likelihood of purchase, are more likely to browse a travel review site, and are also more likely to be recorded browsing the internet in general. This suggests the large gains for retargeting that are reported in many industry studies may be because these are people who are already more likely to purchase as they have already sought the product out. In any case, our results should be interpreted as only reflecting the behavior of consumers who visit the firm’s website. However, since a necessary condition for dynamic retargeting is that a consumer has visited the website, this is the local average treatment effect of interest.

[Table 3 about here.]

[Table 4 about here.]

[Figure 2 about here.]

Figure 2 presents average daily conversion rates by whether someone who was part of the field experiment was exposed to a particular type of ad on that day. This shows that consumers who were not browsing other websites served by any of the ad networks were unlikely to purchase. However, their lack of exposure to ads could simply reflect that they were not online that day and consequently were not making online purchases (Lewis et al., 2011). In general, Figure 2 emphasizes the difficulty in ascribing causality between different types of online advertising and purchases in this kind of data, given that ad exposure is a function of a consumer’s browsing behavior which in turn may reflect other unobservable characteristics. For example, in our data it would appear that for individuals who have previously visited the firm’s website, contextual ads are extremely successful and that re-targeted ads are unsuccessful. However, this correlation may simply reflect that consumers who are browsing travel content are more likely to purchase travel products in general. It is this type of endogeneity which leads us to focus on the field test in our analysis.

## 4 Results

### 4.1 Generic Retargeting Performs Better on Average

We first explore whether generic and dynamic retargeted ads differ in their effectiveness in converting a consumer to purchase. Figure 3 plots the average daily purchase probability for a consumer by whether they had been exposed to a generic or a dynamic retargeted ad that day. This initial evidence suggests that a generic ad is more likely to induce consumers to purchase than a dynamic retargeted ad.

In limiting an ad’s effect to the day it is shown, we follow current industry marketing practice in terms of how online ad networks award commissions to their affiliates (Weiman, 2010). This echoes Tellis and Franses (2006), who suggest that econometricians should use the most disaggregated unit of ad exposure available, to avoid the upward bias inherent in aggregate advertising data.<sup>8</sup>

[Figure 3 about here.]

There are obviously important factors that this simple analysis in Figure 3 does not control for. For example, this analysis does not control for the effect of covariates, such as whether a consumer had been exposed to contextually or behaviorally targeted ads, or the cumulative effect of any of the four types of ads employed by the firm. Additionally, we need to control for a consumer’s underlying probability to purchase that may change over time. We therefore turn to a proportional hazard model (Cox, 1972; Jain and Vilcassim, 1991; Seetharaman and Chintagunta, 2003). This approach is also in line with earlier work that tracks the effect of banner advertising on purchasing (Manchanda et al., 2006).<sup>9</sup>

---

<sup>8</sup>We find that our results are robust to allowing ad exposure to affect purchases within a two-day and a four-day window.

<sup>9</sup>Our results are robust to parametric specifications where the baseline hazard is modeled using the Weibull and Exponential distribution.

Empirically, we measure whether exposure to dynamic ad content indeed increased the likelihood of purchase on the day the consumer was exposed to the ad, relative to the control condition, controlling for covariates and the time elapsed since we first observe the consumer being exposed to one of the ads in the test.<sup>10</sup>

The dependent variable in a proportional hazard model is  $T$ , which captures the time to purchase. We also show robustness to a specification with a binary dependent variable that captures whether or not someone purchases a product on that day.

There are two types of explanatory variables in a proportional hazard models: The baseline hazard,  $h_0(t)$ , and the vector of covariates,  $(X_{it})$ . The baseline hazard captures the effect of how long it is since we first observed a consumer being exposed to an ad in our data.<sup>11</sup> Once the consumer has purchased from the travel firm, they exit the data. To increase flexibility, we estimate the baseline hazard non-parametrically (Seetharaman and Chintagunta, 2003). The vector of covariates,  $X_{it}$ , captures the effect of different types of ads a consumer was exposed to on the probability to purchase on any given day. The hazard rate for consumer  $i$   $h_i(t, X_t)$  is therefore:

$$h_i(t, X_t) = h_0(t) \times \exp(X_{it}\beta) \quad (1)$$

We specify the vector of covariates for consumer  $i$  as

$$\begin{aligned} \exp(X_{it}\beta) = & \exp(\beta_1 \text{DynamicRetargetedAd}_{it} + \beta_2 \text{AnyRetargetedAd}_{it} \quad (2) \\ & + \beta_3 \text{OtherBehavioralAd}_{it} + \beta_4 \text{ContextualAd}_{it} + \beta_5 \text{CumRetargetedDynamicAds}_{it} \\ & + \beta_6 \text{CumRetargetedAds}_{it} + \beta_7 \text{CumOtherBehavioralAds}_{it} + \beta_8 \text{CumContextualAds}_{it}) \end{aligned}$$

---

<sup>10</sup>Hazard models allow for censoring and so are appropriate to model events that, for a subset of the population, may never occur.

<sup>11</sup>We are forced to use days since the first date of ad exposure as we do not have data on when a consumer first contemplated purchasing the product. The randomization inherent in our field experiment means, however, that any error this introduces will at least be orthogonal to the main effect of interest.

$\beta_1$  measures the effect of the consumer being exposed to a dynamic retargeted ad, that is, an ad which had information content that was specific to the previous products they were browsing on the website.  $\beta_2$  measures the effect of the baseline control condition where the consumer was shown a generic retargeted ad.  $\beta_3$  controls for whether the person had seen another form of behavioral targeted ad and  $\beta_4$  measures response to a contextual targeted ad.  $\beta_5$  measures response to the cumulative number of retargeted ads with specific content that the person has seen so far. These allow us to control for the ‘stock’ of advertising a consumer has seen before. Similarly,  $\beta_6$  measures response to the cumulative number of generic retargeted ads.  $\beta_7$  and  $\beta_8$  measure response to the cumulative number of behavioral and cumulative number of contextual ads.

[Table 5 about here.]

In Table 5, Column (1), we initially show our results hold when we use as our dependent variable a straightforward measure of purchase incidence. In this discrete choice specification, the dependent variable is whether or not the consumer purchases that day. Correspondingly, all our explanatory variables are on a consumer-day level. They include the type of ads the user was exposed to that day as well as a vector of controls for each different day in our data.<sup>12</sup> Even with additional controls, these results reinforce the basic pattern in the data shown in Figure 3 that dynamic retargeting is on average less effective. The small size of the coefficients emphasizes that though relative effects of different types of ads are large, absolute effects are small. This is in line with previous evidence such as Manchanda et al. (2006), which suggests that banner ads have low effectiveness, but that their low cost means they still offer a reasonable return on investment. Our controls proxy not only for different types of targeting but also for whether or not someone is seeking travel-category content that day. The cumulative ad controls measure the effect of the stock of previous online ads

---

<sup>12</sup>As we include controls for each day, this specification is equivalent to a discrete-time hazard model (Allison, 1982).

that the person has been exposed to. The estimates for the controls do not have a clear causal interpretation. All results hold when we exclude cumulative ad totals.

Though our use of a linear probability model in Column (1) facilitates the interpretation of interactions (Ai and Norton, 2003), we show that our results also hold when using a probit specification as shown in Column (2).

Column (3) of Table 5 reports results for our main specification where we use a Cox proportional hazards model with time to purchase as the dependent variable as represented by equation (2). Again we see that in general retargeted ads are positively correlated with purchase, but the addition of personalized dynamic content on average reduces their effectiveness sizeably.

We then check robustness to different specifications for the baseline hazard and timing assumptions. Column (4) checks that our results are robust to excluding observations where a consumer saw more than one impression of either a generic or a dynamic ad that day. The results are similar. Column (5) allows our effect to vary with the number of ads that a consumer saw that day. The results are similar but less precise, partly because multiple impressions could be driven by user behavior such as repeated reloading of a page, making the impact of advertising not necessarily always increase with the number of impressions.

So far, the results have assumed that the incremental effect of online advertising is limited to the day that consumers are exposed to it. Column (6) measures the same-day effect of advertising on purchasing, while also controlling for the effect of a one-day lag of exposure to retargeted ads and the lagged values of each of our cumulative counts of ad exposure. The estimates are again robust and the estimates for the one-day lag of the effect of a retargeted ad are not significant. This provides evidence for the validity of our approach, that focuses on purchases that occur the same day as ad exposures.

In sum, the empirical evidence presented in Table 5 confirms the insight of Figure 3, that on average generic retargeting is more effective than consumer-specific dynamic retargeting.

## 4.2 The Performance of Dynamic Retargeting Varies with Browsing

### 4.2.1 Theoretical Grounding

The result that on average dynamically retargeted ads underperform is surprising. We suggest that the effectiveness of a retargeted ad depends on whether the concreteness of its message matches how narrowly consumers’ construe their preferences (Trope et al., 2007; Lee et al., 2010).

Consumers often initially have only a broad idea of what they like. At this stage, when their preferences are construed at an abstract, high level, consumers focus on the general desirability of an activity. Later, they develop narrowly construed preferences and focus on details. When consumers only have a broad idea of what they like, they do not know the specific trade-offs they would like to make in order to satisfy their needs best. They may know, for example, they are looking for a relaxing vacation but not whether they prefer a large hotel with a large pool or a small and more intimate hotel that may not feature a pool. Over time, however, consumers are more likely to construe their preferences on a concrete and specific level.<sup>13</sup> They also learn the weights to place on different attributes (Hoeffler and Ariely, 1999). They may, for example, learn their preference for access to a pool outweighs the cost of choosing a less intimate hotel. In using the term ‘narrowly construed preferences,’ we therefore refer to consumers who have a detailed viewpoint on what kind of products they wish to purchase.

We build on the insight by Simonson (2005) that a consumer’s stage of preference development may significantly affect the effectiveness of personalized ad content. We propose that ads that convey information on high-level characteristics are more effective when consumers have a broad idea of what they want. Such generic ads deliver a broad message about the product, for example that a travel firm offers relaxing vacations, and addresses consumers’

---

<sup>13</sup>This may be linked to a consumer’s stage in their decision process (Lavidge and Steiner, 1961; Hauser, 1990; Häubl and Trifts, 2000; Wu and Rangaswamy, 2003).

higher-level goals, like relaxation. They can enhance brand preferences at a point where the consumer does not yet have a clear picture of which attributes they value. By contrast, ads that focus on specific products will be more effective for consumers with narrowly construed preferences who have shifted their focus to specific product attribute. Therefore, we propose that dynamic ads may be ineffective when shown to consumers who have a broad idea of what they want but more effective when shown to consumers with narrowly construed preferences.

What causes this shift from broad ideas to narrowly construed preferences? As a consumer's uncertainty about making a category purchase decreases, the psychological distance to the event diminishes (Trope et al., 2007). For example, a consumer may be fixing a specific date for their vacation, reducing uncertainty about whether and when to travel. Such consumers are more likely to narrowly construe their preferences and will explore specific choices instead of focusing on their higher level goal. When searching for detailed product information, they will start trading-off how much they value certain attributes and learn the weights to place on them (Hoeffler and Ariely, 1999).

#### **4.2.2 Empirical Results on Browsing**

Empirically identifying an indicator of whether a consumer has yet developed narrow preferences and is more positively disposed towards a dynamic ad is challenging. We suggest that in an online environment, a visit to a site that provides detailed information about specific products, such as a product review site, signals that a consumer worries about specific product attributes, is comparing and contrasting product features and confronting the trade-offs inherent in a product choice, and so is very likely to be developing, or have developed, narrowly construed preferences.

We use data collected by the travel firm on whether a consumer had visited a travel review website such as TripAdvisor.<sup>14</sup> These review websites provide large numbers of detailed

---

<sup>14</sup>The focal firm only displayed standard generic ads on these review sites.

traveler reviews about hotels and travel products. For example, TripAdvisor has nearly 25 million reviews and opinions on more than 490,000 hotels and attractions, has more than 11 million registered members, and operates in 14 countries and 10 different languages. We recognize that a review site visit may be a manifestation of many other different phenomena and explore alternate explanations in detail in our empirical analysis. We also recognize that consumers may have means of obtaining detailed product information that we do not observe in our data. Such misclassification would introduce classification error into our specification.

[Figure 4 about here.]

In our data, 54% of consumers who both purchased a product and visited a review site, visited the review site before making their purchase and 46% visited after making a purchase. Typically, consumers who visit a review site after making a purchase hope to learn more about the destination they have chosen (e.g., places to visit or restaurants nearby).

Figure 4 provides exploratory graphical evidence where we stratify the purchase probability conditional on being exposed to a generic or a dynamic ad by whether or not the consumer had yet visited a product review site. This figure uses data only on consumers who at some point in our data visit a review site. It suggests that after a consumer has visited such a review site, the comparative advantages of the different types of advertising shift so that dynamic retargeted ads no longer under-perform. After viewing a review site, generic brand ads become relatively less effective while dynamic ads become relatively more effective.

As in Section 4.1, we estimate the probability for a consumer to purchase in a hazard model. We interact the key variables of equation (2) with a binary indicator variable for whether or not the person had yet visited a product review website. Table 6 reports the results. Column (1) displays results for a proportional hazard model for all consumers that were eligible for the field experiment. As before, we find that the coefficient

of *DynamicRetargetedAd* is negative, that is, the dynamic retargeted ad performs worse on average than the generic ad. However, *DynamicRetargetedAd*  $\times$  *AfterReviewSite* is positive and significant. This result means that the effectiveness of the dynamic retargeted ad improves after someone has visited a review site.

In sum, our results suggest that dynamic retargeting is ineffective when targeted at consumers who only have a broad idea of what they like. It is an effective form of advertising, however, when addressing consumers who have narrowly construed preferences and so are more likely to focus on specific product details.

[Table 6 about here.]

In column (1) we restrict our data to consumers who at some point visit a review site. This is because consumer characteristics which may be correlated with the decision to visit a review site might likewise be correlated with a partiality for dynamic retargeted ads. For example, consumers who visit reviews sites may be more experienced, have seen more competitive ads, be more familiar with the travel category, or have a preference for drawing independent conclusions. For comparison, in Column (2) of Table 6 we report the results for the entire sample. The coefficients are similar in precision and direction, though those for the unrestricted sample are slightly smaller. This reassures us that the sample selection issues inherent in comparing visitors and non-visitors to review sites is less important empirically than might have been supposed.

A second concern is that even when we exclude users who never visited a review site, our results could still be an artifact of differences in intensity of exposure. Suppose, for example, that less technology-able consumers are more likely to prefer generic ads and are also more likely to cease all web activity after visiting a review site. This would mean that our result would be an artifact of the fact that, post-review site visit, we only observe consumers who are predisposed to dynamic retargeted ads. To address this concern, we report in Column (3)

of Table 6 results restricted to individuals who were exposed at least once to an ad following their review site visit . The results are again similar.

A third possible set of concerns is centered around the fact that a review site visit might in itself directly provide new information altering consumers' choices. For example, there could be a direct effect of reinforcing quality information. To tackle this, we excluded observations of consumers who purchased the product within two days of visiting the review site, reasoning that if such direct effects of the information provided on a review site were present, the effect we measure should weaken substantially. The results reported in Column (4) of Table 6 show that our results hold. Columns (5)-(7) show that our results are robust to our earlier robustness tests for a discrete choice specification and the inclusion of lags.

### **4.2.3 Additional Evidence of Mechanism**

So far, our robustness checks rule out selection or changes in the environment as alternative explanations for our result. We next aim to provide positive evidence that the effect we document was driven by the changing appeal of the dynamic ad rather than other factors. To do this, we turn to a factor that has been documented to shift the appeal of an ad. Specifically, Petty et al. (1983) show that the argument quality of an ad has greater effect under high than under low involvement. If we find that under higher involvement, the appeal of a dynamic ad to a consumer who has narrowly construed preferences increases further, then this is indirect evidence that the effect we document is driven by the changing effectiveness of advertising content rather than external factors. We investigate empirically how such involvement changes the relative effectiveness of generic and dynamic retargeted ads.

We suggest that consumers browsing travel content sites is a good proxy for them being involved in the travel category that day. Travel content websites, such as Conde Nast Traveler, provide a wide range of information about vacation destinations but not about

specific hotels. Since in our data we observe whether a consumer was exposed to an ad by the travel firm on a travel content site, we use this as an indicator for browsing a travel content site and the consumer being involved in the category.

We again acknowledge that there may be alternative interpretations of a visit to a travel content website in addition to it being a proxy for involvement. For example, travel content websites are also more likely than other categories of websites to show ads for competitive products. We have no data, however, on whether or not such ads were present. Therefore our estimates should be considered as reflecting and not controlling for this change in competition.

In Figure 5 we stratify our data by whether a consumer browses the travel category on that day. It illustrates that, on average, browsing the travel category is an important predictor of the likelihood of conversion. When consumers are not involved in the category, they are much less likely to make a purchase that day.

In Figure 6, we decompose Figure 4 by whether or not the consumer visited a travel content site that day. We restrict our analysis to consumers who had both visited a travel content website and viewed a review site, to control for possible issues of sample selection.

[Figure 5 about here.]

Figure 6 illustrates that browsing of travel content websites narrows the proportional gap in effectiveness between generic and dynamic advertising for consumers who have not yet visited a review site. Importantly, the dynamic retargeted ad is clearly more effective than the generic ad for consumers who visited a travel review site (which implies that the consumer is confronting trade-offs in product features) and are also browsing travel content websites that day (which implies they are involved in the category). One interpretation of this finding is that the consumer involvement proxied for by the travel content website browsing is enhancing the argument quality of the dynamic ad for consumers with narrowly

construed preferences.<sup>15</sup>

[Figure 6 about here.]

We repeat this analysis in a regression framework in Table 7. We limit our sample to those consumers who visited both travel and review sites. Column (1) reports the results of interacting our basic specification summarized by equation (2) with an indicator variable for whether or not that person was observed visiting a website devoted to category-related information that day. The negative and significant coefficient for  $RetargetedAd \times DynamicAdContent$  suggests that dynamic ads are less effective than generic ads. However, this is mediated by the positive and significant coefficient for  $RetargetedAd \times DynamicAdContent \times BrowsingTravelthatDay$  which suggests that dynamic retargeted ads perform relatively better on days that consumers browse travel.<sup>16</sup> We echo the analysis of Table 6 and stratify the results by whether the consumer has yet visited a review site in Columns (2) and (3). The baseline measure of  $RetargetedAd \times BrowsingTravelthatDay$  is more negative after the consumer visits the review site. Therefore, the performance of generic retargeted ads gets relatively worse on days where a consumer browses travel after they visit a review site. However, the increasing size of the coefficient  $RetargetedAd \times DynamicAdContent \times BrowsingTravelthatDay$  after a consumer visits a review site suggests that by contrast dynamic retargeted ads perform relatively better after a consumer visits a travel review site and they are browsing the category that day. In general, these results suggest that the most effective time to use dynamic retargeting rather than generic retargeting is after a consumer visits a review site and appears to be actively involved in the category. This is, again, in line with Figure 6.

---

<sup>15</sup>In a robustness check, we find that these results hold when we exclude observations where the consumer was exposed to an ad during or after their browsing of the travel category on that particular day. This means that our results are not driven by reverse causation, where the ad provokes people to browse the travel category, or by a contextual effect of the ad. Likewise, our results hold when we exclude observations where the consumer was exposed to an ad before they browsed the travel category.

<sup>16</sup>As before, these results are robust to different definitions of the baseline hazard or a discrete choice model.

[Table 7 about here.]

These results support the interpretation of our previous findings that dynamic retargeted ads perform well only when preferences are narrowly construed. We recognize that this evidence is partly based on natural variation in the browsing data which may be endogenous in ways we are unable to control for. We turn to the lab to explicitly rule out alternative explanations.

## 5 Confirming the Results in the Lab

*Objective.* We aim to show that the interpretation of our results hold in a controlled lab environment. Specifically, we directly test whether the fact that consumers have formed narrowly construed preferences indeed affects how they react to a generic versus a dynamic ad. Second, we aim to rule out privacy concerns, reactance, competitive effects and consumer experience as alternative explanations of our results. Our experimental set-up additionally provides evidence that neither social validation nor increased access to quality information that are inherent in visiting a review site are driving our results.

*Design and procedure.* The study has a  $2 \times 2$  design. We vary whether a consumer has refined their preferences (only broad idea of what they want versus narrowly construed preferences) and the type of ad they are exposed to (generic ad vs. dynamic ad). In the *broad-idea condition*, participants are asked to imagine they would like to go on a beach vacation. They do not yet know where they would like to go and are still exploring different destinations. They have not yet thought about specific hotels since they would like to first choose their destination. They are then asked to imagine that on the website of a travel company they broadly looked at hotels in many different regions. Participants in the *narrow-preferences condition* are asked to imagine they would like to go on a beach vacation in Hawaii. They are specifically looking for a hotel with a very large pool. They are then asked to imagine that on the website of a travel company they evaluated hotels with large

pools in Hawaii.

All participants are told that, as they now browse the Internet, they see an ad on another website for this travel company. Participants in the *generic-ad condition* are told it shows a picture generally relating to beach vacations. They are shown an ad for a fictitious travel company that displays a sun lounge on the beach. Participants in the *dynamic-ad condition* are shown an ad for a fictitious travel company with pictures of four hotels with pools. They are told the ad shows one of the hotels they looked at plus three other hotels.<sup>17</sup> Participants are then asked how likely they are to visit the firm's website and book a vacation (1: "very unlikely" - 7: "very likely").

Finally, participants answer several additional questions that relate to the scenario. First, we ask how likely they are at this stage of their travel planning to have already visited a travel review site, such as Tripadvisor. Second, we measure whether either ad causes privacy concerns and whether this possibly varies across conditions. We use the scale developed by Xu (2007) that captures how much consumers are concerned about the privacy of their data in online environments. We then measure reactance to the ad using the scale developed by Edwards and Li (2002). We also ask participants how likely they are to buy from a competitor at this stage in their travel planning, how often they had booked a vacation package, hotel or vacation rental in the past three years, and how often they had booked such services online.

We expect participants in the *broad-idea condition* to be more likely to visit the firm's website and book a vacation when they view a generic ad than when they view a dynamic retargeted ad. By contrast, participants in the *narrow preferences condition* should be more likely to visit the firm's website and book a vacation when they view a dynamic retargeted ad. We expect this effect to persist when controlling for other factors.

---

<sup>17</sup>We pretested the ads with 85 participants in a survey on Mechanical Turk. There was no significant difference in how much participants liked the two ads (5.558 vs. 5.286,  $p=0.208$ ).

162 participants were recruited online through Mechanical Turk<sup>18</sup> and randomly assigned to conditions (82 responses for the *broad-idea condition*, 80 responses for the *narrow preferences condition*).<sup>19</sup>

*Results.* As hypothesized, in the *broad-idea condition*, participants react significantly more favorably to the generic ad than to the dynamic ad (5.214 vs. 4.574,  $p=0.014$ ). However, in the *narrow preferences condition*, participants react significantly less favorably to the generic ad than to the dynamic ad (4.227 vs. 5.583,  $p<0.001$ ). A regression model reported in Column (1) of Table 8 confirms that overall consumers are less likely to buy after a dynamic ad unless consumers already have narrowly construed preferences.

We check whether these results are robust to controls. Participants who view the generic ad are less concerned about privacy than participants who view the dynamic ad (3.706 vs. 4.247,  $p<0.001$ ). Privacy concerns do not differ between participants in the *broad-idea condition* and participants in the *narrow-preferences condition* (3.931 vs. 3.809,  $p=0.650$ ). Reactance to the ad is generally low and does not differ across conditions (*generic* versus *dynamic*: 2.199 vs. 2.466,  $p=0.134$ ; *broad idea* versus *narrow preferences*: 2.290 vs. 2.361,  $p=0.688$ ). Including these variables in a regression model does not change the main effect of interest (Column (2)).

Similarly, we find that a consumer’s propensity to buy from a competitor does not differ depending on whether they have narrowly construed preferences (4.244 vs. 4.275,  $p=0.826$ ) or across different ad conditions (4.267 vs. 4.250,  $p=0.902$ ). Including this variable in a regression model again does not change our main effect of interest (Column (3)).

To check whether experience with travel booking in general or online increases the effectiveness of a dynamic ad, we estimate alternative models in which we include how often

---

<sup>18</sup>As discussed by Buhrmester et al. (2011), though Mechanical Turk has disadvantages, there is also evidence that using it may lead to more diverse and hence representative samples than traditional samples of American college students.

<sup>19</sup>We dropped 5 outliers out of 167 participants. Including these outliers lead to results that are directionally consistent and significant, though less precise.

a participant has booked a vacation product (Column (4)) or, alternatively, how often they have booked such services online (Column (5)). We also interact these variables with the type of ad that is displayed in the survey. None of these variables significantly affects whether a consumer will book a vacation.

Last, we confirm that consumers are indeed less likely to have yet visited a travel review site when they only have a broad idea of what they are looking for compared to when they have narrowly construed preferences (5.171 vs. 5.663,  $p=0.025$ ). This confirms that visits to travel review sites are a good indicator for whether a consumer has narrowly construed preferences.

In this study we are directly testing for the effect in the lab and consumers were not required to visit a review site. The experiment therefore likewise provides evidence that social validation through or access to quality information on a review site is not the primary driver of our results.

The results confirm that whether a consumer has narrowly construed preferences is an important determinant for the effectiveness of generic versus dynamic retargeted ads. They illustrate that only when consumers have refined their preferences is it effective to show dynamic retargeted ads to consumers, while in earlier stages it is more effective to show generic retargeted ads.

The web appendix includes a similar experiment where we replicate our results for a different product category - bathroom fixtures as well as an additional study in the travel industry that rules out social validation as an alternative explanation.<sup>20</sup>

[Table 8 about here.]

---

<sup>20</sup>Social Validation may potentially make personalized advertising more persuasive (Cialdini and Goldstein, 2004; Bakshy et al., 2009).

## 6 Conclusion

The digital revolution has seen advances in the use of data on browsing behavior both inside and outside a firm's website to improve its marketing appeals. Internal browsing data has allowed firms to customize their websites so that when a consumer returns, a firm can show them personalized recommendations based on their previous browsing behavior. External browsing data has allowed firms to target their ads better to consumers who fit a particular profile, such as people who have recently been browsing travel websites.

'Dynamic Retargeting' represents a combination of these two techniques. Dynamic re-targeting allows firms to target consumers who have previously been to the firm's website, on other sites across the Internet, with content that is specific to the product the consumer previously viewed at the firm's website. There is, however, little evidence to show whether tailoring advertising content to an individual's observed preferences is effective.

This paper evaluates whether indeed firms benefit from targeting consumers with information that is highly specific to their prior interest. We use field experiment data from an online travel firm to evaluate whether retargeting consumers with a brand-level ad (generic retargeting) or with information that reflects the specific product the consumer has viewed earlier on the firm's website (dynamic retargeting) is more effective. Surprisingly, we find that advertising content that specifically reflects the product consumers viewed earlier is in general not effective.

We then ask whether there is any time when consumers may find the retargeted ad's emphasis on specific products appealing. We build on a consumer behavior literature which suggests that such a specific emphasis on product features will be most effective when a consumer has established narrowly construed preferences. Consumers who have developed narrowly construed preferences have a greater focus on specific and detailed product information and therefore are more likely to respond positively to ads displaying specific products.

Part of the process of establishing these kind of narrowly construed preferences is the act of comparing and contrasting product features, something which consumers naturally do when consulting review sites. Therefore, we empirically explore whether the effectiveness of dynamic retargeting changes after a consumers' visit to a review site. Our results show that dynamic retargeting is not effective when consumers have not yet visited a review site and so are less likely to have developed narrowly construed preferences. However, when consumers have visited a review site, and have refined their product preferences, they are relatively more likely to respond positively to a dynamic ad. This is further augmented if the consumer is involved in the category as proxied for by browsing other websites in the category. Since the decision to visit a review site is potentially endogenous, we perform a battery of robustness tests and also provide direct evidence for the proposed mechanism in a lab experiment.

There are two major managerial insights from these results. First, one would expect individual-level content for ads based on browsing histories to be highly effective, given the generally positive effect of personalized recommendations. However, we find that on average generic ads are more effective than dynamic retargeted ads. Second, we show that the effectiveness of dynamic retargeted ads changes as consumers define their product preferences better and when browsing related content online. Data on browsing of external websites which is currently available to advertisers, but which is rarely evaluated in detail, can be used to identify whether a consumer has broadly or narrowly construed preferences and their category focus and then used to *time* the targeting of ads for maximum effectiveness.

There are of course limitations to our results. First, in the travel category that we study, firms consolidate products and variety may be important. This could explain why dynamic retargeted ads that tend to focus on similar products are particularly ineffective in this setting. In particular, generalizability to products where there is little consumer product research and purchasing behavior is driven by either impulse or habit may be limited. Second,

we do not explicitly address the question of the specifics of dynamic retargeting ad design and in particular which products should be highlighted and how. Third, we do not have data on competitors' advertising decisions which would allow us to tease apart how the effectiveness of dynamic retargeting is moderated by competitive ads. Notwithstanding these limitations, we believe that our paper, by documenting the general ineffectiveness of dynamic retargeting and the circumstances under which dynamic retargeting becomes relatively more effective than generic ads, does represent a useful contribution to our knowledge about this new form of online advertising.

## References

- Agarwal, N., S. Athey, and D. Yang (2009, May). Skewed bidding in pay-per-action auctions for online advertising. *American Economic Review* 99(2), 441–47.
- Ai, C. and E. C. Norton (2003, July). Interaction terms in logit and probit models. *Economic Letters* 80(1), 123–129.
- Allison, P. (1982). Discrete-time methods for the analysis of event histories. *Sociological methodology*, 61–98.
- Ansari, A. and C. Mela (2003). E-customization. *Journal of Marketing Research* 40(2), 131–145.
- Bakshy, E., B. Karrer, and L. A. Adamic (2009). Social influence and the diffusion of user-created content. In *Proceedings of the 10th ACM conference on Electronic commerce, EC '09*, New York, NY, USA, pp. 325–334. ACM.
- Beales, H. (2011). The value of behavioral targeting. *Mimeo, George Washington University*.
- Buhrmester, M., T. Kwang, and S. D. Gosling (2011). Amazon's Mechanical Turk: A new source of inexpensive, yet high-quality, data? *Perspectives on Psychological Science* 6(1), 3–5.
- Chen, Y., D. Pavlov, and J. F. Canny (2009). Large-scale behavioral targeting. In *Proceedings of the 15th ACM SIGKDD international conference on knowledge discovery and data mining, KDD '09*, New York, NY, USA, pp. 209–218. ACM.
- Cialdini, R. B. and N. J. Goldstein (2004). Social Influence: Compliance and Conformity. *Annual Review of Psychology* 55(1), 591–621.
- Cox, D. (1972). Regression models and life-tables. *Journal of the Royal Statistical Society* 34, 187–200.

- Criteo (2010, August 26). Targeting & retargeting interview with Criteo. <http://behavioraltargeting.biz>.
- Dalessandro, B., R. Hook, C. Perlich, and F. Provost (2012). Evaluating and optimizing online advertising: Forget the click, but there are good proxies. *NYU Working Paper No CBA-12-02*.
- Dias, M. B., D. Locher, M. Li, W. El-Deredy, and P. J. Lisboa (2008). The value of personalised recommender systems to e-business: a case study. In *Proceedings of the 2008 ACM conference on Recommender Systems, RecSys '08*, New York, NY, USA, pp. 291–294. ACM.
- Edwards, S. M. and H. Li (2002). Forced exposure and psychological reactance: Antecedents and consequences of the perceived intrusiveness of pop-up ads. *Journal of Advertising* 31, 83–95.
- Gatarski, R. (2002). Breed better banners: Design automation through on-line interaction. *Journal of Interactive Marketing* 16(1), 2–13.
- Goldfarb, A. and C. E. Tucker (2011). Privacy regulation and online advertising. *Management Science* 57(1), 57–71.
- Hargrave, S. (2011, January). Targeted ads aim for greater accuracy. *Marketing Week*.
- Häubl, G. and V. Trifts (2000). Consumer decision making in online shopping environments: The effects of interactive decision aids. *Marketing Science* 19(1), 4–21.
- Hauser, J. R., G. L. Urban, G. Liberali, and M. Braun (2009). Website morphing. *Marketing Science* 28(2), 202–223.
- Hauser, J. R. B. W. (1990, March). An evaluation cost model of consideration sets. *Journal of Consumer Research* 16, 393–408.
- Hoeffler, S. and D. Ariely (1999). Constructing stable preferences: A look into dimensions of experience and their impact on preference stability. *Journal of Consumer Psychology* 8, 113–139.
- Hoffman, D. L. and T. P. Novak (1996). Marketing in Hypermedia Computer-Mediated Environments: Conceptual Foundations. *Journal of Marketing* 60(3), 50–68.
- Hunter, A. (2010, Sep 22). Identifying the right media-placement strategies for digital display. *comScore and ValueClick Media White Paper*.
- Hunter, A., M. Jacobsen, R. Talens, and T. Winders (2010, September). When money moves to digital, where should it go? Identifying the right media-placement strategies for digital display. *Comscore White Paper*.

- Jain, D. C. and N. J. Vilcassim (1991). Investigating Household Purchase Timing Decisions: A Conditional Hazard Function Approach. *Marketing Science* 10(1), 1–23.
- Joshi, A., A. Bagherjeiran, and A. Ratnaparkhi (2011). User demographic and behavioral targeting for content match advertising. In *Data Mining and Audience Intelligence for Advertising*.
- Komiak, S. Y. X. and I. Benbasat (2006). The effects of personalization and familiarity on trust and adoption of recommendation agents. *MIS Quarterly* 30(4), 941–960.
- Lambrecht, A., K. Seim, and C. Tucker (2011). Stuck in the adoption funnel: The effect of interruptions in the adoption process on usage. *Marketing Science* (2), 355–367.
- Lavidge, R. J. and G. A. Steiner (1961). A model for predictive measurements of advertising effectiveness. *Journal of Marketing* 25(6), 59–62.
- Lee, A. Y., P. A. Keller, and B. Sternthal (2010). Value from regulatory construal fit: The persuasive impact of fit between consumer goals and message concreteness. *Journal of Consumer Research* 36(5), 735–747.
- Lewis, R. A., J. Rao, and D. Reiley (2011). Here, there, everywhere: Correlated online behaviors can lead to overestimates of the effects of advertising. In *International World Wide Web Conference*.
- Linden, G., B. Smith, and J. York (2003, Jan/Feb). Amazon.com recommendations: item-to-item collaborative filtering. *Internet Computing, IEEE* 7(1), 76 – 80.
- Malthouse, E. and R. Elsner (2006). Customisation with crossed-basis sub-segmentation. *The Journal of Database Marketing & Customer Strategy Management* 14(1), 4–50.
- Manchanda, P., J.-P. Dube, K. Y. Goh, and P. K. Chintagunta (2006). The effect of banner advertising on internet purchasing. *Journal of Marketing Research* 43(1), 98 – 108.
- Petty, R. E., J. T. Cacioppo, and D. Schumann (1983, September). Central and peripheral routes to advertising effectiveness: The moderating role of involvement. *Journal of Consumer Research: An Interdisciplinary Quarterly* 10(2), 135–46.
- Rusli, E. M. (2013, April 9). Buy signal: Facebook widens data targeting. *Wall Street Journal*.
- Rutz, O. J. and R. E. Bucklin (2011). From Generic to Branded: A Model of Spillover Dynamics in Paid Search Advertising. *Journal of Marketing Research* 48(1), 87–102.
- Seetharaman, P. B. and P. K. Chintagunta (2003, July). The proportional hazard model for purchase timing: A comparison of alternative specifications. *Journal of Business & Economic Statistics* 21(3), 368–82.

- Simonson, I. (2005). Determinants of customers' responses to customized offers: Conceptual framework and research propositions. *Journal of Marketing* 69, 3245.
- Tellis, G. J. and P. H. Franses (2006). Optimal data interval for estimating advertising response. *Marketing Science* 25(3), 217–229.
- Trope, Y., N. Liberman, and C. Wakslak (2007). Construal levels and psychological distance: Effects on representation, prediction, evaluation, and behavior. *Journal of Consumer Psychology* 17(2), 8395.
- Tucker, C. (2011). Social networks, personalized advertising, and privacy controls. *Mimeo, MIT*.
- Weiman, L. (2010, Nov 10). Is the 'last click wins' model best for your affiliates? *Search Engine Land*.
- Wu, J. and A. Rangaswamy (2003). A fuzzy set model of search and consideration with an application to an online market. *Marketing Science* 22(3), pp. 411–434.
- Xu, H. (2007). The effects of self-construal and perceived control on privacy concerns. In *ICIS*, pp. 125.
- Yan, J., N. Liu, G. Wang, W. Zhang, Y. Jiang, and Z. Chen (2009). How much can behavioral targeting help online advertising? In *Proceedings of the 18th international conference on World wide web, WWW '09*, New York, NY, USA, pp. 261–270. ACM.




Figure 1: Dynamic Retargeting Examples

Figure 2: Conversion Rate with Same-Day Ad Exposure




Figure 3: Comparison of Conversion for Generic vs Dynamic Ad Exposure




Figure 4: Comparison of Conversion for Generic vs Dynamic Ad Exposure: Sample Restricted to Those Who Visited A Review Site




Figure 5: Comparison of Conversion for Generic vs Dynamic Ad Exposure: Sample Restricted to Consumers Who Browsed a Travel Website and a Review Website




Figure 6: Comparison of Conversion for Generic vs Dynamic Ad exposure: Sample Restricted to Consumers Who Browsed a Travel Website and a Review Website






Table 1: Previous Literature

Paper	Setting	Personalization	Targeting	Decision stages	Finding
Personalized Recommendations					
Linden et al. (2003)	Portal	Collaborative filtering	None	No	Collaborative filtering improves recommender systems
Komiak and Benbasat (2006)	Lab	Recommendation Agents	None	No	Perceived personalization significantly increases consumers' intention to adopt by increasing cognitive trust and emotional trust.
Dias et al. (2008)	Grocery website recommendations	Past product purchases and shopping basket content	None	No	Supermarket revenues increased by 0.30%
Tailored Communications					
Ansari and Mela (2003)	Content of email newsletter	Customer content category	None	No	Personalization increases click-throughs
Malthouse and Elsner (2006)	Content of cover letter of mail order catalogue	Recency, Frequency, Monetary	None	No	Segment-based customization is cost-effective
Agarwal et al. (2009)	Web content on a firm's website	Segments defined by demographics and browsing behavior	None	No	Bayesian approach dominates non-personalized content selection
Hauser et al. (2009)	Content of firm's website	Cognitive style segments	None	No	Personalization on basis of cognitive style revealed by browsing behavior improves profitability
Tucker (2011)	Banner Ads	Based on stated celebrity preferences	None	No	Privacy controls improve response to personalized ads
Targeted Advertising					
Chen et al. (2009)	Portal	None	Behavioral	No	Adding more categories of browsing behavior to algorithm makes behavioral targeting more effective
Yan et al. (2009)	Search Engine	None	Search	No	Behavioral data on prior searches makes search-engine ads more effective
Beales (2011)	ad network	None	Behavioral	No	Behaviorally targeted ads cost 100% more
Goldfarb and Tucker (2011)	ad network	None	Behavioral	No	Privacy regulation that restricts behavioral targeting reduces ad effectiveness
Joshi et al. (2011)	Match ads to users and content on firm's website	Based on demographics and website visits, searches, ad views, ad click	Behavioral, contextual	No	Matching ads to the right website content can be improved by integrating user characteristics

For descriptions of the different forms of targeting techniques such as behavioral and contextual targeting, see Table 2.

Table 2: Summary of Online Advertising Methods.

Label	Type of Targeting	Ad Image	Part of Field Test
Contextual Targeting	Firm advertises on website that has travel content	Generic ad displaying brand and evocative vacation image.	No
Behavioral Targeting	Firm advertises to consumers who had previously visited a travel website	Generic ad displaying brand and evocative vacation image.	No
Generic Retargeting	Firm advertises to consumers who had previously visited the <i>firm's</i> website	Generic ad displaying brand and evocative vacation image.	Yes
Dynamic Retargeting	Firm advertises to consumers who had previously visited the <i>firm's</i> website	Ad displays products reflecting consumers' prior product search <sup>a</sup>	Yes

<sup>a</sup>Figure 1(a) shows an example of a dynamic retargeted ad. After browsing a certain style of children's shoe, under dynamic retargeting the consumer would be retargeted with ads displaying the specific shoe the consumer looked at, alongside similar shoes.

Table 3: Consumers Eligible for Dynamic Retargeting

(a) Cross-Sectional Descriptives

	Mean	Std Dev
Purchase	0.100	0.300
Visited Review Site	0.402	0.490
Observations	4542	

(b) Time-varying Covariates

	Mean	Std Dev	Min	Max
Any Retargeted Ad	0.089	0.284	0	1
Dynamic Retargeted Ad	0.047	0.211	0	1
Any Ad	0.214	0.410	0	1
Other Behavioral Ad	0.122	0.328	0	1
Contextual Ad	0.042	0.202	0	1
Cumulative Any Retargeted Ads	8.021	13.300	0	151
Cumulative Dynamic Retargeted Ads	6.772	11.581	0	151
Cumulative Other Behavioral Ads	19.082	39.267	0	881
Cumulative Contextual Ads	9.485	25.948	0	1313
Observations	83214			

Table 4: All Consumers  
(a) Cross-Sectional Descriptives

	Mean	Std Dev
Purchase	0.020	0.139
Eligible for Retargeting	0.069	0.253
Visited Review Site	0.091	0.288
Observations	104846	

(b) Time-varying Covariates

	Mean	Std Dev	Min	Max
Any Retargeted Ad	0.002	0.049	0	1
Dynamic Retargeted Ad	0.001	0.036	0	1
Any Ad	0.028	0.166	0	1
Other Behavioral Ad	0.023	0.150	0	1
Contextual Ad	0.005	0.068	0	1
Cumulative Any Retargeted Ads	0.343	2.763	0	248
Cumulative Dynamic Retargeted Ads	0.309	2.455	0	248
Cumulative Other Behavioral Ads	3.952	23.586	0	5504
Cumulative Contextual Ads	1.297	7.519	0	1313
Observations	2138038			

Table 5: Dynamic Retargeting for those Eligible for the Retargeting Campaign

	(1)	(2)	(3)	(4)	(5)	(6)
	Discrete Choice:OLS	Discrete Choice:Probit	Cox Hazard	Exc. Mult. Impr.	Mult. Impr.	Lags
Dynamic Retargeted Ad	-0.0106*** (0.000733)	-0.467*** (0.0302)	-1.111*** (0.340)	-1.118*** (0.417)		-1.348*** (0.480)
Any Retargeted Ad	0.00873*** (0.000614)	0.430*** (0.0224)	0.695*** (0.250)	0.815*** (0.304)		1.339*** (0.332)
Total Dynamic Retargeted Ads					-0.251** (0.118)	
Total Retargeted Ads					0.0150 (0.0752)	
Other Behavioral Ad	0.0182*** (0.000362)	0.774*** (0.00937)	1.821*** (0.161)	1.959*** (0.170)	1.853*** (0.162)	1.825*** (0.163)
Contextual Ad	0.0502*** (0.000946)	1.150*** (0.0115)	2.560*** (0.176)	2.664*** (0.185)	2.586*** (0.177)	2.518*** (0.172)
Cumulative Dynamic Retargeted Ads	0.000180*** (0.0000179)	0.0121*** (0.00186)	0.0456*** (0.0176)	0.0240 (0.0181)	0.0207 (0.0161)	0.125 (0.179)
Cumulative Any Retargeted Ads	-0.000255*** (0.0000162)	-0.0184*** (0.00177)	-0.0564*** (0.0163)	-0.0309* (0.0169)	-0.0305** (0.0154)	-0.374*** (0.134)
Cumulative Other Behavioral Ads	0.00000339 (0.00000338)	0.000257*** (0.0000774)	0.000758 (0.000729)	0.000718 (0.000737)	0.000717 (0.000742)	0.00445** (0.00196)
Cumulative Contextual Ads	-0.0000729*** (0.00000340)	-0.00233*** (0.000172)	-0.00490** (0.00212)	-0.00507** (0.00218)	-0.00498** (0.00210)	0.0167*** (0.00267)
Day Controls	Yes	Yes	No	No	No	No
Lagged Ad Effects	No	No	No	No	No	Yes
Observations	1502514	1502514	1502514	1419428	1502514	1502514
Log-Likelihood	1819099.5	-40349.9	-70059.8	-62344.9	-70148.6	-69779.8

Dependent variable is time to purchase in all columns except columns (1)-(2), where it is whether or not a purchase was made that day. OLS regression coefficients reported in column (1). Probit regression coefficients reported in column (2). Hazard-model coefficients presented in other columns. Observations for hazard models are the number of 'days at risk'.

Robust standard errors. \*  $p < 0.10$ , \*\*  $p < 0.05$ , \*\*\*  $p < 0.01$ .

Table 6: Interactions with Visit to Review Site

	(1)	(2)	(3)	(4)	(5)	(6)	(7)
	Only Review Site Users	All Users	Saw Ads Post Review:Site	Exclude 2 Days Review Visit	Discrete Choice OLS	Discrete Choice Probit	Lags
Dynamic Retargeted Ad $\times$ After Review Site	0.881*** (0.161)	0.783*** (0.137)	0.962*** (0.189)	0.685*** (0.205)	0.011*** (0.002)	0.493*** (0.068)	0.903*** (0.161)
Dynamic Retargeted Ad	-1.718*** (0.117)	-1.331*** (0.077)	-1.614*** (0.154)	-1.106*** (0.147)	-0.014*** (0.001)	-0.622*** (0.049)	-1.318*** (0.157)
Retargeted Ad	1.036*** (0.073)	0.880*** (0.051)	1.068*** (0.110)	1.662*** (0.098)	0.010*** (0.001)	0.426*** (0.029)	1.229*** (0.091)
After Review Site	0.408*** (0.042)	0.212*** (0.035)	0.557*** (0.058)	0.688*** (0.055)	0.004*** (0.000)	0.244*** (0.020)	0.452*** (0.042)
Retargeted Ad $\times$ After Review Site	-0.340*** (0.118)	-0.733*** (0.106)	-0.305*** (0.144)	-1.176*** (0.160)	-0.009*** (0.001)	-0.389*** (0.047)	-0.278*** (0.117)
Cumulative Any Retargeted Ads	-0.071*** (0.007)	-0.052*** (0.004)	-0.105*** (0.013)	-0.043*** (0.008)	-0.000*** (0.000)	-0.006*** (0.001)	-0.156*** (0.033)
Cumulative Retargeted Ads $\times$ After Review Site	-0.001 (0.010)	-0.022*** (0.009)	0.031** (0.015)	-0.023** (0.011)	-0.023** (0.011)	-0.003*** (0.000)	-0.003 (0.010)
Cumulative Dynamic Ads	0.071*** (0.008)	0.043*** (0.005)	0.125*** (0.014)	0.016* (0.009)	0.000*** (0.000)	0.003*** (0.001)	-0.228*** (0.059)
Cumulative Retargeted Dynamic Ads $\times$ After Review Site	-0.019* (0.011)	0.014 (0.010)	-0.073*** (0.016)	0.034*** (0.012)	-0.000*** (0.000)	-0.000 (0.001)	-0.017 (0.010)
Cumulative Total Ads							
Cumulative Total Ads $\times$ After Review Site							
Constant							
Date Controls	No	No	No	No	Yes	Yes	No
Further Ad Controls	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Lagged Ad Effects	No	No	No	No	No	No	Yes
Observations	601475	1502514	276000	548191	601475	601475	601475
Log-Likelihood	-32391.6	-69998.3	-20103.0	-20033.1	652083.7	-19459.8	-32182.3

Dependent variable is time to purchase in all columns except columns (5)-(6) where it is whether or not a purchase was made that day. OLS regression coefficients reported in column (5). Probit regression coefficients reported in column (6). Hazard coefficients reported in all other columns. All controls and appropriate interactions from Table 5, column (3), included but not reported for readability. Further Ad Controls refers to the full set of controls for Contextual and Behavioral Ads as well as their cumulative totals.

Robust standard errors. \*  $p < 0.10$ , \*\*  $p < 0.05$ , \*\*\*  $p < 0.01$ .

Table 7: Further Evidence of Moderating Effect of Category Activity

	All (1) Survival Time	Before Review Site (2) Survival Time	After Review Site (3) Survival Time
Dynamic Retargeted Ad $\times$ Browsing Travel that Day	1.515*** (0.175)	0.838*** (0.267)	2.512*** (0.279)
Dynamic Retargeted Ad	-2.114*** (0.140)	-2.411*** (0.213)	-1.803*** (0.195)
Retargeted Ad $\times$ Browsing Travel that Day	-0.392*** (0.123)	-0.306* (0.170)	-1.496*** (0.238)
Retargeted Ad	0.585*** (0.090)	0.774*** (0.135)	0.951*** (0.145)
Browsing Travel that Day	1.356*** (0.053)	1.959*** (0.085)	1.479*** (0.104)
Cumulative Any Retargeted Ads	-0.089*** (0.009)	-0.081*** (0.012)	-0.105*** (0.013)
Cumulative Retargeted Ads $\times$ Browsing Travel that Day	0.002 (0.012)	-0.033* (0.017)	0.037** (0.017)
Cumulative Dynamic Ads	0.070*** (0.009)	0.072*** (0.013)	0.076*** (0.015)
Cumulative Retargeted Dynamic Ads $\times$ Browsing Travel that Day	0.003 (0.013)	0.047*** (0.018)	-0.024 (0.019)
Further Ad Controls	No	Yes	Yes
Observations	145452	80581	64871
Log-Likelihood	-24077.7	-12040.4	-9418.0

Proportional Hazard regression coefficients shown. Dependent variable is time to purchase. Sample restricted to consumers who at some point visited a travel website and a review site. Further Ad Controls refers to the inclusion of the full set of controls for Contextual and Behavioral Ads as well their cumulative totals that are reported in column (3) of Table 5, but they are not reported here for reasons of space.

Robust standard errors. \*  $p < 0.10$ , \*\*  $p < 0.05$ , \*\*\*  $p < 0.01$ .

Table 8: Lab Study One - Travel

	(1)	(2)	(3)	(4)	(5)
Dynamic Ad	-0.639** (0.256)	-0.601** (0.259)	-0.599** (0.260)	0.435 (0.276)	0.465* (0.270)
Stable Preferences	-0.987*** (0.249)	-0.951*** (0.249)	-0.943*** (0.250)		
Dynamic Ad $\times$ Stable Preferences	1.995*** (0.364)	1.951*** (0.364)	1.938*** (0.366)		
Privacy		0.0475 (0.0743)	0.0522 (0.0749)	0.0530 (0.0813)	0.0533 (0.0812)
Reactance		-0.155* (0.0856)	-0.151* (0.0861)	-0.185* (0.0947)	-0.180* (0.0958)
Buy Competitor			-0.0580 (0.103)	-0.104 (0.113)	-0.0958 (0.112)
Booked Vacation				0.0229 (0.0259)	
Dynamic Ad $\times$ Booked Vacation				-0.0208 (0.0608)	
Booked Online					0.0290 (0.0285)
Dynamic Ad $\times$ Booked Online					-0.0325 (0.0664)
Constant	5.214*** (0.178)	5.360*** (0.329)	5.577*** (0.508)	5.284*** (0.541)	5.222*** (0.546)
Observations	162	162	162	162	162
R-Squared	0.175	0.192	0.194	0.0525	0.0541

Dependent variable is purchase probability scale. Ordinary Least Square Estimates. \*  $p < 0.10$ , \*\*  $p < 0.05$ , \*\*\*  $p < 0.01$ .