ANJA LAMBRECHT and CATHERINE TUCKER*

Firms can now offer personalized recommendations to consumers who return to their website, using consumers' previous browsing history on that website. In addition, online advertising has greatly improved in its use of external browsing data to target Internet ads. Dynamic retargeting integrates these two advances by using information from the browsing history on the firm's website to improve advertising content on external websites. When surfing the Internet, consumers who previously viewed products on the firm's website are shown ads with images of those same products. To examine whether this is more effective than simply showing generic brand ads, the authors use data from a field experiment conducted by an online travel firm. Surprisingly, the data suggest that dynamic retargeted ads are, on average, less effective than their generic equivalents. However, when consumers exhibit browsing behavior that suggests their product preferences have evolved (e.g., visiting review websites), dynamic retargeted ads no longer underperform. One explanation for this finding is that when consumers begin a product search, their preferences are initially construed at a high level. As a result, they respond best to higher-level product information. Only when they have narrowly construed preferences do they respond positively to ads that display detailed product information. This finding suggests that in evaluating how best to reach consumers through ads, managers should be aware of the multistage nature of consumers' decision processes and vary advertising content along these stages.

*Keywords*: retargeting, online advertising, field experiments, online decision process, construal level theory

*Online Supplement*: http://dx.doi.org/10.1509/jmr.11.0503

# When Does Retargeting Work? Information Specificity in Online Advertising

Innovations in the parsing and processing of individual-level browsing data now enable firms to offer 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 leaving the website, and they may increase sales (Dias et al. 2008; Linden, Smith, and York 2003). 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 the exact product consumers previously browsed.

At first glance, this makes sense: previous marketing literature has emphasized that greater specificity of a firm's

*Anja Lambrecht is Assistant Professor of Marketing, London Business School (e-mail: alambrecht@london.edu). Catherine Tucker is the Mark Hyman Jr. Career Development Professor and Associate Professor of Marketing, MIT Sloan School of Management, Massachusetts Institute of Technology, and Research Associate, National Bureau of Economic Research (e-mail: cetucker@mit.edu). The authors thank Havas Digital, particularly Katrin Ribant for access to data from Artemis and Marco Bertini for facilitating contact. They gratefully acknowledge financial support from the National Science Foundation CAREER Award 1053398. They 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 Summer Institute in Competitive Strategy conference and seminar participants at Cass Business School, European School of Management and Technology, ESSEC Business School, London Business School, and the National University of Singapore. Ron Shachar served as associate editor for this article.

interactions with consumers should increase relevance and consumer response (Dias et al. 2008; Hoffman and Novak 1996; Komiak and Benbasat 2006). Indeed, firms that offer retargeting services point to strong increases in advertising effectiveness. For example, the Behavioral Targeting Blog (2010) reports that personalized retargeted ads are six times more effective than standard banner ads and four times more effective than generic retargeted ads. As a result, dynamic retargeting has engendered much enthusiasm among online advertising practitioners (Hargrave 2011; Hunter et al. 2010). For example, the Next Performance ad network, a firm that sells retargeting solutions, reports that it has served 30 billion retargeted impressions, analyzed one billion products for possible inclusion in a dynamic retargeted ad, and offered dynamic retargeted ads to 500 million unique visitors.[1]

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, in contrast, is designed to engage people who have not yet returned to the firm's website. Despite much enthusiasm about dynamic retargeting, advertisers currently do not know whether this technique is indeed effective, nor do they know what information they can use to determine when to show such ads featuring content that is highly specific to an individual consumer. With this research, we intend to fill these gaps.

We empirically explore these questions using data from an online field experiment by a travel company. After consumers viewed hotel options on the travel company's website, 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 saw either a random ad that contained an image of the specific hotel they had previously browsed plus three similar hotels (dynamic retargeting) or a random generic brand ad for the travel firm (generic retargeting). Surprisingly, we find that, 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 (Lee, Keller, and Sternthal 2010; Trope, Liberman, and Wakslak 2007).

Consumers may initially have only a broad idea what they want. Their preferences are construed at a high level, and they focus on higher-level goals. For example, they may want a "relaxing vacation." Over time, consumers shift their focus to specific product attributes and develop narrowly construed preferences: they may narrow their search to a hotel with a large swimming pool near the beach. Thus, in using the term "narrowly construed preferences," we refer to consumers who have a detailed viewpoint of the kinds of products they want to purchase. We propose that consumers who focus on higher-level goals respond better to advertising messages that address such goals than to messages that display specific products. Only consumers with narrowly construed preferences are 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 that reflects such a shift in goals. To do so, we isolate browsing behavior that may indicate that a consumer has shifted to narrowly construed preferences and may be more receptive to such highly specific advertising. We use whether a consumer has visited a travel review site as a proxy. When searching product-specific information on a review site (e.g., tripadvisor.com), a consumer compares and contrasts product features and confronts 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 consider a visit to a review site a potential proxy measure that a consumer now has narrowly construed preferences.

We find that generic ads are most effective for consumers who have not yet sought out product quality information at a review site. Dynamic retargeting becomes relatively more effective only after consumers have visited a product review site. We observe that the effectiveness of retargeting further increases for consumers who are browsing category-level content simultaneously. This is consistent with previous literature that has suggested that the quality of the advertising message is mediated by consumers' involvement (Petty, Cacioppo, and Schumann 1983). In summary, our results suggest that consumers' response to ads varies as they progress to different stages in their decision process.

We acknowledge that visiting a travel review website, in addition to being a proxy for consumers' 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 explain why consumers react more positively to a specific (vs. generic) ad after visiting a travel review site.

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, however, generic retargeting is more effective. Our findings regarding 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 introduced retargeted ads into their members' news feeds as a central plank in their advertising strategy (Rusli 2013). In addition, our results provide operational insights for managers who are considering embracing this new technology. They suggest that dynamic retargeting is best employed when managers also have access to external browsing data that would help them identify whether preferences have evolved and, thus, when dynamic retargeting will be effective. Otherwise, more generic campaigns may work better. We next discuss how this result, in addition to being managerially useful, adds to the existing literature.

## *RELATIONSHIP TO PRIOR LITERATURE*

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

[1]See http://www.nextperformance.com.

Table 1
PREVIOUS LITERATURE

| Source | Setting | Personalization | Targeting | Decision Stages | Finding |
|---|---|---|---|---|---|
| *Personalized Recommendations* | | | | | |
| Linden, Smith, and York (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 .30%. |
| *Tailored Communications* | | | | | |
| Ansari and Mela (2003) | Content of e-mail newsletter | Customer content category | None | No | Personalization increases click-throughs. |
| Malthouse and Elsner (2006) | Content of cover letter of mail-order catalog | Recency, frequency, monetary | None | No | Segment-based customization is cost-effective. |
| Agarwal, Athey, and Yang (2009) | Content on a firm's website | Segments defined by demographics and browsing behavior | None | No | Bayesian approach dominates nonpersonalized content selection. |
| Hauser et al. (2009) | Content on a firm's website | Cognitive style segments | None | No | Personalization based on 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, Pavlov, and Canny (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 (2011c) | Ad network | None | Behavioral | No | Privacy regulation that restricts behavioral targeting reduces ad effectiveness. |
| Joshi, Bagherjeiran, and Ratnaparkhi (2011) | Match ads to users and content on firm's website | Based on demographics and website visits, searches, ad views, and ad click | Behavioral, contextual | No | Matching ads to the appropriate website content can be improved by integrating user characteristics. |

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

ommendations on a firm's website has focused on both documenting their effectiveness (Dias et al. 2008) and suggesting ways of improving their effectiveness (Linden, Smith, and York 2003). However, by their very nature, these personalized recommendations only appear to consumers who have 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 indicates that tailoring improves communications' performance. Marketers can use consumer characteristics such as cognitive style (Hauser et al. 2009), celebrity affinity (Tucker 2011), browsing behavior (e.g., previous ads clicked; Agarwal, Athey, and Yang 2009), and prior purchases (Malthouse and Elsner 2006) to identify appropriate segments for which they can customize ads. However, this extant literature has focused on identifying consumer segments and then showing the appropriate ad for that segment rather than individualizing each ad.

An increasing 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. These researchers find that data on consumer browsing behavior (Chen, Pavlov, and Canny 2009) and demographics (Joshi, Bagherjeiran, and Ratnaparkhi 2011) can improve targeting. However, they have not examined whether individual advertisers might benefit by incorporating information into the content of their ads that is highly specific to individual consumers, such as their prior product interests. Gartaski (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.

Our study builds on this literature in four ways. First, we study personalized recommendations outside the firm's website. Second, we focus on ad content personalized to individual browsing history rather than segments. Third, we examine the tailoring of ad content rather than the selection of who sees ads based on prior browsing behavior. Fourth, we analyze whether and how the effectiveness of such retargeted ads changes depending on browsing behavior that suggests a consumer's preferences have evolved. Our results show that online data gathering of consumers' behavior outside the firm's boundaries can be used not only to target but also to *time* when ads are shown. This tactic builds on prior work by Lambrecht, Seim, and Tucker (2011) that shows that marketers can use detailed online data to understand how different stages of a consumer's purchase process interconnect. In recognizing that consumers' choice decisions are often the result of multiple decision stages, our work further builds on insights that such decisions can be modeled as multistage decision processes (Lambrecht, Seim, and Tucker 2011) and more broadly relates to a literature that models behavior as an outcome of a two-stage decision process involving both product choice and choice of timing of consumption (Ascarza, Lambrecht, and Vilcassim 2013; Lambrecht, Seim, and Skiera 2007).

## EMPIRICAL SETTING AND DATA

### *Retargeting*

Retargeting is typically organized by an advertising network on behalf of a focal firm. Ad networks aggregate advertising space across multiple web content publishers and then sell this space to advertisers. The result is that advertisers do not need to manage relationships with a large number of web publishers, which significantly increases the efficiency in the market for online ad space. Because retargeting is a new technology, we describe in more detail how it is typically implemented:

1. Product exposure: The consumer visits the focal firm's website and views products. For each product page the consumer views, a pixel tag (i.e., a $1 \times 1$ image) is downloaded automatically, recording that the consumer was viewing a specific product. This information becomes part of the individual user profile held by the ad network on behalf of the focal firm and is typically tracked by cookies.
2. Targeting consumers: The consumer browses the Internet. At some point, he or she 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 to identify people who have visited the focal firm's website 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 its 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 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, Panel A, we display an example of dynamic retargeting. A consumer who browsed a pair of children's shoes at an online retailer (left panel) later viewed an ad that displayed this exact product alongside three other similar products (right panel). Dynamic retargeted ads use standardized designs in which 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 he or she visits the website again. The ad network is usually not given information about the precise product that was purchased.

### *Data*

We use data on a travel website that sells hotel stays and hotel vacation packages to consumers. It advertises its services on external websites using several ad networks. The firm engages in four types of targeted online advertising, which we summarize in Table 2.

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 he or she subsequently visited an external website on which the ad network showed ads on behalf of the firm.[2]

[2]This particular retargeting network did not engage in real-time bidding for the pricing of its ads but instead used a previously agreed-on rate. This strategy reduces the potential for distortion that would result if the allocation of advertising were decided on the basis of an auction network.

Figure 1
DYNAMIC RETARGETING EXAMPLES

*A: Dynamic Retargeting in Apparel Category*


*B: Mock-Up of Travel Ads Used in Field Experiment*

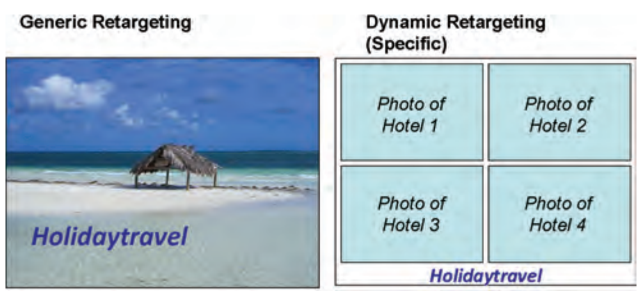


The travel firm ran the field test for 21 days for the hotel category, which is its major product focus. All consumers who viewed a specific hotel on the travel firm's website during the 21-day time period were eligible for the field experiment. Reflecting the beach destination focus of the travel firm, 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 in addition to three others that were similar in location and star rating. We do not have information on which hotel was displayed.[3] In Figure 1, Panel B, due to confidentiality agreements, we can show only an approximation of the design of the dynamic retargeting ads the firm used; the actual ads displayed were more expertly and attractively designed. We use it to illustrate 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 and generic retargeting as they are commonly practiced.

If, on any day, the consumer visited multiple websites that were part of the ad network that implemented the field experiment, he or she would see multiple retargeted ads. However, we designed the randomized trial 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 subsequent regression analysis.

Table 3, Panel A, summarizes the individual-level data. It covers ad exposure data collected by the firm for the 77,937 individual profiles of consumers who were part of the field experiment because they had visited both a page on the firm's website devoted to a specific hotel and, subsequently, websites that were part of the ad network that implemented retargeting.[4] The firm did not store data recording the products consumers initially browsed. Although only one ad network implemented retargeting, our data include all ad exposures across all ad networks with which the firm collaborated.

We do not observe whether the consumer visits the firm's website again. If consumers revisit the firm's website, the firm may better understand their preferences and so may be better able to match products to them in a dynamic targeted ad compared with consumers who do not revisit the firm's website. This is an additional factor that our field data preclude us from measuring, which could contribute 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 made a purchase within the 21-day observation period. Very few consumers clicked on ads. Indeed, click-through rates were, on average, less than .01%. However, as Dalessandro et al.

[3]Typically, 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.

[4]As in prior research (Rutz and Bucklin 2011), we do not have data to address measurement errors introduced by multiple cookies on multiple computers that lead to inaccurate matching of an individual consumer with an individual profile.

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 firm's website | Generic ad displaying brand and evocative vacation image | Yes |
| Dynamic retargeting | Firm advertises to consumers who had previously visited the firm's website | Ad displays products reflecting consumers' prior product search[a] | Yes |

[a]Figure 1, Panel 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 customer viewed alongside similar shoes.

Table 3

CONSUMERS ELIGIBLE FOR DYNAMIC RETARGETING

| A: Cross-Sectional Descriptives | | |
|---|---|---|
| | *M* | *SD* |
| Purchase | .100 | .300 |
| Visited review site | .402 | .490 |
| Observations | 4,542 | |

| B: Time-Varying Covariates | | | | |
|---|---|---|---|---|
| | *M* | *SD* | *Min* | *Max* |
| Any Retargeted Ad | .089 | .284 | 0 | 1 |
| Dynamic Retargeted Ad | .047 | .211 | 0 | 1 |
| Any Ad | .214 | .410 | 0 | 1 |
| Other Behavioral Ad | .122 | .328 | 0 | 1 |
| Contextual Ad | .042 | .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 | 1,313 |
| Observations | 83,214 | | | |

(2012) discuss, 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 purchases made either after the campaign or offline.

Visited Review Site indicates that 40% of consumers were exposed to one of the firm's ads when browsing a travel review site. Because the firm was a major advertiser on the main travel review sites, it is reasonable to believe that most visitors to a travel review site would have been exposed to its advertising. There is a positive correlation between visiting a travel review site and the likelihood of purchase. Approximately 8.6% of consumers who do not visit a travel review site purchase the product, whereas 14.6% of consumers who do visit a travel review site ultimately purchase. This finding suggests that a visit to a review site may not be random, a subject we address directly in our empirical specification.

Table 3, Panel B, describes the data at a daily level over the 21 days, including the types and number of ads to which each consumer was exposed. Any Retargeted Ad 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. Any Retargeted Ad × Dynamic Retargeted Ad reflects that approximately half of these ads were dynamic retargeted ads. In total, more than 149,000 generic retargeted ads and more than 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 dynamic retargeted ad ($p$ = .56) on successive days. In addition, consumers who had viewed a specific type of ad content on a particular day were not more likely to receive either a generic or dynamic ad on that day (viewed travel website: $p$ = .19; viewed news website: $p$ = .21). Importantly, the amount of ads they had previously seen also did not affect the type of retargeted ad they were shown on their next visit ($p$ = .46). This evidence provides further support that generic or dynamic retargeted ads were shown randomly and that there is no systematic variation in websites that showed the different types of ads that could explain our result.

Any Ad captures that, on average, a consumer had a 21.4% probability of being exposed to at least one ad by the travel firm. Contextual Ad captures that a consumer was exposed to a contextual targeted ad on 4.2% of days; similarly, Other Behavioral Ad captures that a consumer was exposed to behavioral targeted ad on 12.2% of days. Neither of the latter categories displayed dynamic content. Similarly, we summarize the cumulative number of ads in each category that a consumer profile viewed before that particular date throughout the 21 days of the field experiment.

Table 4, Panel A, reports the same data as Table 3, Panel A, but it does so for all 2,818,661 consumers to whom the firm served any type of ad during the 21 days of the field experiment rather than just those who were part of the field test. This means that it includes consumers who had not viewed a specific product on the firm's website. The indicator variable Eligible for Retargeting reflects whether 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 greater 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 that the large gains for retargeting that many industry studies report may be because these consumers are already more likely to purchase, because 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, because a necessary condition for dynamic retargeting is that a consumer has visited the website, this is the local average treatment effect of interest.

Figure 2 presents average daily conversion rates by whether a field experiment participant was exposed to a particular type of ad that day. This figure illustrates that con-

Table 4

ALL CONSUMERS

| A: Cross-Sectional Descriptives | | |
|---|---|---|
| | *M* | *SD* |
| Purchase | .020 | .139 |
| Eligible for retargeting | .069 | .253 |
| Visited review site | .091 | .288 |
| Observations | 104,846 | |

| B: Time-Varying Covariates | | | | |
|---|---|---|---|---|
| | *M* | *SD* | *Min* | *Max* |
| Any Retargeted Ad | .002 | .049 | 0 | 1 |
| Dynamic Retargeted Ad | .001 | .036 | 0 | 1 |
| Any Ad | .028 | .166 | 0 | 1 |
| Other Behavioral Ad | .023 | .150 | 0 | 1 |
| Contextual Ad | .005 | .068 | 0 | 1 |
| Cumulative Any Retargeted Ads | .343 | 2.763 | 0 | 248 |
| Cumulative Dynamic Retargeted Ads | .309 | 2.455 | 0 | 248 |
| Cumulative Other Behavioral Ads | 3.952 | 23.586 | 0 | 5,504 |
| Cumulative Contextual Ads | 9.485 | 7.519 | 0 | 1,313 |
| Observations | 2,138,038 | | | |

Figure 2
CONVERSION RATE WITH SAME-DAY AD EXPOSURE

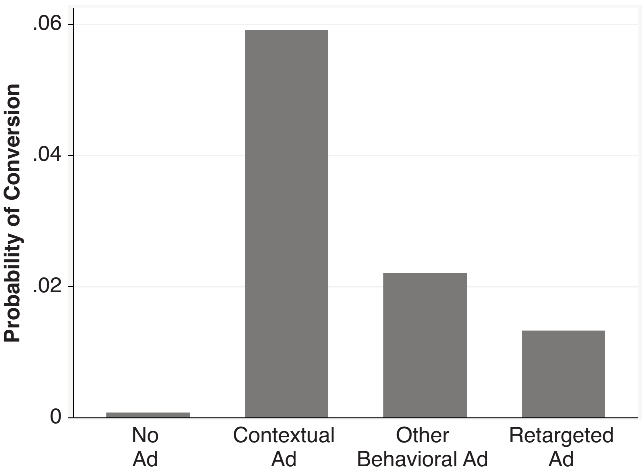


Figure 3
COMPARISON OF CONVERSION FOR GENERIC VERSUS DYNAMIC AD EXPOSURE

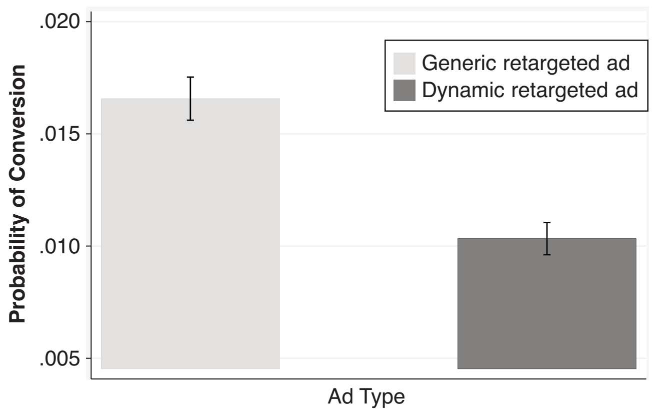


sumers who did not browse other websites served by any of the ad networks were unlikely to purchase. However, their lack of exposure to ads could simply reflect that they were not online that day and, consequently, were not making online purchases (Lewis, Rao, and Reiley 2011). In general, Figure 2 emphasizes 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, from our data, it would seem that for people who previously visited the firm's website, contextual ads are extremely successful and retargeted ads are unsuccessful. However, this correlation may only reflect that consumers who browse travel content are more likely to purchase travel products in general. It is this type of endogeneity that leads us to focus on the field test in our analysis.

## *RESULTS*

### *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 he or she 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 practice 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.[5]

There are important factors for which Figure 3's simple analysis does not control. 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 the firm used. In addition, we need to control for a consumer's underlying probability of purchasing, which 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).

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 observed the consumer being exposed to one of the ads in the test. Because hazard models allow for censoring, they are commonly used to model events that, for a subset of the population, may never occur.

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 someone purchases a product on that day. There are two types of explanatory variables in a proportional hazard model: the baseline hazard, $h_0(t)$, and the vector of covariates, $(X_{it})$. The baseline hazard captures the effect of the time since we first observed a consumer being exposed to an ad in our data.[6] After the consumer purchases from the travel firm, we no longer include the consumer's data. To increase flexibility, we estimate the baseline hazard nonparametrically (Seetharaman and Chintagunta 2003). The vector of covariates, $X_{it}$, captures the effect of different types of ads to which a consumer was exposed on the probability to purchase on any given day. The hazard rate for consumer i $h_i(t, X_t)$ is, therefore,

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

[5]Our results are robust to allowing ad exposure to affect purchases within a two-day and four-day window.

[6]We are forced to use days since the first date of ad exposure because we do not have data recording 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.

We specify the vector of covariates for consumer i as

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

The term $\beta_1$ measures the effect of the consumer being exposed to a dynamic retargeted ad—that is, an ad that had information content that was specific to the previous products he or she was browsing on the website; $\beta_2$ measures the effect of the baseline control condition in which 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. The term $\beta_5$ measures response to the cumulative number of retargeted ads with specific content that the person has seen so far. These terms enable 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, and $\beta_7$ and $\beta_8$ measure response to the cumulative number of behavioral and cumulative number of contextual ads.

In Table 5, Column 1, we show that our results hold when we use a straightforward measure of purchase incidence as our dependent variable. In this discrete choice specification, the dependent variable is whether 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. Because we include controls for each day, this specification is equivalent to a discrete-time hazard model (Allison 1982). Even with additional controls, these results reinforce the basic pattern in the data shown in Figure 3: dynamic retargeting is, on average, less effective. The small size of the coefficients emphasizes that although 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 their low cost means they still offer a reasonable return on investment. Our controls serve as a proxy not only for different types of targeting but also for whether a consumer is seeking travel-category content that day. The cumulative ad controls measure the effect of the stock of previous online ads to which the person has been exposed. The estimates for the controls do not have a clear causal interpretation. All results hold when we exclude cumulative ad totals.

Although our use of a linear probability model in Table 5, Column 1, facilitates the interpretation of interactions (Ai and Norton 2003), Column 2 shows that our results also hold when using a probit specification. Column 3 reports

Table 5

DYNAMIC RETARGETING FOR THOSE ELIGIBLE FOR THE RETARGETING CAMPAIGN

| | *Discrete Choice: Ordinary Least Squares* | *Discrete Choice: Probit* | *Cox Hazard* | *Multiple Impressions Excluded* | *Multiple Impressions Included* | *Lags* |
|---|---|---|---|---|---|---|
| Dynamic Retargeted Ad | –.0106*** | –.467*** | –1.111*** | –1.118*** | | –1.348*** |
| | (.000733) | (.0302) | (.340) | (.417) | | (.480) |
| Any Retargeted Ad | .00873*** | .430*** | .695*** | .815*** | | 1.339*** |
| | (.000614) | (.0224) | (.250) | (.304) | | (.332) |
| Total Dynamic Retargeted Ads | | | | | –.251** | |
| | | | | | (.118) | |
| Total Retargeted Ads | | | | | .0150 | |
| | | | | | (.0752) | |
| Other Behavioral Ad | .0182*** | .774*** | 1.821*** | 1.959*** | 1.853*** | 1.825*** |
| | (.000362) | (.00937) | (.161) | (.170) | (.162) | (.163) |
| Contextual Ad | .0502*** | 1.150*** | 2.560*** | 2.664*** | 2.586*** | 2.518*** |
| | (.000946) | (.0115) | (.176) | (.185) | (.177) | (.172) |
| Cumulative Dynamic Retargeted Ads | .000180*** | .0121*** | .0456*** | .0240 | .0207 | .125 |
| | (.0000179) | (.00186) | (.0176) | (.0181) | (.0161) | (.179) |
| Cumulative Any Retargeted Ads | –.000255*** | –.0184*** | –.0564*** | –.0309* | –.0305** | –.374*** |
| | (.0000162) | (.00177) | (.0163) | (.0169) | (.0154) | (.134) |
| Cumulative Other Behavioral Ads | .00000339 | .000257*** | .000758 | .000718 | .000717 | .00445** |
| | (.00000338) | (.0000774) | (.000729) | (.000737) | (.000742) | (.00196) |
| Cumulative Contextual Ads | –.0000729*** | –.00233*** | –.00490** | –.00507** | –.00498** | .0167*** |
| | (.00000340) | (.000172) | (.00212) | (.00218) | (.00210) | (.00267) |
| Day controls | Yes | Yes | No | No | No | No |
| Lagged ad effects | No | No | No | No | No | Yes |
| Observations | 1,502,514 | 1,502,514 | 1,502,514 | 1,419,428 | 1,502,514 | 1,502,514 |
| Log-likelihood | 1,819,099.5 | –40,349.9 | –70,059.8 | –62,344.9 | –70,148.6 | –69,779.8 |

*$p < .10$.
**$p < .05$.
***$p < .01$.

Notes: Dependent variable is time to purchase in all columns except Columns 1–2, in which the dependent variable is whether a purchase was made that day. Column 1 reports ordinary least squares regression coefficients. Column 2 reports probit regression coefficients. All other columns present hazard-model coefficients. Observations for hazard models are the number of "days at risk." Standard errors are robust.

results for our main specification, in which we use a Cox proportional hazards model with time to purchase as the dependent variable, as represented by Equation 2. Again, we observe that, in general, retargeted ads are positively correlated with purchase, but the addition of personalized dynamic content greatly reduces their effectiveness on average. Note that our results are robust to parametric specifications in which the baseline hazard is modeled using the Weibull and exponential distribution.

We then check robustness to different specifications for the baseline hazard and timing assumptions. Table 5, Column 4, checks that our results are robust to the exclusion of observations in which a consumer saw more than one impression of either a generic or 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 user behavior (e.g., repeated reloading of a page) can drive multiple impressions, meaning that advertising impact does not necessarily increase with the number of impressions.

Thus far, the results have assumed that the incremental effect of online advertising is limited to the day that consumers are exposed to it. Table 5, 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, providing evidence for the validity of our approach, which focuses on purchases that occur on the same day as ad exposures. In summary, the empirical evidence presented in Table 5 confirms the insight derived from Figure 3 that, on average, generic retargeting is more effective than consumer-specific dynamic retargeting.

### *The Performance of Dynamic Retargeting Varies with Browsing*

*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 (Lee, Keller, and Sternthal 2010; Trope, Liberman, and Wakslak 2007).

Consumers often initially have only a broad idea of what they want. 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 general idea of what they like, they do not know the specific trade-offs they would like to make to satisfy their needs best. They may know, for example, that they want a relaxing vacation but not whether they would prefer a large hotel with a large pool or a 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.[7] They also determine how much weight to place on different attributes (Hoeffler and Ariely 1999). They may, for example, learn that their preference for access to a pool outweighs the cost of choosing a less intimate hotel. The term "narrowly construed preferences" refers to consumers who have a detailed viewpoint of the kinds of products they want to purchase.

We build on Simonson's (2005) insight that a consumer's stage of preference development may significantly affect the effectiveness of personalized ad content. We propose that ads conveying information about 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' higher-level goals, such as relaxation. They can enhance brand preferences at a point when consumers do 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 attributes. Therefore, we propose that dynamic ads may be ineffective when shown to consumers with only 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, Liberman, and Wakslak 2007). For example, a consumer may set a specific date for a vacation, reducing uncertainty about whether and when to travel. Such consumers are more likely to construe their preferences narrowly and will explore specific choices instead of focusing on their higher-level goal. When searching for detailed product information, they will begin to make trade-offs on the basis of how much they value and weight certain attributes (Hoeffler and Ariely 1999).

*Empirical results on browsing*. The task of empirically identifying an indicator of a consumer's narrow preferences and positive disposition to a dynamic ad is a challenging one. We suggest that, in an online environment, a visit to a site that provides detailed information about specific products (e.g., a product review site) signals that a consumer is thinking about specific product attributes, comparing and contrasting product features, and confronting the trade-offs inherent in a product choice and so is likely to be developing (or have already developed) narrowly construed preferences.

We use data collected by the travel firm regarding whether a consumer had visited a travel review website, such as tripadvisor.com. The focal firm only displayed standard generic ads on these review sites. These websites provide large numbers of detailed traveler reviews about hotels and travel products. For example, tripadvisor.com has nearly 25 million reviews on more than 490,000 hotels and attractions, boasts more than 11 million registered members, and operates in 14 countries and ten languages. We recognize that a review site visit may be a manifestation of many other different phenomena and explore alternative 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.

In our data, of consumers who both purchased a product and visited a review site, 54% visited the review site before making their purchase, whereas 46% visited after making a

[7]This may be linked to the consumer's stage in his or her decision process (Häubl and Trifts 2000; Hauser and Wernerfelt 1990; Lavidge and Steiner 1961; Wu and Rangaswamy 2003).

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 graphic evidence in which we stratify the purchase probability conditional on being exposed to a generic or a dynamic ad by whether the consumer had already visited a product review site. This figure uses data only on consumers who visit a review site at some point during our observation. 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 underperform. After viewing a review site, generic brand ads become relatively less effective, whereas dynamic ads become relatively more effective.

As before, 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 the person had already visited a product review website. Table 6 reports the results. Column 1 displays results for a proportional hazard model for all consumers who were eligible for the field experiment. As before, we find that the coefficient of Dynamic Retargeted Ad is negative—that is, the dynamic retargeted ad performs worse than the generic ad on average. However, Dynamic Retargeted Ad × After Review Site is positive and significant. This means that the effectiveness of the dynamic retargeted ad improves after the consumer visits a review site.

Figure 4

COMPARISON OF CONVERSION FOR GENERIC VERSUS DYNAMIC AD EXPOSURE: SAMPLE RESTRICTED TO THOSE WHO VISITED A REVIEW SITE

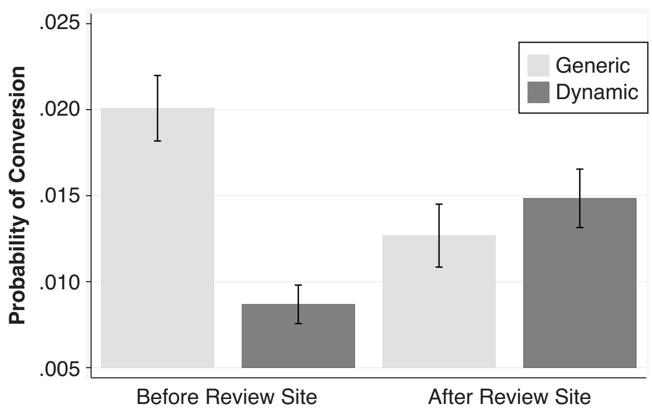


In summary, our results suggest that dynamic retargeting is ineffective when targeted at consumers who have only a broad idea of what they want. It is an effective form of

Table 6

INTERACTIONS WITH VISIT TO REVIEW SITE

| | *All Users* | *Only Review Site Users* | *Saw Ads Post–Review Site* | *Exclude Purchases Two Days After Review Site Visit* | *Discrete Choice: Ordinary Least Squares* | *Discrete Choice: Probit* | *Lags* |
|---|---|---|---|---|---|---|---|
| Dynamic Retargeted Ad | .783*** | .881*** | .962*** | .685*** | .011*** | .493*** | .903*** |
| × After Review Site | (.137) | (.161) | (.189) | (.205) | (.002) | (.068) | (.161) |
| Dynamic Retargeted Ad | –1.331*** | –1.718*** | –1.614*** | –1.106*** | –.014*** | –.622*** | –1.318*** |
| | (.077) | (.117) | (.154) | (.147) | (.001) | (.049) | (.157) |
| Retargeted Ad | .880*** | 1.036*** | 1.068*** | 1.662*** | .010*** | .426*** | 1.229*** |
| | (.051) | (.073) | (.110) | (.098) | (.001) | (.029) | (.091) |
| After Review Site | .212*** | .408*** | .557*** | .688*** | .004*** | .244*** | .452*** |
| | (.035) | (.042) | (.058) | (.055) | (.000) | (.020) | (.042) |
| Retargeted Ad | –.733*** | –.340*** | –.305** | –1.176*** | –.009*** | –.389*** | –.278** |
| × After Review Site | (.106) | (.118) | (.144) | (.160) | (.001) | (.047) | (.117) |
| Cumulative Any Retargeted Ads | –.052*** | –.071*** | –.105*** | –.043*** | | | –.156*** |
| | (.004) | (.007) | (.013) | (.008) | | | (.033) |
| Cumulative Retargeted Ads | –.022** | –.001 | .031** | –.023** | | | –.003 |
| × After Review Site | (.009) | (.010) | (.015) | (.011) | | | (.010) |
| Cumulative Dynamic Ads | .043*** | .071*** | .125*** | .016* | .000*** | .003*** | –.228*** |
| | (.005) | (.008) | (.014) | (.009) | (.000) | (.001) | (.059) |
| Cumulative Retargeted Dynamic Ads | .014 | –.019* | –.073*** | .034*** | –.000*** | –.000 | –.017 |
| × After Review Site | (.010) | (.011) | (.016) | (.012) | (.000) | (.001) | (.010) |
| Cumulative Total Ads | | | | | –.000*** | –.006*** | |
| | | | | | (.000) | (.001) | |
| Cumulative Total Ads | | | | | –.000*** | –.003*** | |
| × After Review Site | | | | | (.000) | (.000) | |
| Constant | | | | | .005*** | –2.806*** | |
| | | | | | (.001) | (.029) | |
| 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 | 1,502,514 | 601,475 | 276,000 | 548,191 | 601,475 | 601,475 | 601,475 |
| Log-likelihood | –69,998.3 | –32,391.6 | –20,103.0 | –20,033.1 | 652,083.7 | –19,459.8 | –32,182.3 |

*p < .10.
**p < .05.
***p < .01.

Notes: Dependent variable is time to purchase in all columns except Columns 5–6, in which the dependent variable is whether a purchase was made that day. Column 5 reports ordinary least squares regression coefficients. Column 6 reports probit regression coefficients. All other columns report hazard coefficients. All controls and appropriate interactions from Table 5, Column 3, are 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. Standard errors are robust.

advertising, however, when addressing consumers who have narrowly construed preferences and so are more likely to focus on specific product details.

In Table 6, Column 2, we restrict our data to consumers who visited a review site at some point during our data period. This is because consumer characteristics that 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 review sites may (1) be more experienced, (2) have seen more competitive ads, (3) be more familiar with the travel category, or (4) have a preference for drawing independent conclusions. For comparison, in Column 1 of Table 6, we report the results for the entire sample. The coefficients are similar in precision and direction, although those for the unrestricted sample are slightly smaller. This reassures us that the sample selection issues inherent in comparing visitors to review sites with nonvisitors are less empirically important than might have been initially 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 technologically "able" consumers are more likely to prefer generic ads and cease all web activity after visiting a review site. This would mean that our result is an artifact capturing that, post–review site visit, we only observe consumers who are predisposed to dynamic retargeted ads. To address this concern, we report results restricted to people who were exposed to an ad at least once following their review site visit in Column 3 of Table 6. Again, the results are similar.

A third possible set of concerns centers on the possibility that a review site visit might in itself directly provide new information, thus altering consumers' choices. For example, there could be a direct effect of reinforcing quality information. To address 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. Column 4 of Table 6 shows that our results hold. Columns 5–7 demonstrate that our results are robust to our earlier robustness tests for a discrete choice specification and the inclusion of lags.

*Additional evidence of mechanism.* Thus 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 by other factors. To do this, we turn to a factor that has been documented to shift the appeal of an ad. Specifically, Petty, Cacioppo, and Schumann (1983) show that the argument quality of an ad has greater effect under high versus 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, this is indirect evidence that the effect we document is driven by the changing effectiveness of advertising content rather than by external factors. We empirically investigate 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 (e.g., Condé Nast Traveler) provide a wide range of information about vacation destinations but not about specific hotels. Because we observe in our data whether a consumer was exposed to an ad by the travel firm on a travel content site, we use this as an indicator for the consumer's browsing a travel content site and 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 website categories to show ads for competitive products. We have no data, however, on whether such ads were present. Therefore, our estimates should be considered to reflect but not control 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 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 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 more effective than the generic ad for consumers who visited a travel review site (which implies that the consumer is considering trade-offs in product features) and also browsed travel content websites that day (which implies they are involved in the category). An interpretation of this finding is that the consumer involvement proxied for by the travel content website browsing enhances the argument quality of the dynamic ad for consumers with narrowly construed preferences. In a robustness check, we find that these results hold when we exclude observations in which consumers were exposed to an ad during or after their travel category browsing on that particular day. This means that our results are not driven by reverse causation (in which the

Figure 5

COMPARISON OF A CONVERSION FOR GENERIC VERSUS DYNAMIC AD EXPOSURE BY BROWSING BEHAVIOR

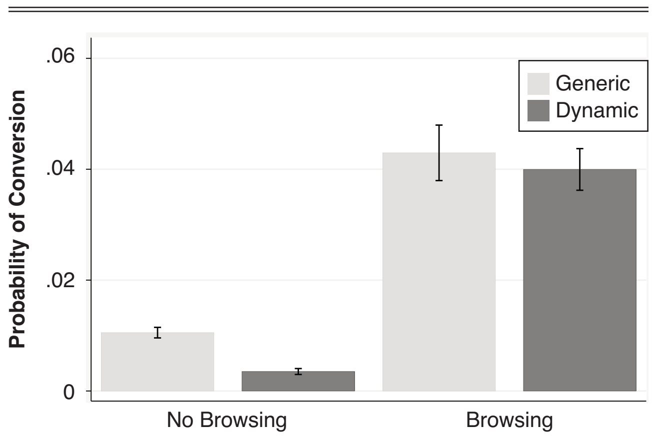


Notes: Sample restricted to consumers who browsed a travel website and a review website.

Figure 6

COMPARISON OF A CONVERSION FOR GENERIC VERSUS DYNAMIC AD EXPOSURE BY BROWSING BEHAVIOR AND BEFORE/AFTER REVIEW SITE VISIT

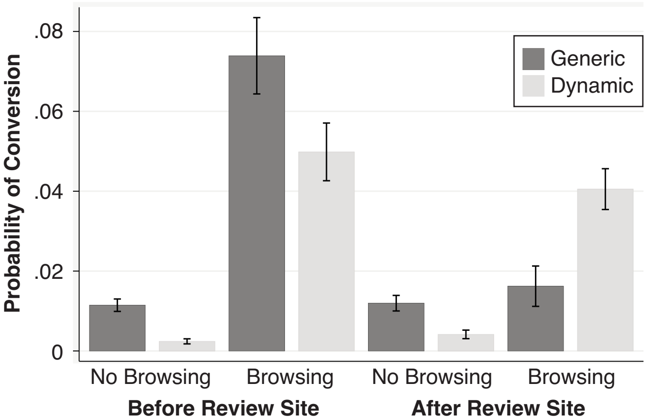


Notes: Sample restricted to consumers who browsed a travel website and a review website.

ad provokes people to browse the travel category) or by a contextual effect of the ad. Likewise, our results hold when we exclude observations in which consumers were exposed to an ad before they browsed the travel category.

We repeat this analysis in a regression framework in Table 7. We limit our sample to 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 that person visited a website devoted to category-related information that day. The negative and significant coefficient for Retargeted Ad × Dynamic Retargeted Ad suggests that dynamic ads are less effective than generic ads. However, this is mediated by the positive and significant coefficient for Retargeted Ad × Dynamic Retargeted Ad × Browsing Travel That Day, which suggests that dynamic retargeted ads perform relatively better on days when consumers browse travel content. As before, these results are robust to different definitions of the baseline hazard or a discrete choice model. We echo the analysis of Table 6 and stratify the results by whether the consumer has visited a review site in Columns 2 and 3. The baseline measure of Retargeted Ad × Browsing Travel That Day is more negative after the consumer visits the review site. Therefore, the performance of generic retargeted ads becomes relatively worse on days when a consumer browses travel content after he or she visits a review site. However, the increasing size of the coefficient Retargeted Ad × Dynamic Retargeted Ad × Browsing Travel That Day after a consumer visits a review site suggests that, in contrast, dynamic retargeted ads perform relatively better after a consumer has visited a travel review site and browsed the category that day. In general, these results suggest that the most effective time to use dynamic (vs. generic) retargeting is after a consumer visits a review site and seems to be actively involved in the category. This is, again, in line with Figure 6.

These results support the interpretation of our previous findings that dynamic retargeted ads perform well only when preferences are narrowly construed. We recognize that

Table 7

FURTHER EVIDENCE OF MODERATING EFFECT OF CATEGORY ACTIVITY

| | *All Survival Time* | *Before Review Site Survival Time* | *After Review Site Survival Time* |
|---|---|---|---|
| Dynamic Retargeted Ad × Browsing Travel That Day | 1.515*** | .838*** | 2.512*** |
| | (.175) | (.267) | (.279) |
| Dynamic Retargeted Ad | –2.114*** | –2.411*** | –1.803*** |
| | (.140) | (.213) | (.195) |
| Retargeted Ad × Browsing Travel That Day | –.392*** | –.306* | –1.496*** |
| | (.123) | (.170) | (.238) |
| Retargeted Ad | .585*** | .774*** | .951*** |
| | (.090) | (.135) | (.145) |
| Browsing Travel That Day | 1.356*** | 1.959*** | 1.479*** |
| | (.053) | (.085) | (.104) |
| Cumulative Any Retargeted Ads | –.089*** | –.081*** | –.105*** |
| | (.009) | (.012) | (.013) |
| Cumulative Retargeted Ads × Browsing Travel That Day | .002 | –.033* | .037** |
| | (.012) | (.017) | (.017) |
| Cumulative Dynamic Ads | .070*** | .072*** | .076*** |
| | (.009) | (.013) | (.015) |
| Cumulative Retargeted Dynamic Ads × Browsing Travel That Day | .003 | .047*** | –.024 |
| | (.013) | (.018) | (.019) |
| Further ad controls | No | Yes | Yes |
| Observations | 145,452 | 80,581 | 64,871 |
| Log-likelihood | –24,077.7 | –12,040.4 | –9,418.0 |

*$p < .10$.
**$p < .05$.
***$p < .01$.

Notes: Proportional hazard regression coefficients shown. Dependent variable is time to purchase. Sample restricted to consumers who visited a travel website and a review site at some point during our observation. 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; however, they are not reported here for reasons of space. Standard errors are robust.

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.

## *CONFIRMING THE RESULTS IN THE LAB*

### *Objective*

We aim to show that the interpretation of our results holds in a controlled lab environment. Specifically, first, we directly test whether consumers' narrowly construed preferences indeed affect how they react to generic versus dynamic ads. Second, we intend to rule out privacy concerns, reactance, competitive effects, and consumer experience as alternative explanations of our results, because prior literature has shown these factors to affect the performance of online advertising (Goldfarb and Tucker 2011a, b). Our experimental setup also provides evidence that neither social validation nor increased access to quality information—both inherent to visiting a review site—drive our results.

### *Design and Procedure*

The study has a 2 × 2 design. We vary whether consumers have refined their preferences (broad preferences vs. narrowly construed preferences) and the type of ad to which they are exposed (generic ad vs. dynamic ad). In the broad-idea condition, participants were asked to imagine that they would like to go on a beach vacation. They do not know where they would like to go and are still exploring different destinations. In addition, they have not thought about specific hotels because they would like to choose their destination first. They were then asked to imagine that they looked at hotels in many different regions on a travel company website. Participants in the narrow-preferences condition were asked to imagine that they would like to go on a beach vacation in Hawaii and are specifically looking for a hotel with a very large pool. They were then asked to imagine that they evaluated Hawaiian hotels with large pools on a travel company website.

All participants were told to imagine that, as they browse the Internet, they see an ad on another website for this travel company. Participants in the generic-ad condition were told that it shows a picture generally relating to beach vacations. They were shown an ad for a fictitious travel company that displays a sun lounger on the beach. Participants in the dynamic-ad condition were shown an ad for a fictitious travel company with pictures of four hotels with pools. They were told that the ad shows one of the hotels they looked at in addition to three other hotels.[8] Participants were then asked how likely they would be to visit the firm's website and book a vacation (1 = "very unlikely," and 7 = "very likely").

Finally, participants answered several additional questions that related to the scenario. First, we asked how likely they were at this stage of their travel planning to have already visited a travel review site, such as tripadvisor.com. Second, we measured whether either ad causes privacy concerns and whether this possibly varies across conditions. We used the scale developed by Xu (2007) that captures how much consumers are concerned about the privacy of their data in online environments. We then measured reactance to the ad using Edwards, Li, and Lee's (2002) scale. We also asked participants how likely they would be 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. In 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.

One hundred sixty-two participants were recruited online through Amazon.com's Mechanical Turk and randomly assigned to conditions (82 responses for the broad-idea condition, 80 responses for the narrow preferences condition).[9] Buhrmester, Kwang, and Gosling (2011) note that although Mechanical Turk has disadvantages, there is also evidence that using it may lead to more diverse and, thus, more representative samples than traditional samples of American college students.

*Results*

As we hypothesized, participants in the broad-idea condition reacted significantly more favorably to the generic ad than to the dynamic ad (5.214 vs. 4.574; $p = .014$). However, in the narrow preferences condition, participants reacted significantly less favorably to the generic ad than to the dynamic ad (4.227 vs. 5.583; $p < .001$). A regression model in Column 1 of Table 8 confirms that, overall, consumers are less likely to buy after seeing a dynamic ad unless they already have narrowly construed preferences.

Next, we determined whether these results are robust to controls. Participants who viewed the generic ad were less concerned about privacy than participants who viewed the dynamic ad (3.706 vs. 4.247; $p < .001$). Privacy concerns did not differ between participants in the broad-idea condition and those in the narrow-preferences condition (3.931 vs. 3.809; $p = .650$). Reactance to the ad is generally low and did not differ across conditions (generic vs. dynamic: 2.199 vs. 2.466; $p = .134$; broad idea vs. narrow preferences: 2.290 vs. 2.361; $p = .688$). Including these variables in a regression model does not change the main effect of interest (Table 8, Column 2).

Similarly, we find that consumers' propensity to buy from a competitor does not differ by narrowly construed preferences (4.244 vs. 4.275; $p = .826$) or ad conditions (4.267 vs. 4.250; $p = .902$). Including this variable in a regression model again does not change our main effect of interest (Table 8, Column 3).

To determine whether experience with travel booking in general or online increases the effectiveness of a dynamic ad, we estimated alternative models in which we include how often participants have booked a vacation product (Table 8, Column 4) or, alternatively, how often they have booked such services online (Column 5). We also interact these variables with the type of ad displayed in the survey. None of these variables significantly affects whether a consumer will book a vacation.

Finally, we confirm that consumers are indeed less likely to have visited a travel review site when they have broad preferences compared with when they have narrowly construed preferences (5.171 vs. 5.663; $p = .025$). This confirms that visits to travel review sites are a good indicator of whether a consumer has narrowly construed preferences.

In this study, we directly tested for the effect in the lab and did not require consumers to visit a review site. Therefore, the experiment likewise provides evidence that social validation through or access to quality information on a review site is not the primary driver of our results. Our study confirms that whether a consumer has narrowly construed preferences is an important determinant for the effectiveness of generic versus dynamic retargeted ads. We illustrate that it is effective to show consumers dynamic retargeted ads only after they have refined their preferences, whereas in earlier stages, it is more effective to expose them to generic retargeted ads. The Web Appendix includes a similar experiment, in which 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.[10]

[8]We pretested the ads with 85 participants in a survey on Amazon.com's Mechanical Turk. There was no significant difference in how much participants liked the two ads (5.558 vs. 5.286; $p = .208$).

[9]We dropped 5 outliers out of 167 participants. Including these outliers led to results that were directionally consistent and significant, but less precise.

[10]Social validation may potentially make personalized advertising more persuasive (Bakshy and Adamic 2009; Cialdini and Goldstein 2004).

Table 8

LAB STUDY: REPLICATING RESULTS OF FIELD STUDY

| | (1) | (2) | (3) | (4) | (5) |
|---|---|---|---|---|---|
| Dynamic Ad | –.639** | –.601** | –.599** | .435 | .465* |
| | (.256) | (.259) | (.260) | (.276) | (.270) |
| Stable Preferences | –.987*** | –.951*** | –.943*** | | |
| | (.249) | (.249) | (.250) | | |
| Dynamic Ad × Stable Preferences | 1.995*** | 1.951*** | 1.938*** | | |
| | (.364) | (.364) | (.366) | | |
| Privacy | | .0475 | .0522 | .0530 | .0533 |
| | | (.0743) | (.0749) | (.0813) | (.0812) |
| Reactance | | –.155* | –.151* | –.185* | –.180* |
| | | (.0856) | (.0861) | (.0947) | (.0958) |
| Buy Competitor | | | –.0580 | –.104 | –.0958 |
| | | | (.103) | (.113) | (.112) |
| Booked Vacation | | | | .0229 | |
| | | | | (.0259) | |
| Dynamic Ad × Booked Vacation | | | | –.0208 | |
| | | | | (.0608) | |
| Booked Online | | | | | .0290 |
| | | | | | (.0285) |
| Dynamic Ad × Booked Online | | | | | –.0325 |
| | | | | | (.0664) |
| Constant | 5.214*** | 5.360*** | 5.577*** | 5.284*** | 5.222*** |
| | (.178) | (.329) | (.508) | (.541) | (.546) |
| Observations | 162 | 162 | 162 | 162 | 162 |
| R-square | .175 | .192 | .194 | .0525 | .0541 |

*$p < .10$.
**$p < .05$.
***$p < .01$.
Notes: Dependent variable is purchase probability scale. Ordinary least square estimates.

## *CONCLUSION*

The digital revolution has undergone advances in the use of data on browsing behavior both within and outside a firm's website to improve marketing appeals. Internal browsing data have enabled firms to customize their websites so that when consumers return, the firm can show them personalized recommendations based on their previous browsing behavior. External browsing data has enabled firms to better target their ads to consumers who fit a particular profile (e.g., people who have recently browsed travel websites).

Dynamic retargeting represents a combination of these two techniques. Firms can now target consumers on other sites across the Internet with content specific to the product the consumer previously viewed on the firm's website. There is, however, little extant evidence to show whether tailoring advertising content to a consumer's observed preferences is effective.

This article evaluates whether firms indeed benefit from targeting consumers with information that is highly specific to their prior interests. 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 previously viewed 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 at which consumers may find the retargeted ad's emphasis on specific products appealing. We build on a consumer behavior literature stream that suggests that such a specific emphasis on product features will be most effective when a consumer has established narrowly construed preferences. Consumers with 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 narrowly construed preferences is the act of comparing and contrasting product features, 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 because they are less likely to have developed narrowly construed preferences. However, when consumers have visited a review site and 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 by browsing other websites in the category. Because the decision to visit a review site is potentially endogenous, we performed a battery of robustness tests and provide direct evidence for the proposed mechanism in a lab experiment.

We discuss three 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 better define their product preferences and browse related content online. Managers can use data on external website browsing—which are currently available to advertisers but rarely evaluated in detail—to (1) identify whether consumers' preferences are broadly or narrowly construed, (2) determine their category focus, and (3) time the targeting of ads for maximum effectiveness. Third, our results should encourage managers to think more broadly about consumers' responsiveness to ads in the context of a multistage decision process in which ad effectiveness may change with their decision stages.

Of course, there are 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 may not extend to products for which there is little consumer product research and purchasing behavior is driven by either impulse or habit. Second, we do not explicitly address the specifics of dynamic retargeting ad design—in particular, which products should be highlighted and how. Third, we do not have data on competitors' advertising decisions, which would enable us to tease apart how competitive ads moderate the effectiveness of dynamic retargeting. Notwithstanding these limitations, we believe that our research, by documenting the general ineffectiveness of dynamic retargeting and the circumstances under which dynamic retargeting becomes relatively more effective than generic ads, represents a useful contribution to knowledge about this new form of online advertising.

## REFERENCES


Agarwal, Nikhil, Susan Athey, and David Yang (2009), "Skewed Bidding in Pay-per-Action Auctions for Online Advertising," *American Economic Review*, 99 (2), 441–47.

Ai, Chunrong and Edward C. Norton (2003), "Interaction Terms in Logit and Probit Models," *Economic Letters*, 80 (1), 123–29.

Allison, Paul D. (1982), "Discrete-Time Methods for the Analysis of Event Histories," *Sociological Methodology*, 13, 61–98.

Ansari, Asim and Carl Mela (2003), "E-Customization," *Journal of Marketing Research*, 40 (May), 131–45.

Ascarza, Eva, Anja Lambrecht, and Naufel Vilcassim (2013), "When Talk Is 'Free': The Effect of Tariff Structure on Usage Under Two- and Three-Part Tariffs," *Journal of Marketing Research*, 49 (December), 882–99.

Bakshy, Eytan, Brian Karrer, and Lada 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: Association for Computing Machinery, 325–34.

Beales, Howard (2011), "The Value of Behavioral Targeting," research study, Network Advertising Initiative, (accessed June 19, 2013), [available at http://www.networkadvertising.org/pdfs/Beales_NAI_Study.pdf].

Behavioral Targeting Blog (2010), "Targeting and Retargeting Interview with Criteo," (August 26), (accessed June 19, 2013), [available at http://behavioraltargeting.biz/targeting-retargeting-interview-with-criteo/].

Buhrmester, Michael, Tracy Kwang, and Samuel 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, Ye, Dmitri Pavlov, and John 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: Association for Computing Machinery, 209–218.

Cialdini, Robert B. and Noah J. Goldstein (2004), "Social Influence: Compliance and Conformity," *Annual Review of Psychology*, 55 (1), 591–621.

Cox, D.R. (1972), "Regression Models and Life-Tables," *Journal of the Royal Statistical Society*, 34 (2), 187–200.

Dalessandro, Brian, Rod Hook, Claudia Perlich, and Foster Provost (2012), "Evaluating and Optimizing Online Advertising: Forget the Click, but There Are Good Proxies," NYU Working Paper No. CBA-12-02.

Dias, M. Benjamin, Dominique Locher, Ming Li, Wael El-Deredy, and Paolo J.G. 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: Association for Computing Machinery, 291–94.

Edwards, Steven M., Hairong Li, and Joo-Hyun Lee (2002), "Forced Exposure and Psychological Reactance: Antecedents and Consequences of the Perceived Intrusiveness of Pop-Up Ads," *Journal of Advertising*, 31 (3), 83–95.

Gatarski, Richard (2002), "Breed Better Banners: Design Automation Through On-Line Interaction," *Journal of Interactive Marketing*, 16 (1), 2–13.

Goldfarb, Avi and Catherine Tucker (2011a), "Advertising Bans and the Substitutability of Online and Offline Advertising," *Journal of Marketing Research*, 48 (April), 207–227.

——— and ——— (2011b), "Online Display Advertising: Targeting and Obtrusiveness," *Marketing Science*, 30 (3), 389–404.

——— and ——— (2011c), "Privacy Regulation and Online Advertising," *Management Science*, 57 (1), 57–71.

Hargrave, Sean (2011), "Targeted Ads Aim for Greater Accuracy," *Marketing Week*, (January), (accessed June 19, 2013), [available at http://www.marketingweek.co.uk/targeted-ads-aim-for-greater-accuracy/3022557.article].

Häubl, Gerald and Valerie Trifts (2000), "Consumer Decision Making in Online Shopping Environments: The Effects of Interactive Decision Aids," *Marketing Science*, 19 (1), 4–21.

Hauser, John R., Glen L. Urban, Guilherme Liberali, and Michael Braun (2009), "Website Morphing," *Marketing Science*, 28 (2), 202–223.

——— and Birger Wernerfelt (1990), "An Evaluation Cost Model of Consideration Sets," *Journal of Consumer Research*, 16 (4), 393–408.

Hoeffler, Steve and Dan Ariely (1999), "Constructing Stable Preferences: A Look into Dimensions of Experience and Their Impact on Preference Stability," *Journal of Consumer Psychology*, 8 (2), 113–39.

Hoffman, Donna L. and Thomas P. Novak (1996), "Marketing in Hypermedia Computer-Mediated Environments: Conceptual Foundations," *Journal of Marketing*, 60 (July), 50–68.

Hunter, Anne, Meredith Jacobsen, Richard Talens, and Tony Winders (2010), "When Money Moves to Digital, Where Should It Go? Identifying the Right Media-Placement Strategies for Digital Display," white paper, comScore and ValueClick Media (September), (accessed June 19, 2013), [available at www.recruemedia.com/downloads/comscore2010.pdf].

Jain, Dipak C. and Naufel J. Vilcassim (1991), "Investigating Household Purchase Timing Decisions: A Conditional Hazard Function Approach," *Marketing Science*, 10 (1), 1–23.

Joshi, Amruta, Abraham Bagherjeiran, and Adwait Ratnaparkhi (2011), "User Demographic and Behavioral Targeting for Content Match Advertising," in *Proceedings of the Fifth International Workshop on Data Mining and Audience Intelligence for Advertising (ADKDD 2011)*. New York: Association for Computing Machinery, 53–60.

Komiak, Sherrie Y.X. and Izak Benbasat (2006), "The Effects of Personalization and Familiarity on Trust and Adoption of Recommendation Agents," *Management Information Systems Quarterly*, 30 (4), 941–60.

Lambrecht, Anja, Katja Seim, and Bernd Skiera (2007), "Does Uncertainty Matter? Consumer Behavior Under Three-Part Tariffs," *Marketing Science*, 26 (5), 698–710.

——, ——, and Catherine Tucker (2011), "Stuck in the Adoption Funnel: The Effect of Interruptions in the Adoption Process on Usage," *Marketing Science*, 30 (2), 355–67.

Lavidge, Robert J. and Gary A. Steiner (1961), "A Model for Predictive Measurements of Advertising Effectiveness," *Journal of Marketing*, 25 (October), 59–62.

Lee, Angela Y., Punam A. Keller, and Brian Sternthal (2010), "Value from Regulatory Construal Fit: The Persuasive Impact of Fit Between Consumer Goals and Message Concreteness," *Journal of Consumer Research*, 36 (5), 735–47.

Lewis, Randall A., Justin M. Rao, and David H. Reiley (2011), "Here, There, Everywhere: Correlated Online Behaviors Can Lead to Overestimates of the Effects of Advertising," in *Proceedings of the International Conference on World Wide Web*. New York: Association for Computing Machinery, 156–66.

Linden, Greg, Brent Smith, and Jeremy York (2003), "Amazon.com Recommendations: Item-to-Item Collaborative Filtering," *IEEE Internet Computing*, 7 (1), 76–80.

Malthouse, Edward C. and Ralf Elsner (2006), "Customisation with Crossed-Basis Sub-Segmentation," *Journal of Database Marketing & Customer Strategy Management*, 14 (1), 4–50.

Manchanda, Puneet, Jean-Pierre Dubé, Khim Yong Goh, and Pradeep K. Chintagunta (2006), "The Effect of Banner Advertising on Internet Purchasing," *Journal of Marketing Research*, 43 (February), 98–108.

Petty, Richard E., John T. Cacioppo, and David Schumann (1983), "Central and Peripheral Routes to Advertising Effectiveness: The Moderating Role of Involvement," *Journal of Consumer Research*, 10 (2), 135–46.

Rusli, Evelyn M. (2013), "Buy Signal: Facebook Widens Data Targeting," *The Wall Street Journal*, (April 9), (accessed June 19, 2013), [available at http://online.wsj.com/article/SB10001424127887324504704578412960951909032.html].

Rutz, Oliver J. and Randolph E. Bucklin (2011), "From Generic to Branded: A Model of Spillover Dynamics in Paid Search Advertising," *Journal of Marketing Research*, 48 (February), 87–102.

Seetharaman, P.B. and Pradeep K. Chintagunta (2003), "The Proportional Hazard Model for Purchase Timing: A Comparison of Alternative Specifications," *Journal of Business & Economic Statistics*, 21 (3), 368–82.

Simonson, Itamar (2005), "Determinants of Customers' Responses to Customized Offers: Conceptual Framework and Research Propositions," *Journal of Marketing*, 69 (January), 32–45.

Tellis, Gerard J. and Philip Hans Franses (2006), "Optimal Data Interval for Estimating Advertising Response," *Marketing Science*, 25 (3), 217–29.

Trope, Yaacov, Nira Liberman, and Cheryl Wakslak (2007), "Construal Levels and Psychological Distance: Effects on Representation, Prediction, Evaluation, and Behavior," *Journal of Consumer Psychology*, 17 (2), 83–95.

Tucker, Catherine (2011), "Social Networks, Personalized Advertising, and Privacy Controls," MIT Sloan Research Paper No. 4851-10.

Weiman, Lori (2010), "Is the 'Last Click Wins' Model Best for Your Affiliates?" SearchEngineLand.com, (November 10), (accessed June 19, 2013), [available at http://searchengineland.com/is-the-last-click-wins-model-best-for-your-affiliates-55104].

Wu, Jianan and Arvind Rangaswamy (2003), "A Fuzzy Set Model of Search and Consideration with an Application to an Online Market," *Marketing Science*, 22 (3), 411–34.

Xu, Heng (2007), "The Effects of Self-Construal and Perceived Control on Privacy Concerns," in *Proceedings of 28th Annual International Conference on Information Systems (ICIS 2007)*. Montreal: Association for Information Systems, 125.

Yan, Jun, Ning Liu, Gang Wang, Wen Zhang, Yun Jiang, and Zheng Chen (2009), "How Much Can Behavioral Targeting Help Online Advertising?" in *Proceedings of the 18th International Conference on World Wide Web*. New York: Association for Computing Machinery, 261–70.