When Does Retargeting Work? Timing Information Specificity*

Anja Lambrecht† and Catherine Tucker‡

December 2, 2011

Abstract

Firms can now serve personalized recommendations to consumers who return to their website, based on their earlier browsing history. At the same time, online advertising has greatly advanced in its use of external browsing data across the web to target internet ads appropriately. 'Dynamic Retargeting' integrates these two advances by using information from *internal* browsing data to improve internet advertising on *external* websites. Consumers who previously visited the firms' website are shown ads that reflect the specific products they have looked at before on the firm's own website when surfing the wider web. 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 increased ad specificity is on average less effective than generic information. We provide evidence that this can be explained by a mismatch between the specificity of this form of advertising and whether a consumer has well-defined product preferences. Only when consumers have well-defined product preferences and are actively engaged in the category are specific ads more effective than generic ads.

*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 participants at the 2011 SICS conference and seminar participants at Cass Business School, ESMT, ESSEC, London Business School and the National University of Singapore for useful comments.

†London Business School, London, UK; alambrecht@london.edu.

‡MIT Sloan School of Management, MIT, Cambridge, MA; cetucker@mit.edu and NBER

1 Introduction

Innovations in how firms can parse and process individual consumer data now enable them to serve individualized recommendations in real time to consumers who return to their website. These recommendations are often for the specific products that the consumer was previously browsing. These techniques have been successful at improving sales (Linden et al., 2003; Dias et al., 2008). As a result, marketers have begun to use individualized recommendations to enhance the content of online advertising external to the firm's website - a practice known as 'dynamic retargeting'.

Dynamic retargeting combines personalized recommendations based on consumer *internal* browsing of a firm's website with the use of *external* browsing data to track consumers across the web. This external browsing data has been commonly used for targeting ads, that is, selecting the group of consumers who see a certain ad. For example, ads for a vacation product may be seen only by consumers who recently visited a travel site. The innovation of dynamic retargeting is that firms can now, in their online advertising campaign, show to all consumers who browsed their website but did not purchase precisely the product they looked at previously on the firm's website. This significantly extends the reach of a firm's consumer-specific communication, which now is no longer limited to consumers who decide to return to the firm's own website.

At face value, the idea of 'dynamic retargeting' 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). Similarly, firms that offer retargeting technology 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).

However, there is little empirical evidence that targeting consumers with personalized recommendations will similarly benefit firms than when they use the two techniques separately. It is unclear whether a technique designed to engage consumers who are already engaged enough to return to the firm's website will be similarly successful when used to address consumers who may not yet have returned to the firm's website and who may be less aware of what product they are looking to buy. Advertisers currently do not know either whether consumers are at all times similarly receptive to these highly-specific ads. And, if the effectiveness of specific ads varies, what information they can use to time the ads correctly. This research seeks to fill this gap. We ask *whether* and *when* firms benefit from using ads that are highly specific to an individual consumer's prior product search relative to showing ads that display only a generic brand message.

We use data from an online field experiment by a travel firm. The firm tracked consumers who visited their website and the hotels they looked at. When these consumers visited external websites that the travel firm advertised on, the travel firm randomized whether they used dynamic retargeting (showing an ad that contained an image of the specific hotel the consumer had previously browsed plus three similar hotels) or generic retargeting (showing a generic brand ad for the travel firm). This random variation identifies whether highly-specific ads are more effective than generic ads in converting consumers to purchase a travel product. Surprisingly, we find that highly-specific ads are *less* effective than generic ads at convincing consumers to purchase. This suggests that, on average, firms do *not* benefit from targeting consumers with personalized ads that reflect that specific consumer's prior product search.

To explain why specific ads are often less effective than generic ads, we turn to a literature that highlights that consumers may not necessarily have well-defined preferences when they start searching for products (Bettman et al., 1998). Instead, consumers often start their search with a general notion of what they want. During the search, they learn about the available product options and their attributes as well as about their own preferences (Griffin and Broniarczyk, 2010). As a result, consumers may initially focus on broad product benefits and only later, as they refine their preferences, turn to evaluating attributes in more detail.[1]

Building on these behavioral insights, we suggest that the average ineffectiveness of dynamic retargeting may be explained if consumers still lack well-defined product preferences when seeing an ad. When consumers only have a broad notion of what they want and are still constructing their exact preferences, generic advertising may be more effective, since it appeals broadly to their needs. At this point, consumers are not yet evaluating product alternatives in much detail and have little interest in detailed product ads. For example, if a consumer is still unsure about whether to vacation in Florida or in Greece, highlighting that a specific Greek hotel has a large pool may be ineffective. By contrast, specific ads may be more effective when a consumer has further refined their product preferences. Such a consumer is more likely to focus on specific product attributes than on broad category information. For example, they could be evaluating whether the Greek hotel has a large pool or is close to the beach. A firm therefore has to make sure that the level of information specificity in an ad matches to whether a consumer has already developed well-defined product preferences and so searches a category broadly or evaluates product attributes in greater detail.

Insights from consumer's external browsing behavior may help firms to establish whether a consumer's preferences are well-defined. One indicator for whether a consumer is establishing more specific preferences is whether they seek out more specific information on individual products instead of broadly researching a category. Product review sites, such as TripAdvi-

[1]Brucks (1985), for example, discusses that consumers may search for information differently depending on their level of prior product category knowledge.

sor for travel products, allow consumers to evaluate products in depth, and hence support the formation of well-defined preferences.[2] We therefore suggest that a visit to such a site is a good indicator that a consumer is moving from being broadly interested in the category to evaluating specific options in detail.

We explore how whether a consumer has visited a review site affects our results. We find that generic retargeting is most effective before a consumer seeks out product quality information at a review site, that is, before they develop well-defined preferences. Dynamic retargeting becomes relatively more effective only after a consumer has visited a product review site, at which point generic retargeting becomes strikingly ineffective.

We then extend this analysis to account for whether a consumer is actively engaged in the product category on a specific day. We find that dynamic retargeting of consumers with information specific to their prior interests is only effective at encouraging consumers to purchase under very limited circumstances: When consumers have formed well-defined preferences and are engaged in the category. In all other settings, generic retargeting is more effective.

This suggests that firms should be careful in too readily applying insights on the effectiveness of personalized marketing techniques from within their website to consumer behavior outside their own website. More broadly, our results indicate that a firm that aims to advertise with highly relevant information to specific consumers can benefit greatly from using detailed information on browsing behavior across the web rather than limiting itself to information collected on its own website.

2 Relationship to Prior Literature

Our research relates to previous work on personalized recommendations, tailored communications and targeting in online markets. Table 1 summarizes the literature in these fields.

[2]We present survey evidence for how consumers use such websites.

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 customers' 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)	Advertising Network	None	Behavioral	No	Behaviorally targeted ads cost 100% more
Goldfarb and Tucker (2011c)	Advertising 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.

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). It concludes that firms typically benefit from offering their customers personalized recommendations. However, by their very nature these personalized recommendations are only shown to customers who already decided to return to the firm's website. They do not reach consumers who do not return to their 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 on segments rather than individuals means that this kind of communication is not individualized.

The literature on both personalized recommendations and tailored communications has focused on optimizing communications within the confines of a firm's website or direct marketing appeals, limiting the scope of customers the firm can address. Targeted advertising techniques, by contrast, allow firms to connect with customers outside of their own website. A growing body of online targeting literature has attempted to qualify 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 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.[3]

[3] 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.

Though dynamic retargeting builds on elements of personalized recommendations, tailored communications and targeted advertising, it is unclear whether combining these techniques would be as successful as using each of them separately. The fact that consumers' preferences develop over time (Bettman et al., 1998)[4] and that a consumer's stage of preference development may significantly affect the effectiveness of personalized messages (Simonson, 2005), can create a significant challenge for one-to-one marketers who wish to address consumers across the web with highly relevant messages. To do so, advertisers need to empirically identify the stage of preference development for each consumer. However, there is currently little guidance on how this can be done.

Our study fills this gap and is unique in four different ways. First, we focus on advertising messages personalized to individuals, not segments. Second, the messages are highly specific to an individual's interest, since they are not based on demographics or broad browsing behavior but on the specific product a consumer has looked at before but not purchased. Third, these messages address consumers outside of the firm's website. Fourth, we analyze whether and how the effectiveness of such highly-specific messages changes depending on whether a customer has yet developed well-defined preferences and whether they are engaged in the category. We propose that data on external browsing of websites which is now available to advertisers but rarely evaluated in detail, can be used to identify whether a consumer has well-defined preferences. 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* the targeting of their ads. 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.

[4]This may possibly but not necessarily 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).

3 Data

We use data from a travel website that sold hotel stays and hotel vacation packages to consumers. It advertised its services on external websites using several advertising networks.[5] When a consumer viewed a travel product at the firm's website, the firm set a cookie on the consumer's computer to collect data about the consumer's subsequent browsing behavior across the internet. Each time the consumer visited an external website the firm advertised on, the firm used a 'pixel tag' (a small 1x1 pixel image) embedded in the ad to match this exposure to the consumer's cookie. As a result, the firm was able to collect detailed data on advertising exposure and match this with their data on consumers' purchases. The firm engaged in four types of targeted online advertising. These are summarized in Table 2. As discussed in Section 2, the literature has so far focused on behavioral and contextual targeting.

The firm conducted a field experiment, in cooperation with a major advertising network, that allows us to evaluate the relative effectiveness of generic and dynamic retargeting. In this field experiment, the consumer randomly was exposed to a generic or a dynamic retargeted ad when they subsequently visited an external website where the firm advertised.[6]

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. The generic ad focused on an image that evoked vacations alongside the brand logo. As the firm focused on selling beach vacations, this generic image evoked a beach-type holiday. The dynamic retargeted ad displayed one hotel the consumer had browsed on the focal firm's website, alongside three

[5] Advertising networks aggregate advertising space across publishers of web content and sell this space to advertisers. They significantly increase the efficiency in the market of selling ad content, as an advertiser does not have to manage multiple relationships with often very small web publishers.

[6] 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.

Table 2: Summary of different online advertising methods.

Label	**Type of Targeting**	**Ad Image**	**Part of Field Test**
Contextual Targeting	Firm advertises on travel websites[a]	Generic brand-awareness-building ad displaying brand and evocative vacation image.	No
Behavioral Targeting	Firm advertises to consumers who had previously visited a travel website[b]	Generic brand-awareness building ad displaying brand and evocative vacation image.	No
Generic Retargeting	Firm advertises to consumers who had previously visited the *firm's* website	Generic brand-awareness building 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[c]	Yes

[a] This is similar to search advertising where the ad displayed may depend on the keyword used (Goldfarb and Tucker, 2011b). However, our data is limited to contextual banner advertising.

[b] Retargeting is, strictly speaking, a form of behavioral targeting, since it targets ads based on the previously observed behavior of consumers. However, because of its high specificity and different underlying technology it is usually referred to as retargeting, or sometimes as remarketing.

[c] Figure 2(a) shows an example of a dynamically 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.

others that were similar in terms of location and star rating. We do not have information on which hotel was displayed.[7] Due to confidentiality agreements, we are unable to reveal the exact ads the firm showed. Instead, in Figure 2(b), we include an approximation of the design of the travel ads the firm used, though the real ad displayed online were more expertly and attractively designed.

Dynamic retargeted ads use standardized designs where a predefined space is subdivided into multiple areas for images of specific products (see also the right side of Figure 2(a)). 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. This standardized design means that as well as being personalized, dynamic retargeted ads are also more complex in design than most banner ads. Therefore, since the dynamic retargeted ad differed from the generic retargeted ad in many dimensions, the correct way to interpret the results of the field experiment is as a comparison of dynamic retargeting as commonly practiced relative to generic retargeting as commonly practiced.

Figure 1: Dynamic Retargeting Examples

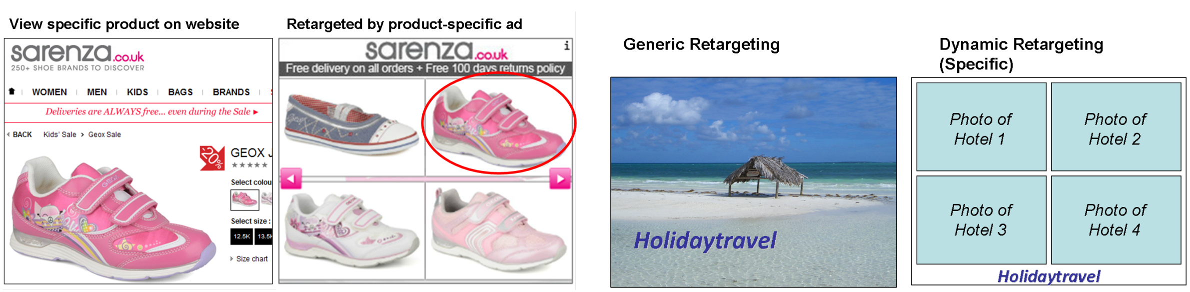

(a) Dynamic retargeting in apparel category (b) Mock-up of travel ads used in field experiment

In our consumer data, we observe each time a consumer was exposed to any type of ad, including the generic and dynamic retargeted ads served during the field experiment as well as any contextual or behavioral targeted ads during the 21 days of the field experiment.

[7] 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.

Importantly, this allows us to 'follow' consumers through any type of website where the firm advertised, across all advertising networks it cooperated with. For each ad exposure we see the time stamp and the name of the site or the advertising network that displayed the ad. We know whether a consumer previously looked at a specific product on the firm's own website because all these consumers were served retargeted ads. Our data also tracks purchases on the firm's own website. One strength of our data is therefore the ability to combine insights on consumers' interest within the firm's website with information on consumers' behavior externally. This data reflects the level of detail that advertising networks are willing to reveal to their clients. It contrasts with clickstream data that, while including greater detail on a consumer's activity, is usually limited to consumer behavior within the firm's own website.

Table 3(a) summarizes consumer-level data for the 77,937 consumers 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 advertising network that implemented retargeting.

Purchase reflects whether that person made a purchase online within the time-frame of the study. In our data, 10% of consumers made a purchase online. Purchase or conversion was measured by whether a person with the same anonymous cookie profile booked or purchased a travel product through their website on a particular day within the time period of the field experiment. We do not know the type 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 customers' purchases after the end of the campaign. We also do not know when a consumer initially browsed the focal firm's website or what specific product they viewed there. It is further possible that there were offline or telephone sales that we do not measure since the firm had no way of linking such offline activity with online advertising activity. However, since a substantial proportion of the firm's travel bookings are now made online, we are confident that we capture a large proportion of sales that relates to the firm's online activities. Last, it is possible that a consumer ultimately bought using a different computer than the one used when first visiting the travel firm's website and so was tracked by a separate cookie, but, as is also the case in prior research, we do not have data to investigate how this influences our results (Rutz and Bucklin, 2011).

$VisitedReviewSite$ indicates that 40% of users visited a travel review site. 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. None of the ads served on any of the travel review sites or travel content sites were retargeted.

Table 3(b) describes the data at a daily level over the 21 days, including the types and number of ads captured by each cookie that consumers were 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. $RetargetedAd \times SpecificAdContent$ reflects that roughly half of these ads were dynamic retargeted ads. $AnyAd$ 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 they saw a contextual targeted ad, and similarly, $OtherBehavioralAd$ captures that on 12.2% of days they saw a behavioral targeted ad. Similarly, we summarize the cumulative number of ads in each category that a consumer viewed prior to that particular date across the 21 days of the field experiment.

We check the validity of the randomization between generic and dynamic retargeted ads. There was no statistically significant relationship between whether an individual was shown a generic or a dynamic retargeted ad (p=0.56) on successive days. Also, individuals who had viewed a specific type of ad content on a 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 provides further support that generic or dynamic retargeted ads were shown randomly.

If on any day the consumer visited multiple websites that were part of the advertising 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 specific retargeted ads. This means that the same individual can be in different treatment groups in different days. This is one of our motivations for including a stock of previous ads the individual is exposed to in our regression analysis.

For comparison, 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. 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. It is noticeable that 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 means that 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.

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. There are three immediate insights. First, it appears that browsing behavior is heavily linked to conversions. This is similar to the activity bias reported by Lewis et al. (2011). People who were not browsing other websites within any of the advertising networks were unlikely to purchase. However, their lack of exposure to ads could simply reflect that they were not online that

Table 3: Consumers Eligible for Dynamic Retargeting

(a) Cross-Sectional Descriptives

	Mean	Std Dev	Min	Max	Observations
Purchase	0.100	0.300	0	1	77937
Visited Review Site	0.402	0.490	0	1	77937

(b) Time-varying Covariates

	Mean	Std Dev	Min	Max	Observations
Retargeted Ad	0.089	0.284	0	1	1502514
Retargeted Ad × Specific Ad Content	0.047	0.211	0	1	1502514
Any Ad	0.214	0.410	0	1	1502514
Other Behavioral Ad	0.122	0.328	0	1	1502514
Contextual Ad	0.042	0.202	0	1	1502514
Cumulative Retargeted Ads	8.021	13.300	0	151	1502514
Cumulative Retargeted Specific Ads	6.772	11.581	0	151	1502514
Cumulative Other Behavioral Ads	19.082	39.267	0	881	1502514
Cumulative Contextual Ads	9.485	25.948	0	1313	1502514

Table 4: All Consumers

(a) Cross-Sectional Descriptives

	Mean	Std Dev	Min	Max	Observations
Purchase	0.020	0.139	0.00	1	2818661
Eligible for Retargeting	0.069	0.253	0.00	1	2818661
Visited Review Site	0.091	0.288	0.00	1	2818661

(b) Time-varying Covariates

	Mean	Std Dev	Min	Max	Observations
Retargeted Ad	0.002	0.049	0	1	59128153
Retargeted Ad × Specific Ad Content	0.001	0.036	0	1	59128153
Any Ad	0.028	0.166	0	1	59128153
Other Behavioral Ad	0.023	0.150	0	1	59128153
Contextual Ad	0.005	0.068	0	1	59128153
Cumulative Retargeted Ads	0.343	2.763	0	248	59128153
Cumulative Retargeted Specific Ads	0.309	2.455	0	248	59128153
Cumulative Other Behavioral Ads	3.952	23.586	0	5504	59128153
Cumulative Contextual Ads	1.297	7.519	0	1313	59128153

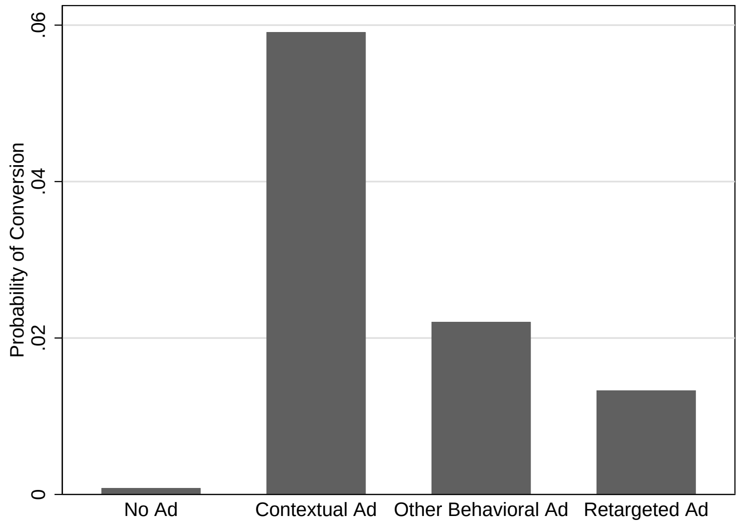

Figure 2: Conversion Rate with Same-day Ad Exposure

day and consequently were not making online purchases. Second, of the different types of ad-exposures, it appears that retargeted ads were the least likely to be correlated to purchase on that particular day. This is striking, as it goes strongly against industry wisdom that has made claims about the high effectiveness of retargeted ad campaigns. For example, Hunter et al. (2010) argued that retargeting increased website visits by 726%, almost double the measured effectiveness of other digital targeting techniques. One explanation is that these industry studies fail to account for sample selection. The baseline tendency to purchase appears far higher in our data in Table 3(a) for people who were eligible to be retargeted because they had visited the website, compared with the people in Table 4(a) who were not eligible because they did not visit the website. Therefore the measured gains to many retargeting campaigns may be because these are people who are already more likely to purchase as they have already sought the product out. Claims about the attractiveness of retargeting may be skewed by self-selection.

The last insight from Figure 2 is 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, it would appear that contextual ads are extremely successful and that retargeted 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. By contrast, the retargeted ads were more likely to be shown on websites that have content unrelated to travel. It is that type of endogeneity which leads us to focus in our analysis on the field test. The fact that otherwise identical consumers who are visiting identical websites are randomly shown different ads allows us to ascribe any differences between the two conditions to the different types of ads.

4 Results

4.1 Information Specificity of Ad Content

We first explore whether generic retargeted ads 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 either a generic ad or a specific ad that day. This initial evidence suggests that a generic ad is more likely to induce consumers to purchase than a specific ad. In limiting an ad's effect to the day it is shown, we follow current industry marketing practice in terms of how online advertising networks award commissions to their affiliates (Weiman, 2010). We also follow 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. We later check that our results are robust to allowing ad exposure to affect purchase within more aggregated intervals, such as a two-day and a four-day window.

There are obviously important factors that this simple analysis in Figure 3 does not control for. For example, the propensity to purchase may vary with how much time has elapsed since the consumer had initially viewed the product on the focal firm's website.

Likewise, this analysis does not control for the effect of other covariates, such as whether a consumer had been exposed to contextual or behavioral targeted ads, or the cumulative effect of any of the four types of ads employed by the firm. We then check whether our results hold when adding further controls. To flexibly control for such factors, we turn to a hazard or survival-time framework. This allows us to identify whether exposure to highly-specific ad content actually increased the likelihood to purchase on the day the customer was exposed to the ad, relative to the control condition, controlling for covariates and the time elapsed since initially visiting the firm's website.

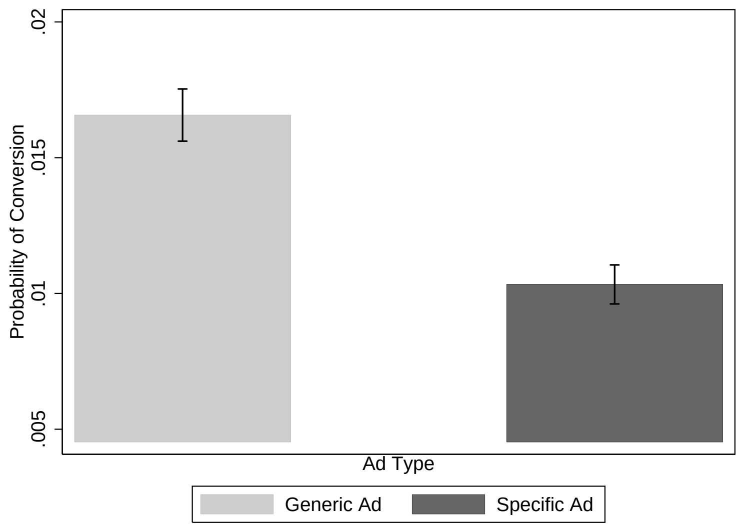

Figure 3: Comparison of Conversion for Generic vs Specific Ad Exposure

Our primary model is a proportional hazards model (Cox, 1972; Jain and Vilcassim, 1991; Seetharaman and Chintagunta, 2003). Such a model had previously been used to study online advertising by Manchanda et al. (2006). Hazard models allow for censoring to account for the fact that not all events, in our case purchases, are observed. Though originally designed to model events that will at some point occur for every individual in the population, they are used to model many other events that, for a subset of the population, may never happen. This includes time to first marriage, time to first child, or time to exit from unemployment.

In hazard models, the key dependent variable is T, a random variable that represents the time to purchase. The empirical model estimates the hazard function of T that captures the instantaneous probability of purchase given that no purchase has been made up to time t. The model has two components: The baseline hazard, $h_0(t)$, and the vector of covariates, (X_{it}). The baseline hazard captures the effect of the time elapsed since we first observe an individual being exposed to an ad in our data. Ideally, we would like to capture the effect of the time elapsed since a consumer first contemplated purchasing the product. However, we do not observe this date in our data. 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. 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 individual i $h_i(t, X_t)$ is therefore:

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

We specify the vector of covariates for person i as

$$\begin{aligned} exp(X_{it}\beta) = exp(&\beta_1 RetargetedAd_{it} \times SpecificAdContent + \beta_2 RetargetedAd_{it} \\ &+\beta_3 OtherBehavioralAd_{it} + \beta_4 ContextualAd_{it} + \beta_5 CumRetargetedSpecificAds_{it} \\ &+\beta_6 CumRetargetedAds_{it} + \beta_7 CumOtherBehavioralAds_{it} + \beta_8 CumContextualAds_{it}) \end{aligned} \tag{2}$$

β_1 measures the effect of the person 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. β_2 measures the effect of the baseline control condition where the consumer was shown a generic retargeted ad. β_3 controls for whether the person had seen another form of behavioral targeted ad and β_4 measures response to a contextual targeted ad. β_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 any effects from the 'stock' of advertising a consumer has seen before. Similarly, β_6 measures response to the cumulative number of generic retargeted ads. β_7 and β_8 measure response to the cumulative number of behavioral and cumulative number of contextual ads.

Column (1) of Table 5 reports a simple model which reflects the findings of Figure 3 in a survival-time framework. It confirms that increased specificity in advertising is, on average, less effective. Column (2) add the full set of controls suggested by equation (2). Again it indicates that on average non-specific ads work better than specific ads. These additional controls proxy not only for different types of targeting but also for whether or not someone is seeking travel-category content that day. Therefore, similarly to Figure 2, a possible interpretation of the smaller coefficient for retargeted ads relative to coefficients for behavioral and contextual ads is simply that people who are seeking travel-category content are more likely to purchase a travel product. The cumulative ad controls measure the effect of the stock of previous online ads that the person has been exposed to. They suggest a possibly lower marginal effect of seeing an additional generic or dynamic retargeted ad when consumers have already viewed many other generic or dynamic retargeted ads. However, the estimates do not have a clear causal interpretation. The negative effects could also result from heavy browsers that are, for exogenous reasons, unlikely to buy in the category. Column (3) presents the estimates from column (2) as hazard ratios to allow an interpretation of the magnitude of the effects. The effects appear economically significant. Exposure to regular generic retargeted ad doubles the probability to purchase for that day, but adding personalized content to this ad reduces the purchase probability by 67%.

Table 6 presents robustness checks for our main specification. In Column (1), we confirm

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

	Coefficients		Hazard Ratio
	(1)	(2)	(3)
	Survival Time	Survival Time	Survival Time
Retargeted Ad × Specific Ad Content	-0.575**	-1.111***	0.329***
	(0.252)	(0.340)	(0.112)
Retargeted Ad	0.984***	0.695***	2.004***
	(0.184)	(0.250)	(0.501)
Other Behavioral Ad		1.821***	6.178***
		(0.161)	(0.998)
Contextual Ad		2.560***	12.942***
		(0.176)	(2.273)
Cumulative Retargeted Specific Ads		0.046***	1.047***
		(0.018)	(0.018)
Cumulative Retargeted Ads		-0.056***	0.945***
		(0.016)	(0.015)
Cumulative Other Behavioral Ads		0.001	1.001
		(0.001)	(0.001)
Cumulative Contextual Ads		-0.005**	0.995**
		(0.002)	(0.002)
Observations	1502514	1502514	1502514
Log-Likelihood	-78158.7	-70059.8	-70059.8

Proportional hazard regression coefficients shown in columns (1)-(2). Column (3) reports hazard ratios for identical specification to that in column (2). Robust standard errors.
* $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

Table 6: Robustness Checks: Dynamic Retargeting for those Eligible for the Retargeting Campaign

	(1) Discrete Time	(2) Weibull	(3) Exponential	(4) Excluding Multiple Impressions	(5) No. Impressions
Retargeted Ad × Specific Ad Content	-0.011***	-1.073***	-1.112***	-1.118***	
	(0.001)	(0.313)	(0.322)	(0.417)	
Retargeted Ad	0.009***	0.750***	0.951***	0.815***	
	(0.001)	(0.243)	(0.250)	(0.304)	
Total Specific Retargeted Ads					-0.251**
					(0.118)
Total Retargeted Ads					0.015
					(0.075)
Other Behavioral Ad	0.018***	1.779***	1.781***	1.959***	1.853***
	(0.000)	(0.160)	(0.159)	(0.170)	(0.162)
Contextual Ad	0.050***	2.529***	2.575***	2.664***	2.586***
	(0.001)	(0.179)	(0.176)	(0.185)	(0.177)
Cumulative Retargeted Specific Ads	0.000***	0.038**	0.037**	0.024	0.021
	(0.000)	(0.018)	(0.018)	(0.018)	(0.016)
Cumulative Retargeted Ads	-0.000***	-0.049***	-0.052***	-0.031*	-0.031**
	(0.000)	(0.016)	(0.017)	(0.017)	(0.015)
Cumulative Other Behavioral Ads	0.000	0.001	0.001	0.001	0.001
	(0.000)	(0.001)	(0.001)	(0.001)	(0.001)
Cumulative Contextual Ads	-0.000***	-0.005**	-0.005**	-0.005**	-0.005**
	(0.000)	(0.002)	(0.002)	(0.002)	(0.002)
Constant	0.001**	-5.113***	-6.074***		
	(0.000)	(0.242)	(0.075)		
ln_p		-0.221***			
		(0.062)			
Day Controls	Yes	No	No	No	No
Observations	1502514	1502514	1502514	1419428	1502514
Log-Likelihood	1819099.5	-12434.7	-12602.9	-62344.9	-70148.6

Dependent variable is time to purchase in all columns except columns (1), where it is whether or not a purchase was made that day. OLS regression coefficients reported in column (1). Hazard-model coefficients presented in other columns. The parameter p in the Weibull model indicates whether the baseline hazard is flat (p=1), monotonically increasing (p>1), or monotonically decreasing (p<1). Robust standard errors.

* $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

Table 7: Shorter Time Window Specifications: Dynamic Retargeting for those Eligible for the Retargeting Campaign

	(1)	(2)	(3)	(4)
	2-Day: PH	2-Day: DT	4-Day: PH	4-Day: DT
Retargeted Ad × Specific Ad Content	-0.737***	-0.014***	-0.552***	-0.020***
	(0.246)	(0.001)	(0.213)	(0.002)
Retargeted Ad	1.059***	0.019***	1.057***	0.033***
	(0.206)	(0.001)	(0.182)	(0.001)
Other Behavioral Ad	1.527***	0.023***	1.238***	0.027***
	(0.137)	(0.000)	(0.135)	(0.001)
Contextual Ad	2.207***	0.063***	1.807***	0.074***
	(0.150)	(0.001)	(0.140)	(0.001)
Cumulative Retargeted Specific Ads	0.042**	0.000***	0.034*	0.001***
	(0.019)	(0.000)	(0.019)	(0.000)
Cumulative Retargeted Ads	-0.060***	-0.001***	-0.053***	-0.001***
	(0.017)	(0.000)	(0.017)	(0.000)
Cumulative Other Behavioral Ads	0.000	-0.000*	0.001	0.000
	(0.001)	(0.000)	(0.001)	(0.000)
Cumulative Contextual Ads	-0.005*	-0.000***	-0.003	-0.000***
	(0.003)	(0.000)	(0.003)	(0.000)
Constant		0.001**		-0.001
		(0.000)		(0.000)
Day Controls	No	Yes	No	Yes
Observations	812503	812503	443216	443216
Log-Likelihood	-88303.6	715426.5	-90070.0	260044.9

Dependent variable is time to purchase in columns (1) and (3) and whether or not a purchase was made that day in columns (2) and (4). Proportional hazard regression coefficients shown in columns (1) and (3). OLS regression coefficients reported in column (2) and (4). Robust standard errors.

* $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

that our results hold in the corresponding discrete-time hazard model, which uses a linear-probability specification with controls for each day in the search process to evaluate how advertising affects the likelihood of conversion (Allison, 1982). In such specifications, the dependent variable is simply whether a purchase was made that day. The results are similar. We use a linear probability model, instead of a probit specification, so that the interpretation of coefficients of interactions in a non-linear framework is not problematic (Ai and Norton, 2003). Results are similar in a logit or probit.

We then check robustness to different specifications for the baseline hazard and timing assumptions. Columns (2) and (3) confirm that our results are robust to using a more parametric formulation of the baseline hazard, such as a Weibull or Exponential specification.

The final two columns of Table 6 check the robustness of our results to multiple impressions. 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.

The specification in Tables 5 and 6 assume that the incremental effect of online advertising is limited to the day that consumers are exposed to it. However, it is possible that ads have carry-over effects to the next day. Columns (1) and (2) of Table 7 investigates this possibility by using a two-day rather than a one-day window as the basic unit of time.[8] The results are similar, if smaller, for both the proportional hazard and discrete time specifications. Columns (3) and (4) use a four-day window.[9] Again, as might be expected given the evidence

[8]Since there are 21 days the final day is treated as a separate window. However, our results are not sensitive to omitting it or including this extra day in a final 3-day window.

[9]Similarly, though this treats the final day as a separate observation, our results are robust to omitting it or collapsing it into the final window.

about limited online ad-recall (Goldfarb and Tucker, 2011a), the results are smaller and less precisely estimated, though they are directionally consistent. In general, 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.

One interpretation of these results is that the complexity of retargeted advertising's design is uniformly unappealing to consumers. Alternatively, it is possible that the hotel the consumer viewed and did not purchase was a hotel they disliked. If either of these interpretations explain our results, as opposed to being explained by dynamically retargeted ads not matching a consumers' need for information, the generic ad should always dominate the specific ad. We now explore whether this is indeed universally the case.

4.2 Timing Specific Ad Content to Match Preferences

A unique feature of our data is that because of the scope of the travel firm's relationships with many different ad-networks the firm had substantial information about when and what kinds of websites each consumer was browsing. Importantly, they observed when a consumer visited a product review website, such as TripAdvisor. These review websites provide large numbers of detailed 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. As discussed in Section 1, a consumer's visit to a review site indicates that they are moving from having just a broad category interest to establishing more defined product preferences.

A visit to such a review site is distinct from a visit to a site that sells commercial products (such as the own website of the firm we study). Consumers may visit a firm's product website to get general ideas of feasibility and availability, but review sites are by their very nature focused on the details of the product itself. We recognize that consumers may have means of obtaining detailed product information that we do not observe in our data. However, by introducing measurement error, this would bias our results downwards.

Figure 4: Comparison of Conversion for Generic vs Specific Ad Exposure

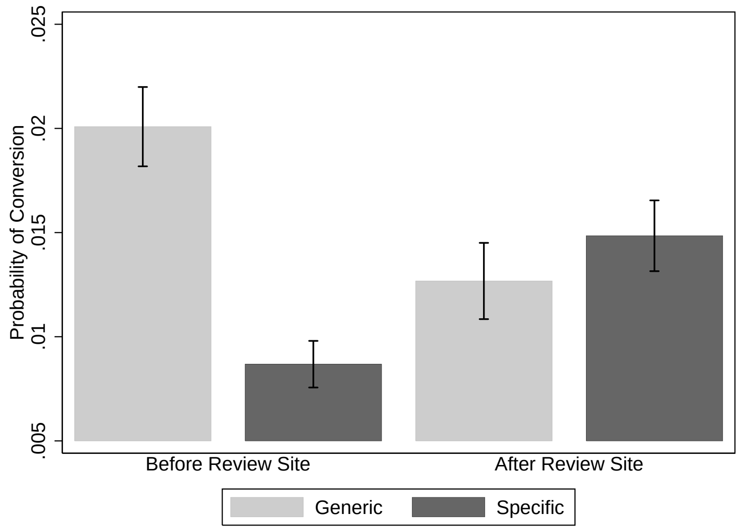

Sample restricted to those who visited a review website.

Of those customers who in our data purchased and visited a review site, 46% purchased prior to visiting the review site, and 54% visited after they purchased.[10] There may be several reasons why a consumer might visit a review site after they have completed a purchase, including learning more about the destination they chose.[11]

Figure 4 provides some 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 an independent product review site. Figure 4 uses data only on consumers who at some point in our data visit a review site. We exclude consumers we do not observe visiting a review site, as they may be different in unobserved ways from

[10]We conducted an online survey on the Mechanical Turk to better understand consumers' use of review sites. Of 55 individuals who had used a travel review site before booking a trip, 73% indicated they typically looked at ratings or comments on hotels or vacation rentals and 92% indicated they viewed either ratings or comments on hotels or they were looking for other information on hotels or vacation rentals, such as availability or location. 82% indicated that typically as a result of a visit to a travel review site they knew more about specific hotels or vacation rentals, including how they liked it, prices or availability.

[11]In a survey on Mechanical Turk of 33 individuals who visited a travel review site after purchasing a travel product, 84% indicate that they did this to get information on restaurants or things to do in the area while 33% indicate they are looking for other information on the hotel or vacation rental they booked.

consumers who ultimately do.[12] Of course the exact motivation for visiting a review site may be different for different consumers. However, this approach allows us to focus on consumers who are similar in their knowledge about the availability of review sites.

Figure 4 illustrates that, after a consumer has visited an independent product review website, the comparative advantages of the different types of advertising change. That is, after consumers visited a review site to seek out detailed product information, highly-specific advertising is relatively more effective than before they seek out a review site. On the other hand, generic brand ads become less effective after viewing a review site. The difference in performance of the two techniques after visiting the review site is not precisely estimated in this raw analysis which is one reason for us to turn to econometric analysis that can control for more factors.

As in the previous section, we use a survival time model. We interact the components of equation (2) with a binary indicator variable for whether or not the person had visited a product review website. Table 8 reports the results. Column (1) displays results for a proportional hazard model for all consumers that were eligible for the field experiment. $RetargetedAd_{it} \times SpecificAdContent$ is negative, that is, the dynamic ad performs worse on average than the generic ad, as found previously. However, $RetargetedAd_{it} \times SpecificAdContent \times AfterReviewSite$ is positive and economically significant. The results suggest that after someone has visited a review site, dynamic retargeted ads are more effective than generic ads. Column (2) restricts the analysis to only consumers whom we observe in our data visiting a review site. The results are similar.

Column (3) shows that the result is robust to a discrete-time specification. Column (4) shows robustness to a Weibull distribution of the baseline hazard. Column (5) shows robustness to an exponential distribution of the baseline hazard. The final two columns

[12]The conversion probabilities of consumers who do not visit a review site from a generic or a specific ad are not significantly different from those of customers in Figure 4 before they visit a review site.

Table 8: Survival Time: Interactions with having visited review site

	All Users: PH (1)	Only Review Site Users: PH (2)	Discrete-time (3)	Weibull (4)	Exponential (5)	2-Day: PH (6)	4-Day: PH (7)
Retargeted Ad × Specific Ad Content × After Review Site	0.783***	0.881***	0.006***	1.030***	1.059***	0.850***	0.948***
	(0.137)	(0.161)	(0.001)	(0.119)	(0.120)	(0.115)	(0.117)
Retargeted Ad × Specific Ad Content	-1.331***	-1.718***	-0.010***	-1.191***	-1.244***	-0.864***	-0.613***
	(0.077)	(0.117)	(0.001)	(0.069)	(0.069)	(0.055)	(0.044)
Retargeted Ad	0.880***	1.036***	0.008***	0.891***	1.036***	1.142***	1.078***
	(0.051)	(0.073)	(0.001)	(0.043)	(0.041)	(0.040)	(0.034)
After Review Site	0.212***	0.408***	0.005***	0.517***	0.629***	0.251***	0.328***
	(0.035)	(0.042)	(0.000)	(0.040)	(0.039)	(0.039)	(0.042)
Retargeted Ad × After Review Site	-0.733***	-0.340***	-0.007***	-1.174***	-1.191***	-0.861***	-1.028***
	(0.106)	(0.118)	(0.001)	(0.088)	(0.088)	(0.098)	(0.106)
Cumulative Retargeted Ads	-0.052***	-0.071***				-0.054***	-0.048***
	(0.004)	(0.007)				(0.004)	(0.004)
Cumulative Retargeted Ads × After Review Site	-0.022**	-0.001				-0.003	0.014*
	(0.009)	(0.010)				(0.009)	(0.009)
Cumulative Specific Ads	0.043***	0.071***	0.000***	0.006***	0.006***	0.039***	0.031***
	(0.005)	(0.008)	(0.000)	(0.002)	(0.002)	(0.005)	(0.005)
Cumulative Retargeted Specific Ads × After Review Site	0.014	-0.019*	-0.000	0.010***	0.010***	-0.008	-0.025***
	(0.010)	(0.011)	(0.000)	(0.003)	(0.003)	(0.009)	(0.009)
Cumulative Total Ads			-0.000***	-0.017***	-0.019***		
			(0.000)	(0.001)	(0.001)		
Cumulative Total Ads × After Review Site			-0.000***	-0.007***	-0.008***		
			(0.000)	(0.001)	(0.001)		
Constant			0.001***	-5.411***	-6.121***		
			(0.000)	(0.056)	(0.020)		
ln_p				-0.157***			
				(0.012)			
Date Controls	No	No	Yes	No	No	No	No
Further Ad Controls	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Observations	1502514	601475	1502514	1502514	1502514	812503	443216
Log-Likelihood	-69998.3	-32391.6	1819646.7	-12204.1	-12288.8	-88239.2	-90041.7

Dependent variable is time to purchase in all columns except column (3). Dependent variable in column (3) is whether or not a purchase was made that day. Proportional hazard regression coefficients shown in Columns (1), (2), (6) and (7). OLS regression coefficients reported in column (3). Weibull coefficients reported in column (4), Exponential coefficients reported in column (5). 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 their cumulative totals. The parameter p in the Weibull model indicates whether the baseline hazard is flat (p=1), monotonically increasing (p>1), or monotonically decreasing (p<1).

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

explore how robust our results are to different time windows. In Column (6) we explore whether our results hold if, rather than using a one-day time window for an ad to potentially have an effect, we use a two-day time window. Our results are robust to the longer time window. In Column (7), we explore what happens if we allow the effect of an ad to persist over four days. As before, the results hold.

The results of Table 8 confirm that a firm should match the level of how specific the information is with a consumer's broader actions online and specifically whether they are, at that point in time, seeking out specific product information. When consumers appear to not yet have well-defined product preferences and are still broadly evaluating their different options, ads that contain broad information on a product line or a brand are more effective than ads that focus on detailed information about specific products or product attributes. However, when consumers appear to have specific product preferences and are actively seeking detailed product information, then using information about their previous product search to tailor ads can be an effective marketing technique.

Though our data does not allow us to control directly for ad design or for consumer liking or disliking of the hotels highlighted in specific ads, the fact that the effectiveness of ads varies with a consumer's stage in their deliberation process lends support to the interpretation that the effectiveness of highly specific ad content is linked to how well consumers have refined their preferences.

4.3 Timing Specific Ad Content to Match Engagement

Our results in Table 8 suggest that highly-specific advertising is more effective once consumers actively seek product information but not when consumers show no signs of actively seeking product information and so are significantly less likely to have well-defined preferences. Prior research on advertising effectiveness suggests that the argument quality of an advertising appeal has a greater effect under high than under low involvement but does not link involvement to whether consumers have well-defined preferences (Petty et al., 1983). Interestingly, findings on the success of personalized recommendations within a firm's own website suggest that customers who return to the firm's own website and so are clearly engaged, are likely to respond positively to personalized recommendations. It is not clear, however, whether greater engagement would also increase the success of highly specific advertising and, if so, how this would interact with a consumer's level of preference development.

We proxy for whether a consumer is engaged in the travel category on a particular day by whether on that day they are browsing websites that offer travel-specific content. We identify this by whether someone has been exposed to a contextual ad, because such ads were served only on travel content sites. People visited such sites before and after they visited a review site. Likewise, a consumer who has visited a review site may on any subsequent day be either involved or not involved in the category.

Figure 5 clearly illustrates that, on average, category engagement is an important predictor of likelihood of conversion. It is also related to the relative performance of dynamic and generic ads in that when consumers are not engaged the generic ads perform better, but when consumers are engaged these ads do not significantly differ in performance. We then examine whether the effect of preference development on the effectiveness of highly specific ads varies with browsing behavior. In Figure 6 we decompose Figure 4 by whether or not the consumer visited a travel website that day. Here we restrict our analysis to consumers who had both visited a travel website and viewed a review site.

Figure 6 illustrates that consumer category engagement increases the effectiveness of generic advertising for consumers who have not yet visited a review site. Category engagement similarly increases the relative effectiveness of the highly-specific ad, especially for those consumers who have visited a review site. Strikingly, we find that the specific ad is only more effective than the generic ad when a consumer has developed well-defined product preferences *and* is exposed to the ad on a day that they are engaged in the category. A

Figure 5: Comparison of Conversion for Generic vs Specific ad exposure

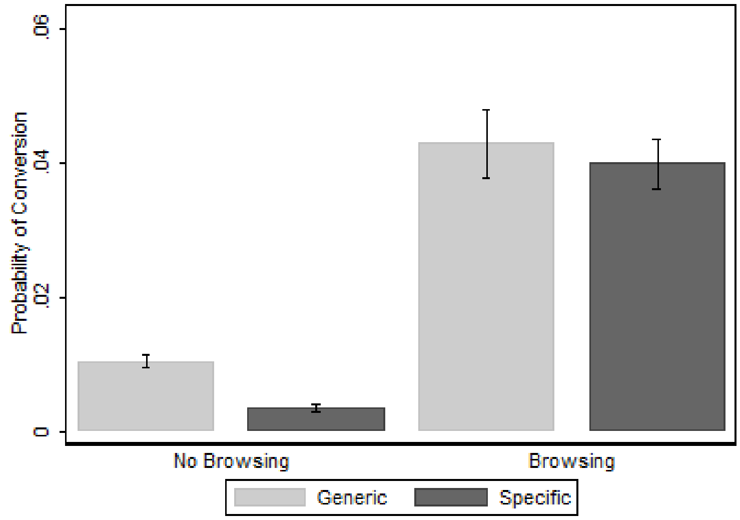

Sample restricted to consumers who at some point browsed a travel website and visited a review website.

consumer who has well-defined preferences but who that day is not engaged in the category would still react more favorably to a generic ad.[13]

To confirm the robustness of these insights, we again turn to a survival-time framework. Table 9 reports the results for the sample of 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 $RetargetedAd \times SpecificAdContent$ suggests that specific ads are less effective than generic ads. However, this is mediated by the positive and significant coefficient for $RetargetedAd \times SpecificAdContent \times BrowsingTravelthatDay$ which suggests that dynamically retargeted ads perform relatively better on days that consumers

[13]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 causality 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.

Figure 6: Comparison of Conversion for Generic vs Specific ad exposure

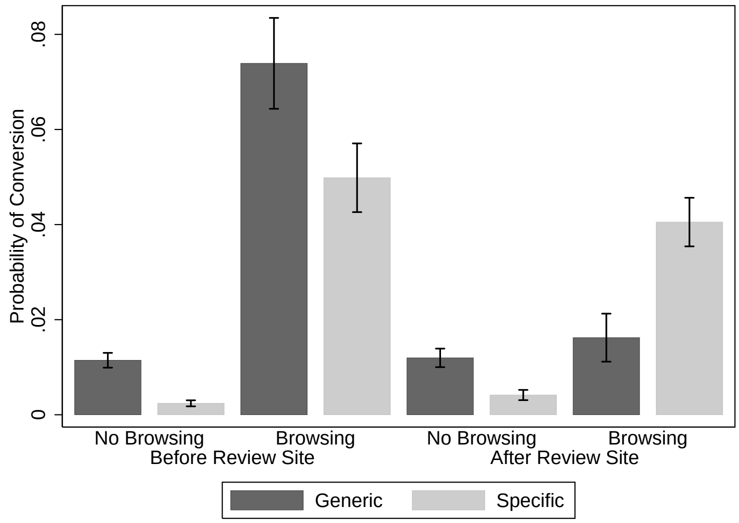

Sample restricted to consumers who at some point browsed a travel website and visited a review website.

browse travel.[14] We echo the analysis of Table 8 and stratify the results by whether the consumer has visited a review site yet or not 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 $RetargetedAd \times SpecificAdContent \times BrowsingTravelthatDay$ after a consumer visits a review site suggests that by contrast specific 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 highly engaged in the category. This is, again, in line with Figure 6.

In sum, our results indicate that the effectiveness of highly specific advertising messages depends on a consumer's engagement and level of preference development at the time of

[14]As before, these results are robust to different definitions of the baseline hazard or a discrete time hazard model.

Table 9: Survival Time: Further interactions with browsing behavior

	(1) All	(2) Before Review Site	(3) After Review Site
Retargeted Ad × Specific Ad Content × Browsing Travel that Day	1.515***	0.838***	2.512***
	(0.175)	(0.267)	(0.279)
Retargeted Ad × Specific Ad Content	-2.114***	-2.411***	-1.803***
	(0.140)	(0.213)	(0.195)
Retargeted Ad × Browsing Travel that Day	-0.392***	-0.306*	-1.496***
	(0.123)	(0.170)	(0.238)
Retargeted Ad	0.585***	0.774***	0.951***
	(0.090)	(0.135)	(0.145)
Browsing Travel that Day	1.356***	1.959***	1.479***
	(0.053)	(0.085)	(0.104)
Cumulative Retargeted Ads	-0.089***	-0.081***	-0.105***
	(0.009)	(0.012)	(0.013)
Cumulative Retargeted Ads × Browsing Travel that Day	0.002	-0.033*	0.037**
	(0.012)	(0.017)	(0.017)
Cumulative Specific Ads	0.070***	0.072***	0.076***
	(0.009)	(0.013)	(0.015)
Cumulative Retargeted Specific Ads × Browsing Travel that Day	0.003	0.047***	-0.024
	(0.013)	(0.018)	(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. Robust standard errors. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.
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) Table 5, but are not reported here for reasons of space.

exposure. Even after consumers have engaged in more detailed product search and so appear to have well-defined preferences, high information specificity is effective only when they are engaged in the category. At all other times, generic advertising is consistently more effective.

5 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 retargeting 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. Industry experts claim that personalized retargeted ads are six times more effective than standard banner ads, and four times more effective than generic retargeted ads (Criteo, 2010). There is, however, little evidence to support whether tailoring advertising content to an individual's observed preferences is effective.

In this paper, we evaluate 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 products 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 what drives the effectiveness of generic versus dynamic retargeting. We suggest that the effectiveness of different forms of retargeting depends on whether consumers have already well-defined product preferences and actively seek detailed product information. We use the visit to a travel review site as a proxy for whether consumers have well-defined product preferences. Our results show that retargeting with individually-tailored ad content is not effective when consumers have not yet visited a review site and are more likely to have poorly-defined preferences. However, when consumers have visited a review site, and so have refined their product preferences, they are more susceptible to higher specificity in ad content. As a result, dynamic retargeting is more effective than generic retargeting.

We then turn to whether on a particular day a consumer is engaged in the product category and ask whether once consumers have refined their preferences, the effectiveness of specific ads holds independently of their level of engagement. We find that targeting consumers with highly specific information is only effective under very limited circumstances: when consumers are engaged and have developed well-defined preferences. Otherwise, generic messages are more effective.

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 content is more effective than highly specific content.

Second, we show that the effectiveness of highly-specific advertising messages changes as customers define their product preferences better and with the level of category engagement. Data on browsing of external websites which is currently available to advertisers, but which is rarely evaluated in detail, can be used to identify a consumer's preference development and engagement and then used to *time* the targeting of ads for maximum effectiveness.

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.

Beales, H. (2011). The value of behavioral targeting. *Mimeo, George Washington University*.

Bettman, J. R., M. F. Luce, and J. W. Payne (1998). Constructive consumer choice processes. *Journal of Consumer Research 25*, 187–217.

Brucks, M. (1985). The effects of product class knowledge on information search behavior. *Journal of Consumer Research 12*(1), 1–16.

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.

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.

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.

Gatarski, R. (2002). Breed better banners: Design automation through on-line interaction. *Journal of Interactive Marketing 16*(1), 2–13.

Goldfarb, A. and C. Tucker (2011a). Online display advertising: Targeting and obtrusiveness. *Marketing Science 30*(3), 389–404.

Goldfarb, A. and C. Tucker (2011b). Search engine advertising: Channel substitution when pricing ads to context. *Management Science 57*(3), 458–470.

Goldfarb, A. and C. E. Tucker (2011c). Privacy regulation and online advertising. *Management Science 57*(1), 57–71.

Griffin, J. G. and S. M. Broniarczyk (2010). The slippery slope: The impact of feature alignability on search and satisfaction. *Journal of Marketing Research 47*(2), 323–334.

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.

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.

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.

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.

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.

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.