Chapter 10: Targeting Display Advertising[1]

Wendy W. Moe, Associate Professor of Marketing
Robert H. Smith School of Business, University of Maryland

Before sponsored search advertising, pop-up ads and email campaigns, there was display advertising. From the very beginning, display ads have been an important source of revenues for online content providers and as such have played a critical role in supporting the free content provided on the Internet. Figure 1 illustrates the growth in online advertising expenditures over the past 8 years. Even with the development of a wide array of new online advertising vehicles, expenditures on displays ads continue to grow and accounted for 24% of the $26.04 billion online advertising industry in 2010.

Figure 1. Online advertising expenditures

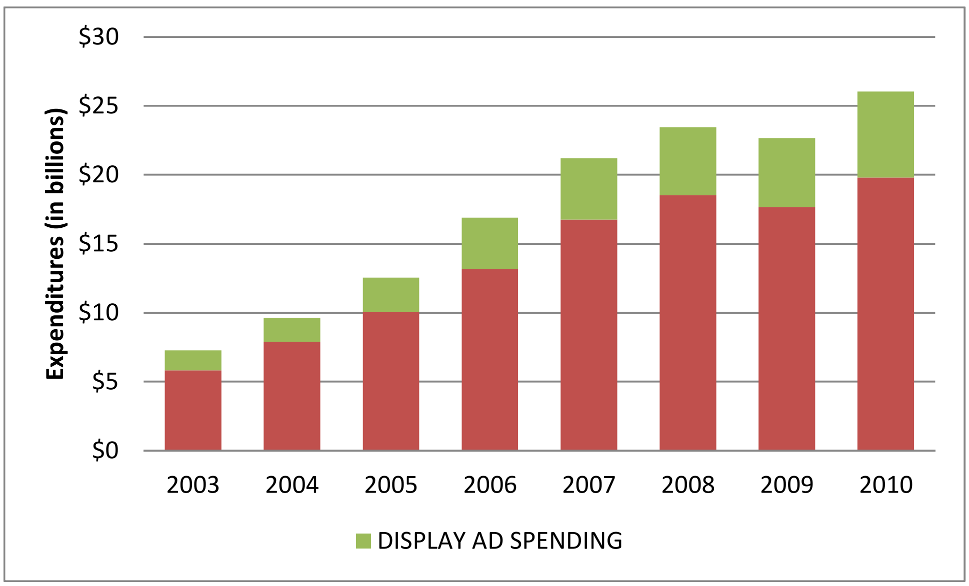

DATA SOURCE: Interactive Advertising Bureau (2003-2010)

The term display ad refers to a family of online advertisements that include text, images and/or animation presented on content sites. Large display ads appearing across the top of a content page are typically referred to as banner ads whereas ads presented along the side of a content page are referred to as skyscraper ads. Regardless of the location of the ad, display ads expose the website visitor to the brand featured in the ad and provide interested individuals the ability to click on the ad for more information.

---

[1] *Moe, Wendy W., "Targeting Display Advertising" in K. Coussement, K.W. De Bock and Scott A. Neslin (Eds.), Advanced Database Marketing: Innovative Methodologies & Applications for Managing Customer Relationships, Gower Publishing, London (United Kingdom) (2013, forthcoming).*

Over the years, as online marketers gather more data and establish better metrics and analytics, advertisers have developed and refined their ability to target display ads to the most promising prospects. This chapter will provide an overview of some of the metrics and analytics used to evaluate display ads and then describe how they are leveraged for targeting policies.

Measuring the effectiveness of online display advertising

While online marketers have more data available, it can still be a challenge to accurately measure advertising effectiveness. The ideal metric would provide some measure of return on investment (ROI) reflecting the added profitability generated by advertising expenditures.

(1)
$$ROI = \frac{Incremental\ profits\ generated\ by\ advertising}{Advertising\ expenditure}$$

While the denominator for any ROI calculation can be relatively straight forward to obtain, the numerator presents a greater challenge. This is true for both online and offline marketers. The advantage that online marketers have is the abundance of performance metrics available for online ads. Offline advertisers must rely on response models that link aggregate sales to advertising activities (which can be difficult) or surveys that gauge consumer awareness of and attitudes toward advertisements. In contrast, online advertisers are able to track a variety of consumer behaviors that can be directly related to a specific online ad.

Table 1 describes common metrics used to evaluate online advertising. Ad impression measures are comparable to offline metrics of audience exposure and reach and represent the number of website visitors exposed to the ad. This metric is determined by the level of visitor traffic at the website on which the ad space is purchased. Unique to the online environment is the ability to measure behavioral responses to the ad exposure. Two common metrics are click-through rates (CTR) and purchase conversion. Click-through rates are computed as the percent of exposures associated with a click on the ad and are a measure of the consumer's immediate behavioral response to the ad. Online advertisers can also link a consumer's click-through to subsequent purchasing behavior at the advertiser's website, an endeavor that has been a challenge for offline advertisers. However, online web retailers can easily compute purchase conversion rates (number of visitors who purchase / total number of visitors) associated with visitors who enter the site through an ad click-through and compare it to their regular visitors. The aforementioned metrics provide marketers with the ability to directly measure an ad's effect on sales and thus allows for more accurate calculations of advertising ROI (if the selection effects resulting from ad targeting are adequately controlled for).

Websites selling ad space also try to capitalize on the availability of these metrics by linking their price to many of the performance metrics mentioned above. Webpages that benefit from high visitor traffic often price based on the number of impressions they deliver. Their proposition is that simply exposing individuals to an ad can be beneficial to the advertiser's brand and long term sales. This is referred to as CPM (or cost per thousand impressions) pricing. Other sites (especially those with very targeted audiences) price according to the number of click-throughs they deliver with the presumption that the visitors to their site are more valuable consumers than the untargeted audience delivered by high traffic web pages. This pricing model is referred to as cost-per-click (or CPC) pricing. Some efforts have also been made to price ads according to the subsequent purchases that may result. However, this can be difficult since it requires that the web retailer track and share data pertaining to purchasing activity at their site, which can raise potential privacy problems thus limiting this practice.

CPM models capitalize on exposure effects while CPC models capitalize on measureable behavioral responses. However, advertising is likely to provide both long term brand-building effects from exposure (Dreze and Hussherr 2003) as well as immediate behavioral effects as measured by click-throughs (Chatterjee, Hoffman and Novak 2003). Thus, increasingly websites are instituting hybrid models where the cost of advertising on their site is function of both the impressions and the clicks delivered.

Table 1. Online Advertising Performance Metrics

	Description	Pricing Model	
Impressions	Number of exposures to ad	CPM (cost per thousand impressiosn)	Hybrid model
Click-Through Rates (CTR)	% of expsoures associated with a click on the ad	CPC (cost per click)	
Purchase Conversion	% of web retailer visitors who purchased	Difficult for ad seller measure	
Return on Investment (ROI)	Incremental profits from ad expenditure		

Performance metrics based on immediate ad effects

The online advertising industry has focused heavily on CTR metrics. While offline ads have been priced almost exclusively according to the level of exposure they provide, pricing for online ads has been based largely on click-through rates. In a 2010 report by the Interactive Advertising Bureau (IAB), 62 % of all online ads were priced using performance metrics only (e.g., CTR) while the remaining considered exposure measures (33% were priced exclusively based on exposure while 5% considered a hybrid pricing model).

Because of this focus on click-through rates, several researchers have directed their attention to the factors that impact click-through performance of display ads. A study by Chatterjee, Hoffman and Novak (2003) examined intrasession versus intersession effects on click-through rates. They found that, within a browsing session, individuals tend to click on the first exposure of any given ad and are far less likely to click on repeat exposures. Across sessions, individuals become more responsive to an ad as time since the last exposure increases. In other words, the researchers argue that ad performance depends on the location of the ad in the individual's navigational clickstream. The study also found high variation in "click proneness" across people, highlighting the potential opportunity to individually target online display advertisements.

Rutz and Bucklin (2009) considered a different performance metric. Rather than examining click-through rates, they investigated how display ads affect browsing paths. Across industries, click-through rates are quite low. In 2009, a study by the National Advertising Initiative showed that the average click-through rate for an untargeted display ad was 2.8% (Beales 2010). However, display ads can affect behaviors other than ad click-through. Rutz and Bucklin (2009) show how consumers exposed to a display ad (even if they don't click on it) will change their online shopping behavior to include searches for the brand featured in the ad.

Long-term performance of online display advertising

Long term effects of advertising can be very difficult to measure, whether online or offline. The theory is that advertising builds awareness and preference for a brand and that this ultimately translates to increased sales. Whiles performance metrics such as CTR measure the immediate effect of display advertising, they fail to capture long term benefits for the brand.

To highlight the drawbacks of relying exclusively on click-through rates, Dreze and Hussherr (2003) use eye tracking technology in a controlled experiment to examine how individuals interact with online display ads. Their findings suggest that low click-through rates are driven in large part by the fact that many individuals avoid looking at display ads altogether (rather than looking but not clicking), a finding supported by Cho and Cheon (2004) who found that consumers avoid looking at display ads because they are perceived as goal impediments. However, Dreze and Hussherr (2003) further find that although website visitors avoid looking at display ads, the presence of the ad is still able to influence traditional brand equity measures such as awareness, recall and recognition. The implication is that click-through rates underestimate the returns on display ads as these ads can benefit a brand even when consumers are not actively engaged with the ad itself.

The long term impact of display ads on the brand is further demonstrated by Manchanda et al (2006). In their study, display ads were shown to influence future repurchasing behavior for a frequently purchased consumer product. Rather than examining click-through rates, Manchanda et al (2006) modeled an individual's inter-purchase time as a hazard process where exposure to display ads can speed up purchase frequency. In other words, they show that display ads can increase customer lifetime value. This benefit of display advertising is ignored by simple performance metrics such as click-through rates, and thus any return on investment (ROI) calculations based on CTR would undervalue display advertising activities.

Advertising Response Models

The alternative to using CTR is to construct an advertising response model that allows for both immediate and short term effects. Decades of offline research has resulted in a family of models that employ a construct that has often been referred to as advertising goodwill (Nerlove and Arrow 1962). The theory is that advertising effects accumulate as consumers are repeatedly exposed to advertising. At the same time, consumers also forget, and the effects of previously seen advertisements decay as time passes. Marketing researchers have operationalized this accumulation and decay of advertising effects by defining advertising goodwill as follows:

(2) $$Goodwill_t = \alpha \ Goodwill_{t-1} + A_t$$

where $Goodwill_t$ represents the inventory of advertising goodwill at time t, α is the decay rate of goodwill from period to period, and A_t is the effect of the advertisement presented at time t. This model is sometimes also referred to as a "leaky bucket" model as the mechanism driving the goodwill measure can be likened to attempts to fill a leaky bucket with water.

The above model has also been extended to allow for advertising wear-out and restoration effects. In a study of TV advertising effects on product sales, Naik, Mantrala and Sawyer (1998) propose a model that allows ads in a campaign to decrease in effectiveness with repetition (wear-out) but regain effectiveness as time passes (restoration). As a result, the effectiveness of any given ad campaign would depend on the ad schedule and the timing of each ad exposure (e.g., *when* should an ad air?), a key component of the media buying decision.

However, advertisers buying online display ads cannot pre-determine the schedule of ad exposures. With offline media, the audience passively absorbs the schedule of ads that is pushed out to them. In contrast, the online consumer interacts with the media and controls what they see as well as the sequence in which they see it. As a result, the schedule of ad exposures is dependent on an individual's sequence of pageviews and therefore unique across individuals. While it is possible for an advertiser to design an online campaign to generally increase or decrease exposures, notable difference will still exist across individuals due to differences in their browsing behaviors.

Braun and Moe (2010) extend the response models developed for offline advertising to explicitly measure the effects of individually varying ad schedules. Their model attributes the effects of an online display ad campaign to (1) the various advertising copy utilized in the campaign, (2) the schedule of ad impressions (e.g., wear-out and restoration effects), and (3) individual differences across consumers. Since the nature of online data provide detailed behavioral data at the individual-level, Braun and Moe (2010) examine how individual visits to the advertiser's website and subsequent conversion at the website are affected by that individual's unique schedule of ad impressions.

Specifically, they employ a zero-inflated Poisson model for the number of visits (v_{it}) an individual makes to the advertiser's website in time *t*, and a zero-inflated Binomial model for that individual's conversion rate at the site (where s_{it} represents the number of successful conversions resulting from the v_{it} visits). In both models, r_v and r_s captures the inflated observations at zero.

(3) $$f(v_{it}) = (1 - r_v)\mathrm{I}(\textstyle\sum_t v_{it} = 0) + r_v \frac{e^{-\mu_{it}} {\mu_{it}}^{v_{it}}}{v_{it}!}$$

(4) $$f(s_{it}) = (1 - r_s)\mathrm{I}(\textstyle\sum_t s_{it} = 0) + r_s \binom{v_{it}}{s_{it}} p_{it}^{s_{it}} (1 - p_{it})^{v_{it} - s_{it}}$$

Both model components include covariate effects that capture the variation in impression schedules across individuals. Methodologically, they employ an advertising goodwill construct with creative-specific, wear-out and restoration effects as a covariate for the individual's Poisson visit rate (λ_i) and Binomial conversion rate (p_i). Bayesian methods are used to allow for heterogeneity across individual baseline visit and conversion rates. The resulting model measures the effects of a display ad campaign on visits to and successful conversions at the advertiser's website as ad exposures vary across individuals, both in terms of the creative content as well as the schedule of ad impressions presented to the target consumer.

(5) $$log\ \mu_{it} = log\ \mu_{0i} + \beta_{\mu i} GOODWILL_{it} + \boldsymbol{\gamma} SEASONALITY_t$$

(6) $$\mathrm{logit}\, p_{it} = \mathrm{logit}\, p_{0i} + \beta_{pi} GOODWILL_{it}$$

where advertising goodwill (GOODWILL) decays over time at rate α and increases with each new ad impression. The magnitude of the increase depends on the creative-specific effect associated with the impression, repetition wear-out effects after *u* repeat exposures of creative *j*, and restoration effects after a period of τ has elapsed since the last exposure of creative *j*.

(7)

$$GOODWILL_{it} = \alpha GOODWILL_{i,t-1} + \sum_{j=1}^{J} \begin{matrix} Effect\ of \\ Creative\ j \end{matrix} \sum_{u=1}^{\substack{Cumulative \\ Ad\ Exposures_{ijt}}} \begin{pmatrix} repetition \\ wearout \end{pmatrix}^{u-1} \begin{pmatrix} restoration \\ effect \end{pmatrix}^{\tau_{ij}}$$

Braun and Moe (2011) also control for potential correlations between how frequently an individual is exposed to an ad and their underlying visitation and purchasing behavior. They do so by explicitly modeling each individual's impression schedule as a separate zero-inflated Poisson process (with impression rate λ) and correlate each individual's rate of impressions with the rate of visits and purchases. This allows for variation across individuals as well as any correlation between ad exposure and visitation and purchasing behavior that may be a result of simply increasing online browsing activity.

(8) $$\log \lambda_i, \log \mu_{0i}, \operatorname{logit} p_{0i}, \beta_{\mu i}, \beta_{pi} \sim MVN(\phi, \Sigma)$$

Their empirical results show substantial ad accumulation, wear-out and restoration effects and have implications for how advertisers should schedule ad impressions for a given individual. For example, in some cases, it may be beneficial to serve an ad with less effective creative content if ads with more effective content are sufficiently "worn-out." A more in-depth discussion of how this study impacts behavioral targeting policies appears later in this chapter.

Cross Channel Effects

Increasingly, consumers are multi-channel shoppers (Inman, Shankar and Ferraro 2004, Neslin and Shankar 2009). Many consumers research a purchase in one channel while transacting the purchase itself in a different channel. However, when evaluating the effectiveness of online ads, researchers tend to consider only the online channel and assume channel/marketing congruency. That is, they assume that shoppers respond to online ads with online purchases (Blattberg et al 2008).

However, Dinner, van Heerde and Neslin (2011) found evidence of multi-channel advertising effects. They decompose sales revenue into online and offline revenues and then distinguish between customer count and customer spend as two separate sources of revenue, allowing them to examine how various forms of advertising affects each.

(9) $$\text{REVENUES}_{mt} = \text{OfflineCustomerCount}_{mt} \text{ x } \text{OfflineCustomerSpend}_{mt} + \text{OnlineCustomerCount}_{mt} * \text{OnlineCustomerSpend}_{mt}$$

While previous research found that advertising and promotional efforts tend to affect customer count but have limited effects on customer spend (Ansari et al 2008), Dinner et al found empirical evidence that advertising affects both customer count and customer spend. But perhaps their most interesting finding is that online advertising can affect offline sales (more so than how much offline advertising affects online sales). This cross-channel effect can be substantial with online advertising affecting offline sales as much as it affects online sales. These results imply that measures of online advertising effectiveness focused exclusively on online sales will underestimate the overall impact of the advertising expenditures for those companies with both an online and offline presence.

Summary of measurement methods for online display advertising

What the above mentioned studies have shown us is that online display advertising can have both immediate effects, which can be captured by click-through rates, as well as longer term effects, which tend to be overlooked by such performance metrics. As a result, a number of studies have set forth to document and measure the long term effects of display advertising on the brand, navigational path, repeat purchasing rate and future visitation and conversion behavior. Table 2 provides a summary of these studies and the effects they found.

Table 2. Summary of online display ad effects

	Effects of Display Ads	Research Study
Immediately observable effects	Click-through rate	Chatterjee et al (2003) Beales (2010)
	Navigational/search behavior	Rutz and Bucklin (2009)
Long-term effects	Repurchase frequency	Manchanda et al (2006)
	Brand equity metrics (awareness, recall, recognition)	Dreze and Hussherr (2003)
	Advertising goodwill Visits to advertiser's website Conversion at advertiser's website	Braun and Moe (2011)
Multi-channel effects	Customer counts (online vs. offline) Customer Spend (online vs. offline)	Dinner, van Heerde and Neslin (2011)

In the process, these studies have also documented notable differences across individual consumers in terms of how they respond to online display advertising. This compels us to think about how we can target online display advertising to improve our returns on online advertising investments.

Targeting Strategies

The goal with target advertising is to selectively expose your advertisements to those individuals who are most likely to respond positively. The challenge is identifying these individuals. A number of different approaches have been used to target display ads (see Table 3).

Table 3. Overview of targeting approaches

Targeting Strategy	Description
Targeting based on user characteristics	Identify individuals for ad targeting based on demographic and/or interest profiles which they have self-reported
Geotargeting	Present display ads to individuals in certain geographic locations based on IP address of computer
Contextual targeting	Identify webpages for ad placement based on content
Behavioral targeting	Identify individuals for ad targeting based on online browsing behavior

Targeting based on user characteristics

Traditional approaches such as demographic and/or interest based targeting have been transported to the online environment. Many online communities (e.g., Facebook) allow advertisers to serve ads to community members who self-report certain demographic characteristic and/or interests in their member profiles. For example, an advertiser interested in reaching college students in the mid-Atlantic region may design a Facebook ad campaign that serves their ads to University of Maryland students when they log in.

Additionally, researchers have also developed methods that use website visitors' clickstream behavior to infer demographic characteristics for the purposes of online ad targeting. In a study by De Bock and Van den Poel (2010), Random Forest classifiers were used to predict the demographic attributes of anonymous website visitors using their Internet clickstream behavior. Their methodology was applied to clickstream data that tracked user behavior across 260 associated web sites. This data was coupled with data from an online survey that solicited demographic information from a random sample of web site visitors, allowing the researchers to both train and validate their method. Their approach provides additional demographic information that can be useful for online targeting, both in terms of the individualized targeting of ads as well as the profiling of web sites for ad placement decisions.

Advertisers can also geotarget their ads by delivering their ad content to only individuals in certain geographic regions based on their computer's IP address (while the precise location of the computer is difficult to ascertain since most home computers are assigned rotating IP addresses by their broadband provider, the geographic region to which the IP belongs can still be informative for targeting purposes). Geotargeting can be effective in limiting advertising efforts to the advertiser's geographic reach, making the advertising spending more efficient. However, it is less effective in matching the ad content to individual interests.

Contextual targeting

Contextual targeting, which relies on matching the content of the ad to the context in which the ad is seen, provides another targeting approach. For example, a brand like Nike may choose to purchase ad space in the sports section of the WashingtonPost.com website while a brand like Nordstrom may purchase ad space in the fashion section. The theory is that an individual viewing the webpage on which the ad appears has an interest in the topic featured on that page and therefore potentially has an interest in the advertiser if the content of the ad is matched to the content of the page.

A number of studies have investigated the impact of contextually targeted ads on behavior. In an early study conducted by Shamdasani, Stanaland and Tan (2001), controlled experiments showed that relevance between website content and the product category featured in the display ad can induce: (1) more favorable attitude toward the ad, (2) more favorable attitude toward the product, (3) higher intention to click on the ad and (4) higher purchase intentions. In other words, advertisers that match the content of their ad to the content of the webpage on which the ad appears will benefit from more successful ad campaigns. This study by Shamdsani, Stanaland and Tan (2001) provided early evidence that contextual targeting improves ad performance.

Other studies found similar effects. Moore, Stammerjohan and Coulter (2005) also found that context congruity between the ad and the website resulted in more favorable attitudes toward the advertised brand. However, they also found that context congruity decreased recall and recognition of the ad since congruent information is easily overlooked (in contrast, incongruent information is more likely to attract the consumer's attention). Thus, an advertiser must trade off the benefits achieved in terms of improved consumer attitudes resulting from contextual targeting with the loss in recall and recognition associated with the practice.

Finally, two other studies examined how context congruity affects the intrusiveness of online ads (Edwards, Li and Less 2002; Goldfarb and Tucker 2010). In a controlled experiment, Edwards, Li and Less (2002) demonstrated how context congruity decreases ad intrusiveness, which can improve advertising performance as consumers tend to avoid more intrusive ads. Edwards, Li and Less (2002) argue that context congruent ads are *more* effective since they are *more* likely to attract the attention of the consumer, contradicting the conclusions of Moore, Stammerjohan and Coulter (2005), and suggesting that context congruent ads are *less* effective as they are *less* likely to attract the attention of the consumer. Thus, further research to explore this relationship is needed. A formal scale to measure ad intrusiveness was developed in a separate study by Li, Edwards and Lee (2002). Like the abovementioned researchers, Goldfarb and Tucker (2010) also examined the effects of context congruity and ad intrusiveness. In a large-scale field experiment, they found that both context congruity and ad intrusiveness can increase purchasing intent. Specifically, when ads are both contextually targeted and perceived to be intrusive, ad performance decreases as the practice raises privacy concerns that the consumer may already have. This effect is strongest among people in their study who choose not to share sensitive information (e.g., income, age, etc.).

The practice of matching an advertisement to website content is analogous to the offline practice of matching ads with the appropriate newspaper section, TV show, or radio station. What is unique about contextually targeting online is the opportunity to identify very specific individual interests. For example, Google serves contextually targeted ads in user's Gmail accounts by scanning the content of emails and serving ads that are contextually congruent. If an individual is viewing an email from a friend in which they are talking about planning a ski vacation, an ad for a local ski resort may appear on the page.

The above mentioned studies and practices employ contextual targeting simply to match the ad content with the website content. In a study by Sherman and Deighton (2001), contextual targeting strategies were extended to help identify websites for ad placement based not just on the content of the websites but based on whether targeted customers were likely to visit these websites. In their analysis, the websites visited by an advertiser's customers were cluster analyzed. The resulting clusters effectively identified website genres that were popular amongst their target customers, providing target websites for ad placement. When banner ads were offered on these targeted websites, conversion rates (purchases per impression) were 10 times higher that when the ads were offered on non-targeted websites.

The study by Sherman and Deighton (2001) began to leverage customer behaviors for the purpose of ad targeting. In the next section, we further discuss behavioral targeting as a strategy for online display advertising.

Behavioral Targeting

An increasingly popular targeting strategy online is behavioral targeting. Behavioral targeting leverages observed behavioral histories to infer an individual's interest in and likelihood of purchasing from a particular product category. Ads are then targeted at individuals whose interests are aligned with the product being advertised. For example, if an individual visits a number of automobile related website, it would be fair to assume that this individual has an interest in cars and may be in the market for a new car, making this individual a prime target for car ads.

A study conducted by the National Advertising Initiative (2010) estimated that spending on behaviorally targeted display ads account for 17.9% of all spending for online display advertising, and it is expected to grow. The popularity and growth of behavioral targeting can be easily explained by the dramatic improvements in performance metrics associated with behaviorally targeted ads when compared to untargeted, or run-of-network, advertisements. In the same study by the NAI, behaviorally targeted ads achieved a click-through rate of 6.8% compared to just 2.8% for untargeted ads. As a result of this drastic difference in CTR, advertising networks are able to sell behaviorally targeted ad space at a price of 2.68 times more than what they charge for untargeted ad space.

Clearly, there are tremendous opportunities to improve advertising ROI by turning to behavioral targeting. Therefore, we will review how the technology driving behavioral targeting works and discuss some of the methodological approaches available for identifying target individuals and customizing advertising for them. We discuss both behavioral targeting within a website and across a network of websites below.

Behavioral targeting within a website

Within site behavioral targeting relies on individual clickstream and purchase history data, both of which are easily obtainable from website server logs. These behavioral histories can reveal consumer interests without the need to ask individuals to self-report their interests. For example, an individual who has viewed or purchased a product from a particular category can be assumed to have an interest in that category.

A common practice is for multi-category web retailers to target category-specific promotional offers to those customers who have previously purchased or shopped for products in that category. For example, Amazon may target a frequent book buyer with a display ad on their site that offers 20% off of her next book purchase. In theory, these targeted customers are more likely to respond positively to promotional messages that are matched to their interests (as indicated by their behavioral histories).

To this end, researchers have developed advanced methodologies that utilize pageview and purchase history data to decipher visitor preferences. For example, Moe (2006) developed a two-staged choice model where both product pageviews and purchase choices at an online retailer were modeled to identify shoppers' search criteria and attribute preferences. In stage one of the model, a product is included in a shopper's search set if its attributes meet or exceed the individual's decision criteria for that attribute. A multinomial logit choice model (with a no-choice option) is employed where the utility of viewing one product out of all the products available in the category is determined by product *j*'s *K* attributes, x_{jk}, and how that attribute would contribute to the variety of attributes in the search set (D). The latter component of the utility function accommodates behaviors where consumer *i* actively seeks out variety along a specific attribute for comparison shopping purposes.

(10) $$U_i(j) = \sum_{k=1}^{K}\left[\beta_k \cdot x_{jk} + \gamma_k \cdot D_{ijk}\right]$$

To operationalize the use of simplified decision criteria in the search stage, not all attributes contribute to the utility function specified in (10). Instead, the inclusion of an attribute is dependent on a latent indicator variable which is estimated. For example, if price is used as a screening criterion in the search stage while color is not, the indicator for price is 1 and the indicator for color is 0. This variable is estimated with heterogeneity across individuals.

The above conceptualization of screening criteria follows established research on how different decision rules are used in various stages of the decision process (Gensch 1987; Bettman, Johnson and Payne

1987; Fader and McAlister 1990; Gilbride and Allenby 2004). The resulting model provided estimates that describe both the search criteria in stage one as well as the underlying preference structure driving both search and purchasing decisions.

In stage two, the shopper is assumed to choose a product from the resulting search set using a more effortful compensatory decision rule. Again, a multinomial logit choice model (with a no-choice option) is estimated. However, unlike the formulation used in the search stage where only some attributes factor into the utility function, all attributes contribute to the purchasing utility. Additionally, since the objective in this stage is to choose the one utility maximizing option, attribute variety (D) is not considered.

(11) $$V(j) = \sum_{k=1}^{K} \beta_k \cdot x_{jk}$$

Moe (2006) applied this model to clickstream data pertaining to two separate product categories at an online retailer. Empirical results indicated that a smaller set of product attributes are considered in stage one (i.e., search stage) of the model than in stage two (i.e., purchasing stage). However, this finding should not be surprising given the extant research on two-staged decision processes. Instead, the model provides a tool for online retailers interested in micro-targeting individuals based on their search criteria and attribute preferences. For example, an individual interested in searching across a wide variety of price points may respond more positively to ads that feature products of varying prices than to ads that feature only products within a single price tier. Additionally, such a model would also allow marketers to better design promotional ads that feature the consumer's preferred product based on estimated attribute preferences in order to encourage a purchase.

While a web retailer has abundant customer data on search and purchasing behavior at their site, this data provides only a partial view of an individual's overall online shopping. Shopper behaviors at competing retail sites or third-party information sites are not available to the web retailer. Furthermore, analysis of within-site data can only inform targeting decisions involving shoppers who have already visited the site. Advertising to prospective visitors requires different data and different methodologies, which we discuss next.

Behavioral targeting across an ad network

Content providers across the Internet sell ad space on their webpages in order to support their ability to provide free content. Ad networks purchase an inventory of ad space from a wide variety of content providers and resell it to advertisers. This allows the ad network to create both an inventory of websites that will host ads as well as an inventory of varied advertisements to serve. The resulting network of ads and ad space is what enables behavioral targeting across the Internet by (1) providing the ability to observe behavior across websites and (2) providing the opportunity to present targeted ads based on the visitor's observed behavior.

One common approach to behavioral targeting is to differentiate between users according to their online navigational clickstream and advertise only to those whose behavioral history indicates an interest in the advertised product (Ha 2004). Figure 2 provides a simple illustration of such a process and highlights the role that ad networks have in enabling this process. Technologically, when an individual visits a website on which an ad network has purchased space, the ad network saves a cookie file on the individual's computer. On this cookie, they record an interest category based on the content available on the page. For example, an individual who visits a technology review site is assumed to be interested in technology. If an ad network has purchased space on this site, they will then record this interest on a cookie stored on the user's computer. Later, when the same computer visits a news site on which the ad network has also purchased space, the ad network will read the cookie and show the user a targeted ad that features a technology brand (rather than an auto brand, for example). In the absence of an ad network, a mass market (untargeted) ad would be served instead.

Figure 2. Illustration of clickstream-based behavioral targeting across an ad network

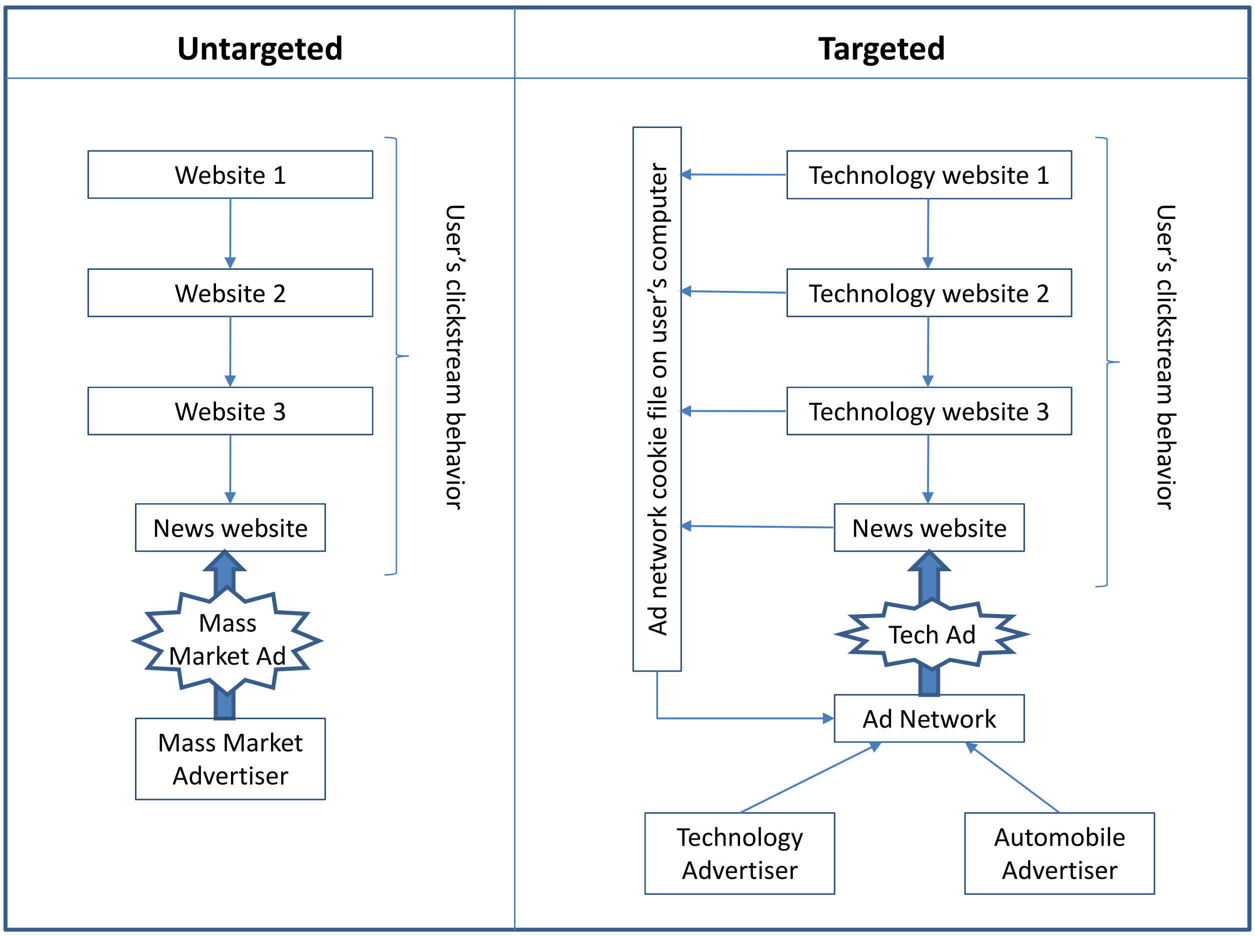

While the above illustration highlights the use of observed clickstream behavior to identify individuals to target, other behaviors can also be tracked and leveraged using similar cookie technologies and analysis methods. For example, Yan et al (2006) examine the value of using browsing/clickstream, search query and click-through histories for targeting purposes. The key question they address is whether or not these behavioral histories can actually improve advertising effectiveness. First, they compared the behavior of users who clicked on an ad to those who did not. They represented each individual with a large scale matrix that indicates either the pages they have viewed or the words they have searched. Using these matrices, they computed classical cosine similarity measures between individuals. On average, those who clicked on an ad were similar to each other but notably different from those who did not click, in terms of their pageview and search behavior. Then, to quantify the potential benefits of behavioral targeted, they cluster analyzed individuals using only their browsing or search behavior and compared CTRs across the resulting segments. They found that an ad's CTR can vary greatly across segments, highlighting the potential benefits of matching ads to appropriately targeted user segments. Further examination of potential segmentation schemes showed that short behavioral histories (1 day) are more effective than long histories (7 days) and search queries are more effective than pageview histories in identifying the variance in CTR across segments. Overall, Yan et al (2006) demonstrate how short click-through histories can be used to target individuals based on what pages they have viewed or what words they have searched to increase future advertising click-through rates (CTR).

In a later study of how behavioral history can inform ad targeting, Chen, Pavlov and Canny (2009) propose a linear Poisson regression model to examine how various types of behaviors can predict ad clicks. Let y=the number of ad clicks, w=a weighting vector and x=behavioral covariates:

$$P(y) = \frac{\lambda^y \exp\{-\lambda\}}{y!} \qquad where\ \ \lambda = w^{\intercal} x \tag{12}$$

In their model, six types of covariates that described an individual's behavioral history were considered: ad clicks, ad views, pageviews, search queries, algorithmic search result clicks, and sponsored search result clicks. Their approach provided a parsimonious and highly efficient method for identifying individuals for targeting purposes. In a set of large-scale experiments using Yahoo's behavioral data, the authors document a 20% lift in CTR using their approach.

As both computing and analytical technologies develop, advertisers are increasingly incorporating more information into their targeting models and integrating multiple targeting approaches. For example, Kazienko and Adamski (2007) propose a system that they named AdROSA. This system utilizes behavioral history (i.e., clickstream and click-through) and web page content to personalize advertisements for specific individuals and to match the advertisements to the web page on which it is offered.

Finally, while most of the extant behavioral targeting research focuses on leveraging behavioral data related to pageviews, ad clicks and search queries, there are opportunities to further refine behavioral targeting algorithms by bringing in ad impression histories (e.g., Braun and Moe 2011). Consider an individual with an interest in buying a new computer. Figure 3 illustrates how repeated exposure to the same ad creative affects advertising goodwill in the presence of wear-out effects. For comparison purposes, Figure 3 also illustrates how exposures to two unique ad creatives affect advertising goodwill. The figure highlights that in the presence of advertising wear-out effects, repeated exposures to the same computer ad will decrease the effectiveness of that ad and may be an inefficient use of advertising resources, even when the alternative ad is of lower (baseline) quality. Consequently, it can be beneficial for an ad network to introduce some variety, either in terms of the creative copy presented or in advertising another product category for which the user has also exhibited an interest, in order to optimize the value of that advertising opportunity.

Figure 3. Illustration of how ad exposures affect advertising goodwill

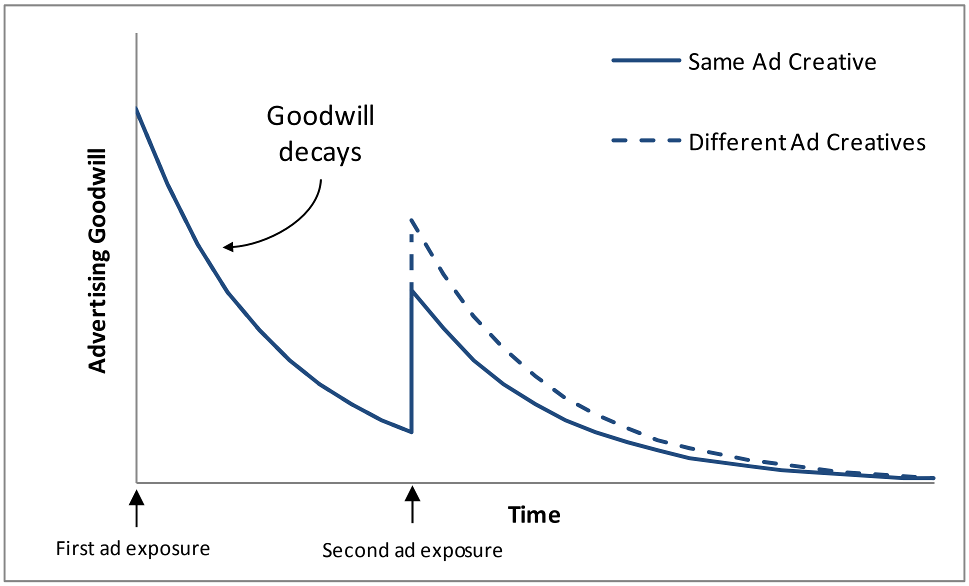

The model proposed by Braun and Moe (2011) provides a measure of how effective an ad will be given an individual's ad impression history. For example, consider a campaign consisting of three ads (A, B and E) where A is the ad with the most effective creative content and B and E have similar but less effective content. Based on a simulation, Tables 4 and 5 show how the next ad (served in week 5) will affect visits and purchase conversions for various individuals (each with a unique ad impression history) depending on the ad's creative content. Notice that while ad A can be considered a more effective ad (sans wear-out and restoration effects), there are instances where the advertiser would be better off serving ads B or E. For example, impression histories 1, 3 and 5 consist exclusively of ad A in the first 4 weeks. As a result, ad A is less effective in week 5 due to ad wear-out and ads B and E would be more effective. Impression histories 2, 4 and 6 lead to a worn-out ad B (in addition to ad A). For individuals with those histories, ad E would be more effective as the week 5 ad.

While Tables 4 and 5 provide stylized examples to illustrate the role of impression histories, the model provided by Braun and Moe (2011) can be applied to any impression history observed in the data. This allows the advertiser to determine which ad to serve next given the individual being targeted. Overall, the research and its findings suggest that advertisers should consider impression histories (and include this information in their cookies to inform future ad impressions), in addition to behavioral histories, to design more effective targeting policies.

Table 4. Expected Visits for Various Ad Impression Histories

	Wk1	Wk2	Wk3	Wk4	Week 5 ad A	Week 5 ad B	Week 5 ad E
Impression History 1	A	A	A	A	.820	**.921**	**.922**
Impression History 2	A	B	A	B	1.049	1.014	**1.117**
Impression History 3	AA		AA		.977	**1.058**	**1.067**
Impression History 4	AB		AB		1.253	1.242	**1.306**
Impression History 5	AAAA				1.390	**1.479**	**1.471**
Impression History 6	ABAB				2.305	2.237	**2.425**

Table 5. Expected Conversion for Various Ad Impression Histories

	Wk1	Wk2	Wk3	Wk4	Week 5 ad A	Week 5 ad B	Week 5 ad E
Impression History 1	A	A	A	A	.308	**.345**	**.341**
Impression History 2	A	B	A	B	.397	.383	**.427**
Impression History 3	AA		AA		.364	**.407**	**.400**
Impression History 4	AB		AB		.477	.467	**.499**
Impression History 5	AAAA				.524	**.567**	**.562**
Impression History 6	ABAB				.898	.838	**.937**

Risks of Targeting Display Ads

The targeting of display ads has greatly improved the ROI of online display advertising. However, this practice is not without its risks. These risks range from those associated with ignoring sentiment to missed cross-selling opportunities to privacy concerns.

In December 2006, Reuters.com ran a story with the headline, “Over 250 sick after eating at Indiana Olive Garden.” On the left side of the page was a contextually targeted ad that read, “Free Dinner for Two at Olive Garden” (Aker 2008). While the targeting algorithm correctly identified the topic featured on the webpage, the sentiment in the content was ignored. Often times the text analysis conducted to categorize webpage content focuses exclusive on identifying the interest category without any consideration of the sentiment expressed in the content. This can lead to inappropriate, distasteful and sometimes comical results that may draw unwanted attention to the advertiser.

Figure 4. Contextually targeted ad for Olive Garden

SOURCE: Aker (2008)

Additionally, efforts expended in matching advertising content with an individual's past browsing behavior tends to overlook cross-selling opportunities. For example, an Amazon shopper who has a long history of purchasing books would likely be targeted with book-related promotions and ads. However, there is an opportunity cost. Using the advertising opportunities to deepen the interest in books does so at the expense of potentially cross-selling another product category and broadening the overall relationship with the retailer.

Finally, many consumer advocates have been very concerned with the privacy implications of online targeting. While many consumers prefer targeted ads which they consider to be more relevant to them individually, some are concerned about personal privacy. This issue is particularly salient for behaviorally targeted advertising practices that rely on tracking individual's web browsing behavior across the Internet. Data that tracks strictly behavioral data is anonymous, that is no personally identifiable information (PII) is recorded. However, it is possible to collect PII or infer sensitive user characteristics from one's behavioral data. Currently, these practices are self-regulated which leaves the onus of protecting consumers' identities and data on the ad networks.

Future Research

The current state of research has documented effects of display ads, identified a number of performance metrics, explored the effects of various forms of targeting and considered privacy concerns. However, there are still several questions that remain unanswered.

First, research in academia and practice has favored click-through rates as a measure of ad performance. Some recent research has begun to examine non-immediate effects of ads on future behavior. However, there is still a need to better understanding how online ads affect ROI and the overall brand. Additionally, consumers are increasingly engaged in multi-channel and mobile commerce. Future research on the effects of advertising across platforms is also needed.

Second, behavioral targeting across websites is a relatively nascent practice. Further research is needed to better understand what works and what doesn't. While computer scientists at ad networks have been diligently data mining and refining adaptive algorithms, marketers have little understanding as to the theoretical mechanisms that are driving the consumer's response to behavioral targeting. This is a critical gap in our knowledge that needs to be filled if we want to improve how we design our advertising campaigns, identify target consumers and structure targeting policies. Critical in this pursuit is a careful study of how to measure advertising ROI, especially in the presence of behavioral targeting. When advertisers pursue behavioral targeting measures, the causal effect of the ad impression must be carefully separated from the elevated CTR and purchasing conversion rate associated with the targeted consumer, who should theoretically be more likely than the average consumer to click-through and purchase.

In assessing the value of behavioral targeting, the practice should be compared to simpler targeting methods based on user characteristics. While some ad networks employ algorithms that consider both behavior and inferred user characteristics, the value of each in terms of how it might improve advertising ROI is unclear. Additionally, a significant risk is associated with attempts to identify user characteristics. Consumers typically view their behavior online as anonymous. Efforts to associate an identity to these consumers may raise privacy concerns (Wall Street Journal 2010).

Overall, as behavioral targeting becomes increasingly popular and the profiling of website visitors becomes more accurate, privacy issues begin to concern consumer advocates. Many have advocated government regulation of the industry. Researchers need to begin considering how possible regulation would affect online advertising. In the absence of regulation, advertisers need to be mindful of the ethical responsibility of securing consumer data, masking individual consumer identities, and protecting consumer privacy. Therefore, we encourage additional research into how advertising practices affects these responsibilities.

REFERENCES

Aker, J. Brooke (2008), “How to Avoid Contextual Tragedies,” *iMedia Connection*, December 9, http://www.imediaconnection.com/article_full.aspx?id=21333 (accessed November 2, 2011).

Angwin, Julia (2010), “The Web’s New Gold Mine: Your Secrets,” *The Wall Street Journal*, July 30.

Ansari, Asim, Carl Mela, and Scott Neslin (2008), “Customer Channel Migration,” *Journal of Marketing Research*, 45 (1), 60-76.

Beales, Howard (2010), “The Value of Behavioral Targeting,” http://www.networkadvertising.org/pdfs/Beales_NAI_Study.pdf

Bettman, James R., Eric J. Johnson, and John W. Payne (1990), “A Componential Analysis of Cognitive Effort in Choice,” *Organizational Behavior and Human Decision Processes*, 45 (1), 111-139.

Blattberg, Robert C., Byung-Do Kim, and Scott A. Neslin (2008), *Database Marketing: Analyzing and Managing Customers,* New York: Springer Science+Business Media LLC.

Braun, Michael and Wendy W. Moe (2011), "Online Advertising Response Models: Incorporating Multiple Creatives and Impression Histories," available at SSRN: http://ssrn.com/abstract=1896486.

Chatterjee, Patrali, Donna L. Hoffman and Thomas P. Novak (2003), “Modeling the Clickstream: Implications for Web-Based Advertising Efforts,” *Marketing Science,* 22 (4), 520-541.

Chen, Ye, Dmitry Pavlov and John F. Canny (2009), “Large-Scale Behavioral Targeting,” *Proceedings of the 15th ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*, Paris, France.

Cho, Chang-Hoan and Hongsik John Cheon (2004), “Why Do People Avoid Advertising on the Internet,” *Journal of Advertising*, 33 (4), 89-97.

De Bock, K. W. and Van den Poel, D. (2010), "Predicting website audience demographics for web advertising targeting using multi-website clickstream data," Fundamenta Informaticae, 98 (1), 49-67.

Dinner, Isaac M., Harald J. van Heerde and Scott A. Neslin (2011), “Driving Online and Offline Sales: The Cross-Channel Effects of Digital versus Traditional Advertising,” *Working Paper.*

Drèze, Xavier, and François-Xavier Hussherr (2003), "Internet Advertising: Is Anybody Watching?" *Journal of Interactive Marketing,* 17 (4), 8-23.

Edwards, S. M., Li, H., and Less, J.-H. (2002). Forced Exposure and Psychological Reactance: Antecedents and Consequences of the Perceived Intrusiveness of Pop-Up Ads. *Journal of Advertising, 31* (3), 83-95.

Fader, Peter S. and Leigh McAlister (1990), "An Elimination by Aspects Model of Consumer Response to Promotion Calibrated on UPC Scanner Data," Journal of Marketing Research, 27:3 (August), 322-332.

Gensch, Dennis H. (1987), "A Two Stage Disaggregate Attribute Choice Model," Marketing Science, 6 (Summer), 223-31.

Gilbride, Timothy J. and Greg M. Allenby (2004), "A Choice Model with Conjunctive, Disjunctive, and Compensatory Screening Rules," *Marketing Science*, 23 (3), 391-406.

Goldfarb, Avi and Catherine Tucker (2011), "Online Display Advertising: Targeting and Obtrusiveness," *Marketing Science*, 30 (3), 389-404.

Ha, Sung Ho (2004), "An Intelligent System for Personalized Advertising on the Internet," *Proceedings of the 5th International Conference on E-Commerce and Web Technology*, K. Bauknecht, M. Bichler and B. Proll, Ed., LNCS 3182.

Inman, J. Jeffrey, Venkatesh Shankar, and Rosellina Ferraro (2004), "The Roles of Channel-Category Associations and Geodemographics in Channel Patronage" *Journal of Marketing*, 68 (April), 51-71.

Interactive Advertising Bureau (2010), "IAD Internet Advertising Revenue Report (April), http://www.iab.net/media/file/IAB_Full_year_2010_0413_Final.pdf.

Interactive Advertising Bureau (2008), "IAD Internet Advertising Revenue Report (March), http://www.iab.net/media/file/IAB_PwC_2008_full_year.pdf.

Interactive Advertising Bureau (2006), "IAD Internet Advertising Revenue Report (May), http://www.iab.net/media/file/resources_adrevenue_pdf_IAB_PwC_2006_Final.pdf.

Interactive Advertising Bureau (2004), "IAD Internet Advertising Revenue Report (April), http://www.iab.net/media/file/resources_adrevenue_pdf_IAB_PwC_2004full.pdf.

Kazienko, Przemyslaw and Michal Adamski (2007), "AdROSA – Adaptive Personalization of Web Advertising," *Information Sciences*, 177 (11), 2269-2295.

Kumar, Subodha, Milind Dawande and Vijay S. Mookerjee (2007), "Optimal Scheduling and Placement of Internet Banner Advertisements," *IEEE Transactions on Knowledge and Data Engineering*, 19 (11), 1571-1584.

Lambrecht, Anja and Catherine Tucker (2011), "When Does Reargeting Work? Timing Information Specificity," *Marketing Science Institute Working Paper Series 2011 Report No. 11-105.*

Li, Hairong, Steven M. Edwards, and Joo-Hyun Lee (2002), "Measuring the Intrusiveness of Advertisements: Scale Development and Validation," *Journal of Advertising*, 31 (2), 37-47.

Manchanda, P; Dube, JP; Goh, KY; et al. (2006), "The effect of banner advertising on Internet purchasing," *Journal of Marketing Research*, 43 (1): 98-108.

Moe, Wendy (2006), "An Empirical Two-Stage Choice Model with Decision Rules Applied to Internet Clickstream Data," *Journal of Marketing Research*, 43 (4), 680-692.

Moore, Robert S., Claire Allison Stammerjohan, Robin A. Coulter (2005), "Banner Advertiser-Web Site Context Congruity and Color Effects on Attention and Attitudes," *Journal of Advertising,* 34 (2), 71-84.

Naik, PA; Mantrala, MK; Sawyer, AG (1998), "Planning media schedules in the presence of dynamic advertising quality," *Marketing Science*, 17 (3): 214-235.

Neslin, Scott A. and Venkatesh Shankar (2009), "Key Issues in Multichannel Customer Management: Current Knowledge and Future Directions," *Journal of Interactive Marketing*, 23, 70-81.

Nerlove, M; Arrow, Kj (1962), "Optimal Advertising Policy under Dynamic Conditions," *Economica*, 29 (114): 129-142.

Rutz, OJ; Bucklin, RE (2011), "From Generic to Branded: A Model of Spillover in Paid Search Advertising," *Journal of Marketing Research*, 48 (1): 87-102.

Sherman, Lee and John Deighton (2001), "Banner Advertising: Measuring Effectiveness and Optimizing Placement," *Journal of Interactive Marketing*, 15 (2), 60-64.

Shamdasani, Prem N., Andrea J. S. Stanaland and Juliana Tan (2001), "Location, Location, Location: Insights for Advertising Placement on the Web," *Journal of Advertising Research*, 41 (4), 7-21.

Yan, Jun, Ning Liu, Gang Wang, Wen Zhang, Yun Jiang, Zheng Chen (2009), "How much can Behavioral Targeting Help Online Advertising," *Proceedings of the 18th international conference on World wide web*, Madrid, Spain.