PUNEET MANCHANDA, JEAN-PIERRE DUBÉ, KHIM YONG GOH, and
PRADEEP K. CHINTAGUNTA*

This article focuses on whether banner advertising affects purchasing patterns on the Internet. Using a behavioral database that consists of customer purchases at a Web site along with individual advertising exposure, the authors measure the impact of banner advertising on current customers' probabilities of repurchase, while accounting for duration dependence. The authors model the probability of a current customer making a purchase in any given week (since the last purchase) with a survival model that uses a flexible, piecewise exponential hazard function. The advertising covariates are purely advertising variables and advertising/individual browsing variables. The model is cast in a hierarchical Bayesian framework, which enables the authors to obtain individual advertising response parameters. The results show that the number of exposures, number of Web sites, and number of pages all have a positive effect on repeat purchase probabilities, whereas the number of unique creatives has a negative effect. Returns from targeting are the highest for the number of advertising exposures. The findings also add to the general advertising literature by showing that advertising affects the purchase behavior of current (versus new) customers.

The Effect of Banner Advertising on Internet Purchasing

Online advertising expenditures were expected to rise 29% in 2004 to approximately $9.3 billion (*BusinessWeek* 2004). Although Internet advertising is beginning to emerge as a viable medium (Silk, Klein, and Bernt 2001), its role and effectiveness has been the source of much debate. Prior research has shown that exposure to banner advertising leads to increased advertisement awareness, brand awareness, purchase intention, and site visits (Dreze and Hussherr 2003; Ilfeld and Winer 2002; Internet Advertising Bureau 1997; Sherman and Deighton 2001), but the relationship between advertising exposure and actual purchasing on the Internet has not been investigated.[1]

Since the early days of Internet commerce, there has been a lot of discussion about how the effectiveness of banner advertising should be measured. Web sites hosting online advertisements have been pushing for traditional exposure-based metrics, such as "impressions" served, to allow them to charge for each banner exposure. However, difficulties in measuring online impressions precisely have caused much dissatisfaction among managers, resulting in a reluctance to commit funds to banner advertising (Hoffman and Novak 2000). Moreover, advertisers, which prefer to pay on the basis of the performance of their advertisements, believe that, in general, impressions overstate advertising effectiveness. Instead, advertisers have been pushing for heuristic metrics of performance, such as "click-through," which indicates when a Web surfer clicks through to the advertiser's URL from the banner (for an analysis of click-

*Puneet Manchanda is Associate Professor of Marketing (e-mail: puneet.manchanda@ChicagoGSB.edu), Jean-Pierre Dubé is Associate Professor of Marketing (e-mail: jdube@gsb.uchicago.edu), and Pradeep K. Chintagunta is Robert Law Professor of Marketing (e-mail: pradeep.chintagunta@ChicagoGSB.edu), Graduate School of Business, University of Chicago. Khim Yong Goh is an assistant professor, Department of Information Systems, National University of Singapore (e-mail: gohky@comp.nus.edu.sg). The authors thank France Leclerc, Stijn M.J. van Osselaer, Wendy Moe, and Young-Bean Song for feedback; John Mracek, Steve Findley, Nancy Niu, and Susan Gertzis for help on data issues; two anonymous firms for providing the data; and the Kilts Center for Marketing at the Graduate School of Business, University of Chicago, for research support. Puneet Manchanda and Jean-Pierre Dubé also acknowledge research support from the Beatrice Foods Faculty Fund at the University of Chicago.

[1]A banner advertisement is a section of online advertising space that is typically 480 × 60 pixels in size. In general, it consists of a combination of graphic and textual content and contains a link to the advertiser's Web site by means of a click-through URL, which acts as a Web address.

ISSN: 0022-2437 (print), 1547-7193 (electronic)

Journal of Marketing Research
Vol. XLIII (February 2006), 98–108

through behavior, see Chatterjee, Hoffman, and Novak 2002). However, the effectiveness of click-through as a valid measure is also being called into question (Briggs 2001; *BusinessWeek Online* 2001; Dreze and Hussherr 2003; Song 2001). Typical click-through rates are quite small in magnitude, .5% on average (Dahlen 2001; Sherman and Deighton 2001; Warren 2001), which has led practitioners to believe that banners are ineffective. Moreover, click-through is a measure of a *visit* to the Web site. Because there is considerable evidence that only a small proportion of visits translate into final purchase (Moe and Fader 2003), click-through may be too imprecise to measure the effectiveness of banners served to the mass market. Therefore, these studies underscore the importance of investigating the impact of banner advertising on actual purchase behavior.

In this research, we focus on a previously unexplored question: Does banner advertising affect purchasing patterns on the Internet? In particular, using a behavioral database that consists of customer purchases at a Web site along with individual advertising exposure, we measure the impact of banner advertising on current customers' probabilities of buying again, while accounting for duration dependence. In particular, we examine whether, given a temporal interval since the last purchase, a customer makes a purchase at the Web site of interest and how this decision is influenced by exposure to banner advertising. We formulate a model of individual purchase timing behavior as a function of advertising exposure. We model the probability of a current customer making a purchase in any given week (since the last purchase) with a survival model. Effectively, a purchase represents "failure," whereas no purchase represents "survival." We capture the duration dependence in the customers' purchase behavior through a flexible, piecewise exponential hazard function (Wedel et. al. 1995). The advertising covariates enter through a proportional hazards specification. We use a much richer set of covariates than has typically been used in prior research (in which advertising is measured only as the amount of exposure). Specifically, the covariates we use consist of strictly advertising variables, such as weight and "diversity" (number of creative treatments), as well as advertising/individual browsing variables represented by how many and which pages expose customers to advertising. Our proposed model also controls for unobserved individual differences by specifying a distribution over the individual customer advertising response parameters. We do this by formulating our model in a hierarchical Bayesian framework. This also enables us to provide some insights into where the returns from targeted banner advertising are the highest and the extent to which the returns are higher than no targeting.

In terms of the broader area of research on the effects of (any type of) advertising on individual consumers, our work adds to the studies that investigate the effects of advertising on purchase timing and incidence behavior in at least two ways: First, it documents the effect of more facets of advertising than has been in studies with individual data (as we described previously). Second, a banner advertisement is a different form of advertising from a standard advertisement in terms of visual quality, attention-getting ability, and creative execution. Thus, our findings complement the findings of advertising's effect at the individual level that previous research has documented. Our main finding is that contrary to popular belief, exposure to banner advertising has a significant effect on Internet purchase behavior. This is reflected in our model as an increase in purchase probability (after we control for duration dependence) as a function of banner advertising exposure. From a managerial perspective, banner advertising has a positive effect on purchase probabilities in any given week (since the last purchase) beyond the duration-dependence effects. These results also suggest indirectly that click-through is a relatively poor measure of advertising effectiveness because it results in a very small proportion of overall purchases.

We find that the number of exposures, number of Web sites, and number of pages on which a customer is exposed to advertising all have a significant effect on customer purchase probabilities. Notably, increasing the number of unique creatives to which a customer is exposed lowers the purchase probability. In general, the effect sizes of banner advertising on purchase are in the same order of magnitude as the effects sizes of traditional advertising. We also find evidence of considerable heterogeneity across consumers in response to various aspects of banner advertising. The extent of heterogeneity shows that the returns from targeting individual customers are likely to be the highest for the weight of advertising (the number of advertisements that customers were exposed to in a given week), followed by the number of sites on which customers are exposed to advertising. Using individual response parameters, we conduct an experiment that demonstrates that even under very simple targeting approaches, there are significant increases in the effectiveness of banner advertising in terms of changing purchase probabilities and, thus, profitability. Finally, in terms of the broader area of research on the effects of (any type of) advertising, we provide somewhat unique evidence that advertising affects the purchase behavior of current (versus new) customers.

The structure of the article is as follows: We briefly discuss prior work in this and related areas. We then give an overview of the data. We present the details of the models and then discuss the results and the managerial implications of our findings. We conclude with a discussion of the limitations of the study and provide directions for further research.

LITERATURE REVIEW

Our specific focus in this article is the role of banner advertising in a digital environment, such as the Internet. However, our study also builds on a long tradition in marketing of estimating (conventional) advertising response models with individual-level data. Therefore, we discuss the relationship between our study and previous studies in both domains.

Most of the academic (e.g., Cho, Lee, and Tharp 2001; Dahlen 2001; Dreze and Hussherr 2003; Gallagher, Foster, and Parsons 2001) and industry research on advertising in digital environments has focused on measuring changes in brand awareness, brand attitudes, and purchase intentions as a function of exposure (as against the effects of banner advertising on actual purchase behavior). Such research is usually done with field surveys or laboratory experiments that use individual- (or cookie-) level data. Thus, the focus has been on understanding the effect of banner advertising on the awareness stage.

In contrast to studies that use experimental data, Sherman and Deighton (2001) describe the process of serving banner advertisements and collecting response data in detail. They also report the results of an experiment carried out by a Web advertising agency and an online merchant that show that targeting advertising to specific customers and on Web sites increases response rates and drives down the average cost per action (because of confidentiality restrictions, they report only broad, aggregate-level findings). Using aggregate data, Ilfeld and Winer (2002) show that increased online advertising leads to more site visits.

As we mentioned previously, there is a long tradition of research in marketing that models response to advertising with conventional scanner panel data (see Lodish et al. 1995; Vakratsas and Ambler 1999). Our research builds on this tradition by estimating a purchase incidence advertising response model with individual-level response parameters after controlling for unobserved heterogeneity. Thus, our research complements other research that has used individual-level data but has estimated only brand choice models (Deighton, Henderson, and Neslin 1994; Tellis 1988). Other researchers (e.g., Pedrick and Zufryden 1991, p. 112) have also questioned the managerial usefulness of brand choice models that ignore purchase incidence. Note that given our data, we cannot model brand choice. Our work extends previous research that models purchase incidence with a more detailed treatment of unobserved heterogeneity (e.g., Zufryden [1987] uses a summary measure) and the explicit incorporation of advertising covariates (in contrast to Pedrick and Zufryden 1991). Finally, in contrast to other studies that measure (individual) exposure to advertising with aggregate advertising dollars (e.g., Ilfeld and Winer 2002; Mela, Gupta, and Jedidi 1998), we use individual banner advertising exposure.

Findings from industry research (Black 2001; Briggs 2001; *BusinessWeek Online* 2001; DoubleClick Press Release 2001; Song 2001; Tran 2001; Warren 2001) show that banner advertising has attitudinal effects and that click-through is a poor measure of advertising response. In general, these findings are consistent with the findings of the academic research we discussed previously. Notably, in addition to the attitudinal effects of banner advertising, we find a few studies that provide some informal evidence of its behavioral effects as well. In this article, we use a formal model to investigate such behavioral effects for current customers.

The key differentiating managerial issue on the Internet is that firms and customers can build and manage relationships with individual customers in a much more cost-effective manner than is possible in other domains. Our research examines the influence of one marketing instrument (i.e., banner advertising) on a specific aspect of this relationship (i.e., purchase probability). To this end, our research uses banner advertising exposure and purchase data at the individual consumer (i.e., cookie) level and calibrates advertising response parameters at the individual level. This also distinguishes our work from previous research on advertising response that uses conventional panel data. Our research is also distinct from extant banner advertising research because this prior research has largely been limited to the influence of banner advertisements on attitudes, whereas the current study examines the influence of banner advertisements on behavior.

DATA

We obtained the data for this study from an Internet-only firm that sells health care and beauty products as well as nonprescription drugs to consumers. The data were processed and made available to us by the advertising agency that was responsible for serving the advertisements for the firm in question. Because of the nature of the data-sharing agreement between us and the two firms, we are unable to reveal the name of either firm. The data span all purchasers at the site during a period of three months in the third quarter of 2000, specifically, from June 11 to September 16. The data are available at the individual cookie level. As we mentioned previously, most data sets used to investigate online environments usually comprise only browsing behavior. Our data are unique in that we have individual-level stimulus (advertising) and response (purchase incidence). The data are contained in two databases: the CAMPAIGN database and the TRACER database.

The CAMPAIGN database comprises the online advertisement banner exposure and click-through response originating from promotional campaigns that were run on Web sites. The data fields in the CAMPAIGN database consist of consumer data—a unique cookie (we use the terms "cookie" and "consumer" interchangeably herein) that identifies the individual computer;[2] an indicator variable that denotes consumer response to the banner advertisement (view or click);[3] and the date and time the person viewed or clicked on the banner advertisement, the portal or alliance site's Web page on which the banner advertisement view or click occurred, and a unique key that identifies the specific banner advertisement.

In terms of the Web sites on which the advertising was delivered, the database contains records of the company's advertising on portal and alliance Web sites, such as Yahoo, America Online, Women.com, iVillage.com, Healthcentral.com, and E*Trade, among others. These sites accounted for more than 80% of the firm's advertising activity during this period. Note that though we have a unique identifier for each site on which the banner advertisement was served, we do not know the specific identity of each site. Advertising activity typically consisted of a specific creative that oper-

[2]Our data only allow us to identify a unique computer, not a unique consumer. Our assumption of equivalency between a consumer and a computer could be strong in certain environments (for details, see Dreze and Zufryden 1998).

[3]Note that because we are working with behavioral data, we are unable to control for the exact nature of exposure. In other words, we are assuming that if the consumer was on a specific page and the banner appeared on that same page, he or she actually viewed the advertisement. This assumption is consistent with prior research that has examined the effect of advertising exposure on sales for individual consumers (for a detailed discussion, see Deighton, Henderson, and Neslin 1994, p. 34) and experimental studies that have used banner advertisements (e.g., Dreze and Hussherr 2003). Even given this assumption, banner advertisements are probably even lower involvement than television advertisements (small size, exposure in the presence of competing information), which argues against our finding any effects of banner advertisement exposure.

ated over several weeks. In terms of the advertising message contained in the various creatives, we know that the majority of the messages were of the brand-building type for the Web site (i.e., the message consisted of the name of the Web site and a line describing the benefits of purchasing from the Web site). A limitation of the data is that we do not have information on the specific message in each banner (even though we have an indicator that tells us that one creative was different from another). This creative was delivered to Web sites in the form of a digital graphic, generally referred to as a GIF. These GIFs were of the usual size for banner advertisements (480 × 60 pixels). New GIFs were typically released at the beginning of a calendar week (i.e., on Sunday and/or Monday, reflecting media buying patterns). During the period covered by our data, there were 100 total GIFs spread over 15 major sites. However, the majority of exposures came from a small number of GIFs: 7 GIFs accounted for approximately 55% of all exposures.

The TRACER database contains the date and time of the purchase transaction for each unique cookie identifier. Note that we do not have information on visits to the site that did not result in a purchase. We merged the CAMPAIGN database with the TRACER database using unique cookie identifiers. This resulted in 14,370 unique cookies. We then examined the purchasing patterns of these cookies in the context of our discrete-time formulation. Given that banner advertising activity was planned by the firm for each week, we chose the time interval to be a single week. Thus, our unit of observation is a "cookie week." However, if there were a significant number of cookies who purchased multiple times in a single week, our model would be inappropriate. An examination of the data revealed that 99% of the 14,370 cookies did not purchase multiple times in any given week. We then deleted all the cookies for which we could construct only one observation (i.e., if their purchase occurred in the last six calendar days of our data) because we would be unable to obtain individual-level parameters for these cookies. (We describe how we construct the weekly data for each cookie in the "Model Specification" section.) Finally, we deleted purchase transactions with blank cookies, repeat transactions (identical transactions at identical times), and observations with obvious data entry errors. This resulted in a panel of 12,748 cookies with a total of 97,805 observations. The number of observations in the data is the sum (over the 12,748 cookies) of the total number of weeks for each cookie after the cookie's first purchase.

Of these 97,805 observations, a purchase is made on 14.3% (13,955) of the observations; there is no purchase on the remaining 85.7% (by no purchase, we mean that there was no purchase from the online store that provided us with the data). It is instructive to compare this proportion with purchases based on click-through. On the basis of the sample click-through and purchase rate of .25% (consistent with rates documented in prior studies and validated by the firm's advertising agency), we find that click-through-only purchases are fewer than purchases in our data (1134 purchases versus 13,955 purchases across all purchasers). Combined with feedback from the firm's executives, this leads us to conclude that click-through is not an important path to purchase (for customers of this Web site). This is also consistent with findings from experimental research (Dreze and Hussherr 2003).

MODEL

We investigate the purchase behavior of customers who are exposed to banner advertising by the Web site. We model the potentially duration-dependent purchase incidence decision—that is, whether and when to buy from the Web site—with a semiparametric survival model (for a comparison of alternative specifications, see Seetharaman and Chintagunta 2003). Specifically, we estimate a constant piecewise exponential hazard model in discrete time (Wedel et al. 1995). This allows the intrinsic purchase incidence probabilities to vary over time in the absence of covariates. We use this formulation rather than a standard purchase incidence model because we believe that relative to the cross-sectional effects, duration-dependence effects are likely to be large and significant in our data. We also model the decisions of when and whether to purchase as a function of the advertising exposure and browsing behavior variables at the individual customer level. To capture variability in individual choices, we allow the individual response parameters to be distributed across customers.

As noted previously, our model formulation focuses on the weekly purchase decision; that is, consumers decide every week whether they plan to purchase as a function of the timing of their last purchase, marketing and behavioral variables, and unobserved heterogeneity. Our model falls into the class of semiparametric survival models (Meyer 1990). The main advantage of the semiparametric specification is that it does not impose a specific distributional assumption or a shape on duration dependence (i.e., the baseline hazard). This model treats the no-purchase weeks for each customer as the survival weeks, whereas it treats the purchase weeks as the failure weeks. Prior modeling research that uses customer browsing data has found evidence of heterogeneity (Bucklin and Sismeiro 2003; Moe and Fader 2003). Therefore, we account for heterogeneity using a continuous distribution over the individual customer response parameters. We cast our model in a hierarchical Bayesian framework and estimate it using Markov chain Monte Carlo methods (for a detailed review of such models, see Rossi and Allenby 2003). In general, with a few notable exceptions (Allenby, Leone, and Jen 1999; Lee, Boatwright, and Kamakura 2003), the use of proportional hazard models under the hierarchical Bayesian framework has been somewhat limited in the marketing literature. Note that our model form is analogous to purchase incidence models, such as the binary logit/probit with temporal fixed effects. We now describe the specific model, the prior distribution of the unknowns, the likelihood function, and the resultant posterior distributions.

THE SEMIPARAMETRIC SURVIVAL MODEL

Let t_{ij} denote the interpurchase time for consumer i's spell j, and let h(u) denote a hazard function. The survivor function corresponding to this time is as follows:

$$S(t_{ij}) = \exp\left[-\int_0^{t_{ij}} h(u)du\right]. \tag{1}$$

Note that because our data are discrete survival data, we use a discrete-time model to predict the probability of purchase.

We first split the time axis into a finite number of intervals, $0 < s_1 < s_2 < \ldots < s_J$, where $s_J > y_{it}$ for all i = 1, 2, …, I and t = 1, 2, …, T_i and where y_{it} represents the survival time for customer i's tth observation. Thus, we have J intervals, $(0, s_1]$, $(s_1, s_2]$, …, $(s_{J-1}, s_J]$. Following the convention in the discrete-time semiparametric hazard function literature (e.g., Meyer 1990), we replace the integral in Equation 1 for each of the J intervals with the following expression:

$$\int_{(t-1)_{ij}}^{t_{ij}} h(u)du = \exp(\lambda_j). \tag{2}$$

This represents a piecewise exponential hazard model in which we assume a constant baseline hazard, $h_{0j}(y) = \log(\lambda_j)$, for $t_{ij} \in I_j = (s_{j-1}, s_j]$, where I_j is the indicator function. The $\log(\lambda_j)$ parameters do not correspond to calendar time but rather to the time interval following the last purchase. They enable us to assess whether the data indicate duration dependence when the parameters are different for different time intervals or durations. Note that for most cookies, we have only one observation for each of the J intervals. Thus, the data support the inference about the baseline hazard only at the pooled level; that is, we cannot specify a heterogeneity distribution across customers for any of the $\log(\lambda_j)$ parameters.

We then let the effect of the covariates enter multiplicatively; that is, we use a proportional hazard formulation. Let x_{pij} represent the pth covariate for customer i in the time interval j. Because we have repeated measures across customers (after we control for the pooled baseline hazard), the response parameters can be customer specific. Thus, Equation 2 becomes the following:

$$\int_{(t-1)_{ij}}^{t_{ij}} h(u)du = \exp\left[\lambda_j + \sum_{p=1}^{P}(x_{pij} \times \beta_{pi})\right]. \tag{3}$$

The piecewise exponential model is general in the sense that it is sufficiently flexible to accommodate a wide variety of shapes of the baseline hazard. Note that if J = 1, the model reduces to a parametric exponential model with a failure rate $\lambda = \lambda_1$. It is also parsimonious in the sense that there is only one unknown parameter per time period. Given Equation 3, the probability of purchase (failure) in any of the j time intervals for a customer i is given as follows:

$$Pr_{ij}(\text{purchase}) = 1 - \exp\left[-\exp(u_{ij})\right], \tag{4}$$

where $u_{ij} = \Sigma_{j=1}^{J}(\lambda_j \times I_j) + \Sigma_{p=1}^{P}(x_{pij} \times \beta_{pi})$ and where $I_j = 1$ in time interval j and equals 0 otherwise. Thus, the overall log-likelihood for all the customers in the sample is as follows:

$$LL_i(\beta_i, \lambda|\phi, x) = \sum_{i=1}^{I}\sum_{j=1}^{J}[\log(Pr_{ij}) \times \phi_{ij} + \log(1 - Pr_{ij}) \times (1 - \phi_{ij})], \tag{5}$$

where ϕ_{ij} is an indicator function that is equal to 1 if customer i purchases in time interval j and 0 if otherwise and where β_i and λ are vectors of β_{pi} and λ_j.

THE BAYESIAN HIERARCHY AND INFERENCE

We cast our model in a hierarchical Bayesian framework. Given that we want to obtain simultaneously the cross-sectional parameters for the discrete-time hazards and the individual-level parameters for the response coefficients, this framework is particularly appealing. Under this framework, to complete the model, we need to specify the prior distribution of the unknowns and derive the full conditional distributions.

Let $\psi_j = \log(\lambda_j)$ for j = 1, 2, …, J. We assume that ψ_j are distributed multivariate normal with mean ψ_0 and variance V_ψ. We capture unobserved heterogeneity with the distribution of β_i (where β_i is the vector of the response parameters) by allowing it to be distributed multivariate normal with mean β_0 and variance V_β; that is,

$$\beta_i = \beta_0 + \nu_i, \tag{6}$$

where $\nu_i \sim N(0, V_\beta)$. The hyperparameters β_0 and V_β are distributed multivariate normal and inverse Wishart, respectively.

We derive the full conditional distributions of the unknowns, $(\psi, \beta_i, \beta_0, V_\beta)$, using the joint density (Equation 5) and the specified prior distributions. We then draw sequentially from this series of full conditional distributions until convergence is achieved. Both the full conditional distributions and the inference process are standard. Thus, we do not describe them here (a detailed Appendix is available from the authors on request).

MODEL SPECIFICATION

We first discuss how we specify the baseline hazard. We have 13 calendar weeks in our data. We created the spell variables for each cookie in the following manner: We initialized the first spell for each cookie to the calendar week that corresponded to the first purchase occasion. We then created purchase indicators, ϕ_{ij} (Equation 5), for each week following this initial week for each cookie. If there was no purchase in a subsequent week, we incremented the spell counter by one and set the purchase indicator to zero. If there was a purchase, we set the indicator variable to one. We restarted the spell counter at one for the week following the purchase week. Then, we created 13 indicator variables, I_1, …, I_{13} (Equation 4), and set them to one, corresponding to the spell counter for that week for that cookie. As we mentioned previously, the indicator variables do not represent the calendar week but rather the number of weeks elapsed since the last purchase. The coefficients of each of these variables, $\psi_j = \log(\lambda_j)$, represent the constant hazard for that week (except for ψ_{13}, which represents the hazard of weeks 13 and higher). We then constructed the four advertising covariates that we use (described subsequently) for each cookie week.

We now consider the advertising variables. We postulate that the decision of whether to purchase in each week will be affected by advertising exposure (weight and diversity) and individual differences (both observed and unobserved). We first discuss the advertising variables. We expect that banner advertisements act as reminder tools and/or brand builders for current customers. Thus, exposure to banner advertising is likely to increase the probability of purchase (Cho, Lee, and Tharp 2001). Therefore, we construct the following variables: VIEWNUM represents the total number of advertising exposures in each week for each customer, and ADNUM represents the number of creatives (GIFs) that the consumer was exposed to each week.

Prior research has shown that repeated exposures to an advertisement prevent the early decay of advertising effects (Cacioppo and Petty 1985; Dreze and Hussherr 2003; Pechmann and Stewart 1988). Therefore, we expect that increased exposure to advertising (VIEWNUM) should increase the probability of purchase in a given week. However, at some point, the response to advertising should provide diminishing returns. Because the empirical evidence supports a concave response to advertising weight in general (Lilien, Kotler, and Moorthy 1992, p. 267), we use log(1 + VIEWNUM), or LVIEWNUM, in our specification (to accommodate weeks when VIEWNUM = 0). In terms of the variety of creative execution, prior research has indicated that response to different creatives can be quite different (Lodish et al. 1995). It has also been shown that recall is enhanced if consumers are exposed to different creatives in the same campaign (Rao and Burnkrant 1991). However, in our case, whereas all the creatives essentially advertise the Web site, they are not part of a single campaign. Therefore, we have no prediction about the effect of ADNUM on interpurchase times.

We also need to control for differences across consumers in terms of prior purchase behavior and browsing behavior. These differences could be both observed and unobserved. Observed differences may arise as a result of two kinds of variation: purely cross-sectional variation (e.g., demographics) and cross-sectional combined with longitudinal variation (e.g., usage and browsing behavior). Usage variables capture systematic differences in customers' use of digital environments. For example, some customers may spend more time on the Internet and therefore may be more prone to buying from Web merchants. Thus, the individual probability of buying for such a consumer could be affected by individual browsing behavior and individual advertising exposure.[4] Our data do not contain any direct measures of Internet usage and browsing behavior. However, we use the data available to us and develop two proxy variables that potentially reflect individual differences in Internet usage: SITENUM represents the total number of unique Web sites on which the consumer was exposed to advertising each week, and PAGENUM represents the number of unique Web pages on which the consumer was exposed to advertising each week.

Note that the use of these variables controls for both across-customer (in that the means of these variables are likely to differ across customers) and within-customer (there may be differences in these variables for the same customer across weeks) differences. From the usage-based arguments discussed previously, we expect these variables to have a positive effect on the decision to purchase each week. From an advertising perspective, prior research has shown that viewing a series of advertisements leads to higher recall and more positive attitudes (Pechmann and Stewart 1988; Zielske and Henry 1980). Therefore, we expect that the probability of purchase is higher for consumers who are exposed to advertising on many different Web sites (SITENUM) and pages (PAGENUM).

Taken together, these four covariates provide a richer description of individual exposure to advertising than has typically been studied in the literature.[5] Specifically, we have data on quantity, diversity, and the location of exposure in contrast to just quantity. In addition, a banner advertisement is a different form of advertising than is a standard advertisement in terms of visual quality, attention-getting ability, and creative execution. To summarize, we expect to observe a positive sign for the coefficient of LVIEWNUM, SITENUM, and PAGENUM, whereas the sign for ADNUM could be either positive or negative. Note that a positive coefficient increases the purchase probability and a negative coefficient decreases the purchase probability.

The temporal sequence of events for a typical customer in our data is as follows: Every week, the consumer is exposed to some advertising that spans (possibly) different creatives. These exposures occur at (possibly) different Web pages on (possibly) different Web sites. As a result, each week the consumer decides whether to purchase (given a purchase in the past). We identify the three sets of parameters in the following manner: First, we identify the time dummies from the aggregate temporal purchase patterns (we pool them across consumers) after controlling for the effect of the covariates. Second, we identify the mean response parameters by variation across consumers. Third, we identify the individual response parameters by variation within consumers. In terms of the data, the means (standard deviations) of LVIEWNUM, ADNUM, SITENUM, and PAGENUM are .25 (.66), .23 (.64), .07 (.30), and .09 (.33), respectively.

RESULTS

Model Estimates: Duration Dependence

The mean (posterior standard deviation) baseline hazard parameters, ψ, for weeks 1–13 are .57 (.06), –2.80 (.04), –2.54 (.03), –2.74 (.04), –2.48 (.04), –2.49 (.04), –1.78 (.03), –1.97 (.04), –1.98 (.04), –.96 (.02), –1.30 (.03), –.89 (.03), and –.31 (.03), respectively. All the parameters have posterior distributions that are massed at a considerable distance from zero. These parameters map directly onto the purchase probability in a given week j. The higher the magnitude of ψ_j, the higher is the probability of purchase. The estimates show that there is some nonmonotonicity in the probability of purchase as the number of weeks since the last purchase increases. As can be observed from the estimates, the probability of purchase in a given week increases somewhat for the first three weeks and then remains flat

[4]Note that individual exposure to advertising may be systematically different across consumers if the firm and its advertising agency are strategically targeting advertising to individual cookies on the basis of prior browsing and/or purchase behavior. However, our discussions with the firm revealed that during the time period of our data, this was not the case, because the technology to do this was still not available for commercial use.

[5]Note that some of our measures may be perfectly confounded with non-advertising-related behavior; for example, ADNUM could be confounded with site content if different creatives appear on different sites systematically. Given our data, we cannot disentangle these effects. We thank an anonymous reviewer for pointing this out.

until Week 6. After Week 6, we observe two peaks in Week 7 and Week 10, followed by a dip in Week 11 and then an increase for the following weeks. This is consistent with the mean time between purchases in our data (approximately seven weeks). The estimated survival pattern does not conform to any well-known parametric survival function formulation. This provides some support for our choice of a piecewise exponential hazard formulation.

Model Estimates: Advertising Covariates

We now examine the effect of covariates at the mean level (i.e., the β_0 vector; see Table 1). The overall pattern of the results indicates that advertising weight, diversity, and the individual browsing variables have an effect on the decision to purchase in any given week (all the posterior means are massed away from zero). These effects are as we predicted. As advertising weight (the log of the number of advertising exposures every week – LVIEWNUM) increases, the survival probability is lowered. In other words, greater exposure to advertising (numbers) has a positive effect on the purchase probability, albeit in a manner consistent with diminishing returns. Notably, the two main studies that have investigated the effect of advertising on repeat purchasers using individual exposure data are those of Deighton, Henderson, and Neslin (1994) and Tellis (1988). Neither study finds any effects of advertising on repeat customers (operationalized as the interaction between exposure and the last brand chosen). Thus, these studies find no effects on repeat brand choice behavior. To the best of our knowledge, there are no other studies that have found a positive effect of advertising on current customers. Thus, our finding is unique in this regard.

However, advertising diversity (the number of creatives that a customer is exposed to in every week, or ADNUM) has a positive effect on the survival probability. Thus, exposure to more creatives in a week decreases the probability of purchase. This result is not surprising given that previous research has not hypothesized or documented a specific direction of the relationship. This may be because repetition (of the same message) aids consumers in learning and retaining the ad content. This idea is especially relevant given the plethora of competing banner messages to which a Web surfer may be simultaneously exposed. This result is also more consistent with Lodish and colleagues' (1995) finding that we cited previously. In addition, given that the creative content of a banner advertisement is relatively constrained, exposure to more creatives can lead to more fragmentation rather than to reinforcement.

The effect of exposure to advertisements on many different Web sites (SITENUM) and many different pages (PAGENUM) is to lower the survival probability. Thus, broadly speaking, the more locations (sites and pages) at which the consumer is exposed to advertising, the higher is the probability of purchase. Cross-sectional differences in browsing behavior across the total number of consumers could account for this effect. However, as we discussed previously, individual response coefficients are all positive, which implies that even within consumers, exposure at a greater number of locations increases the purchase probabilities.

In terms of the differences across customers, the means (posterior standard deviations) of the V_β diagonal elements of LVIEWNUM, ADNUM, SITENUM, and PAGENUM are .11 (.03), .10 (.01), 1.86 (.53), and .15 (.03), respectively. This implies that there is considerable heterogeneity across customers (a detailed discussion follows). The correlation patterns across the response parameters are also noteworthy. First, only two of the six correlations are massed away from zero. This suggests that the response coefficients of four advertising variables are not highly correlated across consumers (i.e., the four variables measure different facets of responsiveness to advertising). Second, the correlation in response parameters across LVIEWNUM and SITENUM is negative and massed away from zero (–.54). This implies that responsiveness to advertising is lower for consumers who are more responsive to being exposed to banner advertising on many Web sites. Because the mean effect of both LVIEWNUM and SITENUM is positive, there are tradeoffs in the development of individual-level targeting based on these two variables. Third, the correlation in response parameters across ADNUM and SITENUM is negative and massed away from zero (–.41); that is, customers who are more responsive to different creatives are less responsive to being exposed to advertising on many Web sites. However, because the main effect of ADNUM seems to be negative, this correlation implies that it may be better to expose consumers to the same creative at a small number of sites for maximal response.

In summary, we find that in our data, advertising weight and copy affect consumers' decisions to visit the Web site and make purchases. In addition, we also find that cross-sectional differences in browsing behavior have an effect on purchase probabilities. Finally, exposure on distinct locations (sites and pages) for the same consumer also tends to increase the purchase probability. We also find that these response parameters vary across consumers and that there are some interesting correlations across these parameters.[6]

MANAGERIAL IMPLICATIONS

We now use our results to explore their implications on managerial practice. First, we compute the average effect sizes of the various advertising variables. Second, we investigate the variation in consumer responsiveness to advertising to obtain an understanding of the returns to targeting.

Table 1
EXPECTED AND ESTIMATED EFFECTS (COVARIATES)

Variable	*Expected Sign*	*M*	*Posterior SD*
LVIEWNUM	+	.10	.01
ADNUM	?	–.25	.02
SITENUM	+	1.55	.07
PAGENUM	+	.84	.05

[6]We also compared our model with a series of null models in which we assumed (1) only duration dependence, (2) only advertising covariates (simple purchase incidence model), and (3) duration dependence (time since last purchase and time since last purchase squared) and advertising covariates using a purchase incidence model. Our proposed model performed better than these models both within and out of sample.

Elasticities

To understand the extent of the effect sizes, we compute the change in probability of purchase for a 10% change in the advertising variable for each observation and then compute the mean elasticity across observations and draws post-burn-in. The mean (standard deviation) elasticity magnitudes for VIEWNUM, ADNUM, SITENUM, and PAGENUM are .02 (.01), –.03 (.01), .05 (.02), and .04 (.02), respectively. Note that though the mean effects are small, the standard deviations indicate that they are massed away from zero. In addition, these elasticities are in the same order of magnitude, that is, within one-tenth the estimate reported for conventional advertising in the literature (e.g., Sethuraman and Tellis [1991] report an average advertising elasticity of demand of .10).

These findings have some interesting implications. First, it seems that firms need to cut back on exposing customers to different creatives and stick with a smaller set of creatives. This result may also be a reflection of banner advertisements being limited in terms of the creative that they can deliver. Second, it seems that both the weight and the diversity of advertising have smaller effects than the sites on which consumers are exposed to banner advertising. This reinforces the belief in the industry that delivering a consistent message across many different sites and pages is the most effective method of marketing communication on the Internet, where there are many distractions on the same Web page. Third, the number of Web sites that expose a customer to advertising is about as important as the number of Web pages. However, because the average effect of Web sites is higher than the effect of pages, for advertising decisions on the margin, firms should locate themselves on the high-traffic pages across more Web sites rather than spread exposures across many, possibly unrelated, pages.

Returns to Targeting

Given that the average effect sizes of banner advertising are significant, managers may be interested in exploiting the one-to-one targeted marketing potential of the Internet. To analyze this, we need to compute the returns to targeting across the four advertising/individual difference variables in our analysis. We accomplish this through an examination of heterogeneity in response parameters across the individual customers. We compute the coefficient of variation (the standard deviation divided by the mean) for the distribution of the individual response parameters to describe the size of the variation in response across the four parameters. They are .92 (LVIEWNUM), .22 (ADNUM), .34 (SITENUM), and .07 (PAGENUM). These numbers imply that people differ the most in terms of their response to the number of advertisements to which they are exposed, followed by the number of sites on which they are exposed, the number of creatives to which they are exposed, and then the number of pages to which they are exposed. Therefore, the returns to targeting follow the same order.

We formalize our findings on the returns to targeting through a stylized revenue (profitability) experiment. We classify customers as "high sensitives" (H) and "low sensitives" (L) by a median split on each of the four individual-level parameters. We then classify them in four groups—HH, HL, LH, and LL—on the basis of their sensitivity to the two stimuli we identified as having the highest likely returns from targeting: LVIEWNUM and SITENUM. In other words, an HL customer is a high sensitive on the amount of exposed advertising but a low sensitive on the number of sites on which the exposure occurs. We then choose a week sufficiently far from the previous purchase occasion (Week 9) to contrast the effect of targeted and untargeted advertising purchase behavior and profitability. We assume that the firm can expose customers (who do not purchase in Week 9) to one, two, or three banner advertisements. These advertisements can be distributed on one, two, or three sites. In the untargeted strategy, we expose all the customers to the identical number of banner advertisements on the same number of sites. In the targeted strategy, we choose different levels of exposure and sites, depending on the classification of the customer. We then compute the new probability of purchase (in each condition) and take the product of that probability with the average historical dollar expenditure by that customer (i.e., the expected revenue) on the Web site. We sum across all the chosen customers to obtain the total revenue in each condition. On the basis of industry feedback, we assume that every additional exposure costs the firm \$.05 and that there is a \$.02 charge for every additional Web site on which the advertising is placed. We compute the return for each strategy as (total revenue – total cost)/total cost of each strategy. The results appear in Table 2. Note that the absolute values that we report in the table may or may not be meaningful. For the purpose of our stylized experiment, it is only important to focus on the relative values.

There are three interesting findings in the table. First, the return is always greater for the targeted advertising strategy, even though we use a simple targeting rule, such as a median split, and then assign the number of exposures and the number of Web sites in a fairly simple manner. Second, the general magnitude of the return becomes smaller as more exposures are provided; the maximum return is 108, 141, and 159 for three, two, and one exposure, respectively. This is likely due to the diminishing return nature of the response to exposure. Third, the return becomes greater as the options for targeting become larger. To clarify, the incremental return for an additional three, two, and one exposure is 19% (108/91), 14% (141/123), and 5% (159/152), respectively. With only one additional exposure, firms are limited in how they can target (e.g., they can decide which customers to expose and on which site). In contrast, with three exposures, finer targeting is possible, leading to higher (relative) returns. In conclusion, even with a simple targeting strategy, the firm can reap significant benefits. Thus, this experiment demonstrates the value of obtaining individual-level parameters to create profitable targeting strategies.[7]

[7]We also carried out a separate experiment in which we targeted banner advertising with a median split on the LVIEWNUM response coefficient. This targeting exercise was subject to the constraint that the advertising expenditure was identical across the targeted and untargeted groups. Even under this extremely simplistic targeting strategy (which provides a lower bound), we found the gain from targeting to be between 1.2% and 2.1% for up to three additional advertisements.

Table 2
RETURNS TO TARGETING

Group/Stimulus	*Optimal Exposure/Site*	*Revenue ($)*	*Cost ($)*	*Return*[a]
Additional Exposures = 3				
Untargeted[b]	3/2	71,989	781	91
Targeted[c]		67,376	616	108
HH, HL, LH, LL	3/3, 3/1, 2/2, 0/0			
Additional Exposures = 2				
Untargeted[b]	2/2	71,574	575	123
Targeted[c]		54,074	380	141
HH, HL, LH, LL	2/2, 2/2, 1/1, 0/0			
Additional Exposures = 1				
Untargeted[b]	1/1	43,945	288	152
Targeted[c]		42,087	263	159
HH, HL, LH, LL	1/1, 1/1, 1/1, 0/0			

[a]We computed return as (total revenue – total cost)/(total cost).
[b]For the untargeted scenario, there could be different strategies (e.g., 3 exposures on 3 sites versus 2 sites). We simulated all possible scenarios and picked the one with the highest return (for 3 exposures, it was 2 sites, and for 2 exposures, it was also 2 sites).
[c]This represents the banner ad placement for each group (e.g., 3/2 represents 3 exposures on 2 sites).

Discussion

Our findings have several implications for managers. First, we find unique evidence that exposure to banner advertising has a significant effect on Internet purchasing. Specifically, all else being equal, we find that exposure to banner advertising increases the purchase probabilities for current customers. Second, the elasticity estimates are in the same order of magnitude as those documented for conventional advertising, suggesting that managers should expect to observe effect sizes that are consistent with other forms of advertising. To the best of our knowledge, this is the first documentation of such effect sizes. Third, our data and results show that ad exposure is likely to lead to an increase in purchase probabilities after exposure. This implies that managers may be focusing on the wrong metric when they use instantaneous metrics, such as click-through, to measure advertising effectiveness. Fourth, our results have implications for the design and execution of banner ad campaigns. Broadly speaking, campaigns should be designed such that customers are exposed to fewer (and more consistent) creatives across many pages and Web sites. Fifth, given a fixed number of exposures and creatives, returns to exposure are somewhat higher for sites first and then pages. Finally, although the mean response to the number of exposures is somewhat lower, the returns to targeting on this measure are likely to be the highest. A stylized experiment shows that the returns to targeting are higher than if there is no targeting (even when the number of exposures is constant across the two groups) and that these returns become relatively higher as the targeting options become larger.

CONCLUSION

Our research fits into the small but growing subfield of empirical research that is dedicated to measuring the impact of the Internet on marketing policy areas, such as pricing, product assortment, and advertising. To the best of our knowledge, this article is the first attempt to model the effects of banner advertising on customers' Internet purchasing. We use a unique data set to investigate the effects of banner advertising on the weekly purchase probability of existing customers. Our main finding is that contrary to popular belief, banner advertising does affect purchase probabilities. We find this result because our modeling approach allows for temporal separation between advertising and purchase behavior. We speculate that the temporal separation exists because advertising acts as a brand-building tool and/or a reminder. The corollary to this finding is that measures of instantaneous behavior, such as click-through, may be poor measures of advertising effectiveness. We find that both the weight and the diversity of advertising have an effect on customers' purchase probabilities. We also find that the more creatives a customer is exposed to in a given week, the lower is the purchase probability. Given a relatively simple medium, such as banner advertising, and the amount of competing information on Web pages, different messages may be diluting the impact of advertising. Our findings also show that exposure to banner advertising on more (unique) Web sites and Web pages has a slightly greater effect on the individual purchase probabilities than the weight and diversity of advertising. This may be because these two covariates contain information about both advertising exposure and individual differences in browsing behavior.

We also find evidence of considerable heterogeneity across consumers in response to advertising. Heterogeneity is highest for the ad weight response, followed by the number of unique exposure sites. Thus, for maximal response, targeted communication should be based on these two response variables. We show the managerial benefits of our approach through a stylized experiment that demonstrates that targeted banner advertising provides relatively higher returns. Finally, in terms of the broader area of research on the effects of (any type of) advertising, we provide somewhat unique evidence that advertising affects the purchase behavior of current customers.

There are also some limitations of our research that arise primarily from the lack of information in our data. First, our results may not apply to customers who have not purchased items at least once at this Web site. Second, we do not have any demographic information and other relevant behavioral metrics (e.g., Internet usage) on the cookies. This information may have been useful in explaining a larger part of the unobserved heterogeneity. Third, our results would have been richer if we had information on the actual message contained in each advertisement and the identity of the referral sites. Fourth, we do not have any knowledge of the other marketing variables, such as price and promotion, during consumers' purchase visits. We were also unable to detect any effects of lagged advertising in our data. Finally, the correct targeting experiment to carry out is to allocate advertising in the targeted case such that the marginal profit (revenue less cost of goods sold) for each customer week is equal to the marginal cost of delivering the banner advertisement. Because our data do not contain information on the profit per customer and cost of delivering an advertisement to a specific Web site on a specific day, we are unable to do this. Further research could address these limitations by running formal field experiments (see Lodish et al. 1995) or by obtaining richer data sets that provide natural variation on these dimensions.

REFERENCES

Allenby, G.M., R.P. Leone, and L. Jen (1999), "A Dynamic Model of Purchase Timing with Application to Direct Marketing," *Journal of the American Statistical Association*, 94 (446), 365–74.

Black, Jane (2001), "Online Advertising: It's Just the Beginning," *BusinessWeek Online*, (July 12), (accessed October 10, 2005), [available at http://www.businessweek.com/technology/content/jul2001/tc20010712_790.htm].

Briggs, R. (2001), "Abolish Clickthrough Now!" (accessed October 16, 2005), [available at http://www.intelliquest.com/resources/whitepapers/abolishclick.pdf].

Bucklin, Randolph E. and Catarina Sismeiro (2003), "A Model of Web Site Browsing Behavior Estimated on Clickstream Data," *Journal of Marketing Research*, 40 (August), 249–67.

BusinessWeek (2004), "The Online Ad Surge," (November 22), 76–80.

BusinessWeek Online (2001), "A Chat with a True Believer," (July 12), (accessed October 10, 2005), [available at http://www.businessweek.com/technology/content/jul2001/tc20010712_184.htm].

Cacioppo, J.T. and R.E. Petty (1985), "Central and Peripheral Routes to Persuasion: The Role of Message Repetition," in *Psychological Processes and Advertising Effects*, A. Mitchell and L.A. Alwitt, eds. Hilsdale, NJ: Lawrence Erlbaum Associates, 9–11.

Chatterjee, P., D.L. Hoffman, and T.P. Novak (2002), "Modeling the Clickstream: Implications for Web-Based Advertising Efforts," *Marketing Science*, 22 (4), 520–41.

Cho, C., J. Lee, and M. Tharp (2001), "Different Forced-Exposure Levels to Banner Advertisements," *Journal of Advertising Research*, 41 (4), 45–56.

Dahlen, M. (2001), "Banner Advertisements Through a New Lens," *Journal of Advertising Research*, 41 (4), 23–30.

Deighton, John, Caroline M. Henderson, and Scott A. Neslin (1994), "The Effects of Advertising on Brand Switching and Repeat Purchasing," *Journal of Marketing Research*, 31 (February), 28–43.

DoubleClick (2001), "Diameter's Campaign Evaluation Illustrates Direct Correlation Between Exposure to On-Line Ads and Brand Impact," press release, (July 30), (accessed October 12, 2005), [available at http://www.doubleclick.com/us/about_doubleclick/press_releases/default.asp?p=216.]

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

——— and F. Zufryden (1998), "Is Internet Advertising Ready for Prime Time?" *Journal of Advertising Research*, 38 (3), 7–18.

Gallagher, K., K.D. Foster, and J. Parsons (2001), "The Medium Is Not the Message: Advertising Effectiveness and Content Evaluation in Print and on the Web," *Journal of Advertising Research*, 41 (4), 57–70.

Hoffman, D. and T. Novak (2000), "When Exposure-Based Advertising Stops Making Sense (and What CDNOW Did About It)," working paper, Owen Graduate School of Management, Vanderbilt University.

Ilfeld, J.S. and R.S. Winer (2002), "Generating Website Traffic," *Journal of Advertising Research*, 42 (5), 49–61.

Internet Advertising Bureau (1997), *IAB Online Advertising Effectiveness Study*, (accessed October 12, 2005), [available at http://www.intelliquest.com/resources/reports/mbi_report02.asp].

Lee, J., P. Boatwright, and W.A. Kamakura (2003), "A Bayesian Model for Prelaunch Sales Forecasting on Recorded Music," *Management Science*, 49 (2), 179–96.

Lilien, G.L., P. Kotler, and K.S. Moorthy (1992), *Marketing Models*. Englewood Cliffs, NJ: Prentice Hall.

Lodish, Leonard M., Magid Abraham, Stuart Kalmenson, Jeanne Livelsberger, Beth Lubetkin, Bruce Richardson, Mary Ellen Stevens (1995), "How T.V. Advertising Works: A Meta-Analysis of 389 Real World Split Cable T.V. Advertising Experiments," *Journal of Marketing Research*, 32 (May), 125–39.

Mela, C.F., S. Gupta, and K. Jedidi (1998), "Assessing Long Term Promotional Influences on Market Structure," *International Journal of Research in Marketing*, 15 (2), 89–107.

Meyer, B. (1990), "Unemployment Insurance and Unemployment Spells," *Econometrica*, 58 (4), 757–82.

Moe, W.W. and P.S. Fader (2003), "Dynamic Conversion Behavior at E-Commerce Sites," *Management Science*, 50 (3), 326–35.

Pechmann, C. and D.W. Stewart (1988), "Advertising Repetition: A Critical Review of Wearin and Wearout," in *Current Issues and Research in Advertising*, Vol. 11 (1–2), J.H. Leigh and C.R. Martin Jr., eds. Ann Arbor: Ross School of Business, University of Michigan, 285–329.

Pedrick, J.H. and F. Zufryden (1991), "Evaluating the Impact of Advertising Media Plans: A Model of Consumer Purchase Dynamics Using Single-Source Data," *Marketing Science*, 10 (2), 111–30.

Rossi, P.E. and G.M. Allenby (2003), "Bayesian Statistics and Marketing," *Marketing Science*, 22 (3), 304–329.

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

Sethuraman, Raj and Gerard J. Tellis (1991), "An Analysis of the Tradeoff Between Advertising and Price Discounting," *Journal of Marketing Research*, 28 (May), 160–74.

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

Silk, A.J., L.R. Klein, and E.R. Bernt (2001), "The Emerging Position of the Internet as an Advertising Medium," *Netnomics*, 3 (2), 129–48.

Song, Y. (2001), "Proof That On-line Advertising Works," Atlas Institute, (accessed October 12, 2005), [available at https://iab.net/resources/pdf/OnlineAdvertisingWorks.pdf].

Tellis, Gerard J. (1988), "Advertising Exposure, Loyalty, and Brand Purchase: A Two-Stage Model of Choice," *Journal of Marketing Research*, 25 (May), 134–44.

Tran, K.T.L. (2001), "Its Own Language: A Dictionary of Online-Marketing Jargon," *The Wall Street Journal*, (April 23), (accessed October 10, 2005), [available at http://online.wsj.com/public/us].

Unnava, H. Rao and Robert E. Burnkrant (1991), "Effects of Repeating Varied Ad Executions on Brand Name Memory," *Journal of Marketing Research*, 28 (November), 406–417.

Vakratsas, Demetrios and Tim Ambler (1999), "How Advertising Works: What Do We Really Know," *Journal of Marketing*, 63 (January), 26–43.

Warren, Susan (2001), "Winners and Losers," *The Wall Street Journal*, (April 23), (accessed October 10, 2005), [available at http://online.wsj.com/public/us].

Wedel, Michel, Wagner A. Kamakura, Wagner S. DeSarbo, and Frenkel T. Hofstede (1995), "Implications for Asymmetry, Nonproportionality, and Heterogeneity in Brand Switching from Piece-Wise Exponential Mixture Hazard Models," *Journal of Marketing Research*, 32 (November), 457–62.

Zielske, H. and W. Henry (1980), "Remembering and Forgetting Television Advertisements," *Journal of Advertising Research*, 20 (2), 7–13.

Zufryden, F. (1987), "A Model for Relating Advertising Media Exposures to Purchase Incidence Behavior Patterns," *Management Science*, 33 (10), 1253–66.