CLICKABLE WORLD WIDE WEB BANNER ADS AND CONTENT SITES

Charles F. Hofacker
Jamie Murphy

CHARLES F. HOFACKER is a Professor of Marketing in the College of Business, Florida State University, in Tallahassee. JAMIE MURPHY is a visiting Fellow in the Department of Information Management and Marketing at the University of Western Australia.

ABSTRACT

Some businesses entice visitors to their Websites by offering entertaining or useful information, or *content*. Similar to radio and television models, advertising revenue supports the firms providing this free information (Hoffman & Novak, 1996). When a visitor clicks on a content site's advertising banner, the visitor leaves the content site for the advertiser's site. While clicking on a banner can create revenue for the content site, it can also reduce additional revenue had the visitor not clicked on the ad. The webmaster therefore faces a dilemma as to how many clickable banner ads to have, or even if such ads should appear on the site. This paper explores these decisions, the optimal answers to which naturally depend on the clicking behavior of consumers arriving at the site. Our empirical results suggest that a webmaster with a single ad on a page can safely add a second banner to that page without impacting the clickthrough rate on the first ad.

JOURNAL OF INTERACTIVE MARKETING
VOLUME 14 / NUMBER 1 / WINTER 2000

Almost overnight, the Internet has spawned a marvelous variety of marketing activity. Companies marketing products or services in the physical world use the virtual world to directly sell, provide product information, project a corporate image, generate qualified leads, dispense electronic coupons, and handle a wide range of post-purchase tasks (Berthon, Pitt, & Watson 1996).

In this paper, we focus on companies whose product or service is predominantly virtual. These companies use computer-mediated communication to create and display *content*; a phenomenon somewhat analogous to broadcast media such as radio or TV in that content is free and hosting paid advertising provides revenue (Hoffman & Novak, 1996). Although Novak and Hoffman (1996) have proposed additional revenue mechanisms, content sites usually generate revenue in one of two ways. The advertiser pays by the impression, which by analogy to other media we refer to as *cost per thousand (cpm) pricing*, or the advertiser pays for each time a content site visitor clicks on the ad and subsequently visits the advertiser's site. We refer to this payment mode as *clickthrough pricing*. Of course, combinations of the two are also possible.

Clickthrough pricing creates an unusual dilemma for the content site manager. In effect, management profits by losing its visitor. By analogy, advertisers would pay TV stations when the viewer switched channels. A clickthrough may provide advertising revenue, but it may preclude additional advertising revenue if the visitor does not return to view additional ads. This paper presents a model for two simple content sites and offers empirical evidence as to how the content site manager should deal with this situation.

It is perhaps not surprising that the World Wide Web presents a conundrum unique in the field of advertising decision-making. A number of authors, including Turkle (1995), Stewart and Ward (1994), and Novak and Hoffman (1996), argue that existing approaches may not generalize well from old media to this new medium.

Advertising Banners

Companies use traditional media to drive traffic to their site by incorporating their Website address (URL or Uniform Resource Locator) in print, television, outdoor and radio ads. Advertising in traditional media, however, promotes an online site via offline media to an offline audience. The Web drives online consumers to a particular online Website. Online promotion techniques include search engine registration, including the URL in email signature files, seeking and obtaining hypertext links with relevant external Websites, and purchasing the banner advertisements that create the managerial dilemma addressed in this paper.

Issues that advertisers, who want to maximize their Web advertising dollars, might contemplate include defining the communication purpose or goal for the banner, choosing between generating many clicks with an untargeted audience or less clicks with a targeted audience, selecting the host site type and content, and deciding between cpm, clickthrough, or other types of pricing.

Advertising banners themselves are small, generally rectangular displays on a webpage ranging in size from 120 to 500 pixels wide by 45 to 120 pixels high (Novak & Hoffman, 1996). Web banner advertising revenue projections for the year 2000 range from $1.7 to $5 billion (Murphy & Forrest, 1996).

Interestingly, a large number of sponsored content sites both buy and sell Web ad banners. For example, Yahoo! buys banners to drive consumers from other sites to its site. Once these consumers arrive, Yahoo! then tries to send them off to still other sites to generate positive clickthrough banner revenue of its own. We focus on the tension that underlies this dual role in our model development later in the paper. A sponsored content site's webmaster wants to keep visitors so as to maintain high traffic on as many pages as possible. Yet cashing in on that traffic by sending visitors off site, out into the virtual world, motivates that same webmaster to sell clickable banner ads. The webmaster's banner ad sales decision is an important and a subtle one. Fortunately, a webserver

logs every click visitors make to a site, helping the webmaster decide to sell by the click or sell by the impression.

Practitioner analysis of webserver log data suggests that about 2 out of every 100 visitors actually click on an ad banner (DoubleClick and I/PRO, 1996) but clickthrough rates vary greatly. For example, DoubleClick and I/PRO observed that ads with animation have a higher average clickthrough rate than ads without animation. We note from the academic literature that Hofacker and Murphy (1998) used a randomized Web experiment to demonstrate that efficacious ad copy could improve the click response to banner ads and that Drèze and Zufreyden (1997) found that background, image size, use of java, and operating system influenced clicking rates on a webpage.

In addition to these results, we assume that the webpage's context drastically impacts clickthrough behavior. Since the Web is primarily a *pull* medium as compared to a *push* medium, the page's content powerfully self-selects visitors to that page. Beyond that, it seems likely to us that the nature and number of other links on the page will influence the probability of a clickthrough on a particular link, like an ad banner. Conversely, the presence of a clickable ad will tend to influence the visitor's behavior towards the other links on a page.

The goal of a clickable ad banner is to move the consumer through a compressed hierarchy of effects (HOE), ending with clicking behavior. Our conceptualization of the classic HOE theory (Barry & Howard, 1990), as observed on the Web, resembles that of Robertson (1976) in his discussion of the "low commitment" audience. While Robertson dealt with a passive TV audience, he refers to a "collapsed" hierarchy consisting of two effects: awareness and trial. When the banner is successful, there may be two or three stages: the consumer becomes aware of the banner, experiences interest or curiosity about the sponsoring site, and then clicks on the banner.

Human information processing brings key limitations to this rapid sequence (Bettman, 1979), not the least of which is a limited capacity to attend to and focus on links, banners, and other information on a page. Further, the limited capacity of short-term memory can subvert the desire to click on a link. The limitations of our short-term memories makes it quite possible to find oneself lost in a series of links that traverse sites and multiple levels within sites. Visitors may forget wanting to visit a link or how to return to a desired link, or even seeing that link. The upshot is that the specific environment created by a page or a site may have an impact on the clicking probability of any item on that page. The more links on a page, the better the capacity of that page to lose or distract visitors clicking on those links.

The interaction of two or more links can be seen as the impact of a visitor contemplating a competitive set. We would therefore emphasize that an ad banner cannot be looked at in isolation. The banner's environment, specifically the set of other links or banners on the hosting page competing for clicks, affects the banner's performance. The Web user ventures into a constantly changing collection of link choice sets. The binary outcome of clicking or not clicking is reminiscent of discrete choice processes associated with branded goods retailing. The price paid for, and the benefit accrued from, clicking on a link, compares on a reduced scale to low-involvement convenience goods. The "price" is paid in opportunity cost, wasted time, connect time charges, and frustration.

On one hand, we assume that many content site visitors are hedonically consuming (Holbrook & Hirschman, 1982) the site with essentially no plan of action. The perceived or expected utility associated with a given link—that is, a link that has been attended to but not clicked on—is no doubt highly stochastic with high variance across short time intervals and possesses modest valences. There may be minimal systematic variability of utility from one attended link to another as compared to the random component. On the other hand, some visitors engage in goal-directed, utilitarian seeking (Babin, Darden, & Griffin, 1994), the converse of hedonic Website consumption. These individuals would have much less variance for the same link over a short time period and much more systematic variance between inter-

esting and less interesting links. We conceptualize these utilitarian individuals as foraging in an information environment, expending time on the links most likely to yield useful informational nuggets (Pirolli & Card, 1995). For these visitors, the information-processing effort (e.g., Lohse & Johnson, 1996) is more thoughtfully traded off against possibly rewarding discoveries at the other end of the banner hyperlink. But for both the hedonic surfer and the utilitarian searcher, the set of links on a page can be thought of as a competitive set whose members pull at visitors with a variety of strengths.

We now explore ad decisions for two basic link environments or Web spaces. In the first, the webmaster manages a single webpage with a preexisting banner ad. The decision is whether or not to add a second ad banner to the page. This decision is a difficult one since revenue depends on the click choice probabilities for both ads. In the second link environment we assume a Website with two cross-linked pages. One page has an ad, and the decision is whether or not to add an ad on the second page. While the situations we describe are relatively simple, they offer a beginning for modeling Web clicking behavior.

A MODEL FOR ONE PAGE WITH ONE OR TWO BANNERS

The general research question that concerns us is as follows: Is it possible to know a priori how visitors to a page will behave if links on the page change in certain systematic ways? In what follows we assume a homogeneous population of visitors arriving at the site at a steady rate. To specify the situation further, consider a single webpage presenting entertaining content that contains a single banner ad, called banner 1. That ad generates revenue r_1 each time a visitor clicks on it, with the probability p_1 that a visitor clicks on this banner. Revenue is therefore proportional to r_1p_1, the constant of proportionality equal to the rate of traffic arriving at the site. Assume further that the webmaster faces the decision to add or not add a second banner paying r_2 with clickthrough probability p_2^*. We use an asterisk in our notation to refer to clickthrough probabilities of pages that have been changed. In this case the change stems from the addition of the second banner.

We cannot in general assume that the clickthrough rate for the first ad will remain p_1 after the addition of another banner changes the page. Instead, we propose the following axiom:

> *Axiom 1.* The probability of clicking on a link will either decrease or stay the same after adding an additional link to the same page.

According to *Axiom 1*, the new clickthrough probability is $p_i^* \leq p_i$.

This hyperlink axiom is inspired by Bell, Keeney, and Little (1975) and the well-known Luce (1959) choice axiom. In *Axiom 1*, we are making explicit our description of the clickable links on a page as comprising a competitive set from which a visitor chooses zero, one, or more links to visit.

Now the webmaster should accept the second banner if

$$p_1^* r_1 + p_2^* r_2 > p_1 r_1 \tag{1a}$$

$$p_1 - p_1^* < p_2^*(r_2/r_1) \tag{1b}$$

These equations specify the relationship between revenue and clicking probabilities. The left-hand side of equation (1a) represents the revenue expectation with two ads while the right-hand side shows the revenue with a single ad. If $r_1 = r_2$, then the webmaster should accept the second banner if the clickthrough rate of the new banner is greater than the reduction in clickthroughs for the first banner. We suppose that this is likely and express this supposition as *Axiom 2*:

> *Axiom 2.* The total number of links visited from a given page will either increase or stay the same after adding an additional link to the page.

Consequently we have $p_i^* + p_j^* \geq p_i$.

For a one-page site, impression pricing means that the site generates revenue simply from the number of visitors who arrive and the number of ads on the page. The various clicking

probabilities of the ads are irrelevant and revenue is proportional to visitor frequency. If impression pricing is in effect already, and the advertiser wants to negotiate clickthrough pricing, the webmaster would be in a position to compare revenue expectations using the clickthrough formulas $p_1^* r_1 + p_2^* r_2$ in the case of two ads or simply $p_1 r_1$ in the case of one ad.

We started this section discussing a page with a single ad. Assuming the webmaster has diligently studied the log files, the value p_1, which is the clickthrough rate for the original ad, should be known. It would be a fairly simple matter to also determine p_2, the clickthrough rate for the second ad, by substituting ad 2 for ad 1 for a time on the page. In that case we may well ask what we can say a priori about p_1^* and p_2^* from knowing p_1 and p_2?

> *Axiom 3.* If two ads which are alternated on a page have clickthrough rates $p_i > p_j$, when they are combined on that page they will have clickthrough rates $p_i^* > p_j^*$.

As is the case with *Axiom 1*, this axiom can be derived from the Luce choice axiom but is a weaker statement than the Luce axiom. Note that to test this axiom fairly, the positions of the two ads when they are combined on a page must be experimentally balanced.

What happens when a preexisting banner is replaced with a new banner? Assume that a page containing banner i is replaced with a new version indicated with an asterisk such that the clickthrough probability $p_i^* > p_i$. We propose that the other unchanged links on the page will behave as follows:

> *Axiom 4.* The probability of clicking on a given link will either decrease or stay the same after the clickthrough probability for some other preexisting link on the same page is made to increase.

Axiom 4 is analogous to *Axiom 1* except that here we have described changing a preexisting link rather than adding an additional link. Analogous to *Axiom 2* we have

> *Axiom 5.* The total number of links visited from a given page will either increase or stay the same after the clickthrough probability for some preexisting link on that page is made to increase.

We now turn to two simple experiments that test the first four of these axioms.

Empirical Results

We selected two banners from Yahoo! (www.yahoo.com), a popular content site, for the experiment. One banner was for Saturn automobiles, and the other advertised a site called "Women's Wire," which offers women's content. Both banners were 460 by 55 pixels, and both presented two scenes in succession, a technique achieved using the animated GIF format. We chose a shopping-oriented webpage, which has been on the Web since the fall of 1994, as the home page for these banners.

The page, which contains three direct-marketing Website indices, receives approximately 1,000 visitors per week. This page contains four links, a three-item main menu and a clickable company logo. Four versions of the page were created: one with the Saturn banner at the top of the page, one with the Women's Wire banner at the top, and two versions with both banners. One of the two-banner versions had the Saturn ad on top with the second ad just below the first, while the other page reversed the layout putting the Women's Wire ad on top. The UNIX cron facility, which is used to automatically run computer programs, changed the live page, substituting a different experimental condition every 15 minutes. Thus visitors might find themselves in any one of the four experimental conditions depending on when they retrieved the page. A single pixel transparent gif file with a unique name was added to each experimental page to create a foolproof method to match visitors' records to their experimental condition.

The analyst of Web log data must make certain decisions due to the peculiarities of the HyperText Transfer Protocol (HTTP). For one, we limited our analysis to HTTP GET requests where the return code was 200, an indication of a successful operation (Stout, 1997). A second consideration is that visitors do not always retrieve images, including banners, that are em-

bedded in a webpage. We limited our analyses to those visitors who retrieved the banner or banners implied by their experimental condition. Identifying unique webpage visitors is another complication. As the HTTP is a connectionless protocol, the server maintains a record of domain names and IP addresses for GET requests but keeps no unique visitor identification per se. An HTTP feature known as "cookies" can identify unique computers, but cookies are controversial and not widely applied at this time. We defined a visitor as one or more log records from the same IP address within a 30-min period (Novak & Hoffman, 1996).

The study ran during 4 weeks late in the spring of 1997. During that time period, 3,335 visitors, as defined above, retrieved the page. As the Saturn ad had the extremely low click-through rate of .0132 (1.32%), the study was replicated during another 3-week period with a replacement banner of the same size for a personal services company called "Single Search." The second study saw 2,629 visitors. Results from the two studies appear in Tables 1 and 2, respectively.

Axiom 1 states that the number of clicks on a banner should decrease or remain invariant when a second banner is added to the page. We therefore tested the hypothesis:

$$H_0: p_i^* \leq p_i$$

$$p_i^* - p_i \leq 0,$$

where the asterisk again signifies the click-through probability after adding an additional banner. The axiom and the hypothesis predict a negative or zero Z score. Since both p_i^* and p_i come from independent samples, we used a Z score for this test of the difference between two independent binomial populations[1].

TABLE 1
Study One

Condition	*Saturn Banner*	*Women Banner*	*N*
Saturn Only	.0132	—	836
Women Only	—	.0531	772
Women/Sat	.0075	.0574	802
Sat/Women	.0119	.0670	925

TABLE 2
Study Two

Condition	*Singles Banner*	*Women Banner*	*N*
Singles Only	.1330	—	722
Women Only	—	.0630	635
Women/Singles	.1207	.0620	613
Singles/Women	.0956	.0668	659

The estimates for the p_i^* were pooled across two experimental groups: one in which banner i appeared on top and one where it was on the bottom. We failed to reject the hypothesis of *Axiom 1* for both ads in the first study (Saturn, $Z = -.74$, $p = .77$; Women, $Z = .94$, $p = .17$) and for both in the second study (Singles, $Z = -4.16$, $p = .99$; Women, $Z = .13$, $p = .45$) although there was weak (but nonsignificant) evidence in both samples that the Women's Wire ad actually became more popular when paired with another ad.

Axiom 2 holds that the combined number of clicks for two banners is greater than or equal to the number of clicks on either banner when they appear by themselves on the same page. The hypothesis in this case is that

$$H_0: p_i^* + p_j^* \geq p_i$$

$$p_i^* + p_j^* - p_i \geq 0 \text{ for } i \neq j.$$

For both banners in both studies, the Z score[2] was in the correct direction (positive) predicted

[1] The Z test for the difference between proportions from two independent populations takes on the form:

$$Z = p_i^* - p_i / \sqrt{[V(p_i^*) + V(p_i)]},$$

where $V(p_i) = p_i(1 - p_i)/N$ and a similar statement holds for $V(p_i^*)$, which comes from a second independent sample.

[2] In this case, the values p_i^* and p_j^* come from the same sample and so cannot be considered independent for the purposes of calculating the variance of $p_i^* + p_j^* - p_i$. That variance is needed for the

by Axiom 2 (Study 1, Saturn, $Z = 7.78$, $p = .99$; Study 1, Women, $Z = 1.85$, $p = .97$; Study 2, Singles, $Z = 2.27$, $p = .99$; Study 2, Women, $Z = 7.18$, $p = .99$). We therefore failed to reject the hypothesis embodied by *Axiom 2* in all four cases and in fact confirmed it.

Looking at the data, we were tempted to conclude post hoc that *Axiom 1* holds as an equality. Therefore the combined clickthroughs when two banners share a page can be predicted by the sum of the individual clickthroughs when the banners appear by themselves on that page. We looked at this possibility, post hoc, testing the hypothesis that

$$H_0: p_i^* + p_j^* = p_i + p_j$$

and found that for the first study we could not reject the hypothesis ($Z = .54$, $p = .58$), nor for the second ($Z = -1.20$, $p = .23$).[3] To a first approximation, banner clickthroughs were simply additive on this page. We refer to this post hoc conclusion as the *simple additivity* hypothesis.

In the preceding post hoc analysis it appeared that the marginal clickthrough probabilities act very similarly whether the banner is alone or combined on a page. But by looking at additional information, namely the joint clicking probabilities where both banners were present on the page, we found for the first study that the null hypothesis of independence could not be rejected, with a chi-square on one degree of freedom of 1.136, $p = .248$. For the second study, however, the hypothesis of independence was rejected with a chi-square of 20.084, $p = .001$. It is tempting to conclude that the lack of similarity between an automotive product and a women's service led to independence in the first case and that an overlap of similarity led to dependence in the second. Such a result suggests that clickthroughs, which can be thought of as pick-any choice data, are subject to violations of the Independence of Irrelevant Alternatives property (Currim, 1982).

Axiom 3 claims that if a banner tested by itself on a certain page outperforms a second banner, it will also outperform that banner when combined with it on a page. In the first study, the Women's Wire banner had a significantly higher click rate than the Saturn banner, $Z = 4.44$, $p = .99$. This also proved to be the case when those two banners were combined on the same page, $Z = 8.68$, $p = .99$. Similarly, in the second study the Singles banner outperformed the Women's Wire banner separately ($Z = 4.40$, $p = .99$) and together ($Z = 4.16$, $p = .99$). The Z scores calculated for the combined banner conditions took into account the covariance between the proportions as described above.

Axiom 4 implies that by improving the clickthrough probability of a banner, the clickthrough probabilities of other links will be nonincreasing. We tested this in Study 1 by comparing two groups, the group that saw the Women's Wire banner and the group that saw the less popular Saturn banner. Recall that the clickthrough rates on these two banners were significantly different. As has been described, the experimental page contained four additional links. For each of these four links, we compared the Saturn-only and Women's Wire-only groups to see if there was an association between the banner presented and the probability of clicking on the link. In effect, we created a series of two by two contingency tables each with two rows; the Saturn-only group and the Women's-only group; and two columns that

---

denominator of the Z score to test the hypothesis. That Z is given by

$$Z = p_i^* + p_j^* - p_i / \sqrt{[V(p_i^*) + V(p_j^*) + 2\,\mathrm{COV}(p_i^*, p_j^*) + V(p_i)]}.$$

The covariance between two proportions is provided by Christofferson (1975, Appendix 2) as

$$\mathrm{COV}(p_i^*, p_j^*) = [p_{ij}^* - p_i^* p_j^*]/N,$$

where p_{ij}^* is the joint probability that a visitor clicks on banners i and j. This joint probability was .0029 in the first study and .0165 in the second.

[3] Here we used the formula

$$Z = p_i^* + p_j^* - p_i - p_j / \sqrt{[V(p_i^*) + V(p_j^*) + 2\,\mathrm{COV}(p_i^*, p_j^*) + V(p_i) + V(p_j)]}$$

and used a two-tailed test since the hypothesis was an equality.

tabulated how many visitors clicked on a particular additional link and how many did not. *Axiom 4* would predict that the Women's Wire group would be associated with the same or reduced clicking on the additional link. There was no association for any of the four links, nor was there any analogous association in Study 2 when we compared the Singles-only and Women's Wire-only banner groups. Thus replacing a poorly performing banner with a better performing banner did not impact clicking on the other links on the page. The simple additivity hypothesis held for *Axiom 4* much the same way it held for *Axiom 1*: A link in a modified page context, in this case provided by a improved banner, behaved about the same as it had previously.

We next turn our attention to a slightly more complex Website; namely a two-page site with revenue possibilities on each page.

A MODEL FOR TWO PAGES

Consider a Website with two cross-linked pages; h_1 and h_2. Revenue from page h_1 accrues either deterministically or probabilistically as visitors land on h_1, but page h_2 has no ad and hence accrues no revenue. Revenue sources for h_1 include either a banner ad or a selling opportunity. The revenue outcome of a hit on h_1 is r_1. The value r_1 is deterministic in the case of an ad banner and impression-based pricing. In the case of clickthrough pricing, r_1 represents expected revenue based on the clickthrough rate of the banner appearing on page h_1. For a Website employed in direct marketing, r_1 reflects an expectation based on the percentage of visits converted to purchase events. In all three cases, the HTTP log sequences and their (expected) payouts are:

Page Hits	Payout
h_1	r_1
$h_1,\ h_2$	r_1
h_2	0
$h_2,\ h_1$	r_1.

Define $\Pr(h_i)$ as the proportion of site visits beginning with page h_i. Also define $\Pr[h_j \mid h_i]$ as the proportion of visitors who click on the link to page j, given that they began their visit to the site on page i. There are two ways for visitors to generate revenue for the webmaster; they can begin their site visit on page 1 or get to page 1 from page 2. The two components in the curly brackets in the revenue function below reflect these two ways:

$$f = \{\Pr(h_1) + \Pr(h_2)\Pr[h_1|h_2]\}r_1, \qquad (2)$$

if we add the second revenue source to page h_2, what will happen to total revenue? Enumerating the log sequences and their new payouts, we have:

Page Hits	Payout
h_1	r_1
$h_1,\ h_2$	$r_1 + r_2$
h_2	r_2
$h_2,\ h_1$	$r_1 + r_2$

which makes it clear that there are more payout possibilities than before. The new revenue function is

$$f^* = \Pr(h_1)r_1 + \Pr(h_2)r_2 + \Pr(h_1)\Pr[h_2|h_1]r_1 + \Pr(h_2)\Pr^*[h_1|h_2]r_2. \qquad (3)$$

But *Axiom 1* suggests a danger since $\Pr^*[h_1 \mid h_2] \leq \Pr[h_1 \mid h_2]$ because some other link has been added to page 2. The second revenue source should be added if

$$f^* - f > 0 \qquad (4)$$

so substituting (2) and (3) into (4) and canceling terms leads to

$$\Pr(h_2)r_2 + \Pr(h_1)\Pr[h_2|h_1]r_2 + \Pr(h_2)\Pr^*[h_1|h_2]r_1 - \Pr(h_2)\Pr[h_1|h_2]r_1 > 0$$

$$r_2\{\Pr(h_2) + \Pr(h_1)\Pr[h_2|h_1]\} > r_1\Pr(h_2)\{\Pr[h_1|h_2] - \Pr^*[h_1|h_2]\} \qquad (5)$$

According to (5), you compare the left hand side, which is the increase in revenue due to the new ad on page 2, with the right hand side,

which is the loss of revenue due to the Axiom 1 effect operating on page 2.

The new banner on h_2 is now competing with the cross-link to h_1. We looked at our data to see whether, and by how much, cross-linking might be reduced on our experimental page in Study 1 by comparing the group that saw only the Saturn ad to groups that saw both the Saturn and Women's Wire ads. We performed four chi square tests of independence in a two-way table to see if those who saw the additional banner were less likely to click on each of the four links to other parts of the site.

In all four cases we failed to reject the hypothesis of no association, and that was also true for Study 2 when we compared the group that saw the Women's Wire banner to the groups that saw both banners. Once again *Axiom 1* held as an equality and no penalty was paid for the additional banner.

CONCLUSIONS AND SUGGESTIONS FOR FURTHER RESEARCH

Context Effects

In a word, the results surprised us. Specifically, the finding of simple additivity—under two separate banner conditions—was not expected. While we have some doubts about the viability of that hypothesis, it is tempting to propose that the more general Luce (1959) choice axiom could hold in the milieu of the Web. In our notation, we would have

$$\frac{p_i}{p_j} = \frac{p_i^*}{p_j^*}.$$

Our studies used a simple page with a total of four links in addition to one or two banners. Results might have been different in a more complex page environment. As mentioned above, we are suspicious of the additivity of clickthroughs exhibited by the two sets of two banners. Would clicks keep simply adding for the webmaster if there were 4, 8, or 16 additional links?

In other words, if I have a revenue-producing link and 50% of my visitors click on it, can I get away with selling multiple banners on the page without reducing the percentage of clicking visitors to below 50%? We uncovered no evidence to the contrary after one additional banner, but the limited capacity of the human information processor suggests otherwise (e.g., Lohse & Johnson, 1996). It would be a simple, if slightly cumbersome, experimental procedure to determine how long *Axiom 1* would hold as a strict equality. Surely at some point a page would be so cluttered with links that the probability of clicking on some link on the page would dip. While the exact point at which that would happen awaits future discovery, for the time being, our advice to the webmaster would be to include as many advertising banners on a page as will fit.

Web Research Generalizations

The computer-mediated nature of the Web makes harvesting large sample sizes easy for webpage market response research, given that the webmaster has the requisite skills and available time. If so, the results can then be generalized to (potential) visitors to those pages. What is more difficult is to generalize to the universe of possible pages. It would be useful to explore different types of Websites including sites featuring different content topics, direct shopping sites, search engines and other network utilities, or corporate promotional sites. Then, these early results might be applicable across a variety of Websites. This generalizability would be useful to a webmaster who lacks the time and skills to perform these experiments themselves.

Visitors to any given page self-select and so we take it as given that different pages have different audiences. We assume that the behavior of any individual link on a page will not be reproduced exactly when that element is transferred to another page and consequently viewed by a different visitor population. Of course content sites such as Jumbo leverage this sort of whole-page context effect. The Jumbo site offers over one hundred thousand webpages featuring free software. As the user navigates his way deeper into the site, the ads become more targeted.

Further research could look at the interactions between targeted ads and different hosting pages.

Modeling Considerations

The ease with which data can be collected on Web choice processes makes it tempting to ignore marketing theory. We strongly believe that a strictly empirical approach will shortchange both practitioners and academics. The role of theory has long been assigned a high importance by marketing scientists (see, e.g., Hunt, 1983; Naert & Leeflang, 1978). We believe that in the long run, theory leads to better decisions, improved understanding, and a superior framework for discussion. Further, ultimately only theory can suggest what experiments and conditions it makes most sense to test.

In this paper we suggested an analogy between visits to a World Wide Web site and a supermarket. Gingerly proceeding with this analogy, the clicking behavior of both the hedonic surfer and the utilitarian searcher leaves behind a pick-any data set (Holbrook, Moore, & Winer, 1982) in the content site's server records. At this early date it is unclear whether it might be best to model individuals or to model links (DeSarbo & Hoffman, 1987). But the analogy makes us wonder whether other tools used in stochastic choice modeling might not be applicable to server log data.

Drèze and Zufryden (1997) modeled clicking incidence using the Poisson distribution, sometimes used for purchase incidence. It makes sense to us to think of visitor navigation through the pages of a Website as the realization of stochastic transition matrix (see, e.g., Lilien, Kotler, & Moorthy, 1992, Ch. 2). Unfortunately for the webmaster, leaving the site is the absorbing state, the condition all visitors ultimately end up in.

The Advertiser's Decisions

The decision to add a second banner to a page, assuming that the first and second banner generate revenue by the clickthrough, is a simple one for the advertiser if not for the webmaster. The advertiser of either the original or second banner will be indifferent to the relative click probabilities of the two banners and how they change as a result of page structure. The advertiser in this case pays only for visitors who click through the banner regardless of the rate at which that occurs.

If the advertiser, however, pays for impressions, issues of the visitor's limited attention capacity come into play and the advertiser might prefer to be the sole sponsor of a page. The advertiser should be similarly motivated to be sole sponsor if they pay a flat fee for an appearance on the page. If the ad is clickable, at best *Axiom 1* implies no better results for the same flat fee. If the ad is not clickable, the advertiser again runs into considerations of limited attention.

Ad banner frequency and placement—not to mention executional elements such as size, copy, animation and color—depend upon the banner's objective. Ad banners designed for impressions benefit from repeated exposures, while banners designed for clicking seem to suffer from repeated exposures according to practitioners (DoubleClick and I/PRO, 1996; Modahl & MacQuiddy, 1996). Yet a clickable banner also produces an impression of corporate logos, slogans, or messages. Novak and Hoffman (1996) have argued for transaction-based metrics of banner effectiveness. Surely the clickthrough is not the end all and be all; further research is needed to show why people click on a banner, how additional banners or links affect clickthroughs, and how those that clicked navigate the advertiser's site.

REFERENCES

Babin, B.J., Darden, W.R., & Griffin, M. (1994). Work and/or Fun: Measuring Hedonic and Utilitarian Shopping Value. Journal of Consumer Research, 20 (March), 644–656.

Barry, T.E., & Howard, D.J. (1990). A Review and Critique of the Hierarchy of Effects in Advertising. International Journal of Advertising, 9, 121–135.

Bell, D.E., Keeney, R.L., & Little, J.D.C. (1975). A Market Share Theorem. Journal of Marketing Research, 12, 136–141.

Berthon, P., Pitt, L.F., & Watson, R.T. (1996). The World Wide Web as an Advertising Medium: Toward

an Understanding of Conversion Efficiency. Journal of Advertising Research, 36 (1), 43–53.

Bettman, J.R. (1979). An Information Processing Theory of Consumer Choice. Reading, MA: Addison–Wesley.

Christofferson, A. (1975). Factor Analysis of Dichotomized Variables. Psychometrika, 40 (March), 5–32.

Currim, I.S. (1982). Predictive Testing of Consumer Choice Models Not Subject to Independence of Irrelevant Alternatives. Journal of Marketing Research, 19 (May), 208–222.

DeSarbo, W.S., & Hoffman, D.L. (1987). Constructing MDS Joint Spaces from Binary Choice Data: A Multidimensional Unfolding Threshold Model for Marketing Research. Journal of Marketing Research, 24 (February), 40–54.

DoubleClick and I/PRO. (1996, October 10). Banner Effectiveness Tips. Web in Perspective: Ad Response research report.

Drèze, X., & Zufryden, F. (1997). Testing Web Site Design and Promotional Content. Journal of Advertising Research, March, 77–91.

Hofacker, C.F., & Murphy, J. (1998). World Wide Web Banner Ad Copy Testing. European Journal of Marketing, 32 (7/8), 703–712.

Hoffman, D.L., & Novak, T.P. (1996). Marketing in Hypermedia Computer-Mediated Environments: Conceptual Foundations. Journal of Marketing, 60 (July), 50–68.

Holbrook, M.B., & Hirschman, E.C. (1982). The Experiential Aspects of Consumption: Consumer Fantasies, Feelings, and Fun. Journal of Consumer Research, 9 (September), 132–140.

Holbrook, M.B., Moore, W.L., & Winer, R.S. (1982). Constructing Joint Spaces from Pick-Any Data: A New Tool for Consumer Analysis. Journal of Consumer Research, 9 (June), 99–105.

Hunt, S.D. (1983). Marketing Theory: The Philosophy of Marketing Science. Homewood, IL: Richard Irwin.

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

Lohse, G.L., & Johnson, E.J. (1996). A Comparison of Two Process Tracing Methods for Choice Tasks. Organizational Behavior and Human Decision Processes, 68 (1), 28–43.

Luce, R.D. (1959). Individual Choice Behavior: A Theoretical Analysis. New York: Wiley.

Modahl, M.A., & MacQuiddy, R. (1996). Internet Advertising. The Forrester Report, Media and Technology Strategies (1)(September) 1, corporate research report.

Murphy, J., & Forrest, F. (1996, May 26). Hits, Views, Clicks and Visits: Web Advertisers Face Data Jungle. New York Times on the Web.

Naert, P.A., & Leeflang, P.S.H. (1978). Building Implementable Marketing Models. Boston: Leiden.

Novak, T.P., & Hoffman, D.L. (1996). New Metrics for New Media: Toward the Development of Web Measurement Standards. Project 2000 White Paper, http://www2000.ogsm.vanderbilt.edu/novak/web.standards/webstand.html.

Pirolli, P., & Card, S. (1995). Information Foraging in Information Access Environments. SIGCHI Proceedings 1995, http://www.acm.org/sigchi/chi95/proceedings/papers/ppp_bdy.htm.

Robertson, T.S. (1976). Low Commitment Consumer Behavior. Journal of Advertising Research, 16 (April), 19–27.

Stewart, D.W., & Ward, S. (1994). Media Effects on Advertising. In J. Bryant & D. Zillmann (Eds.), Media Effects: Advances in Theory and Research. Hillsdale, NJ: Lawrence Erlbaum Associates.

Stout, R. (1997). Web Site Stats: Tracking Hits and Analyzing Traffic. Berkeley: Osborne McGraw–Hill.

Turkle, S. (1995). Life on the Screen: Identity in the Age of the Internet. New York: Simon and Schuster.