

Internet Advertising: Is Anybody Watching?

Xavier Drèze
University of Southern California

François-Xavier Husserr*
Ecole Nationale Supérieure des Télécommunications

August 1999

This is a Draft, do not circulate or quote without prior consent from authors.

*Xavier Drèze is an Assistant Professor at the Marshall School of Business, University of Southern California, Los Angeles, CA 90089 (xdm@sbaxdm.usc.edu). François-Xavier Husserr is a researcher at the Centre Nationale d'Etude des Télécommunications (CNET) and at the Ecole National Supérieure de Télécommunication (ENST), Paris, France (francois.husserr@francetelecom.fr).

The authors would like to thank Voilà and France Telecom for funding this research and Philippe Taupin for his help with the Eye-Tracker device.

Internet Advertising: Is Anybody Watching?

Abstract

Click-through rates have emerged as the de facto measure of Internet advertising effectiveness. Unfortunately, click-through rates are plummeting. This decline prompts four critical questions: (1) why do banner ads seem to be ineffective; (2) what can advertisers do to improve their effectiveness; (3) does an immediate measure such as click-through rate undervalue online advertising; and, (4) are memory-based measures such as recall or awareness more appropriate? To address these questions, we utilized an eye-tracking device to investigate online surfers' attention to online advertising. Then we conducted a large-scale survey of Internet users' recall, recognition, and awareness of banner advertising.

Our research suggests that the reason why click-through rates are low is that surfers actually avoid looking at banner ads during their online activities. This suggests that the larger part of a surfer's processing of banners will be done at the pre-attentive level. If such is the case, click-through rate is an ineffective measure of banner ad performance. Our research also shows that banner ads do have an impact on traditional memory-based measure of effectiveness. Thus, we claim that advertisers should rely more on traditional brand equity measures such as brand awareness and advertising recall.

Finally, our study shows that although repetition leads to lower click-through rates, it has a beneficial impact on brand awareness and advertising recall.

Keywords: Internet, Advertising, Awareness, Brand Equity, Eye-tracking, Click-through.

Introduction

As the Internet matures into a viable commercial medium, many web sites (e.g., Lycos, Go Network, Yahoo!) rely on advertising to finance their operations. The lure of advertising is such that some companies provide users with free Internet access (e.g., NetZero.com, FreeI.com) and even free computers (e.g., Free-PC.com) in exchange for their eyeballs (Berst 1999). This should not come as a surprise as advertisers have long used every conceivable vehicle to display their messages in front of the gazing eyes of potential customers, be it magazines, television, or racecars.

As the Internet becomes more mainstream, many companies are budgeting significant dollar amounts for online advertising. The Internet Advertising Bureau (1999) reports 1998 online advertising expenditure of \$1.92 billion, more than double 1997 revenues. The bulk of this expenditure is allocated to banner ads. Banner ads typically consist of rectangular images displayed at the top of web pages and contain the message that the advertiser wants to send to web surfers.

The most widely used measure of online advertising effectiveness is the percentage of the total number of ad exposures that induce the surfer to actually click on the banner in response to the advertised message. This measure is known as the click-through rate (Novak and Hoffman 1997). Click-through rate has become such a dominant measure that in 1996 *Procter and Gamble* made a deal with *Yahoo!* in which P&G would pay only for click-throughs and not for exposures (Associated Press 1996). The ability of a site to generate click-throughs also affects the advertising rates it can command (Hamilton 1998).

Click-through rates started in 1996 at around seven percent. However, they have declined steadily to around 0.6% in 1999 (Nielsen//Netratings 1999). This is problematic

because advertisers typically do not knowingly allocate budgets to media that are not effective. Does this trend suggest that the online advertising community is going to fritter along with the decline in click-through rates? Should one sell his or her stocks in *Yahoo!*? Not necessarily. Many authors (Ambler 1998, Batra, Lehmann, Burke, and Pae 1995) argue that good advertising affects long-term brand equity, not necessarily short-term sales. They contend that equity variables such as brand or ad awareness are better gauges of advertising effectiveness. In this spirit, Briggs and Hollis (1997) have shown using Milward Brown's Brand Dynamics™ system (Dyson, Farr, and Hollis, 1996), that banner ads can have an impact on consumers' attitudes toward a brand independent of click-through.

Briggs and Hollis' study, combined with the decline in click-through rate, begs the questions (1) why do banner ads seem to be ineffective; (2) what can advertisers do to improve their effectiveness; (3) does an immediate measure such as click-through rate under-value online advertising; and, (4) are more traditional measures such as recall or awareness more appropriate? The purpose of this paper is to answer these four questions. We intend to show that because banner ads operate mostly at the pre-attentive processing level (Shapiro, MacInnis, and Hoyer 1997), traditional effectiveness measures are more appropriate than click-through rates. We will then use these measures to study some of the factors that might impact banner ad effectiveness.

The remainder of the paper is organized as follows. The first section discusses the results of a study which utilized an eye-tracking device to determine whether web surfers see banner ads and which factors increase or decrease the probability that a banner ad is seen. We use the results from this first study to generate hypotheses about the characteristics of banner ads that might increase or decrease viewers' attention. The following section relates the results of the follow-up study that tested the hypotheses generated in the first study on a broader sample of

web surfers (807 respondents). The study also explored the effects of Internet advertising on recall, recognition, and awareness. We then consider the results of both studies and discuss their managerial relevance. Finally, we close with concluding remarks, a discussion of the limitations of our methodology and results, and directions for future research.

Study 1: Eye-tracking

The Internet differs from traditional media in at least one significant way. When an advertiser uses Television or Radio to deliver his messages, he preempts the program being broadcast (e.g., a sitcom or song) and uses all the bandwidth of the medium to transmit his message. This means that by default, the viewer or listener is paying attention to the advertisers, and the message is only interrupted if the listener zaps away. Zapping, however, is quite infrequent. Siddarth (1999) reports zapping rates for commercials of less than 3%. By contrast, online banner ads share their bandwidth with other elements of the pages in which they are being displayed. A banner ad typically occupies less than 10% of the area of a web page on a standard VGA computer screen (640x480 pixels). Therefore, the attention of the web surfer is generally focused on other elements of the page. The task of the banner ad is to first grab a surfer's attention and second to induce the surfer to click on the ad. If surfers never look at a banner, they cannot click on it!

Shared bandwidth might explain why click-through rates are low, but not why they are declining. There is some evidence that some online surfers dislike banner ads (Bass 1999). This dislike is widespread enough that various software exist that actually prevent browsers from downloading ads (AdsOff!, @Guard, JunkBuster...). One can thus hypothesize that, as surfers gain more familiarity with the medium, they learn to differentiate informational content from advertising. Ultimately, this would give them the ability to disregard banner ads.

Given this possible learning and avoidance behavior, we start our investigation by measuring the extent to which surfers pay attention to banner ads. We begin by formulating the following two hypotheses:

H₁: Internet users avoid looking at banner ads.

H₂: The more time users have spent on the Internet, the less they pay attention to banner ads.

To test these hypotheses, we asked a group of subjects to look at various web pages while hooked up to an eye-tracking device that records their eye movements and fixations. Eye-tracking studies are not new. Javal (1878) used eye-tracking to study reading patterns more than 100 years ago. Although reading studies are still being conducted through eye-tracking (Hyönä 1995), a growing number of eye-tracking studies have recently addressed marketing problems. For instance, Russo and Leclerc (1994) studied in-store brand choice, Fischer et al (1989) studied warning labels on tobacco ads, Janiszewski (1998) looked at exploratory search behavior with catalogs, Kroeber-Riel (1979) investigated the effect of arousal on advertising copy processing, and Lohse (1997) studied Yellow Page advertising.

Most of the eye-tracking studies, including ours, use Pupil Center/Corneal Reflection (PCCR) monitoring devices to track the eye movement of their subjects. These devices illuminate a subject's eyes with a near-infrared LED while a video camera collects images of the eyes. From these images, a computer calculates the position of the center of the pupil and the specular highlight of the LED (corneal reflection). From the relative position of the pupil and the reflection, the computer can recover the location of subjects' fixation within 1.5 degrees (Young and Sheena 1975).

Study design

Our study was conducted using information portals as a background. The cover study was an ergonomic research on the design for one of the largest French portals: *Voilà* (www.voila.fr). The subjects were asked to perform five searches using three portals: *Voilà*, an alternate layout for *Voilà* (henceforth called *Voilà Bis*), and *Voilà*'s largest competitor. Three of the searches related to general topics (e.g., find information about 'Le Louvre'), the other two related to individuals (e.g., find the phone number of 'Jean Dupont'). Each of the three general-topic searches was made using a different portal. The two other searches were made with *Voilà* and *Voilà Bis*.

To accomplish the task they were assigned, the subjects would click on a link or enter a search string. For the three general-topic searches, this would prompt the display of an answer page (see Figure 1 for the answer page to *Voilà*) containing the information they were asked to look for. They were then asked to indicate with the mouse where the information they were looking for was located. The answer pages were designed to match the look and feel of the question page. In other words, the *Voilà* search page led to a *Voilà* answer page of similar design; the *Voilà Bis* search page led to a *Voilà Bis* answer page, and likewise for the competing portal.

=====

Insert Figure 1 about here

=====

The order in which the pages were shown was rotated across subjects with the only restriction being that the three search/answer page pairs were kept unbroken. Eight banner ads were integrated within the design of the eight web pages (see Figure 1 for the 'Club Internet'

banner of *Voilà Answer*). At no time before or during the search task on these eight pages was any reference to banners ads made.

We collected our data in two steps. First, we used an eye-tracking device to collect eye movements and fixations¹ during the experiment. Second we asked our subjects to fill out a short survey after completing their assigned task on the eight web pages. The survey asked questions about their Internet savvy, the experimental process (e.g., did they encounter any stress), their preferences regarding the various pages to which they were exposed, and a series of questions regarding the banner ads they saw (e.g., do you remember seeing any banner ads).

=====

Insert Figure 2 about here

=====

To simplify the analysis of the eye-tracking data, each page was dissected in a series of mutually exclusive rectangular zones. One zone was created for each paragraph of text, banner ad, or graphical element of the page (see Figure 2 for the zone definition of *Voilà Answer*). The eye-fixation data were then coded by zone. Hence, for every subject we have a list of each zone that they fixated on during the experiment and for how long. Of the 60 subjects that were recruited for the experiment, 11 had to be eliminated because they suffered from heavy nystagmus or because the calibration of the eye-tracking device could not be performed satisfactorily on them. This left us with 49 usable subjects. The experiment was conducted at the CNET from December 16th to the 18th of 1998. The subjects were selected through a street intercept in the center of Paris and paid 100FrF (\cong \$16) to participate in the study.

¹ A fixation was recorded whenever a subject stared at an element of a page for at least 100ms.

Analysis

The primary goal of this first experiment was to measure the extent to which surfers actually look at the banner ads that are embedded within the web pages. Each of our 49 subjects was exposed to 8 banner ads (one per page). Looking at the zones that were focused on by our subjects, we find that every subject looked at one or more banner ads (i.e., nobody managed to avoid every ad). On average they looked at 3.96 of the eight banners during the experiment, which yields a probability of being seen of 0.49 for each individual banner (see Figure 3 for a frequency distribution of the number of banners looked at). This probability is low relative to other media such as television (>90%, per Siddarth 1999) or Yellow Page ads (89 percent for small display ads, 93 percent for large display ads, per Lohse 1997).

=====

Insert Figure 3 about here

=====

To test Hypotheses 1 and 2, we created a data file that lists all of the zones that a subject might focus on (i.e., 111 zones x 49 subjects). Each zone is associated with variables describing their location, shape, and content (see Appendix 1 for a description of the zone description variables). We then ran a logit regression using as the dependent variable an indicator that was set to 1 if the subject fixated on the zone, and 0 otherwise. As test variables we used an Ad dummy (1 if the zone is a banner ad, 0 otherwise), an Expert dummy (1 if the subjects have been on the Internet at least 25 times, 0 otherwise), and an interaction term. To control for page layout as well as possible differences across gender or age, we use the variables described in Appendix 1 as control variables along with two demographic control variables (Age, and Gender).

=====
Insert Table 1 about here
=====

The results shown in Table 4 are very revealing. As one could have expected, a zone's location and size are important. The positive coefficient on Area shows that the bigger a zone is, the more likely it is to capture subjects' attention. The significant interaction between the page and the zone interaction shows that the page layout is important. Similarly, the zone's content is important as evidenced by the significant content dummies.

The negative coefficient on the Ad dummy provides support for Hypothesis 1. It indicates that viewers avoid looking at ads. It also indicates that they are able to recognize that an item is an ad without having to look at it directly. Although the Expert dummy is marginally significant ($p=0.11$), the interaction term between expertise and banner is not significant ($p=0.46$). Hence, we do not find support for Hypothesis 2.

Although we did not find support for Hypothesis 2, it is still interesting to contrast the behavior of our various demographic groups. During the experiment, as well as throughout our analysis, we found significant differences in behavior between novices and experts as well as between young and older surfers (see Figure 4 for an example of eye movements for both an expert and a novice). To illustrate these differences, we ran a series of regressions on the number of fixes, number of zones looked at, and time spent during fixes across these groups. For each page looked at by each respondent, we regressed the three dependent variables against an expertise dummy, an age dummy, and a gender dummy, as well as seven control variables to account for the differences across pages. As the results in Table 2 show, experts tend to process each page by making fewer fixes, looking at fewer zones, and spending less time than novices.

Older people look at the same number of zones as young people, but it takes them longer and they fixate on a larger number of positions. Finally, males and females seem to behave similarly.

=====

Insert Table 2 and Figure 4 about here

=====

As part of the debriefing questionnaire, we asked our subjects if they remembered seeing any banner ads. Only 46.94 percent of the subjects indicated they did. After asking them if they remembered seeing any ad, we showed our subject four banners and asked them if they recalled seeing the ads during the test. Two of the ads were fake ads that had not been part of the test; the other two were real. We did not find significant differences in recognition level between the fake and the real ads ($m=23.5\%$ vs. 18.4 , respectively, $p=0.38$). The number of false positives we encountered is similar to those reported in other studies. Janiszewski (1990) reported 21 percent of false positives (119 subjects, 5 different ads). A forthcoming study jointly conducted by the French division of the Internet Advertising Bureau and SOFRES (1999) reports false positive levels of 17 percent (6,872 subjects, 14 ads).

Take away from Study 1

Our first study was very revealing. It provides us with an answer to the first question motivating this study (Why are banner ads not effective?) and a hint of an answer to the second and third questions (What can advertisers do to improve banner effectiveness? Does click-through rate under-value online advertising?). The study shows that one of the problems hindering banner ad effectiveness is that *half of the banner exposures are not attended to*. The problem is not only that surfers do not look at the banners, but they also seem to *purposefully avoid looking at them* (Hypothesis 1).

There are at least two possible explanations for this apparently clairvoyant behavior. First, site designers have traditionally located banner ads at the top of their web pages. This might lead web surfers to treat as a potential ad every item that is located at the top of the screen. Second, as has been noted by Janiszewski (1998), peripheral vision allows subjects to recognize objects that are located outside their focal point of attention (as measured by the eye-tracking device). This ability, coupled with the fact that most banner ads have the same shape (468x60 pixels) provides web surfers with the ability to train themselves into recognizing banner ads for what they are without having to actually focus on them. Both of these explanations assume that surfers learn over time and develop strategies to avoid devoting attention to advertising. Unfortunately, we did not find support for Hypothesis 2. This probably means that it takes less than 25 Internet sessions to learn to avoid banners.

That only half the banner ads are looked at is highly detrimental to click-through rates. One cannot click on something one does not look at! It does not mean, however, that half of the banner exposures are wasted. Research has shown that consumers do not need to fully process a message in order to be influenced by it. Janiszewski (1990a, 1990b, and 1993) has researched the topic extensively. Among other things, his research shows that incidental exposure to advertising can enhance a consumer's liking for the ad and brand advertised despite the consumer's inability to recognize having previously seen the ad and brand (a situation similar to ours). Other researchers (e.g., Shapiro, MacInnis, and Heckler 1997) have reached similar conclusions. This means that a large part, if not most, of a consumer's processing of banner ads will be done at a pre-attentive level rather than at a full attention level. Further, it implies that click-through rates will not capture the full extent of an ad's effectiveness since pre-attentive processing does not lead to immediate action.

Our study also reveals that experts are more efficient at processing web pages than novices and that young surfers are more efficient than older ones. This does not, however, translate into fewer banners seen by experts or young surfers.

As to what factors might help improve banner effectiveness, we found that location, size, and zone content are important factors when trying to predict whether a zone is attended. We will further investigate these factors along with the relevance of traditional advertising effectiveness measures in Study 2.

Study 2: Impact of Online Advertising

Study 1 provides us with somewhat discouraging results about Internet advertising. Although everybody fixated on at least one ad, only half of the subjects remembered seeing a banner ad during the experiment. Further, the reported recognition rate for ads they have been exposed to is similar to the recognition rate for ads they have never seen before.

The purpose of study 2 was to further explore these findings using a different methodology. Indeed, eye-tracking studies have two major drawbacks. First, they are artificial environments; second their cost and infrastructure prevent their use on large samples. Hence, we sought to validate the results of the first study in a much broader experiment (~1,000 respondents). Study 2 first attempts to measure the effects of online advertising using traditional memory-based measures (i.e., other than click-through), and then establish how the factors uncovered in study 1 (i.e., size, shape, content) affect advertising effectiveness.

Our study of the effects of Internet advertising uses the same benchmarks for online advertising as those that have been used for advertising in traditional media. For television or print ads, advertising agencies use various measures to evaluate an ad campaign (Tellis 1998). They are interested in such constructs as unaided advertising recall, aided advertising recall,

brand recognition, and brand awareness. Consequently, we formulate the following four hypotheses regarding banner ads:

H_{3a}: Banner ads will have a positive impact on aided advertising recall.

H_{3b}: Banner ads will have a positive impact on brand recognition.

H_{3c}: Banner ads will have a positive impact on unaided advertising recall.

H_{3d}: Banner ads will have a positive impact on brand awareness.

These four measures are sorted by increasing level of advertising effectiveness. Brand awareness is the most desired effect; aided advertising recall is the least preferred. Should our analysis reveal that banner ads do indeed have a significant impact beyond immediate click-through, we will be able to use these four measures to analyze the different factors that may ultimately affect banner effectiveness.

Previous studies (e.g., Chatterjee, Hoffman, and Novak 1998) have shown that repetition has a negative impact on click-through rates. Click-through rates are maximum after the first exposure. This is at odds with studies of television advertising (Pechman and Stewart 1989) that shows how a low level of repetition is beneficial. In light of Study 1's result that indicate only 50 percent of the exposures are attended to, we believe that the total number of surfers who pay attention to a banner ad will increase with repeated exposure to the banner. With this increase in audience, we should see an increase in advertising effectiveness. To resolve the issue of the positive or negative impact of repetition we phrase hypothesis two as:

H₄: Advertising effectiveness increases with frequency of exposure.

Our logit analysis performed in Study 1 also gives us some insights as to what attributes of a banner ad will affect its ability to attract attention. The factors included in the logit analysis were size, shape, content, and location. We will not study the impact of location as it is an

attribute over which advertisers have little control. As we have mentioned before, site designers typically use the prime locations to display the informational content of their site, and relegate banner ads to the top of their pages.

As far as size is concerned, our analysis shows that bigger is definitely better, although the S-shaped nature of the logit link function implies that at some level one will see decreasing marginal returns from increases in size. This leads to our next hypothesis:

H_{5a}: Larger banner ads will be more effective than smaller banner ads.

Study 1 also revealed that orientation matters. A vertical zone is more likely to be attended to than a horizontal one. Consequently, we will have the following corollary to Hypothesis 5a:

H_{5b}: Banner ads that are laid-out vertically will be more effective than banner ads that are laid-out horizontally.

Besides location and shape, Study 1 showed that a zone's content impacts its ability to attract attention. The next series of hypotheses will deal with this issue. The size of the CONTRAST coefficient in Table 4 shows that it is an important factor. Hence, we will formulate our next hypothesis as:

H_{6a}: Banner ads that contrast with their environment will be more effective than ads that do not.

One characteristic that sets banner ads apart from other elements on a web page is animation. There is some anecdotal evidence that users are put off by blinking banner ads (Hamilton, 1998). However, animation provides advertisers with the ability to put more content in the same space by spacing the content temporally rather than spatially. Animation also allows

advertisers to build dramatic pauses by delivering the message sequentially. This increased flexibility should translate in more effective ads. Thus, we hypothesize that:

H_{6b}: Animated banners will be more effective than static ones.

As they have gained experience with the online medium, advertisers also have become more and more creative with their use of the limited space available on a banner. They have gone from simple images to fully interactive games, from static words to audible messages. The tools they have at their disposal are now richer. This should, hopefully, translate into more effective ads. This presupposes, of course, that advertisers use these tools effectively and that an ad's execution matters.

While it is hard to evaluate the ability of advertisers, it is relatively easy to figure out whether their ability matters by testing the impact of different executions of the same concept.

We will thus have:

H_{6c}: The effectiveness of a banner ad depends on its message.

In summary, we have generated ten hypotheses that fall into four different types. H_{3a} through H_{3d} deal with the use of traditional measures of advertising effectiveness to evaluate online advertising. H₄ relates to the effect of frequency on online advertising effectiveness. H_{5a} and H_{5b} relate to the physical characteristics of the banner box. Finally, H_{6a} through H_{6c} relate to the content of the banner box.

Study Design

The second study consisted of a two part self-administrated web-based questionnaire. Respondents were recruited through a mass e-mail sent to 6,000 individuals. The e-mail asked them to participate in a study about *Voilà*, the new search engine of France Telecom. The stated purpose of the study was to better understand how people use the Internet. The e-mail contained

an hyperlink that would lead the respondent directly to the web pages containing the survey. As an incentive to participate, 2 digital cameras and 25 CD-Roms were to be awarded to participants through a random drawing.

The questionnaire began with three questions addressing a respondent's current Internet usage and their brand awareness. Once these questions were answered, respondents were led through nine web pages in which they had a specific task to accomplish or a question to answer (e.g., "Enter the word 'Voyage' in the SEARCH string input box"). Seven of the nine pages contained a banner ad. The ads related to five different products. Four of the products were online services offered by France Telecom: *@près l'école* (a subscription based homework help web site), *Le Mél* (a free web e-mail service similar to Hotmail), *Nouba* (a web site listing the time and places of films, plays, and concerts), and *Tout En Ville* (a city guide about France's major cities). The fifth advertised product was *Cortal*, a large French financial institution. Each respondent saw one banner ad for each of the five services plus a second ad for both *Cortal* and *@près l'école*, for a total of seven exposures.

Each of the banner ads was executed with different creative specification purposefully designed to test the hypotheses derived in the previous section. For instance, the ad for *Le Mél* had five different executions. The five banners were identical except for their color (blue, yellow, orange, red, or green). In contrast, *Cortal* was executed in two different sizes: 230x33 pixels and 468x60 pixels. A total of 18 different banners were used. The details of each execution will be discussed along with the experimental results in the next section. The order in which the banners were presented as well as which execution of each banner was used was counter balanced across respondents.

Twenty-four hours after they filled out the first survey and went through the nine test web pages, the respondents who completed the first task were sent a second e-mail asking them to fill-out a follow-up survey. This survey asked them several questions about their advertising recall, and to rate the banner ads to which they were exposed. Of the 6,000 e-mails that were originally sent, 211 came back as undeliverable. 2,297 of the remaining 5,789 prospects successfully completed the first part of the experiment. After being contacted the next day, 807 respondent filled out the second survey, for an overall response rate of 13.5 percent. The data was collected between January 28 and February 4, 1999.

Analysis

Our first order of business was to determine what effects, if any, banner ads have on recall, awareness, and recognition. We will not report any analysis of click-through rate since only three click-throughs were recorded during the experiment. Out of 807 respondents, 460 (57 percent) recalled having seen an ad during their task the previous day. Table 3 reports the brands they recall seeing during the experiments.

=====

Insert Table 3 about here

=====

The five brands with the highest recall (11.4 percent average recall) are those that were advertised during the test. The other brands do not seem to be random brands mentioned just to satisfy the questionnaire. *Voilà* is the search engine of *France Telecom*, the sponsor of the research. The *Voilà* logo was prominently displayed at the top of each page. It is not a banner ad per se, but it could be mistaken for one. *Pages Jaunes* and *Pages Blanches* (Yellow and White pages) are telephone number search services offered by *Voilà*. *Pages Zoom* is the former

name of these services. *Honda* and *123 Achats* are brands that were heavily advertised on the web (and on *Voilà*) at the time of the experiment. Finally, *Internet Explorer*, *Netscape*, and *Linux* were a set of three buttons that were located at the bottom of each page as they are on many web sites. Again, they do not constitute banner ads per se, but could be construed to be advertisements.

To test whether these differences in unaided advertising recall are significant, we ran a logit regression with 14 observations per subject (1 for each brand mentioned except for *Voilà*²). Our dependent variable was set to 1 if the brand was mentioned by the subject, and 0 otherwise. We then created three dummy variables. EXP1 was set to 1 for each of the five banner ads, EXP2 was set to 1 for the two banner ads that had received two exposures (*Cortal* and *@près l'école*), finally BUTTON was set to one for *Internet Explorer*, *Netscape*, and *Linux*. The estimates for the logit regression are shown in Table 4.

=====
 Insert Table 4 about here
 =====

As one can see, the coefficients for both one and two exposures are significant (note that the total effect for two exposures is the sum of EXP1 and EXP2). This supports H_{3c} regarding the positive impact of banner ads on unaided advertising recall. BUTTON is not significant; indicating that the recall of *Netscape*, *Internet Explorer*, or *Linux* cannot be attributed to our experiment.

=====
 Insert Table 5 about here
 =====

² We chose to ignore *Voilà* since some subjects might consider its logo to be a banner ad while others might not.

To investigate hypotheses H_{3a} and H_{3b} , we removed the brand name from each of the five banner ads used in the experiment and showed them to our subjects. We then asked them whether they recalled seeing the ads during the experiment and, if so, could they name the brands that are advertised. Table 5 provides aided advertising recall, brand recognition (conditional on recalling the ad), as well as unconditional brand recognition. We can compare the recall levels to the level obtained in Study 1 for the ads that had never been seen (Recall = 23.45 percent). We find that the overall mean recall of 30.1percent is significantly larger than 23.45 percent ($p=0.0001$), while the individual ads themselves are significantly larger than 23.45 percent for *Tout en Ville* ($p=0.0001$), *Le Mél* ($p=0.0001$), and *@près l'école* ($p=0.0001$). The differences are, however, only marginally significant for *Cortal* (0.1174), and not significantly larger for *Nouba* (0.4635). If we compare these recognition levels to the one obtained by IAB France/SOFRES (17 percent) on its much larger sample, we find that they are all significantly larger at the $p=0.0001$ level. This gives us some support for H_{3a} ; the hypothesis that banner ads have a positive impact on aided advertising recall.

=====

Insert Table 6 about here

=====

In terms of brand recognition, we can turn to Table 6, which shows the brands associated with each ad. As one can see, the overwhelming majority of people make the correct association. The few people that make incorrect brand recognition tend to mistake *Voilà*, *Wanado* or *France Telecom* for the correct brand. A χ^2 test rejects independence at the $p=0.0001$. Hence the data supports the hypothesis that banner ads have a positive impact on brand recognition (H_{3b}).

Our last measure of advertising effectiveness is brand awareness. In the screening questionnaire subjects indicated whether they knew each of ten brands. Five of these brands were the brands advertised later during the experiment; the remainder were Internet-related brands (e.g., *Yahoo!*, *Netscape*). The same questions were asked the next day in the follow up survey. Pre- and Post-awareness measures and the number of exposures for each brand are reported in Table 7. As one can see, the non-advertised brands suffered a slight loss in awareness whereas the advertised brands enjoyed an increase in awareness. For the purpose of this particular analysis, we will consider that the effect of the *Voilà* logo at the top of each page is similar to that of a banner ad for *Voilà*, and include the data in the analysis.

Various models of brand awareness have been proposed over time (see Mahajan, Muller, and Sharma 1984 for an exhaustive review, see also Little 1979). Those models, whether they be TRACKER (Blattberg and Golanty 1978), NEWS (Pringle, Wilson, and Brody 1982), or LITMUS (Blackburn and Clancy 1982), assume that exposure leads to increased awareness. They also assume that people forget about brands in the absence of advertising, and that there is some saturation level to awareness (typically 100 percent). They are usually estimated using OLS on time series data. In our case, we possess data for only two time periods. Hence, we propose an alternative model that takes advantage of the data we have at hand, and still adheres to the spirit of the previous models in that it takes into account buildup due to exposure, decay due to non-exposure, and saturation effects.

If one has individual level data for two consecutive periods (as we do) and one considers that the awareness level of a brand is nothing else than the probability that one remembers the brand, then one can set up a logit regression to represent the awareness building and decay phenomenon. We can say that:

$$A_{it} = \frac{e^{\alpha_i - t \cdot \beta_1 + t \cdot \beta_2 \cdot \text{Exposure}_i}}{1 + e^{\alpha_i - t \cdot \beta_1 + t \cdot \beta_2 \cdot \text{Exposure}_i}}$$

Where: A_{it} is the awareness level of brand i at time t ;

t is a time index taking the values 0 or 1;

Exposure_i is the number of exposures to brand i between time 0 and time 1;

α_i , β_1 , and β_2 are parameters to be estimated.

In this setting, the α_i s are indicators of each brand's pre-test awareness (the bigger α_i , the bigger brand i 's awareness). β_1 represents the decay occurring in the absence of advertising, and β_2 is an indication of the benefits to be gained from advertising exposure. In addition, this functional form accounts for saturation effects when the awareness levels become large.

=====

Insert Table 7 about here

=====

Following the above awareness model, we ran a logit regression on a data set containing 20 observations per subjects (10 pre-awareness, 10 post-awareness). The data set contained 9 dummy variables for the ten brands, one dummy variable (TEST) to indicate the pre-post test, and one variable (EXP) indicating the number of exposures to the brand. The exposure variable is set to 0 for all pre-test data points, and, depending on the brand, 0, 1, 2, or 9 for the post-test data points. In this setup, the brand dummies control for difference in pre-awareness across brands, the test dummy indicates the change in awareness occurring after zero exposure, and the exposure (EXP) variable tells us by how much awareness changes as the number of exposure increases. The parameter estimates and their p-values are shown in Table 8.

=====
 Insert Table 8 about here
 =====

The large variation in brand dummies accurately reflects the differences in brand awareness. The negative sign on the test dummy indicates that there is a slight decrease in awareness after 24 hours in the absence of ad exposure. However, it is not statistically significant in our case. Finally, the exposure variable is positive and significant. This provides support for H_{3d} , the hypothesis that banner exposure has a positive impact on brand awareness. One must remember, however, that the logit link function is inherently S-shaped and thus, as the number of exposure increases, the benefits will decrease. To better illustrate the post-awareness levels as a function of exposure, one can look at Figure 5 that plots the post-awareness level as a function of the pre- level (0, 1, 2, and 9 exposures). One can also look at Figure 6 that plots the post-awareness as a function of the number of exposure for different pre-awareness levels (25 percent, 50 percent, and 75 percent). As one can see, there is a big gain to be earned by having multiple exposures. Further, the higher the current awareness, the less there is to gain from additional exposures.

=====
 Insert Figures 5 and 6 about here
 =====

So far we have found support for H_{3a} through H_{3d} . This gives us a benchmark with which we can test Hypotheses 4 through 6. We, however, already have a partial answer for Hypothesis 4. In our discussion of H_{3c} and H_{3d} we have highlighted the fact that repetition increases both unaided advertising recall and brand awareness. We now only need to look at its effect on aided

advertising recall and brand recognition. To do so, we ran two logit regressions. The first one had five data points for each respondent (one for each of the banner ads) and used advertising recall as dependent variable. As an independent variable, we used a dummy variable that took the value of 0 if the banner was seen once, and 1 if the banner was seen twice. We proceeded in the same fashion for the brand recognition variable except that we limited ourselves to the observations where the recall indicator is set to 1. For each of these regressions, a positive and significant dummy variable indicates that repetition is beneficial.

=====

Insert Table 9 about here

=====

The parameter estimates for each regression as well as their significance level are shown in Table 9. Repetition improves brand recognition, but it does not improve aided advertising recall. That is, repetition improves unaided advertising recall, brand recall, and brand awareness, but does not improve aided advertising recall. This data suggests support for Hypothesis 4 that repeated exposures increase online advertising effectiveness.

The procedure to test the remaining five hypotheses was the same for each hypothesis. In order to be as concise as possible, we will describe it only once, leaving it to the reader to make the necessary adjustments to fit each case. We test each hypothesis by fitting logit regression models on our four advertising effectiveness constructs: unaided advertising recall, aided advertising recall, brand recognition, and brand awareness. In each case, we set up a series of dummy variables that reflect the experimental design used, and we test whether these dummies are significantly different from zero. Rather than report detailed numbers about each of the regression, we report only the significance level of the dummies as a whole.

When conducting these tests, we consider that an ad is more memorable than another if, at the minimum, it leads to higher *brand recognition*. A banner that leads to higher aided advertising recall, but does not improve brand recognition, unaided advertising recall, or brand awareness will not be considered effective as surfers can remember the banner, but not the brand that was advertised.

H_{5a} and H_{5b} relate to the size and orientation of banners. We tested three different banner sizes. We used the traditional 468x60 pixel banner as our null case, and pitted it against both a quarter (230x33³) and a double (468x120) size banner. To test vertical versus horizontal orientations, we pitted a 144x240 pixel banner against the traditional 468*60. The significance levels for the logit regressions are shown in Table 10.

=====
 Insert Table 10 about here
 =====

Overall, we do not find any support for H_{5a} . Small banner ads perform just as well as large ones. This contradicts the findings of Study 1, which showed that bigger is better. The lack of size effect might be due to the decreasing returns from larger size. The small banner ads might be big enough to attract attention and the benefits gained from the larger ads might be too small to be measured accurately given our sample size. We find very weak support for H_{5b} . It seems that banner orientation affects aided advertising recall, but not unaided recall or brand recall. It might also have a small effect on brand awareness, although the effect is only marginally significant (p=0.17). The fact that only aided advertising recall is significant might indicate that our subjects remember the shape of the banner more than its content.

³ 230*33 is slightly larger than a quarter of 468*60, but it is a widely accepted banner size.

H_{6a} to H_{6c} relate to the graphical content of banner ads. To test H_{6a} , we used five banner ads that differ only in their background color (Red, Blue, Yellow, Orange, or Green). Given the background color of the pages and the color of the elements surrounding the ads, the Yellow and Orange ads were coded as not contrasting while the other ads were contrasting. To test H_{6b} , we used a still and an animated banner ad, both displaying the same message. Finally, to test H_{6c} , we compared four different ads for the same product. Our base case ad is a traditional 468x60 pixel banner ad with the brand name of the sponsor. The second ad is an offer to participate in a game associated with the sponsor. The third execution has a drop down menu embedded in the banner ad. A sound file that enunciates the name of the brand being advertised as the banner ad is displayed accompanies the last one.

=====
 Insert Table 11 about here
 =====

As reported in Table 11, we find a weak contrast effect (H_{6a}). However, the effect goes in the opposite direction as predicted. High contrast ads seem to perform worse than non-contrasting one. We do not find any animation effect (H_{6b}), but we find support for H_{6c} . The fact that brand recognition is significant for H_{6c} is important because it shows that the different messages used affect not only the recognition of the visual component of the ads as is the case for H_{6a} and H_{6b} but also the comprehension of the message. Hence, ad content is important. Further, the message is more important than how the message is sent.

Summary of results

We started with an eye-tracking study conducted on 49 subjects. Study 1 revealed that a banner ad typically has a probability of about 50 percent of being seen by a surfer looking at a

page in which the banner is embedded. This number is dramatically lower than the 97 percent reported by Siddarth (1999) for television ads or 93 percent reported by Lohse (1997) for Yellow Pages ads indicating that pre-attentive ad processing might be much more important on the Internet than in more traditional media. We also found that different levels of expertise, age, or gender lead to differences in page processing behavior. However, these differences did not seem to translate into different probabilities of seeing banner ads. A logit analysis of the eye-tracking data revealed that some of the factors that would most likely influence the probability of content being looked at on a web page were a zone's size, shape, content, and location.

We used the results of the eye-tracking experiment to design a large-scale experiment where 807 subjects were exposed to nine pages containing six banner ads. The banner ads were designed to test a series of six hypotheses relating to size, shape, color, animation, message, and repetition. The study was also used to test whether traditional memory-based measures of advertising effectiveness could be applied to the Internet.

We found that traditional memory-based effectiveness measures provide valuable insight into the effects of Internet advertising. These measures outperform the immediate effects measured by the click-through rate. On average, for 100 surfers exposed to a banner ad, 11 recall seeing the ad and can mention the brand name on the ad without any aid 24 hours later (see Table 12 for summary results). Thirty respondents remember seeing the banner when they are shown the same banner but without brand name. Of those 30 surfers, 18.5 (62 percent) can name the banner's brand. In addition, three of the 100 surfers become aware of the brand. We can compare these numbers with the average click-through rate of 0.6 percent (Nielsen//Netratings 1999). The effect on unaided brand awareness is 4.5 times larger than the click-through rate; the effect on unaided brand recall is 19 times larger.

=====

Insert Table 12 about here

=====

In terms of what factors influence online advertising effectiveness we found that frequency is important. Repetition affects unaided advertising recall, brand recognition, and brand awareness. However, execution has little effect. Banner size seems to be unimportant. Contrast, animation content, and shape of the banner influence aided advertising recall but no other dependent measure. Nevertheless, a banner's message influences both aided advertising recall and brand recognition. This indicates that what an ad says is more important than how it says it.

Our finding that repetition increases advertising effectiveness might seem to contradict the results found by Chatterjee et al (1998) in their study of the negative effect of repetition on click-through. We do not believe they do. If we use the framework of Lavidge and Steiner (1961), awareness measures cognition while click-through measures action. Click-through requires the user to see the banner, process it, be convinced by its message, and take action. It is very unlikely that a user would click on a banner ad after two exposures if he or she actually processed and rejected the message during the first exposure. Hence the only users who might click on a banner after multiple exposures are those who have not processed the ad already. As the number of exposures increases, the proportion of 'virgin' users decreases, and thus the click-through rate will decrease too. This might lead Web advertisers to pursue the logic that if short-term results are desired, avoid repetition, but if long-term results are sought, repetition should be encouraged!

Conclusion

We started this research to answer four basic questions regarding online advertising. Given the steady decline in click-through rate over the past three years, we asked (1) why are banner ads not effective; and, (2) what can advertisers do to improve their effectiveness? Further, given the focus of the Internet community of click-through rates as a measure of banner ad effectiveness, we asked (3) does an immediate measure such as click-through rate actually under-value online advertising; and, (4) are more traditional measures such as recall or awareness more appropriate?

Our research shows that traditional memory-based measures of advertising effectiveness should be applied to banner ads as it is to television or newspaper advertising. By focusing on an immediate measure such as click-through rate, one fails to capture the positive effect that banners have on the equity of the brand being advertised. To be sure, click-through rates are not irrelevant. They provide a measure of advertising effectiveness that, along with other measures, will help managers better allocate their advertising budgets. But click-through rates alone are insufficient. Managers who rely solely on them are short-selling their advertising campaigns. Managers must also be aware that click-through rates are still declining and thus inter-temporal comparisons should be treated cautiously.

As to why banner ads seem to be ineffective, we believe that they are actually effective. Banners lead to brand awareness. Click-through rate might be low, but in the long run, awareness is more important than click-through. The low click-through rates are primarily due to the fact that surfers fixate on less than 50 percent of the banners to which they are exposed. Not only do they not see banners, but they actually avoid looking at them.

Finally, advertisers can increase advertising effectiveness by concentrating on the message they send. What is said and how often it is said is more important than how it is said. Bigger ads or animated ads will not compensate for ineffective content.

Putting what we have learned from this study in perspective, we would be inclined to say that the medium Internet advertising resembles most is outdoor billboards. As with banner ads, drivers encounter billboards while engaging in other activities. Billboards occupy only a small portion of their field of vision and typically consist of a simple message and visual. As Donthu, Cherian, and Bhargava (1993) have shown, billboards influence awareness and recall even if they only rarely prompt consumers to take immediate action. Thinking in these terms might help design future research and yield better insight in the mechanisms underlying Internet advertising.

One should also be aware of some of the shortcomings of our research. First and foremost, all of our subjects were engaged in a goal-directed search. As Janiszewski (1998) has shown, goal-directed searches are very different than exploratory searches. Second, we examined memory measures but not behavior measures (e.g., sales or 'store visits').

References

- Ambler, T. (1998), "Myths about the mind: time to end some popular beliefs about how advertising works," *International Journal of Advertising*, Nov, 501-509.
- Associated Press (1996), "Procter & Gamble World Wide Web ad strategy raises online ire," San Francisco, Apr 28.
- Bass, S. (1999), "Unclog the Net: Eliminate Web Ads," *PC World Online*, January.
- Batra, R., D. R. Lehmann, J. Burke, and J. Pae (1995), "When Does Advertising Have an Impact? A Study of Tracking Data," *Journal of Advertising Research*, September/October, 19-32.
- Berst, J. (1999), "Free PCs, ISPs and Software: Are They Worth the Price?," *ZDNet*, February 10.
- Blackburn, J. D. and K. J. Clancy (1982), "LITMUS: A New Product Planning Model," *Studies in the Management Sciences-Marketing Planning Models*, Andy Zoltners (Ed.), 18, 43-62.
- Blattberg, R. and J. Golanty (1978), "Tracker: An Early Test Market Forecasting and Diagnostic Model for New Product Planning," *Journal of Marketing Research*, 15 (May), 192-202.
- Briggs, R. and N. Hollis (1997), "Advertising on the Web: Is There Response before Click-Through?," *Journal of Advertising Research*, 37 (2), 33-45.
- Chatterjee, P, D. L. Hoffman, and T. P. Novak (1998), "Modeling the Clickstream: Implications for Web-Based Advertising Efforts," *Project 2000*, Vanderbilt University.
- Donthu, N., J. Cherian, and M. Bhargava (1993), "Factors Influencing Recall of Outdoor Advertising," *Journal of Advertising Research*, 33 (3), 64-72.

- Dyson, P., A. Farr, and N. Hollis (1996), "Understanding, Measuring, and Using Brand Equity," *Journal of Advertising Research*, 36 (6), 9-21.
- Fischer, P. M., J. W. Richards, E. J. Berman, D. M. Krugman (1989), "Recall and Eye Tracking Study of Adolescents Viewing Tobacco Advertisements," *Journal of the American Medical Association*, 261 (1), 84-89.
- Hamilton, A. (1998), "Secrets of Super-High Web Ad Click-Through," *ZDNet*, August 19.
- Hyönä, J (1995), "Do Irregular Letter Combinations Attract Readers' Attention? Evidence From Fixation Locations in Words," *Journal of Experimental Psychology*, 21 (1), 68-81
- Internet Advertising Bureau (1999), "Internet Advertising Revenues More Than Double in 1998," New York, May 3.
- Internet Advertising Bureau France/SOFRES (1999), "Net.Imp@ct," May.
- Janiszewski, C. (1990a), "The Influence of Nonattended Material on the Processing of Advertising Claims," *Journal of Marketing Research*, August, 263-278.
- Janiszewski, C. (1990b), "The Influence of Print Advertisement Organization on Affect toward a Brand Name," *Journal of Consumer Research*, 17 (June), 53-65.
- Janiszewski, C. (1993), "Preattentive Mere Exposure Effects," *Journal of Consumer Research*, 20 (December), 376-392.
- Janiszewski, C. (1998), "The Influence of Display Characteristics on Visual Exploratory Search Behavior," *Journal of Consumer Research*, 25 (3), 290-301.
- Javal, E. (1878), "Essai sur la physiologie de la lecture," *Anales d'oculistique*, 79, 97-117.
- Kroeber-Riel, W. (1979), "Activation Research: Psychobiological Approaches in Consumer Research," *Journal of Consumer Research*, 5 (March), 240-250.

- Lavidge, R. J. and G. A. Steiner (1961), "A Model for Predictive Measurements of Advertising Effectiveness," *Journal of Marketing*, October, 59-62.
- Little, J. D. C. (1979), "Aggregate Advertising Models: The State of the Art," *Operation Research*, 27 (July-August), 629-667.
- Lohse, G. L. (1997), "Consumer Eye movement Patterns on Yellow Pages Advertising," *Journal of Advertising*, 26 (1), 61-73.
- Mahajan, V., E. Muller, and S. Sharma (1984), "An Empirical Comparison of Awareness Forecasting Models of New Product Introduction," *Marketing Science*, 3 (3), 179-197.
- Nielsen//Netratings (1999), "The Nielsen//Netratings Reporter," www.nielsen-netratings.com, June 17.
- Novak, T. P. and D. L. Hoffman (1997), "New Metrics for New Media: Toward the Development of Web Measurement Standards," *World Wide Web Journal*, 2 (1), 213-246.
- Pechmann, Cornelia and David W. Stewart (1989), "Advertising Repetition: A Critical Review of Wearin and Wearout," *Current Issues and Research in Advertising*, James H. Leigh and Claude R. Martin Jr., Eds. Ann Arbor, MI: University of Michigan, 285-329.
- Pringle, L. G., R. D. Wilson, and E. I. Brody (1982), "NEWS: A Decision Oriented Model for New Product Analysis and Forecasting," *Marketing Science*, 1 (Winter), 1-30.
- Russo, J. E. and F. Leclerc (1994), "An eye-fixation analysis of choice processes for consumer nondurables," *Journal of Consumer Research*, 21 (2), 274-290.
- Shapiro, S., D. J. MacInnis, and S. E. Heckler (1997), "The Effects of Incidental Ad Exposure on the Formation of Consideration Sets," *Journal of Consumer Research*, 24 (June), 94-104

- Siddarth, S. (1999), "Describing the Dynamics of Attention to TV Commercials: A Proportional Hazards Model of the Time to Zap an Ad," *UCLA Working paper*.
- Steinman, R. M. and J. Z. Levinson (1990), "The Role of Eye Movement in the Detection of Contrast and Spatial Details," in *Eye movements and Their Role in Visual and Cognitive Processes*, ed. Elizabeth Kowler, North-Holland, Elsevier, 115-160.
- Tellis, G. J. (1998), "Advertising and sales promotion strategy," *Addison-Wesley Educational Publishers, Inc.*
- Young, L. R. and D. Sheena (1975), "Eye-Movement Measurement Techniques," *American Psychologist*, 30 (March), 315-330.

Appendix 1: Zone description variables

Building on past research (e.g., Janiszewski 1998, Steinman and Levinson 1990) we hypothesize that a zone's ability to attract a subject's gaze will be a function of its location on the page, its size and shape, and its pictorial or textual content. Location was captured by the x and y coordinates of the center of each zone, as well as by the distance to the nearest zone seen by the subject. The top-left corner of the screen was defined as the origin point (0,0), the bottom-right corner was set to (255,255). We also created squared terms for the x and y coordinates to allow for the fact that the center of the pages might be preferred.

Shape was described in terms of the zone's area as well as its orientation⁴ (horizontal vs. vertical) and the general orientation of the page. An interaction dummy was also created to indicate whether the zone and the page follow the same orientation.

The content of each zone was coded using the following dummy variables:

- TEXT (set to 1 if the zone contains text, 0 otherwise),
- IMAGE (1 if the zone contains non-photographic graphical elements),
- PICTURE (1 for photographs),
- INPUT (1 if the zone contains input elements such as buttons),
- CONTRAST (1 if the zone contrasts with its surrounding), and
- AD (1 if the zone contains an advertisement).

Our last control variable is a variable that indicates the number of zones present on each page.

⁴ Remember that all zones are rectangular.

Table 1: Logit regression (Zone Attractiveness)

Variable		Coefficient	Wald χ^2	Pr> χ^2
Intercept		-3.1947	30.22	0.0001
Test	AD	-0.1888	5.10	0.0240
	Expertise	-0.0633	2.50	0.1136
	AD x Expertise	-0.0554	0.54	0.4639
Location	X	0.0203	8.24	0.0041
	Y	0.0211	32.73	0.0001
	X ²	-0.0001	14.52	0.0001
	Y ²	-0.0001	59.24	0.0001
	Distance	-0.0380	326.27	0.0001
Shape	Area	0.0005	348.96	0.0001
	Zone Orientation	0.0482	0.37	0.5437
	Page Orientation	0.0597	1.00	0.3190
	Orientation Int.	0.1135	3.68	0.0550
Content	IMAGE	1.2067	91.52	0.0001
	PICTURE	-3.9477	236.82	0.0001
	CONTRAST	0.2813	12.06	0.0005
	TEXT	1.8351	150.20	0.0001
	INPUT	0.1751	2.23	0.1353
Control	# of Zones/Page	0.0083	0.15	0.7029
	Age	0.0595	2.89	0.0892
	Gender	0.0531	1.60	0.2013
Regression			1234.03	0.0001

Table 2: Differences among groups of surfers

	Number of Fixes		Number of Zones		Fix Time	
	Estimate	p-value	Estimate	p-value	Estimate	p-value
Intercept	26.59	0.0001	6.99	0.0001	12.01	0.0001
Expert	-6.03	0.0006	-0.67	0.0162	-2.52	0.0002
Old	5.36	0.0018	0.21	0.4536	2.65	0.0001
Female	-2.32	0.257	-0.59	0.0753	-0.86	0.2822
Voilà	-0.91	0.7742	0.28	0.5851	-0.23	0.8507
Voilà Answer	4.73	0.1363	1.59	0.0019	2.22	0.0742
Yahoo!	4.62	0.1413	-0.49	0.3309	0.63	0.6078
Yahoo! Answer	12.18	0.0001	-2.27	0.0001	4.39	0.0004
Voilà Bis	4.13	0.1912	0.43	0.3872	1.65	0.1806
Voilà Bis Answer	10.2	0.0014	2.26	0.0001	4.2	0.0008
Voilà Annuaire	6.48	0.0436	2.78	0.0001	0.95	0.4481
N	378		378		378	
R ²	0.11		0.28		0.13	

Table 3: Unaided Advertising Recall (24 hours after surfing)

Brand	Brand Recall*
Cortal	14.3%
@près l'école	12.6%
Le Mél	12.6%
Tout En Ville	10.2%
Nouba	6.9%
Voilà	4.3%
Netscape	3.3%
Honda	3.0%
France Telecom	2.8%
Internet Explorer	2.8%
123 Achats	2.2%
Pages zoom	2.0%
Linux	1.7%
Pages Jaunes	1.5%
Pages blanches	1.1%
Other	4.6%
Cannot Recall Brand	54.1%

*Numbers do not add to 100% due to multiple answers

Table 4: Logit Regression (Unaided Advertising Recall)

Variable	Coefficient	Wald χ^2	Pr> χ^2
INTERCEPT	-3.8413	837.8418	0.0001
EXP1	1.6360	104.0824	0.0001
EXP2	0.3460	6.8699	0.0088
BUTTON	0.2496	1.3723	0.2414
Regression		245.96	0.0001

Table 5: Aided Recall

Banner	Aided Ad Recall	Brand Recognition (Conditional)	Brand Recognition (Unconditional*)
Tout en Ville	35.4%	64.1%	22.7%
Le Mél	36.3%	66.7%	24.2%
Nouba	25.3%	40.8%	10.3%
@près l'école	31.0%	72.1%	22.4%
Cortal	23.6%	60.6%	14.3%
Overall**	30.1%	61.6%	18.5%

*Unconditional Brand Recognition = Ad Recall * Brand Recognition

**Overall numbers are weighted averages of individual ad's numbers

Table 6: Brand Recognition

	Le Mél	Tout-en-Ville	Nouba	Cortal	@près l'école
Le Mél	180	1	0	0	1
Tout-en-Ville	1	134	8	1	1
Nouba	0	4	78	0	1
Cortal	0	0	0	106	1
@près l'école	0	0	1	0	168
Voilà	6	2	15	0	1
Similar brands	4	3	2	2	0
Wanado	1	0	0	0	7
France Telecom	4	2	0	0	1
Honda	0	0	1	0	1
Other	3	1	3	3	2
Don't Remember	71	62	83	63	49
Total	270	209	191	175	233

Table 7: Aided Brand Awareness

Brand	Pre	Post	Δ	$\Delta\%$	# of Exposures
Alta Vista	95.8%	95.2%	-0.6%	-0.6%	0
Nomade	85.0%	83.9%	-1.1%	-1.3%	
Lycos	81.5%	78.9%	-2.6%	-3.2%	
Yahoo!	98.5%	96.7%	-1.8%	-1.8%	
Le Mél	47.6%	50.4%	+2.8%	+5.8%	1
Tout-en-Ville	22.3%	25.5%	+3.2%	+14.3%	
Nouba	11.3%	14.4%	+3.1%	+27.4%	
Cortal	17.2%	18.6%	+1.4%	+8.1%	2
@près l'école	21.4%	25.0%	+3.6%	+16.8%	
Voilà	78.8%	90.5%	+11.7%	+14.8%	9

Table 8: Logit Regression (Aided Brand Awareness)

Variable	Parameter Estimate	Wald χ^2	Pr $>\chi^2$
Intercept	-1.2862	413.70	0.0001
Le Mél	1.2145	243.93	0.0001
Alta Vista	4.3612	1045.96	0.0001
Nomade	3.0034	1054.84	0.0001
Voilà	2.5864	655.96	0.0001
Tout en Ville	0.0966	1.34	0.2472
Lycos	2.7124	950.62	0.0001
Cortal	-0.3280	13.98	0.0002
Nouba	-0.6627	48.19	0.0001
Yahoo!	5.0099	833.75	0.0001
TEST	-0.0499	0.90	0.3435
EXP	0.1139	38.32	0.0001
Regression		7243.60	0.0001

Table 9: Logit Regression (Aided advertising recall and brand recognition)

Variable	Aided Advertising Recall		Brand Recognition	
	Parameter Estimate	Pr > χ^2	Parameter Estimate	Pr > χ^2
Intercept	-0.7503	0.0001	0.3436	0.0001
Dummy	-0.2293	0.0021	0.3716	0.0047
N	3597		1078	

Table 10: Size and Orientation Significance Level

Construct	H_{5a} : Size		H_{5b} : Orientation	
	Pr > χ^2	N	Pr > χ^2	N
Aided Advertising Recall	0.9483	3926	0.0003	918
Brand Recall	0.6397	1018	0.1716	240
Unaided Advertising Recall	0.3056	2360	0.9314	552
Brand Awareness	0.8785	4146	0.1714	962
Overall	No Effect		Weak Effect	

Table 11: Contrast, Execution, and Animation Significance Level

Construct	H_{6a} : Contrast		H_{6b} : Animation		H_{6c} : Message	
	Pr > χ^2	N	Pr > χ^2	N	Pr > χ^2	N
Aided Advertising Recall	0.0012	1500	0.0989	900	0.0001	1182
Brand Recall	0.5407	540	0.6084	216	0.0022	418
Unaided Advertising Recall	0.1194	902	0.2726	540	0.9297	716
Brand Awareness	0.2719	1582	0.6447	944	0.6114	1266
Overall	Weak Effect		No Effect		Effect	

Table 12: Overall Banner Exposure Effects

Advertising Construct	Effect
Exposure	100.0%
Aided Advertising Recall	30.1%
Aided Brand Recognition	18.5%
Unaided Brand Recall	11.4%
Increased Brand Awareness	2.8%
Click-through*	0.6%

* From Nielsen/Netratings 1999

Figure 1: Voilà Answer page

Voilà, les réponses...

Actualité Infoville
 Annuaire Recherche
 Communication Shopping

Club-Internet l'a fait!

Le Salon Nautique prend le vent avec Voilà Achats, Cadeaux avec Voilà Noël

Votre Recherche :
 Réponses : 14.251
Louvre : 14.250
 Louvres : 294
 louvre : 1

cinéma 2
 Affiner la recherche

Autres Recherches :
 Nouveautés
 Les Pages Jaunes®
 Les Pages Blanches®
 Les Pages Marquées
 Sites Commerciaux
 Journal du Jour
 Exclusif & Gratuit !
 Offrez un véritable moteur de recherche à votre site Web.
 Moteur Personnel : toute la puissance de Voilà pour votre Web.

Réponses Chaînes et Web Web uniquement Réponses récentes

Le musée du Louvre
 [1 / 1] Visite virtuelle - Visitez le Louvre en vous promenant dans les salles et en découvrant les photographies des oeuvres d'art exposées.
 www.smartweb.fr/louvre/
 Culture_loisirs/Musees_monuments/Musees_nationaux_grands_musees

Carrousel du Louvre
 [1 / 1] Centre commercial du musée du Louvre - Guide complet des commerces et restaurants sous la Pyramide inversée, à Paris.
 www.smartweb.fr/carrusel/
 Habitat_vie_familiale/Achats_commerce/Gde_surface_gd_magasin_centre.com

Musée du Louvre
 [1 / 1] Découvrez l'histoire du musée, ses collections, les expositions temporaires, l'agenda de l'auditorium, ainsi que les visites-conférences et les ateliers. Une visite virtuelle vous permettra de découvrir les salles d'exposition.
 mistral.culture.fr/louvre/
 Culture_loisirs/Musees_monuments/Musees_nationaux_grands_musees

Restaurants du Carrousel du Louvre
 [1 / 1] La liste des restaurants dans l'enceinte du musée du Louvre, accompagnée d'un plan.
 www.smartweb.fr/carrusel/restaurants/
 Culture_loisirs/Restaurants_bar/Restaurants

Illumination du musée du Louvre
 [1 / 1] Dans le cadre du projet scientifique ISA (Image, Synthèse et Analyse), le laboratoire présente une de ses applications pour le musée du Louvre : la modélisation de la lumière du monument est exposée ici avec les distributions spatiales et spectrales des
 www.lona.fr/texte/eur/equipe/isa/Louvre/light/indexfr.html
 Culture_loisirs/Musees_monuments/Musees_nationaux_grands_musees

Les nouvelles salles des Antiquités égyptiennes
 [1 / 1] Musée du Louvre - Partez à la découverte des nouvelles salles des antiquités grecques, étrusques et romaines et de l'histoire de la construction du Louvre.
 www.smartweb.fr/louvre/special/egypte.htm
 Culture_loisirs/Musees_monuments/Musees_nationaux_grands_musees

Louvre Bible en main
 [1 / 1] Page personnelle - Ce particulier vous propose une visite guidée des départements des Antiquités grecques, égyptiennes et orientales du musée du Louvre (Paris).
 ouworld.computerweb.com/homepages/hgueusquin/
 Culture_loisirs/Musees_monuments/Musees_nationaux_grands_musees

Le Grand Louvre
 [1 / 1] Salles égyptiennes - Parcours thématique et chronologique dans les salles du Département des antiquités égyptiennes du Louvre.
 ouelleweb.stoa.fr/louvre.htm
 Culture_loisirs/Musees_monuments/Musees_nationaux_grands_musees

L'Egypte au Louvre
 [1 / 1] Journal Le Monde - La réouverture des salles égyptiennes du Musée du Louvre est l'occasion de dresser un bilan des connaissances actuelles en matière de chronologie des

Zone de machine locale


Figure 2: Voilà Answer zone definition


Figure 3: Frequency distribution of the number of banners looked at

Figure 4: Experts vs. Novices




Figure 5: Logit Response Function




Figure 6: Logit Response Function

