

A FIELD EXPERIMENT TO ASSESS THE INTERRUPTION EFFECT OF POP-UP PROMOTIONS

WENDY W. MOE

WENDY W. MOE

is Assistant Professor of Marketing,
Robert H. Smith School of Business,
University of Maryland, College Park,
MD 20742-1815; e-mail:
wmoe@rhsmith.umd.edu

This article highlights a large-scale field experiment conducted at an informational Web site where the timing of pop-up promotions being offered was varied. The experiment examines the Web user's reaction to the promotion in terms of (a) a direct response to the promotion (i.e., clickthrough) and (b) any indirect response in terms of the user's Web-site-exit behavior. Factors such as delay in offering the pop-up promotion and the page on which the pop-up appears are identified as variables that can be manipulated to enhance the individual's response. The context, or the page on which the pop-up promotion is offered, also is examined. Overall, this article suggests that characteristics of pop-up promotions, beyond just the message of the promotion, can be utilized to improve the consumer response.

© 2006 Wiley Periodicals, Inc. and Direct Marketing Educational Foundation, Inc.

JOURNAL OF INTERACTIVE MARKETING VOLUME 20 / NUMBER 1 / WINTER 2006

Published online in Wiley InterScience (www.interscience.wiley.com). DOI: 10.1002/dir.20054

INTRODUCTION

As the Internet matures, marketers have been developing new and more creative ways to promote to consumers online. A common practice is to pop-up a promotional window immediately when a user arrives at the home page of a particular site. Web-site visitors may find this practice is annoying and takes away from their experience at the site. As a result, these messages may be ineffective at generating the desired response (i.e., clickthrough) or may lead the user to exit the site sooner than otherwise intended. Rather than completely abandon the use of pop-up promotions, this article explores characteristics of pop-up promotions that can be manipulated to generate a more positive reaction to this marketing tool.

In particular, we conducted a field experiment at an informational site (i.e., no commerce) where the timing of a pop-up promotion, in terms of when it was offered during the user's visit, is varied. We manipulate both across-page timing (i.e., Is the message offered on the first or second page of the visit?) as well as within-page timing (i.e., Is the message offered immediately on a page or after a 15-s delay on that page?). One consequence of manipulating across-page timing is that the context of the pop-up promotion (i.e., the page on which the message appears) will naturally vary. Clickstream data from the site is then collected and analyzed.

Both direct and indirect effects are measured and analyzed as a function of promotional characteristics. Direct effects are defined as the user's response to the promotional message itself where a positive response would be clickthrough. There also may be an indirect effect where the site-exit behavior of the individual presented with the pop-up is altered. A common perspective is that pop-up promotions tend to annoy site visitors and may cause them to exit the site earlier than otherwise intended. Our results show that in several cases, the pop-up promotion can generate a positive indirect response. We find that though few factors have a direct effect on clickthrough rates of the pop-up promotion itself, several factors can lead to a pop-up promotion increasing the number of pages that an individual views at the site.

Several types of promotional tools are available online. The most commonly studied are *banner ads*,

which are ads embedded in the page being viewed. These are similar to print ads in the offline environment. Another class of promotional tools, *interstitial promotions*, has become increasingly popular and includes pop-up messages, pop-under messages, bridge pages, and in-page animations. The focus of this article is on interstitial promotional messages. Specifically, pop-up messages are used in the field experiment. The defining characteristic of this class of promotional messages (especially when compared to banner ads) is the difficulty the Web surfer has ignoring them. These promotional messages are designed to attract attention and interrupt the user's experience at the Web site. Therefore, in the next section, we highlight some research that has been conducted in both the marketing literature with respect to online consumer behavior and promotions as well as in the decision-processes literature with respect to task interruption.

LITERATURE REVIEW

There are several related streams of research that may apply to the research question at hand. The first and most obvious is the research on online consumer behavior.

Online Consumer Behavior

The rapidly growing research on online consumer behavior has focused primarily on search and purchasing behavior across stores (Johnson, Moe, Fader, Bellman, & Lohse, 2004; Lynch & Ariely, 2000; Winer et al., 1997), over time (Moe & Fader, 2004a), and within session (Bucklin & Sismeiro 2003; Li, Montgomery, Srinivasan, & Liechty, 2002; Moe, 2003). Fewer studies have examined the use of and market response to marketing interventions such as pricing (Smith & Brynjolffson, 2000), recommendation engines (Ansari, Essegai, & Kohli, 2000), and banner ads (Dahlen & Bergendahl, 2001). Little to no research has been done on the impact of interstitial promotions. By studying the effect of interstitial promotions on direct promotional response and site exit, we hope to contribute significantly to the online consumer behavior literature.

In terms of methodology and data collection, several studies have used clickstream data which is quickly becoming more widely available to both researchers and practitioners (see Bucklin & Sismeiro, 2003;

Johnson et al., 2004; Li et al., 2002; Moe & Fader, 2004a, 2004b); however, research on clickstream data is often handicapped by the descriptive nature of secondary data. Though clickstream data do contain a lot of information pertaining to the consumer buying process, it is often difficult to draw cause-and-effect conclusions. To overcome this weakness, several researchers have turned to simulating an online environment in an experimental lab setting (e.g., Lynch & Ariely, 2000). But this method is not without its disadvantages. Often, these lab settings are unrealistic and do not reflect the true behavior of consumers. Because of the drawbacks of both clickstream and experimental data, researchers have called on marketing practitioners to conduct experiments on their Web sites; however, many sites are hesitant to do so for fear of annoying site visitors with an inconsistent experience. This article fills that gap and highlights a large-scale field experiment at a high-traffic Web site.

Task Interruption

One unique characteristic of an interstitial promotion is that when it pops up, it interrupts the user's experience at the Web site. Therefore, it is useful to review some of the decision-processes literature that addresses the effect of task interruption. Several studies have asked participants to complete a task during which they are interrupted with other information. These articles aimed to identify the determinants of how individuals respond to task interruptions.

Kirmeyer (1988) identified characteristics of the interruption such as frequency, duration, context, complexity, and timing as factors affecting a participant's response to a task interruption. For the current study, we consider both the effect of timing as well as context. Speier and Valacich (1999), on the other hand, focused on examining the characteristics of the original task and its effect on participants' responses to an interruption. They proposed an information overload perspective, suggesting that participants performing more complicated tasks are more likely to be overloaded by an interruption and thereby react negatively. In contrast, Zijlstra, Roe, Leonora, and Krediet (1999) found that many people when faced with a task will overcompensate when interrupted; that is, they tend to concentrate even more on the task at hand. This suggests that under some circumstances, an interruption can lead to a positive response.

These task-interruption studies are analogous to the experience that Web users face when presented with pop-up promotions. Users visit a particular Web site oftentimes with a task in mind, be it browsing for entertainment purposes or searching for specific information. Pop-up messages effectively interrupt the process. How site visitors respond to this interruption will vary depending on characteristics of the pop-up promotion such as timing and context, two characteristics previously addressed in the task-interruption literature.

DATA

Description of Site

The Web site studied in this article is a purely informational site; that is, no commerce is transacted directly through the site at the time of the study. It is a well-trafficked site (i.e., over 2.7 million readers per month in 2002 were claimed) that provides information about movies both in the theaters and on DVD, including critic reviews, trailers, actor biographies, and so on.

The organization of the site is best understood by categorizing its pages as either *content* pages or *gateway* pages. Content pages are pages that provide movie-specific information. Each movie has multiple unique content pages that the user can drilldown and view. Gateway pages, on the other hand, are navigational pages that link the user to the content pages. For example, the homepage is considered a gateway page since it provides no in-depth discussion of any specific movie but instead highlights a few new releases, upcoming movies, and/or actors. To read more about the featured information, the user must click through on the hyperlink, which redirects the user to a movie-specific content page. The site also provides pages that summarize the Top 10 movies for the week, a full listing of new releases for the month, and so on. These pages also would be considered gateway pages since their purpose is purely navigational and are supposed to redirect the Web user to a product-specific page. This structure of gateway and content pages is common across informational Web sites. News sites, for example, also follow this format. When visiting the CNN.com homepage, the user is faced with a number of headlines—each of which is a hyperlink that redirects the user to the full story. In that case, the

homepage is a gateway page, and the news story itself is a content page. The site also has other gateway pages that summarize world headlines or sports headlines that also redirect the browser to specific stories of interest.

Experimental Design

Clickstream data were collected for 5 nonholiday weekdays in December 2000.¹ The experiment itself ran over a span of 4 days. For the experiment, a pop-up message was designed to solicit site visitors to subscribe to a weekly newsletter. There was no monetary value associated with this offer, only the promise of receiving weekly e-mail messages about movies. The pop-up was offered only to those visitors who had not previously registered with the site, signed up for the newsletter, or received a similar offer in the past. A fifth day during which no promotions were offered was used as a control sample.

The experiment was designed to manipulate the timing of pop-up messages. Timing was varied along two dimensions: across-page delay and within-page delay. Three levels of across-page delay were used. Browsers were presented with the pop-up either on the first page, the second page, or the fourth page of their session. Within-page delay also had three levels. The pop-up appeared either immediately on a given page, after a 15-s delay, or after a 30-s delay. The within-page delay allows for individuals to process some of the information on the page before being interrupted by the pop-up. Our analysis showed that there were no significant differences between a 15-s or a 30-s delay. Therefore, the within-page delay variable was recoded to be an indicator variable where 0 represented no within-page delay and 1 represented a 15-s or a 30-s delay.

Another factor of interest and one that has been shown to have a significant effect on a user's response to task interruption is context. In our case, we tested whether the effect of the pop-up differs depending on the type of page (i.e., gateway vs. content page) during which it is offered. This factor will not be actively manipulated as part of the experimental design but will naturally vary across the experimental delay conditions since each

¹No significant day of week or time of day effects were apparent from the analysis of the data.

TABLE 1

Experimental Design

	0-S DELAY	15-S DELAY	30-S DELAY
1 page viewed	Night <i>n</i> = 6,226	Day <i>n</i> = 8,826	Evening <i>n</i> = 5,788
2 pages viewed	Evening <i>n</i> = 3,426	Night <i>n</i> = 1,765	Day <i>n</i> = 5,978
4 pages viewed	Day <i>n</i> = 2,857	Evening <i>n</i> = 1,583	Night <i>n</i> = 2,385

site visit follows a different navigational path. For example, if the experimental condition is one where the pop-up is offered on the second page of the visit, that could mean that it is offered on a content page for one visitor versus a gateway page for another visitor.

There also is reason to believe that online behavior varies by time of day (Telang, Boatwright, & Mukhopadhyay, 2002). Someone searching for information online during the day is likely to exhibit very different behavior from someone searching in the middle of the night. Therefore, time of day is another factor that we must balance in our experiment. Ideally, we should randomly assign visitors regardless of time of day into one of our page-depth/page-delay conditions; however, due to technology constraints at the Web site, we had to define experimental conditions by time of day. Each calendar day was divided into daytime (9 a.m.–5 p.m.), evening (5 p.m.–midnight), and nighttime (midnight–9 a.m.) as requested by the Web site.² Effectively, we had three factors to manipulate (across-page delay, within-page delay, and time of day), each with three levels. Therefore, a Latin square design was implemented. Table 1 describes each experimental condition and the number of observations collected in each condition.

Data Collection and Measures

For the duration of the experiment, clickstream data were collected via a cookie on the visitor's PC. The

²Ultimately, time of day had no significant effect on either direct or indirect response and therefore will not be discussed as a factor in the analysis presented in this article.

TABLE 2

Descriptive Statistics

	CONTROL	EXPERIMENTAL	RECEIVED POP-UP	DID NOT RECEIVE POP-UP
Number of sessions	44,954	83,136	38,834	44,302
Average no. of pages viewed	3.45	3.47	5.86	1.38
Percent of content pages	67.52%	68.26%	69.04%	67.58%

clickstream data recorded each and every page requested by visitors. That is, each line of data represented a page view. It provided the date and time of the page view and contained information regarding the type of page (content vs. gateway). Each page view was associated with a unique user number specific to the PC used to browse the site and a unique session number. Sessions were defined by the Web site as consecutive pages requested no more than 1 hr apart. If a user is idle for 1 hr or longer, any subsequent page requested by that user would be considered a new session.

During the 4-day experiment, 83,136 nonregistered users visited the Web site and were eligible for the experiment (i.e., they had not previously received a pop-up offer to register for the newsletter). On average, these visitors viewed 3.47 pages during a given session. Of these 83,136 visitors, 38,834 received a pop-up message during their visit. Others were eligible for a pop-up, but were assigned to an experimental condition that did not trigger a promotion during their stay at the site. For example, a visitor may be assigned to the condition where a pop-up is offered on the fourth page, but he or she exits after the second page. Though this visit was assigned to one of the experimental conditions, the pop-up itself was never triggered and therefore never seen by the visitor. As a consequence, sessions affected by the experiment (with the exception of those in the one page/0-s delay condition) are in effect a conditional sample. We will address this issue later in the article when we conduct our hypotheses testing. Finally, a total of 44,954 individuals were in the control group, with an average of 3.45 pages viewed per session.

Table 2 provides some descriptive statistics of the data. A comparison of the sessions conducted during the control period versus those conducted during the experimental period reveal no significant differences in either the average number of pages viewed or the proportion of pages that are content versus gateway pages. The last two columns of the table decompose the sessions conducted during the experimental period into those that received a pop-up and those that did not. The difference in the number of pages viewed is a direct result of how pop-up promotions are triggered during the experimental period. The percent of content pages indicates no significant difference in the nature of the sessions.

Two dependent measures were obtained from the data. The direct effect is represented by the click-through rate for the promotion. The indirect effect is represented by the number of pages viewed per session. Additionally, two categories of independent measures were used to represent promotion characteristics: (a) timing measures, which include across-page delay (1, 2, or 4 pages) and within-page delay (0 or 1); and (b) a measure of context. Specifically, the type of page on which the promotion is offered is coded as either a content page or a gateway page.

Hypotheses

Previous research in task interruption has shown that context and timing of the interruption (or in our case, the pop-up message) can have an effect on how the individual responds to the interruption (Kirmeyer, 1988). Therefore, in this section, we consider how both direct and indirect responses to pop-up

messages are affected by these two characteristics. In developing hypotheses, two perspectives are discussed. The first perspective is the idea that interruptions have a negative effect on the individual. The context and timing of the pop-up determines the likelihood of that pop-up being perceived as an interruption rather than just additional information. The second perspective is the information-overload perspective presented by Speier and Valacich (1999). The argument here is that a pop-up resulting in information overload will likely generate a negative response.

We begin by examining the direct effect of context by considering whether the pop-up promotion appears on a content or a gateway page. These two types of pages have very distinct roles in the Web-site visitor's navigational process. Gateway pages are typically used to help navigate the site to arrive at the desired content page. Therefore, by definition, any pop-up appearing on a gateway page is more likely to present itself as an interruption, and an individual is more likely to ignore a pop-up if it is an interruption so that he or she may continue with the task at hand, leading to a more negative direct response, or lower clickthrough rate. Thus, we hypothesize that:

H1a: Pop-up messages appearing on gateway pages will lead to lower clickthrough rates than pop-up messages appearing on content pages.

However, if we consider the information-overload perspective, an alternative hypothesis results. Consistent with this framework, one could argue that individuals will respond more negatively to a pop-up message that adds to their current information load to the point of overload. Since the amount of information provided by the pop-up message itself remains constant across individuals in our experiment, the determinant of whether the individual becomes overloaded by information lies in the content of the page being viewed when the pop-up appears. When the viewed page contains more content, the pop-up is more likely to overload the individual. As a result, the individual is more likely to ignore the pop-up message itself, leading to a lower clickthrough rate.

H1b: Pop-up messages appearing on content pages will lead to lower clickthrough rates than pop-up messages appearing on gateway pages.

In addition to measuring the direct clickthrough response to pop-up promotions, we also are interested in the indirect effects. That is, how does the pop-up message, regardless of whether the individual clicks through on the pop-up itself, affect the user's exit behavior at the Web site? Web sites are continually trying to devise ways to keep visitors at their sites, creating more opportunities to engage and promote to the visitor. Can pop-up promotions affect the visitors length of stay at the site? Additionally, do context and timing make a difference?

Again, if we assume that a pop-up promotion appearing on a content page is less likely perceived as an interruption than when appearing on a gateway page, then it should have a more positive effect on the individual's overall experience at the site, leading to a longer stay at the site (or a higher page count).

H2a: Individuals receiving a pop-up promotion on a content page will view more pages than those receiving a pop-up promotion on a gateway page.

Alternatively, if a pop-up appears on a content page, it is more likely to overload an individual. Not only will this have a negative impact on the direct response to the pop-up itself, the information overload also may negatively affect the individual's overall experience at the site, which may lead to the individual exiting the site earlier than otherwise intended.

H2b: Individuals receiving a pop-up promotion on a content page will view fewer pages than those receiving a pop-up promotion on a gateway page.

The next characteristic of the promotion we consider is timing. Common practice is to present the pop-up message immediately on the home page, typically the first page viewed by the Web-site visitor; however, recent practices also include offering pop-up promotions on later pages in an individual's visit. Therefore, we consider the effect of across-page delay. In other words, is a pop-up promotion appearing on first page more or less effective than one appearing on the fourth page, for example? Additionally, the Web site can vary the timing of a pop-up message even within a given page. Again, common practice is to present the pop-up message immediately once a page is requested; however, a slight delay would give the individual an opportunity to view the contents of the page before redirecting his

or her attention to the pop-up message. We will refer to this type of delay as a within-page delay.

We begin by discussing the effects of a within-page delay. One could argue that delaying the pop-up message within a page is more likely to interrupt the site visitor in the process of reading the contents of the page. As a result, the delayed pop-up would be more likely perceived as an interruption in the individual's information-processing efforts, thereby causing a negative response to the pop-up. Alternatively, if time were to lapse before a pop-up appears, that could give the visitor more time to process the information on that page, decreasing the likelihood of information overload and causing a more positive response to the pop-up. These arguments lead to the following competing hypotheses.

H3a: Within-page delay increases the click-through rate of a pop-up promotion.

H3b: Within-page delay decreases the click-through rate of a pop-up promotion.

Similar arguments can be made for across-page delay effects on clickthrough response:

H4a: Across-page delay increases the clickthrough rate of a pop-up promotion.

H4b: Across-page delay decreases the clickthrough rate of a pop-up promotion.

In terms of timing effects on site-exit behavior, we consider how the individual is likely to react to an interruption with respect to the original task at hand. The research by Zijlstra et al. (1999) is particularly relevant. This research showed that when interrupted, many people will overcompensate and concentrate even more on the task at hand regardless of their direct response to the interruption itself. This suggests the following hypotheses with respect to delay.

H5a: A pop-up message appearing after a within-page delay will increase the number of pages viewed when compared to one without a delay.

H6a: A pop-up message appearing on a later page (across-page delay) will increase the number of pages viewed when compared to one without a delay.

In addition to the task interruption research, we also have to consider some relevant research on Internet browsing behavior. Specifically, Hoffman and Novak

(1996) proposed that online browsers enter a state of flow in which they are heavily engaged with their Internet experience. One characteristic of this state of flow is that the individual is not aware of the passage of time as he or she browses from page to page. A pop-up message that disrupts this flow may lead to the individual exiting the site earlier than otherwise intended. Any delay (either across-page or within-page) is more likely to disrupt a state of flow than no delay. This generates the following counterhypotheses.

H5b: A pop-up message appearing after a within-page delay will decrease the number of pages viewed when compared to one without a delay.

H6b: A pop-up message appearing on a later page (across-page delay) will decrease the number of pages viewed when compared to one without a delay.

Finally, we hypothesize that context and across-page delay will have an interaction effect on the individual's response to the pop-up message. Thus far, we have argued that an across-page delay in offering the pop-up message can have both a positive or a negative effect on the individual's response to the pop-up depending on the degree to which it is perceived as an interruption. A delayed pop-up may be more likely perceived as an interruption of a task already in process; however, task objectives at an informational-content site like the one in this study tend to involve viewing a content page. Therefore, if a pop-up is delayed to a point where it appears on a content page, it becomes less likely that the pop-up is interrupting a task and more likely that the pop-up appears at the end of a task.

H7: A pop-up message that is delayed to appear on a content page will generate a more positive direct response than one that is delayed to appear on a gateway page.

H8: A pop-up message that is delayed to appear on a content page will generate a more positive indirect response than one that is delayed to appear on a gateway page.

RESULTS

We first examined any direct responses to the pop-up message as measured by clickthrough rates. We found that the only significant main effect is that of within-page delay. Across-page delay and context had

no significant main effect. Overall, clickthrough rates for the pop-up message were 0.81%. Although this appears low, it is not out of line with typical online clickthrough rates (DoubleClick reported an average clickthrough rate of .41% in the third quarter of 2004.) By offering the pop-up message immediately within a page, clickthrough rates increase to 1.14% compared to a clickthrough rate of .65% with a slight delay ($p < .0001$). This supports Hypothesis H3a and suggests that a delayed pop-up is more likely to act as an interruption. This implies that efforts should be made to shorten the delay time when a pop-up message appears. There is no evidence that context (H1) or across-page delays (H4) has any effect, positive or negative, on clickthrough rates nor is there evidence of any significant interaction effects (H7).

With respect to indirect effects, we examined site-exit behavior by modeling the effects of delay and context on page count (i.e., the number of pages viewed by those in the experimental group). This analysis will highlight the differential effects of each of the experimental conditions on page count.

Consistent with H2a, a pop-up that appears on a content page versus one that appears on a gateway page tends to increase the number of pages from 4.86 to 6.31 pages ($p < .0001$). This would suggest that a pop-up appearing on a content page is less likely to be perceived as an interruption and more likely to enhance the online browsing experience than one that appears on a gateway page. This also would suggest that at least in this case, there is minimal risk of information overload by presenting the given pop-up message on a content-heavy page.

Interestingly, delay affects page count differently depending on whether it is an across-page delay or a within-page delay. Delay within a page tends to accelerate the individual's site exit from an average of 6.39 pages viewed when presented with an immediate pop-up to 5.61 pages viewed when presented with a delayed pop-up ($p = .0295$). This supports Hypothesis H5b and suggests that providing time for an individual to process the information on any given page before presenting the pop-up is viewed more as an interruption that negatively affects the Web-site experience. Again, this suggests that there is minimal risk of overloading the individual with the added information from the pop-up message since providing

the extra time to process the information does not improve the experience. Instead, the driving factor as to how the pop-up affects the Web-site experience is the degree to which it is perceived as an interruption. The results show that a within-page delay increases the likelihood of the pop-up acting as an interruption.

A pop-up that is delayed to a later page appears to increase the number of pages viewed to 7.53 with a two-page delay and 11.66 with a four-page delay when compared to a pop-up offered on the first page (3.06 pages); however, there is a statistical effect that is contributing to this apparent pattern. On average, those sessions in the two-page delay condition will by definition have more pages viewed than those sessions in the one-page condition simply because an individual must view at least two pages to be offered the pop-up in the two-page delay condition, and likewise for the four-page delay condition. Therefore, when examining the effect of delay on page count, we must consider a conditional control sample.³ Therefore, for the remainder of this section, we discuss comparisons to a conditional control group in addition to differential effects of the experimental conditions.

As mentioned earlier, the control group in our study consists of sessions occurring on a day when no pop-up messages appeared on the site. When compared to this control group, promotions, regardless of the experimental condition, increased the page count from an average of 3.45 to 5.86 pages ($p < .001$). While there are main effects of context and within-page delay that indicate differential effects of the experimental conditions, all experimental conditions for these two factors tended to increase the page count over the control group. This has significant implications, suggesting that pop-up promotions are not necessarily viewed as nuisances that interrupt the user's experience at the site. Instead, under certain circumstances, they actually can enhance the user's experience and encourage a longer visit.

Note that this increase in page count is mostly due to an increase in viewing content pages as opposed to gateway pages. Table 3 illustrates the main effect of

³ When examining direct responses, we compared only between experimental conditions and did not compare to a control sample. Since no promotions were offered to the control group, we cannot measure clickthrough rates.

TABLE 3 Effect of Context on Pagecount and the Percent of Content Pages

	CONTENT PAGE	GATEWAY PAGE
Number of pages viewed	6.31	4.86
Percent of content pages	36.50%	21.33%

Note: The percent of content pages is calculated by dividing the number of content pages by the total number of pages views. Excluded from the calculation is the page on which the pop-up appeared.

TABLE 4 Main Effect of Across-Page Delay on Pagecount

PROMOTION APPEARS ON:	MEAN	CONTROL MEAN	p-VALUE
Page 1	3.06	3.45	<0.0001
Page 2	7.53	7.07	0.0001
Page 4	11.66	11.59	0.3578

context on both page count and the percentage of pages viewed that are content pages. The implication is that not only are the visitors staying at the site longer but the clickstream behavior reveals a deeper level of involvement where the visitors are engaged in content at the site.

Table 4 provides the mean page count for each of the across-page delay conditions and their conditional control samples. Since the across-page delay condition affects only those site visitors who view at least as many pages as those in the delay condition, we need to construct a condition control sample for each across-page delay condition. For example, the one-page delay condition will affect all sessions in which visitors view at least one page during that phase of experiment. Therefore, the appropriate comparison would be against all sessions in our control group. However, the two-page delay condition will affect only sessions in which visitors view at least two pages. Therefore, the appropriate comparison for the two-page delay condition would be all sessions in which visitors view at least two pages in the control sample, and likewise for the four-page delay condition. Across-page delay appears to increase page count; however, the increase,

though significant, is very slight. This provides some degree of support for Hypothesis H6a.

Although we have not hypothesized an interaction effect between within-page delay and context, there is a slight, but still statistically significant, effect. Figure 1 illustrates the effects of within-page delay and context on page count. From this figure, we see the main effect of both context and within-page delay mentioned earlier. Though not obvious from the figure, the interaction effect between context and within-page delay is significant ($p = .0314$). The results show that a pop-up message that appears with no within-page delay on a gateway page has an added positive effect on page counts beyond the additive main effects. Stated another way, a delayed pop-up presented on a content page is considered more of an interruption than one presented on a gateway page.

Figure 2 plots the effects of across-page delay and context on page count against a conditional control sample. Again, Figure 2 illustrates the main effect of context. The effect of across-page delay is less




FIGURE 1 Effects of Within-Page Delay and Context on the Number of Page Count




FIGURE 2 Effects of Across-Page Delay and Context on Page Count

straightforward. Having already established that there is a main effect of across-page delay, we turn our attention to the interaction effect between across-page delay and context ($p < .0001$). When the pop-up is delayed and appears on a gateway page, the average page count is lower compared to the control group; however, when a pop-up is delayed and appears on a content page, the page count actually increases relative to the control group, suggesting that not only is the pop-up not perceived as a negative interruption but it can encourage continued page views. Both differences are significant and provide support for Hypothesis H8.

DISCUSSION

The field experiment presented in this article studied how timing and context of a pop-up promotion affects the Web-site-visitors' direct clickthrough response as well as their indirect exit response. Typically, experiments and studies aimed at assessing the effectiveness of promotions focus on direct responses only. In this article, we showed that direct clickthrough response is only part of the picture and that a substantial indirect effect can exist. Specifically, we tested the effects of three factors: the page on which the promotion appears, across-page delay, and within-page delay. The results showed that within-page delay is the only one of our three factors that had an effect on clickthrough rates. In fact, within-page delay had a negative effect on both clickthrough rates and exit behavior. Consequently, Web sites should make every effort to minimize any delay in delivery of a pop-up promotion on a given page. From a behavioral perspective, it seems that a within-page delay is more likely to result in information overload, leading to not only a lower clickthrough response but also a more negative experience at the site.

Across-page delay and context, according to the results in our experiment, had no effect on clickthrough rates but do affect the number of pages viewed. In fact, delaying the pop-up promotion to a later page and/or presenting the pop-up promotion on a content page increased the number of pages viewed by the Web-site visitor. Note that not only does a pop-up that is delayed to a later page and/or presented on a content page generate a more positive response than one that appears on an earlier page and/or on a

gateway page but it actually induced more page views than the those in the control group. In other words, a well-timed pop-up promotion can actually *enhance* the visitor's experience at the site and encourage him or her to stay a little longer.

From a behavioral perspective, there are two possible explanations for this response. First, individuals become more involved when faced with an interruption that possesses these two characteristics (across-page delay and appearing on a content page). This would be consistent with the results of Zijlstra et al. (1999), which showed that participants concentrated even more on a task when they were interrupted. Second, pop-up promotions with these two characteristics are less likely to be perceived as interruptions, thereby avoiding any of the negative reactions that arise from being interrupted.

Regardless of the behavioral explanation, the managerial implication is clear. Pop-up promotions that are delayed to appear later during an individual's visit and/or on a content page are better received. Though this has no impact on clickthrough rates, it does seem to have a positive effect on the visitor's overall experience at the site as demonstrated by a longer visit in terms of page count. However, when considering a delay in presenting pop-up promotions, the Web site must weigh the aforementioned benefits against the attrition of visitors from page to page. That is, if the pop-up is delayed to the fourth page, for example, how many fewer visitors will see the pop-up when compared to one that appears on the first page? If the pop-up appears on the first page, the benefit to the Web site is the clickthrough rate applied to all visitors to the site. If the pop-up is delayed, the benefit to the Web site is the clickthrough rate applied to a subset of the site's visitors plus any benefit to the site resulting from the added page views from those visitors.

The results of the experiment presented in this article show that a well-timed pop-up promotion can enhance a visitor's Web-site experience and keep him or her at the site for more page views. Contrary to popular belief that pop-ups are at best ineffective (because of the low clickthrough rates) and annoying at worse, the results of this experiment suggest that there are positive indirect effects that may be underestimated. Little research has been conducted to explore the effect of

pop-up promotions beyond simply looking at click-through rates. Further research along these lines could improve our understanding of how individuals respond to promotional interruptions and help marketers design more effective promotional tactics.

REFERENCES

- Ansari, A., Essegai, S., & Kohli, R. (2000). Internet Recommendation Systems. *Journal of Marketing Research*, 37(3), 363–375.
- Bucklin, R.E., & Sismeiro, C. (2003). A Model of Web Site Browsing Behavior Estimated on Clickstream Data. *Journal of Marketing Research*, 40(3), 249–267.
- Dahlen, M., & Bergendahl, J. (2001). Informing and Transforming on the Web: An Empirical Study of Response to Banner Ads for Functional and Expressive Products. *International Journal of Advertising*, 20(2), 189–205.
- DoubleClick. (2004). Ad Serving Trend Report. Retrieved November, 200x, from http://www.doubleclick.com/us/knowledge_central/documents/trend_reports/dc_q304_adservingtrends_0411.pdf
- Hoffman, D.L., & Novak, T.P. (1996). Marketing in Hypermedia Computer-Mediated Environments: Conceptual Foundations. *Journal of Marketing*, 60(3), 50–68.
- Johnson, E.J., Moe, W.W., Fader, P.S., Bellman, S., & Lohse, J. (2004). On the Depth and Dynamics of World Wide Web Shopping Behavior. *Management Science*, 50(3), 299–308.
- Kirmeyer, S.L. (1988) Coping With Competing Demands: Interruptions and the Type A Pattern. *Journal of Applied Psychology*, 73(4), 621–629.
- Li, S., Montgomery, A., Srinivasan, K., & Liechty, J.C. (2002). Modeling Online Browsing and Path Analysis Using Clickstream Data. *Marketing Science*, 23(4), 579–595.
- Lynch, J.G., & Ariely, D. (2000). Wine Online: Search Cost Affect Competition on Price, Quality, and Distribution. *Marketing Science*, 19(1), 83–103.
- Moe, W.W. (2003). Buying, Searching, or Browsing: Differentiating Between Online Shoppers Using In-Store Navigational Clickstream. *Journal of Consumer Psychology*, 13(1&2), 29–40.
- Moe, W.W., & Fader, P.S. (2004a). Dynamic Conversion Behavior at e-Commerce Sites. *Management Science*, 50(3), 326–335.
- Moe, W.W., & Fader, P.S. (2004b). Capturing Evolving Visit Behavior in Clickstream Data. *Journal of Interactive Marketing*, 18(1), 5–19.
- Smith, M., & Brynjolfsson, E. (2000). Frictionless Commerce? A Comparison of Internet and Conventional Retailers. *Management Science*, 46(4), 563–585.
- Speier, C., & Valacich, J.S. (1999). The Influence of Task Interruption on Individual Decision Making: An Information Overload Perspective. *Decision Sciences*, 30(2), 337–360.
- Telang, R., Boatwright, P., & Mukhopadhyay, T. (2004). A Mixture Model for Internet Search Engine Visits. *Journal of Marketing Research*, 41(2), 206–214.
- Winer, R.S., Deighton, J., Gupta, S., Johnson, E.J., Mellers, B., Morwitz, V.G., et al. (1997). Choice in Computer-Mediated Environments. *Marketing Letters*, 8(July), 287–296.
- Zijlstra, F.R.H., Roe, R.A., Leonora, A.B., & Krediet, I. (1999). Temporal Factors in Mental Work: Effects of Interrupted Activities. *Journal of Occupational and Organizational Psychology*, 72, 163–185.