

Improving the Effectiveness of Time-Based Display Advertising

DANIEL G. GOLDSTEIN, Microsoft Research, New York City
 R. PRESTON MCAFEE, Microsoft Corporation, Redmond
 SIDDHARTH SURI, Microsoft Research, New York City

Display advertisements are typically sold by the impression, where one impression is simply one download of an ad. Previous work has shown that the longer an ad is in view, the more likely a user is to remember it and that there are diminishing returns to increased exposure time [Goldstein et al. 2011]. Since a pricing scheme that is at least partially based on time is more exact than one based solely on impressions, time-based advertising may become an industry standard. We answer an open question concerning time-based pricing schemes: how should time slots for advertisements be divided? We provide evidence that ads can be scheduled in a way that leads to greater total recollection, which advertisers value, and increased revenue, which publishers value. We document two main findings. First, we show that displaying two shorter ads results in more total recollection than displaying one longer ad of twice the duration. Second, we show that this effect disappears as the duration of these ads increases. We conclude with a theoretical prediction regarding the circumstances under which the display advertising industry would benefit if it moved to a partially or fully time-based standard.

Categories and Subject Descriptors: J.4 [Social and Behavioral Sciences]: Economics

General Terms: Economics, Experimentation

Additional Key Words and Phrases: Display, advertising, memory, recall, recognition, exposure, time

ACM Reference Format:

Goldstein, D. G., McAfee, R. P., and Suri, S. 2015. Improving the effectiveness of time-based display advertising. *ACM Trans. Econ. Comp.* 3, 2, Article 7 (April 2015), 20 pages.
 DOI: <http://dx.doi.org/10.1145/2716323>

1. INTRODUCTION

Display advertising is a \$10 billion dollar per year industry [PricewaterhouseCoopers 2011] in which advertisers pay publishers to place ads alongside content on publisher websites. Display ads are typically sold on a per impression basis, where one impression is simply one download of an ad. Furthermore, agreements between publishers and advertisers are often made through guaranteed contracts. For example, a publisher might agree to deliver 10 million impressions to men age 50–70 on finance-related pages, or 8 million impressions to people interested in sports. Display ads are also sold on exchanges like Google’s DoubleClick exchange (GDC) or Yahoo!’s Right Media Exchange (RMX). In this context, a publisher would auction the right to show a display ad targeted to a specific user in real time. Whether display ads are sold via contract or via an exchange, they are usually sold by the impression.

A previous version of this article appeared in the *Proceedings of the 13th ACM Conference on Electronic Commerce (EC’12)*.

Authors’ addresses: dgg@microsoft.com, preston@mcafee.cc, suri@microsoft.com.

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than ACM must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from permissions@acm.org.

2015 Copyright held by the owner/author(s). Publication rights licensed to ACM.

2167-8375/2015/04-ART7 \$15.00

DOI: <http://dx.doi.org/10.1145/2716323>

In addition to increasing short-term sales [Lewis and Reiley 2011; Manchanda et al. 2006], advertisers seek to increase brand recognition and brand awareness with display ads [Dr ze and Hussherr 2003]. Accordingly, advertisers often measure the effectiveness of brand advertising using memory metrics [Wells 2000], which are proxies for ad effectiveness that have been in use for nearly a century [Starch 1923]. Our previous work established a causal link between the amount of time an ad is in view and the probability that a user will recall or recognize it [Goldstein et al. 2011]. In addition, this research found that there are diminishing returns to increased exposure time. That is, there was a steep increase in the probability of remembering for exposure times up to roughly 40 seconds, followed by a less steep, although still increasing, effect of time beyond that. These results suggest that time of exposure combined with the number of impressions delivered may be a better standard for pricing ads since it causally influences and more exactly captures the recognition and recall that display advertisers seek.

In the present work, we investigate how a publisher might sell time-based ads to increase the total effect on memory per unit of time, a topic of interest to both advertisers and publishers. The central question we address is how advertisers and publishers should set the duration and scheduling of display ads. Specifically, when scheduling a fixed amount of time, we ask if it is better to show multiple ads of shorter duration or fewer ads of longer duration. Longer duration ads might increase the likelihood that a user sees each ad, however, a sequence of shorter ads gives users more ads to notice. It is not *a priori* clear which would result in more overall recollection.

To answer this question we conduct an online behavioral experiment in which participants are asked to read an online news article while ads are displayed according to one of two randomly assigned schedules. In one condition, one display ad is shown for t seconds and then replaced by another display ad for t seconds. In the other condition, a sole display ad is shown for $2t$ seconds. Two phenomena emerge. First, we find that two shorter ads result in greater recollection than one ad of twice the duration. Second, we find that this effect disappears as the durations of the ads, t , increases. This is due, at least in part, to the effect of an ad's onset time, which is defined as the amount of time between the page loading and the ad loading. We will show the longer an ad's onset time, the less likely it is to be remembered. We conclude by providing theoretical conditions under which the display advertising industry as a whole would benefit if time-based pricing were adopted.

1.1. Problem Definition

We want to assess the value of switching from an impression-based advertising system to a time-based advertising system. Because impressions last for a variable length of time, there is an inherent randomness in impression-based systems, in which some impressions last a few minutes while others last a few seconds; it may be useful to divide longer impressions across several ads.

To assess the value of doing so, we compare an ad of long duration to two shorter ads. Consider two impressions, shown to two users u_1 and u_2 , which last for $2t$ seconds each. Suppose that, under impression-based advertising, u_1 sees ad A and u_2 sees ad B . An alternative that becomes feasible with time-based advertising is for u_1 to see ad A for t seconds followed by ad B for t seconds, and u_2 to see ad B for t seconds followed by ad A for t seconds. The purpose of this article is to discover which of these two schedules results in a higher probability of remembering the ad.

More formally, we will compare:

$$\begin{aligned} & \Pr(u_1 \text{ remembers } A \mid A \text{ shown for } 2t \text{ seconds}) + \\ & \Pr(u_2 \text{ remembers } A \mid B \text{ shown for } 2t \text{ seconds}) \end{aligned} \tag{1}$$




Fig. 1. Distribution of time spent, and therefore impression lengths, on a popular Yahoo! site.

with:

$$\Pr(u_1 \text{ remembers } A \mid A \text{ shown for } t \text{ seconds followed by } B \text{ for } t \text{ seconds}) + \Pr(u_2 \text{ remembers } A \mid B \text{ shown for } t \text{ seconds followed by } A \text{ for } t \text{ seconds}). \quad (2)$$

In the next section we precisely define the standard industry metrics that we use to measure memory. Observe that this comparison holds the total time the two ads are in view constant. It also holds the within-impression timing constant. That is, in both schedules each ad is shown from 0 to t seconds, and from t to $2t$ seconds exactly once. The second term in Equation (1) may be considered the false positive rate where one claims to remember ads that were not displayed perhaps due to having seen the ad elsewhere.

Inherent in the comparison of Equation (1) and Equation (2) is the value of showing shorter duration ads to more users, which becomes feasible under time-based scheduling. For example, consider two campaigns, each showing ten million impressions. These impressions can be approximately divided in half where one half of the impressions show ad A first and ad B second, and the other half of the impressions show the ads in the reverse order. Some of the impressions will be of too short a duration for the user to see the second ad, thus no value will be created. But the short impressions, under a time-based system, are the same as under an impression-based system, so the comparison is moot. For the longer impressions, the comparison of Equation (1) and Equation (2) accurately captures the increase, if any, from a greater number of shorter exposures. Figure 1, replicated from Goldstein et al. [2011], gives the distribution of impression lengths on a popular Yahoo! site. Note that it is not obvious whether the schedule underlying Equation (1) or (2) is superior. Longer durations with a single ad increase the probability that one ad is remembered, while shorter durations with two ads provides two opportunities to form memory traces.

We will also investigate, theoretically, whether a time-based system increases or decreases industry revenue. We posit two relevant effects arising from using a time-based system. First, the value to advertisers is more accurately monitored, so that high value advertisers buy more time, and low value advertisers buy less. Second, there is a change in the total awareness or recall provided to advertisers, working like a quantity increase. We provide sufficient conditions when these two effects work together to increase industry revenue.

2. RELATED WORK

Before we discuss how our work relates to the literature, we define and motivate the metrics we use to gauge the effectiveness of a display advertisement. *Unaided recall* is the proportion of site viewers who report remembering an advertiser with only a minimal prompt such as “Which advertisers, if any, did you remember being present on the website?” *Recognition* metrics use probes. *Text recognition* uses the name of the advertiser as a probe, for example “Did you see a Netflix ad on the previous page?” *Visual recognition* uses an image of an ad as the probe, asking, for example, “Did you see this ad on the previous page?” followed by an image of an ad. A vast literature in Marketing studies the effectiveness of television advertising using these and related measures [Lodish et al. 1995].

Drèze and Hussherr [2003] conducted a study on banner advertisements where they advocate the use of these memory metrics in online advertising. When referring to generally affecting the memory of an ad, whether it be recall or recognition, we will use the term “recollection.” Surprisingly few studies, however, have considered the effectiveness of *online* advertising in improving recollection. We will describe the most relevant of these next.

In our previous work we showed that exposure time has a causal effect on memory for an ad [Goldstein et al. 2011], whereas prior work had established only a correlation (see Danaher and Mullarkey [2003] and commentary in Goldstein et al. [2011]). Moreover, we showed that there are diminishing returns to this effect. The first seconds of exposure caused a steep increase in the memory for an ad, and further exposure time had a smaller, albeit still increasing impact on recollection. The implication is that, given advertisers who value memory of their ads, time of exposure is a more exact measure and thus a more efficient basis for pricing: charging based on what advertisers value allows for price discrimination and efficient allocation of advertising slots. While that work laid the groundwork for this research, it gave no guidance to advertisers and publishers as to how ads should be scheduled, which is the topic of this work. More specifically, we seek to understand how to allocate time slots to influence overall recollection.

Sahni [2011] conducted a field experiment on a restaurant search website. The design of the experiment allowed for exogenous variation in the number and frequency of sponsored search ads users saw on the site over the course of two and a half months. The author did not study the exposure time of ads, but rather the amount of time between exposures. The key result of Sahni’s work is that increasing the time between exposures, up to two weeks, increases the probability of a purchasing event. Thus, this result may be viewed as a complement to ours. We investigate how the length and timing of exposures influences recollection, whereas Sahni showed how the amount of time *between* exposures influences purchasing. Braun and Moe [2011] used data on impressions and conversions at the individual level to construct a model of, among other phenomena, ad wear-out and how different creatives should be rotated. Lewis [2010] studied wear-out due to repeated exposures to display ads via a natural experiment. If the industry were to adopt a scheme utilizing multiple shorter impressions as opposed to one longer impression, the literature on advertiser wear-out will shed light on when repeated impressions would no longer be effective.

Next we will see how manipulations conducted in the domain of television advertising impact the memory metrics we use in this work. Pieters and Bijmolt [1997] showed via a field study that longer duration television ads are significantly more likely to be remembered than shorter ads. Singh and Cole [1993] conducted a lab experiment that showed longer television ads have a higher probability of correct brand recall than shorter ones. As mentioned earlier, Goldstein et al. [2011] showed that these

results also hold in the display advertising domain. Interestingly, Singh and Cole [1993] also showed that after repeated exposures the effect of increased duration on memory disappears.

Webb [1979] endeavored to show that the first and last television ad in a series are most likely to be remembered. His results were suggestive but not significant. Pieters and Bijmolt [1997], Newell and Wu [2003], and Terry [2005] later found this effect to be significant. Moreover, these works show that the first ad is more likely to be remembered than the last ad. Pieters and Bijmolt [1997] also showed that the longer the amount of time that has elapsed between a commercial break starting and an ad starting, the less likely people are to recall the ad. If these results were to translate to internet advertising they would all suggest that the first display ad shown will have a higher recall rate than the second.

Finally, Jeong et al. [2011] conducted an observational study of people who watched a SuperBowl and the ads therein. They showed that increased duration of a television ad and increased frequency of the ad both resulted in higher recall. Interestingly, they also showed that two 15-second ads had a higher recall rate than one 30-second ad. In what follows we will show that two 10-second display ads have a higher recall rate than one 20-second display ad.

3. METHODS

The experiments reported here were conducted on Amazon Mechanical Turk.¹ Mechanical Turk is an online labor market where requesters can post jobs and workers can choose which jobs to do for pay. After a worker submits a job, the requester can either accept or reject the work based on its quality. The fraction of jobs that a worker submits that are accepted is that worker's approval rating, which functions as a reputation mechanism used to help ensure work quality. Mechanical Turk was originally built to accomplish tasks that are easy for humans but hard for machines like image recognition, audio transcription and adult content classification. Hence jobs on Mechanical Turk are called "Human Intelligence Tasks" or "HITs." There is a burgeoning literature in the academic community around using Mechanical Turk as a platform for online behavioral experiments [Mason and Watts 2009; Mason and Suri 2012; Rand 2011; Shaw et al. 2011]. In this setting, experimenters take on the role of requesters and post their experiment as a HIT and workers are the paid participants in the experiment. Recent studies show behavior observed in Mechanical Turk experiments matches behavior observed in university lab experiments extremely well [Horton et al. 2011; Paolacci et al. 2010; Suri and Watts 2011].

We used the Mechanical Turk API to restrict our participant pool to workers in the United States to help ensure that they can read and understand English. We also restricted to those workers who have an approval rating of 90% or more. The Amazon API gives each worker account a unique, anonymous identifier. By storing these WorkerIDs we were able to ensure that a worker could only do the experiment one time. In all, 1,100 participants took part in these experiments, which ran over the course of two periods of one week each. We, next describe the format of the experiment and the various treatments to which the participants were randomly assigned.

3.1. Experimental Design

Participants were paid a 50 cent flat rate for the HIT plus 10 cents for each question answered. We chose not to pay based on the correctness of the answers to alleviate incentives for sharing answers between workers. The preview page of the HIT consisted

¹<http://www.mturk.com>



Fig. 2. Screenshot of the article with the Netflix ad.

of a brief consent form along with the instructions indicating that the HIT involved reading a web page and answering questions about it.

After reading and accepting the instructions participants were then shown an image of a webpage from an actual Yahoo! website. An image was used so that if a worker clicked on an ad or what appeared to be a hyperlink, nothing would happen. The article image consisted of text with images along with a display ad. See Figure 2 for a screenshot. Since 99% of screens on the Web can show an image of 600 pixels in height² we chose this to be the height of the article image to ensure that the article and display ad were always in view in their entirety and the user never had to scroll to see any part of them. Even if a small number of users did have to scroll to see part of the page, due to random assignment, they would be evenly spread among our treatments and thus not bias our results.

The goal of this research is to compare the recollection of two short ads to one longer one. Thus, for each memory metric we compared the sum of the metric over the two short ads to the metric measured on one long ad plus the false positive rate. The treatments with two short ads necessarily involve two advertisements, and of course different advertisements may be differentially memorable. To hold all of this constant, the simplest test involves two orderings of the short ads. Denote the treatment that shows ad A followed by ad B as AB. Then the simplest test is to compare the effectiveness of AB for one user and BA for a second user, to AA for a third user and BB for a fourth. The AB/BA treatment shows each ad for the same amount of time and in the portions of an impression as the AA/BB treatment, and thus has dedicated the same amount of resources to each advertiser. If the total effectiveness of AB/BA exceeds AA/BB, then publishers should split time slots between two advertisers.

On the experimental webpage, all objects were static except for the display ads, which were changed in two different ways. In one class of treatments, an ad was displayed for t seconds, then replaced by another ad for t seconds, which was then replaced by whitespace. In a second class of treatments, an ad was displayed for $2t$ seconds and

²<http://www.w3schools.com/browsers/browsers.display.asp>




Fig. 3. The four time treatments in which participants were randomly placed. Each colored rectangle represents an ad with the number of seconds it was in view. The white rectangles on the right side of the figure indicate the absence of an ad.




Fig. 4. The ads used as targets and lures in the experiments. When the Jeep ad (4(a)) was the target, the American Express (4(b)) ad was the lure and *vice versa*. Similarly, when the Netflix ad (4(c)) was the target the Avis ad (4(d)) was the lure and *vice versa*.

then replaced with whitespace. Since participants remained on the page for a variable amount of time, replacing the ad with whitespace was necessary to ensure that the final ad shown was in view for the intended amount of time and not more. This design allowed us to compare the memory of two ads shown for t seconds each to the memory of one ad shown for $2t$ seconds, as described earlier. We used one pair of time treatments where $t = 10$ and another pair of time treatments where $t = 20$ resulting in four time treatments overall. Figure 3 gives a pictorial representation of the time treatments. We chose these treatments as a result of a pilot study [Goldstein et al. 2011] that showed that roughly 80% of participants took over 50 seconds to read this article.

We define the following notational convenience to describe the experimental treatments.

Definition 3.1. If X, Y, Z are (not necessarily different) ads and u is a user, let $\Pr(X|YZ) = \Pr(u \text{ remembers } X \mid Y \text{ shown for } t \text{ seconds followed by } Z \text{ for } t \text{ seconds})$.

If $Y = Z$ then the given definition describes a situation where the same ad is shown continuously for $2t$ seconds. Thus our experiment will measure and compare $\Pr(A|AA) + \Pr(A|BB)$ with $\Pr(A|AB) + \Pr(A|BA)$.

In the interests of external validity we used four different ads as stimuli. The first ad treatment had a Netflix ad shown first and a Jeep ad shown second; the second ad treatment had the opposite order. The third ad treatment had an Avis ad shown first and an American Express ad shown second, and the fourth ad treatment had the opposite order. In “single ad” time treatments, in which only one ad was shown, the second ads in the order described earlier were left out. See Figure 4 for pictures of the ads. In all, the four time treatments and the four ad treatments yielded a 4×4 , between subjects design. Subjects were randomly placed into one of these 16 treatments

at the point of accepting the HIT to avoid any confound between dropping out of the experiment and the treatment assigned.

After participants finished reading the article at their own pace they clicked a link and were taken to a page where they played a game for a fixed amount of time. We chose Tetris as it is a visual game that should not create ad-specific linguistic memory interference. The game was rendered in black and white to avoid interference with the colors in the ads. The game time was chosen such that, on average, the amount of time between the first ad disappearing and the following questionnaire was the same across conditions. This ensures that on average, each participant experienced roughly the same amount of forgetting time between the initial ad exposure and test. After playing for the designated amount of time, participants were automatically directed to a questionnaire.

Once having arrived at the questionnaire, participants were unable to press the “back” button on their browser to return to the article. Participants were asked two multiple choice reading comprehension questions about the article on the previous page, after which they were asked an unaided recall question: “Which advertisements, if any, did you see on the page during this HIT? Type the name of any advertisers here if you can remember seeing their ads on the last page, or indicate that you are unable to remember any.” The next page then consisted of four separate recognition questions with textual cues of the form, “Did you see a ___ ad?” with Netflix, Jeep, Avis, and American Express being the advertisers filling in the blank. After answering these questions participants then went to a page that consisted of four separate recognition questions with cues of the form, “Did you see the following ad?” with a picture of the Netflix, Jeep, Avis, and American Express ads following each question. The ads were chosen such that each had a strong visual resemblance to another ad, exhibited in Figure 4. The Avis ad is primarily red, much like the Netflix ad, and the American Express ad is primarily black, much like the Jeep ad. So when the Netflix ad was shown, the Avis ad was not shown making the Avis ad a “lure” ad in this treatment. Conversely, when the Avis ad was shown the Netflix ad was not shown making the Netflix ad a lure ad in this treatment. The American Express and Jeep ads were used in the analogous way. The lure ads were used to check, for example, if users were simply remembering that there was a red rectangle in the top right part of the screen or if they actually noticed the ad itself and the advertiser depicted in it.

The data from a participant were encoded as 12 binary responses. The first four responses coded mentions of the two target ads and the two lures from the unaided recall question. The next four binary responses coded the recognition questions with textual cues and the final four responses coded the recognition questions with visual cues.

4. RESULTS

From the initial sample of 1,100, we excluded observations on the basis of the following criteria, which were determined before the analysis began. We excluded 42 participants who did not complete the survey, and 22 for incorrectly answering both reading comprehension questions. Recall that the longest time treatment was 40 seconds (see Figure 3), thus participants who took fewer than 40 seconds to read the article were not exposed to the entire intended time treatment and resulted in 110 exclusions. In addition, we excluded participants who took over four minutes to read the article as they may have been interrupted, yielding another 10 exclusions. The remaining 916 participants make up the set we analyze.

Table I addresses the question of whether a greater total probability of remembering an ad is achieved with two ads of length t or one ad of length $2t$. For each of the three memory metrics, the sum of the metric for the two 10 second ads ($\Pr(A|AB) + \Pr(A|BA)$),

Table 1.

Comparison of the sums of the memory rates for two successive ads shown t seconds each ($\Pr(A|AB) + \Pr(A|BA)$) with the sum of the memory rates for one ad shown for $2t$ seconds and the false alarm rate. Standard errors are given in parentheses. The p-values, rounded to two places, are bootstrapped estimates of the probability of the data given the hypothesis that $\Pr(A|AB) + \Pr(A|BA) < \Pr(A|AA) + \Pr(A|BB)$

Condition	Measure	$\Pr(A AB)$	$\Pr(A BA)$	$\Pr(A AB) +$		$\Pr(A AA) +$		p-value
				$\Pr(A BA)$	$\Pr(A AA)$	$\Pr(A BB)$		
10+10 vs. 20+0 ($t = 10$)	Visual	.37(.03)	.21(.03)	.58(.04)	.35(.03)	.05(.01)	.41(.03)	<.001
	Text	.22(.03)	.16(.02)	.38(.04)	.25(.03)	.03(.01)	.28(.03)	.02
	Recall	.12(.02)	.09(.02)	.21(.03)	.15(.02)	0(0)	.16(.02)	.07
20+20 vs. 40+0 ($t = 20$)	Visual	.41(.03)	.17(.02)	.57(.04)	.52(.03)	.09(.02)	.61(.04)	.74
	Text	.28(.03)	.11(.02)	.39(.04)	.37(.03)	.07(.02)	.44(.04)	.83
	Recall	.21(.03)	.04(.01)	.25(.03)	.21(.03)	.01(.01)	.22(.03)	.26




Fig. 5. The x-coordinate of the leftmost endpoint of each horizontal line indicates the time at which the ad appeared (either at 0, 10, or 20 seconds). The length of each line indicates the duration of the ad. The y-axis indicates the probability (or sum of probabilities) of remembering according to the three memory metrics. Vertical line segments are confidence intervals of one standard error. For each memory metric, the dotted green line shows the sum of the metric for the two short ads. The top three panels, the ($t = 10$) condition, compares the sum of two 10 second ads ($P(A|AB) + P(A|BA)$) to the sum of one 20 second ad and the false alarm rate ($P(A|AA) + P(A|BB)$). The bottom three panels, the ($t = 20$) condition, are analogous except that the shorter ads had a duration of 20 seconds each and the longer ads had a duration of 40 seconds each.

$t = 10$) is higher than the sum of the metric for the single 20 second ad and the false alarm rate ($\Pr(A|AA) + \Pr(A|BB)$, $t = 10$). For example, when $t = 10$, the probability of visually recognizing the ad in the AB and BA treatments sums to .58, while that of the AA and BB treatments is only .41. Text recognition and unaided recall show a similar pattern. The p-values listed are the probabilities of these data given the hypothesis that the $\Pr(A|AB) + \Pr(A|BA) < \Pr(A|AA) + \Pr(A|BB)$. Thus, when t is 10 seconds, the total amount of recollection in AB + BA treatments exceeds the total in the AA + BB treatments. The top three panels of Figure 5 graphically represent this result. Here, for all three metrics, the dotted green line indicating the sum of the memory measure for two 10 second ads is significantly higher than the orange line indicating the sum of the memory measure for one 20 second ad plus the false alarm rate.

However, when t is 20 seconds, a very different picture emerges. As can be seen in the bottom rows of Table I and the bottom panels of Figure 5, there is no significant difference between $\Pr(A|AB) + \Pr(A|BA)$ and $\Pr(A|AA) + \Pr(A|BB)$ and the direction of the effect reverses in two of three cases. Surprisingly, the sum of two ads ($\Pr(A|AB) + \Pr(A|BA)$) at $t = 10$ is comparable to the sum of two ads at $t = 20$, differing by at most four percentage points. Thus, not only are two short (10 second) ads better than one ad of twice the duration, they are also roughly equivalent to *two* ads of twice the duration. This result is consistent with the powerful effect of ad onset time, which we address in Section 4.3.

The main conclusion of this analysis is that if A and B are advertisers and ad slots are short (around 10 seconds), it seems that more total impact on memory is created when splitting an impression between two advertisers than giving each advertiser its own full slot. That is, in the terminology established earlier, the memory under AB + BA is greater than memory under AA + BB. However, it is in principle possible for $\Pr(A|AB) + \Pr(A|BA) > \Pr(A|AA) + \Pr(A|BB)$ to hold averaged over many different advertisers, but not for a specific advertiser. For example, advertiser A might benefit greatly from the split impressions, while advertiser B suffers slightly. To check whether this occurs in practice, we take advantage that the experimental design uses four unique advertisers, each of which can be used as a test to see whether two short ads lead to more recall than one ad of twice the duration plus the false alarm rate. The results are shown in Table II. Again, a difference between the $t = 10$ and the $t = 20$ condition emerges. In the former case, in 11 of 12 tests, $\Pr(A|AB) + \Pr(A|BA) > \Pr(A|AA) + \Pr(A|BB)$, while in the latter case, there is no clear pattern. P-values are reported in Table II, though it should be noted that the sample sizes here are four times smaller than in Table I, decreasing statistical power; all else equal, p-values increase as sample size decreases, which explains why the aggregate results for $t = 10$ in Table I have lower p-values than do those in Table II. Nonetheless, the largely consistent sign and at times sizeable absolute magnitude of the differences in the $t = 10$ conditions lend support to the conclusion that the benefits of shorter ads hold within advertisers, especially for the visual recognition task, as shown in Figure 6. However, further experimentation needs to be done with more ads and bigger sample sizes to firmly conclude that no advertiser is hurt by time based advertising.

4.1. Remembering at Least One Ad

An alternate means for comparing the AB + BA schedule with the AA + BB schedule is to calculate, for each schedule, the probability of not remembering either of the ads. More formally, in this analysis we will compare $(1 - \Pr(A|AB))(1 - \Pr(A|BA))$ with $(1 - \Pr(A|AA))(1 - \Pr(A|BB))$. Consistent with the results thus far, Table III shows that on the measure of not remembering any ads, there is an advantage to the AB + BA schedule over the AA + BB schedule when $t = 10$. In other words one is more likely to remember at least one ad in the AB + BA schedule than in the AA + BB schedule. Also consistent with results thus far, when $t = 20$ seconds the effect disappears. This analysis supports the idea that, for relatively short time slots, there is an advantage to showing multiple ads over a single ad on the metric of remembering at least one ad.

4.2. Lure Ads

Across all conditions, the rate of incorrectly indicating memory for one of the visually similar lure ads were low and quite similar to those in Goldstein et al. [2011]: 0% for recall, 6.6% for text recognition and 7.5% for visual recognition questions. These are similar to the false alarm rates, $\Pr(A|BB)$, shown in Table I. Thus, when we asked a user if they remembered an ad that was not actually shown, whether or not that ad was visually similar to the ad that was shown did not have a substantial effect on the false

Table II. Advertiser-Specific Variant of Table I

Condition	Advertiser	Measure	Pr(A AB)			Pr(A AB)+			Pr(A AA)+		p-value
			Pr(A AB)	Pr(A BA)	Pr(A BA)	Pr(A BA)	Pr(A BA)	Pr(A AA)	Pr(A BB)	Pr(A BB)	
10+10 vs. 20+0 ($t = 10$)	Avis	Visual	.24(.06)	.21(.05)	.45(.08)	.29(.06)	.02(.02)	.3(.06)	.07		
		Text	.16(.05)	.12(.04)	.28(.06)	.17(.05)	0(0)	.17(.05)	.1		
	AMEX	Recall	.09(.04)	.1(.04)	.19(.05)	.08(.03)	0(0)	.08(.03)	.04		
		Visual	.47(.07)	.21(.05)	.67(.08)	.38(.06)	.06(.03)	.45(.07)	.02		
	Netflix	Text	.19(.05)	.19(.05)	.38(.07)	.25(.06)	.03(.02)	.28(.06)	.15		
		Recall	.09(.04)	.03(.02)	.12(.04)	.13(.04)	0(0)	.13(.04)	.55		
	Jeep	Visual	.44(.07)	.29(.06)	.72(.09)	.45(.07)	.06(.03)	.51(.07)	.03		
		Text	.22(.06)	.29(.06)	.51(.08)	.38(.06)	.08(.03)	.45(.07)	.32		
	Jeep	Recall	.13(.04)	.23(.06)	.36(.07)	.25(.06)	.02(.02)	.27(.06)	.17		
		Visual	.35(.07)	.13(.04)	.47(.08)	.31(.06)	.07(.03)	.38(.07)	.2		
	20+20 vs. 40+0 ($t = 20$)	Avis	Text	.31(.06)	.05(.03)	.36(.07)	.2(.05)	0(0)	.2(.05)	.03	
			Recall	.17(.05)	.02(.02)	.19(.05)	.16(.05)	0(0)	.16(.05)	.31	
AMEX	Avis	Visual	.42(.06)	.22(.05)	.64(.08)	.5(.07)	.11(.04)	.61(.08)	.41		
		Text	.28(.05)	.12(.04)	4(.07)	.27(.06)	.06(.03)	.33(.07)	.25		
Netflix	AMEX	Recall	.18(.05)	.04(.02)	.22(.05)	.12(.04)	0(0)	.12(.04)	.05		
		Visual	4(.06)	.18(.05)	.58(.07)	.42(.06)	.06(.03)	.48(.07)	.17		
Jeep	Netflix	Text	.26(.05)	.14(.04)	4(.07)	.34(.06)	.08(.04)	.42(.07)	.57		
		Recall	.16(.04)	.04(.02)	2(.05)	.19(.05)	0(0)	.19(.05)	.41		
Jeep	Netflix	Visual	.42(.07)	.15(.06)	.56(.09)	.6(.07)	.13(.04)	.73(.09)	.92		
		Text	.35(.06)	.12(.05)	.47(.08)	.58(.08)	.11(.04)	.69(.09)	.97		
Jeep	Jeep	Recall	.29(.06)	.07(.04)	.36(.07)	.33(.07)	.04(.03)	.36(.08)	.5		
		Visual	.39(.08)	.09(.04)	.48(.08)	.58(.07)	.05(.03)	.63(.07)	.9		
Jeep	Jeep	Text	.22(.06)	.07(.04)	.29(.07)	.35(.06)	.02(.02)	.37(.07)	.78		
		Recall	.22(.06)	.02(.02)	.24(.07)	.25(.06)	0(0)	.25(.06)	.57		

recall rate. We can conclude that in the text and visual recognition numbers, that is, $\Pr(A|AB)$, $\Pr(A|BA)$, and $\Pr(A|AA)$, reported in Table I, users remembered more than just the predominant color of the ad. This is most strongly supported by the unaided recall questions in which false alarms are highly improbable.

4.3. Effect of Onset Time

We define an ad's onset time to be the amount of time between the page loading and the ad appearing. If an ad appears as the page loads, it has an onset time of 0. In Table I if one compares $\Pr(A|AB)$ with $\Pr(A|BA)$ one will see $\Pr(A|AB) > \Pr(A|BA)$ for all time treatments and memory metrics. That is, the second ad presented is at a disadvantage compared to the first. Moreover, the difference $\Pr(A|AB) - \Pr(A|BA)$ is far larger when $t = 20$ seconds than when $t = 10$ seconds. This is also visually apparent in Figures 5 and 6 if one observes that the blue lines (indicating $\Pr(A|AB)$) are always higher than the magenta lines (indicating $\Pr(A|BA)$). In addition, Figure 5 shows the difference in heights is much greater in the bottom three panels ($t = 20$) than the top three panels ($t = 10$). This shows that the longer the onset time of an ad, the less likely it is to be remembered. Accordingly, advertisers should value ads with early onset times.

Onset time can also explain why two short ads are roughly as effective as two longer ads, as mentioned in Section 4. When the ads are longer, the second ad appears at a greater onset time reducing the chance of it being remembered. In addition, this suggests that slotting more than two ads into an impression is not likely to be effective. Given ads of even short durations, such as 10 seconds, ads beyond the second would have onset times so great as to diminish their memory rates. This results also suggests that another way to implement the ideas in this article in the display advertising industry would be to have the first ad display for 10 seconds and then leave the second ad displaying for as long as the user is on the page.

The fact that the first ad has higher recall than the second relates to the primacy effect, which states that people are more likely to remember the first item in a list than subsequent items. Moreover, as mentioned in Section 2 the first ad in a group of television ads is also more likely to be remembered than subsequent ads [Pieters and Bijmolt 1997].

4.4. Do Two Ads Gain an Advantage from Apparent Motion When Loading?

As seen in Tables I and II and Figures 5 and 6, for shorter ad durations, it seems to be the case that two short ads (of 10 seconds) lead to more recollection than one longer ad of 20 seconds. A proposed explanation for this is that apparent motion of the second ad's loading might capture user attention and provide another opportunity for the ad to be noticed. This idea has support in previous research, which found that animated ads are more likely to be noticed than static ones [Yoo and Kim 2005]. To test this explanation, we ran a follow-up study in which we replicated the condition in which the Netflix ad was shown for 20 seconds but introduced apparent motion in the ad. We did so by displaying the ad for 10 seconds, then by covering it with whitespace for 0.125 seconds, and then displaying the ad again for another 9.875 seconds. This provided the user two opportunities to detect a change, one when the whitespace appeared and one when the ad reappeared. To provide comparable statistical power to the original condition, an approximately equal number of participants was assigned to this variant as to the original. Despite this manipulation, the 20 second ad that flashed after 10 seconds was not significantly more memorable than the static version. This relationship held across visual recognition (25/56 participants recognizing in the original vs. 24/54 in the variant, p -value ≈ 1), text recognition (21/56 vs. 25/54, p -value=.46) and recall (14/56 vs. 21/54, p -value=.17). While apparent motion may have an effect at the margin, it




Fig. 6. The visual recognition rate only for each of the four advertisers in the ($t = 10$) condition. The x-coordinate of the leftmost point of each horizontal line indicates the time at which the ad appeared. The length of the line indicates the duration the ad remained in view (i.e. the exposure time). The y-axis indicates the probability of remembering the ad in the visual recognition test. The vertical line segments are confidence intervals of one standard error. The dotted green line shows the sum of the recognition rates of the two 10 second ads ($P(A|AB) + P(A|BA)$). The solid brown line shows the sum of one 20 second ad and the false alarm rate ($P(A|AA) + P(A|BB)$).

Table III.

Comparison of the products of the complements of the recall rates of two successive ads shown t seconds each, $(1 - \Pr(A|AB)) * (1 - \Pr(A|BA))$, compared to the product of complements of the recall rates of one ad shown for $2t$ seconds and the false alarm rate, $(1 - \Pr(A|AA)) * (1 - \Pr(A|BB))$. Standard errors are given in parentheses. The p-values, rounded to two places, are bootstrapped estimates of the probability of the data given the hypothesis that $(1 - \Pr(A|AB)) * (1 - \Pr(A|BA)) > (1 - \Pr(A|AA)) * (1 - \Pr(A|BB))$.

Condition	Measure	$(1 - \Pr(A AB)) * (1 - \Pr(A BA))$	$(1 - \Pr(A AA)) * (1 - \Pr(A BB))$	p-value
10+10 vs. 20+0 ($t = 10$)	Visual	.5(.03)	.61(.03)	<.001
	Text	.66(.03)	.73(.03)	.03
	Recall	.8(.03)	.84(.02)	.10
20+20 vs. 40+0 ($t = 20$)	Visual	.49(.03)	.44(.03)	.90
	Text	.64(.03)	.58(.03)	.90
	Recall	.76(.03)	.78(.03)	.31

does not seem to be the primary reason why two short ads lead to more recollection than one longer ad.

5. EQUILIBRIUM EFFECTS OF TIME-BASED PRICING

Will industry revenue be higher using impression-based pricing, or using time-based pricing? We shed light on this question by analyzing the market equilibrium. Selling advertisements by the page view or impression provides a noisier measure of advertiser value than selling by time, because time more closely correlates with advertiser value (recall and recognition) than impressions. In this section, we consider whether using the more accurate time-based measurement induces higher industry revenue.

Let $v(x)$ be the dollar value for an advertiser created from advertisement of exposure duration x . The function v is assumed to be increasing and concave. To illustrate v , suppose that at each moment in time t , there is a probability r_t that an advertisement goes into a user's memory. Assume that the advertiser obtains a value ρ if the user remembers the ad, and zero otherwise. This functional form is just an illustration, justifying the assumptions. The probability that the ad is not recalled, denoted $(1 - R)$, is just the product of the probabilities of the ad not going into memory at each moment in time, or

$$1 - R = \prod_{t=1}^T (1 - r_t) = \exp \left(\sum_{t=1}^T \log(1 - r_t) \right).$$

Taking the continuous limit, we obtain a value function that is the value of recall ρ times the probability of recall, or

$$v(x) = \rho R = \rho \left(1 - \int_0^x \log(1 - r(s)) ds \right).$$

This particular functional form can accommodate our major findings—that the likelihood of recall is concave in time, that starting a second ad part way through a time interval produces a greater total likelihood of recall compared to showing a single ad, and that the effect of either switching ads or of increasing duration after 20 seconds is small. The needed properties of r for these findings are that it decreases over time and v' becomes small sometime around 20 seconds. For the analysis in this section, we will not need to assume any properties on v beyond the assumption that v is increasing and concave.

There are several ways in which time-based pricing could be implemented, such as switching ads every ten seconds, running one ad for ten seconds and then a second ad indefinitely, or making impressions spill over from page to page to given the requisite number of seconds. Any of these approaches will tend to reduce the variation in impression length. The best fit for the theory is the least practical mechanism, involving the denomination of ads into attention-units, so that an advertisers buys the amount of time to lift awareness by a given target amount.

We model impression-based advertising as producing a random exposure length, compared to a regime of fixed exposure lengths. This is an imperfect fit to the empirical work, since neither impressions nor time is a perfect measure of advertiser value, but the thought experiment is instructive nonetheless because impression-based advertising seems to embody more randomness than time-based advertising.

To model the pricing, we assume price-taking behavior, so that in both cases, all the units sell for a price equal to the marginal value of the buyer. To compare the random (impression-based) to the fixed (time-based) outcomes, first imagine a random allocation X ; an advertiser can buy a quantity s of X at price p . In this case, the advertiser purchasing a quantity s obtains utility $u(s) = E v(sX) - ps$. Thus the advertiser chooses s^* satisfying

$$0 = u'(s^*) = E v'(s^*X)X - p.$$

This outcome compares to the nonrandom allocation where a buyer facing price p per unit will maximize $v(x) - px$, and choose a quantity x^* satisfying $v'(x^*) - p = 0$. This outcome arises from our stochastic case when $X = 1$. It will be useful later to note that, in the nonrandom case, $\frac{dx^*}{dp} = \frac{1}{v''(x^*)}$.

The payment by the advertiser is $R = ps^* = E v'(s^*X)s^*X$. In contrast, when the allocation is not random, the agent will buy $E s^*X$, the total quantity available.

Note, then that utility maximization implies a price of $v'(s^*EX)$, yielding a revenue of $v'(s^*EX)s^*EX$. Whether or not the quantity is random, the average quantity per buyer should be the same with both a random and nonrandom allocation; if in either case there were unsold units on average, the price should drop to clear the market. Thus, revenue is higher under a nonrandom allocation whenever $E\{v'(s^*X)s^*X\} \leq E\{v'(s^*X)\}E\{s^*X\}$.

THEOREM 5.1. *Suppose $v'''(x)x + 2v''(x) < 0$. Then payments by advertisers are higher under deterministic allocations than under random allocations.*

PROOF. Let $f(x) = v'(x)x$. Then $f'(x) = v''(x)x + v'(x)$ and $f''(x) = v'''(x)x + 2v''(x) < 0$ by hypothesis. Thus f is concave and thus $Ef(X) \leq f(EX)$ by Jensen's Inequality. Hence

$$R = Ev'(s^*X)s^*X = Ef(X) < f(EX) = v'(s^*EX)s^*EX. \quad \square$$

In the following corollary, Y is a *garbling* of X if $Y = X + \delta$ where X and δ are independently distributed.

COROLLARY 5.2. *If $v'''(x)x + 2v''(x) < 0$ any garbling of an allocation reduces industry revenue.*

PROOF. Immediate consequence of the concavity of $v'(x)x$ and Blackwell's Theorem (presented in Marschak and Miyasawa [1968]). \square

An implication of the corollary is that, if an impression-based allocation represents a more random allocation in Blackwell's sense than a time-based allocation, and the condition holds, then a time-based allocation will have a higher industry revenue.

How likely is the condition $v'''(x)x + 2v''(x) < 0$? To answer this, we consider the elasticity of demand. Recall the elasticity of demand, which measures the responsiveness of quantity to price, is given by

$$\epsilon = -\frac{p}{x^*} \frac{dx^*}{dp} = -\frac{v'(x^*)}{x^*v''(x^*)} > 0.$$

Here we consider two properties. First, demand is elastic if an increase in price reduces revenue, which occurs when $\epsilon > 1$. Second, Alfred Marshall, the inventor of the supply and demand graph as well as the concept of elasticity, famously codified properties of demand as laws. Marshall's first law is that an increase in price reduces the quantity demanded. His second, much more obscure, law is that the elasticity of demand decreases in quantity x ([Marshall 1920] at p. 102–4). Using these two properties, we can provide sufficient conditions for randomization to reduce revenues.

THEOREM 5.3. *Suppose that demand is elastic and Marshall's Second Law holds. Then $v'''(x)x + 2v''(x) < 0$.*

PROOF. By the definition of elasticity, and omitting *'s for the sake of clarity,

$$\begin{aligned} \frac{d\epsilon}{dx} &= -\left(\frac{v''(x)}{xv''(x)} - \frac{v'(x)}{x^2v''(x)} - \frac{v'(x)v'''(x)}{xv''(x)^2}\right) \\ &= -\frac{v''(x)}{xv''(x)} + \frac{v'(x)}{x^2v''(x)} + \frac{v'(x)v'''(x)}{xv''(x)^2} \\ &= -\frac{1}{x} - \frac{\epsilon}{x} - \epsilon \frac{v'''(x)}{v''(x)}. \end{aligned}$$

Thus,

$$\frac{x}{\epsilon} \frac{d\epsilon}{dx} = -\frac{1}{\epsilon} - 1 - \frac{xv'''(x)}{v''(x)},$$

or,

$$\frac{xv'''(x)}{v''(x)} = -\frac{1}{\epsilon} - 1 - \frac{x}{\epsilon} \frac{d\epsilon}{dx}.$$

Because $v''(x) < 0$, $v'''(x)x + 2v''(x) < 0$ if and only if $\frac{xv'''(x)}{v''(x)} > -2$ if and only if $\frac{1}{\epsilon} + \frac{x}{\epsilon} \frac{d\epsilon}{dx} < 1$. Elastic demand ensures $\frac{1}{\epsilon} < 1$; Marshall's second law guarantees $\frac{d\epsilon}{dx} < 0$. \square

We illustrate the concepts with four functional forms.

Example 1: $v(x) = x^\alpha$ for $0 < \alpha < 1$. Then $2v''(x) + xv'''(x) = 2\alpha^2(\alpha - 1)x^{\alpha-2} < 0$, and thus the addition of noise never improves industry revenue. Note, this is the constant elasticity of demand case with demand elasticity $\frac{1}{1-\alpha} > 1$.

Example 2: $v(x) = 1 - (1 - x)^\alpha$ for $\alpha > 1$. This case is valid for $0 < x < 1$. Then

$$2v''(x) + xv'''(x) = \alpha(\alpha - 1)(1 - x)^{\alpha-3}(-2(1 - x) + x(\alpha - 2)) = \alpha(\alpha - 1)(1 - x)^{\alpha-3}(-2 + x\alpha).$$

so that the addition of a small amount of noise decreases revenue when $\alpha x < 2$. In particular, the condition is met with linear demand ($\alpha = 2$) for any x . As α grows the range of x in which the addition of noise increases revenue also grows.

Example 3: $v(x) = 1 - e^{-\alpha x}$. This case is valid when $0 < x < \infty$. In this case $2v''(x) + xv'''(x) = \alpha^2 e^{-\alpha x}(-2 + \alpha x)$, and thus the addition of noise decreases revenue when $\alpha x < 2$.

Example 4: $v(x) = 1 - x^{-\alpha}$ for $0 < \alpha$. Then $2v''(x) + xv'''(x) = \alpha(\alpha + 1)x^{-\alpha-2}(-2 + 2 + \alpha) > 0$, and thus the addition of noise always increases industry revenue. However, this is a case where the monopoly problem—choose x to maximize revenue—has no solution because revenue is increasing as x decreases toward zero.

In all four examples, Marshall's second law holds. Thus any failure arises because demand is inelastic—so that a destruction of quantity increases overall revenue. In these cases, randomness works like a destruction of quantity—effectively reducing output by making it random.

6. CONCLUSION

Display ads are typically sold by impression. In this pricing scheme, a two-minute impression costs the same as a two-second impression, despite the former having a larger effect on memory. If the total amount of time allocated to a publisher is divided into slots, we have shown that displaying two short ads increases memory per slot over a single, longer duration ad when slot lengths are reasonably small (around 10 seconds). Interestingly, a similar result was found in the domain of television advertising [Jeong et al. 2011]. In addition to increasing the total recollection, this pricing system also increases the probability of remembering at least one ad. We have shown that advertisers should be better off under this policy in the sense their ads are more likely to be remembered. Moreover, publishers may be better off since they will experience increased revenues. In this way, the scheme with two shorter ads benefits both advertisers and publishers. This result strengthens the case for moving from an impression based pricing scheme to one that is either partially or completely based on exposure time.

Given the results presented here, it would be natural to ask how short could the duration of an ad be and how many ads could be placed in succession? An ad needs to be displayed long enough so that it can be seen. Goldstein et al. [2011] showed that

there is a sharp increase in memory for an ad when the ad duration goes from zero to five seconds and there are diminishing returns to increased duration after that. So we suggest five seconds as a lower bound for the duration of an ad. Moreover, as discussed in Section 4.3 there is a dramatic effect of onset time on ad memory. The, greater the amount of time between the page loading and the ad loading, the lower the memory rate will be. Figure 5 shows that ads with an onset time of 20 seconds have fairly low memory rates. Thus, we suggest that as an upper bound on the onset time of an ad.

Increasing online ad effectiveness will likely create additional benefits for online publishers and advertisers at the expense of offline publishers such as television. In a time-based pricing scheme advertisers would be able to improve their metrics while buying the same amount of time on the internet thereby improving advertiser returns. This improvement may induce advertisers to substitute internet advertising for some other forms of advertising, bringing more money online, and increasing online publisher revenues. Advertisers must be better off as a group, because the loss of demand for offline advertising will tend to lower the prices of offline advertising. Since the performance of offline ads is the same and there is less money spent on offline advertising, the value of offline advertising should increase. Furthermore, since online advertising must be competitive with offline advertising, online advertising performance per dollar must rise. Thus some of the value from impression splitting is captured by advertisers, and some is captured by publishers.

7. FUTURE WORK

This work suggests several extensions and future research directions. One natural direction would be to explore how these effects depend on characteristics of the ads. For example, how do these effects depend on whether the ad is animated, informational, emotional etc. One could also look at interactions between ad A and ad B. For example, what are the results if they are different creatives for the same advertiser? What if they are competitors and what if they are synergistic products?

A second natural follow-up study could test how exposure time affects long term advertiser memory. One could conduct a similar experiment where there is a gap of one hour, one day or even one week between the ad exposure and the memory test. In addition, in the experiment reported here, the impact of exposure time was measured while participants read an article, which is only one type of activity done on the Web. One could conduct a similar experiment where different types of content such as video, slideshows, or social media content are consumed.

A third set of experiments would explore the interaction between exposure time and ad relevance. It seems likely that even modest exposures of relevant ads will have a greater impact than long exposures of less relevant ads. If so, a shortened exposure time would be beneficial to advertisers by focusing the impact on users who find the ad relevant and shortening times for other users. This phenomenon is especially important in understanding ads that users can voluntarily turn off, like YouTube's TrueView ads. Along the same lines, the interaction of exposure time and behavioral targeting is important. It might be advantageous to have longer exposures when targeting is effective, for instance, when it is based on good behavioral data. Is behavioral data a substitute for exposure, or a complement? We expect that our results—shorter times are more efficient—would hold when the ad is relevant to the user but this also needs to be verified.

A fourth set of experiments would consider different ad types and the interaction of time and the ad type. Ad type—video, image, or text—may interact with exposure time. Is a picture worth a thousand words? Do video ads justify the sort of premium charged on television channels? Our approach is amenable to addressing such questions. Understanding the relationship between advertising types and the durability of brand

recall/product perception is important for understanding the value of advertising and its ability to influence customer purchases.

Finally, while Mechanical Turk is a more natural setting with a more diverse population than a university lab where similar psychological experiments are often run, it is not as natural and diverse as a live website. Thus, to test the robustness of these effects, a field experiment would be valuable. One challenge with a field experiment is the ability to accurately measure advertising impact.

REFERENCES

- Michael Braun and Wendy Moe. 2011. Online advertising campaigns: Modeling the effects of multiple ad creatives. Working Paper.
- Peter J. Danaher and Guy W. Mullarkey. 2003. Factors affecting online advertising recall: A study of students. *J. Advertis. Res.* 43, 252–267.
- Xavier Drèze and François-Xavier Husherr. 2003. Internet advertising: Is anybody watching? *J. Interactive Marketing* 17, 4, 8–23.
- Daniel G. Goldstein, R. Preston McAfee, and Siddharth Suri. 2011. The effects of exposure time on memory of display advertisements. In *Proceedings of the 12th ACM Conference on Electronic commerce (EC'11)*. ACM, New York, 49–58.
- John J. Horton, David G. Rand, and Richard J. Zeckhauser. 2011. The online laboratory: Conducting experiments in a real labor market. *Experimental Economics* 14, 3, 399–425.
- Yongick Jeong, Meghan Sanders, and Xinshu Zhao. 2011. Bridging the gap between time and space: Examining the impact of commercial length and frequency on advertising effectiveness. *J. Marketing Commun.* 17, 4, 263–279.
- Randall A. Lewis. 2010. Measuring the effects of online advertising on human behavior using natural and field experiments. Ph.D. Dissertation, Massachusetts Institute of Technology.
- Randall A. Lewis and David H. Reiley. 2011. Does retail advertising work? Measuring the effects of advertising on sales via a controlled experiment on Yahoo! Working Paper.
- L. M. Lodish, M. Abraham, S. Kalmenson, J. Livelsberger, B. Lubetkin, B. Richardson, and M. E. Stevens. 1995. How TV advertising works: A meta-analysis of 389 real world split cable TV advertising experiments. *J. Marketing Research* 32, 2, 125–139.
- Puneet Manchanda, Jean-Pierre Dubé, Khim Yong Goh, and Pradeep K. Chintagunta. 2006. The effect of banner advertising on internet purchasing. *J. Marketing Research* 43, 1, 98–108.
- J. Marschak and K. Miyasawa. 1968. Economic comparability of information systems. *Int. Econ. Rev.* 9, 2, 137–174.
- A. Marshall. 1920. *Principles of Economics* 8th Ed. MacMillan and Co., Ltd., London.
- Winter Mason and Siddharth Suri. 2012. Conducting behavioral research on Amazon's Mechanical Turk. *Behavior Research Methods* 44, 1, 1–23.
- Winter A. Mason and Duncan J. Watts. 2009. Financial incentives and the performance of crowds. In *Proceedings of the ACM SIGKDD Workshop on Human Computation*. Paul Bennett, Raman Chandrasekar, and Luis von Ahn, Eds., ACM, 77–85.
- Stephen J. Newell and Bob Wu. 2003. Evaluating the significance of placement on recall of advertisements during the Super Bowl. *J. Current Issues Research in Advertising* 25, 2, 57–67.
- Gabriele Paolacci, Jesse Chandler, and Panagiotis G. Ipeirotis. 2010. Running experiments on Amazon Mechanical Turk. *Judgment and Decision Making* 5, 411–419.
- Rik G. M. Pieters and Tammo H. A. Bijmolt. 1997. Consumer memory for television advertising: A field study of duration, serial position, and competition effects. *J. Consumer Research* 23, 4, 362–372
- PricewaterhouseCoopers. 2011. IAB Internet Advertising Revenue Report: 2010 Full Year Results.
- David G. Rand. 2011. The promise of Mechanical Turk: How Online labor markets can help theorists run behavioral experiments. *J. Theor. Biol.* 299, 172–179
- Navdeep Sahni. 2011. Effect of temporal spacing between advertising exposures: Evidence from an online field experiment. Working Paper.
- Aaron D. Shaw, John J. Horton, and Daniel L. Chen. 2011. Designing incentives for inexpert human raters. In *Proceedings of the ACM Conference on Computer Supported Cooperative Work (CSCW'11)*. ACM, New York, 275–284.
- Surendra N. Singh and Catherine A. Cole. 1993. The effects of length, content, and repetition on television commercial effectiveness. *J. Marketing Research* 30, 91–104.

- D. Starch. 1923. *Principles of Advertising*. A. W. Shaw Company, Chicago.
- Siddharth Suri and Duncan J. Watts. 2011. Cooperation and contagion in web-based, networked public goods experiments. *PLoS One* 6, 3.
- W. Scott Terry. 2005. Serial position effects in recall of television commercials. *J. General Psych.* 132, 2, 151–164.
- Peter H. Webb. 1979. Consumer initial processing in a difficult media environment. *J. Consumer Research* 6, 3, 225–236.
- W. Wells. 2000. Recognition, recall, and rating scales. *J. Advertis. Res.* 40, 6, 14–20.
- Chan Yun Yoo and Kihan Kim. 2005. Processing of animation in online banner advertising: The roles of cognitive and emotional responses. *J. Interactive Market.* 19, 4, 18–34.

Received March 2013; revised May 2014; accepted December 2014