

Web Appendix to “An Experimental Investigation of the Effects of Retargeted Advertising – the Role of Frequency and Timing”

October 18, 2018

A More Background on Retargeting and BuildDirect

A.1 Retargeting

Retargeting refers to advertising targeted to customers based on their past actions at the advertiser’s website. An example of a retargeting campaign is in Figure 1. In the top left panel of the figure, a consumer visits the product page for a specific product at a retailer’s website, and views related information, such as the product price, reviews etc. Then the consumer navigates away from this page, and decides to visit a news webpage on the internet. This action of visiting the product page and navigating away without making a purchase triggers a retargeting campaign paid for by the retailer. Subsequently, if the consumer browses a webpage that is a part of the the retargeting platform’s network, he/she might see a retargeted ad showing the product she saw on the retailer’s website.¹

A number of platforms that enable retargeting have emerged over the last few years. They range from those run by large internet media firms such as Google, Facebook and Twitter, to specialist companies such as Criteo. Google’s Doubleclick platform, which is used by the retailer that we partner with for our experiment, tracks users through a combination

¹This form of retargeting is also referred to as “site retargeting”. Other less-common forms of retargeting have recently emerged. For example, “search retargeting” and “email retargeting” involve showing ads to individuals who searched for a specific product, or engaged with the advertiser’s email marketing campaign respectively.

of cookies and Google user ids. A retargeting campaign on this platform is triggered by a small piece of code that gets executed when the individual visits the retailer’s webpage. This signals to Doubleclick that the consumer is to be included in the retargeting campaign, and also provides the parameters for the campaign. The parameters include the duration of the campaign, and the ceiling on the number of advertising impressions that the consumer can see during any particular day. This latter parameter, the “frequency cap”, is the main variable we vary in our experiment.

Estimates on the size of the retargeting industry vary due to the fact that many players are either startups that do not report revenues, or large multi-product advertising firms such as Google, that do not report numbers separately by product. Nevertheless, it is well accepted that the industry has grown over the last few years at a very rapid rate. For instance, a recent industry study AdRoll (2014) finds that 71% of respondents in a 2014 survey of 1000 marketers in the US reported spending 10-50% of their advertising budget on retargeting. This number was a significant increase from the 53% reported in 2013. The proportion of marketers reporting spending over 50% on retargeting went up from 7% to 14%. One of the few firms in the industry that reported its results and derives most of its revenues from retargeting is Criteo. The firm reported a 70% increase in its earnings in the first quarter of 2015, with annual revenues in the year expected to cross \$1 billion. With Criteo being only one of several players in this market, including several large firms, the retargeting market is expected to be several times this size.

A.2 Experimental Context

In this section we describe the BuildDirect.com website, and describe the observed activity of its users. First, we give a brief description of the firm.² Founded in 1999, BuildDirect is an online marketplace for buying heavyweight home improvement products. The company provides homeowners a wide choice of products in multiple categories such as wood flooring, tile flooring, decking, outdoor living, building materials, landscaping, kitchen and bath and vinyl flooring. Since the home-improvement category involves large purchases (average order

²Some of the material in this section is based on <http://techcrunch.com/2016/02/10/builddirect-wants-to-become-the-amazon-of-the-home-improvement-industry-launches-marketplace/>

on the website is \$1800) shopping cycles can be long. The company allows buyers to obtain samples before making a purchase, so they can touch and feel products for color, texture and quality. The firm delivers these products directly to the consumer’s doorstep. Relative to other online retailing platforms, BuildDirect is highly rated on websites providing seller reviews.³

The BuildDirect.com website allows users to search for and buy products across multiple home-improvement categories. Figure 2 shows the homepage of the website. It allows the user to specify a product category for search, or search using text queries. On searching, the consumer arrives at a search-results page, which looks like the example in Figure 3. A user may browse various product options satisfying her search criteria. Users might face significant uncertainty in purchasing the product online. Therefore, the website allows users to order samples before making actual purchases. Figures 4 and 5 show examples of a product-page and a checkout page respectively.

Advertising It is important to note that BuildDirect engages in marketing via several channels, including email, search advertising and display advertising. It is a major advertiser in its category, with significant online advertising spends in the year 2014. Of this, approximately \$4 million were spent on retargeting, which it conducts on multiple platforms, including DoubleClick, Criteo and Chango. Overall, BuildDirect’s advertising through various online channels is delivered with a high intensity – in our data on average 37% of the impressions delivered occurred within a minute of another BuildDirect impression preceding it, and 9% of the delivered ad impressions had at least one other BuildDirect banner on the same webpage. Our experiment varies the DoubleClick campaign only. A user’s participation in the rest of the campaigns is invariant across our experimental conditions.

Description of user behavior on BuildDirect.com

We describe the observed activity of 234,712 users, identified by DoubleClick ids, who had some interaction recorded in our data (not limiting to users included in our experiment for

³For example, on resellerratings.com BuildDirect is rated 8.8, whereas HomeDepot is 1.0; Lowes’ 1.0; Amazon 4.2. On trustpilot.com BuildDirect is rated 6.3; HomeDepot is 2.6; Lowes’ is 5.1; Amazon is 7.7. We thank an anonymous reviewer for pointing us to this information.

this description).⁴ Table 1 provides descriptive statistics of users on the website. On average, a user interacts with the website for more than two days, but there is large heterogeneity; many users interact with the website more often. These interactions are spread over a large time interval. On average, the time interval between the first and the last interaction is about 16 days. Among individuals that arrived on more than one occasion, this number goes up to 35 days. During this time, users on average browse about 25 product pages and 19 search pages. Since home improvement products are expensive, complicated and not frequently purchased, these searches are likely to correspond to a single purchase occasion. Therefore, these statistics suggest that consumers in our setting spend significant time deliberating on purchase and obtaining information from the website.

Conversion from search to next steps in the purchase process is rare. About 13.5% of individuals who search on the website eventually “create a cart”, which signifies their further interest in the product. About 4% of users order a sample, and 0.4% order a product.⁵ Note that the probability of creating a cart for users who clicked on a retargeted ad is significantly higher than average by 50%; 20% of this selected set create a cart. These statistics indicate a very significant and large *correlation* between clicks on retargeted ads and cart-creation (p-val < 0.01). The rest of the statistics show that there is significant time-lag between the users’ first interaction with the website and their conversion activity.

Competition in this category

BuildDirect faces considerable competition. The data we obtained through the DoubleClick platform records activity on BuildDirect.com only. Therefore, to assess competition, we bring to bear data from comScore MediaMetrix, that inform us about consumer activity across competing retailers in the category. Table 2 shows that in the comScore sample, a significant proportion of individuals who visited BuildDirect.com also visited a competitor’s website during the month. If an individual visited BuildDirect.com, the chance of her visiting HomeDepot.com is 50.5%, which is significantly higher than 13.6%, which is the probability

⁴The Doubleclick id is a user-level identified provided by Google. It is a cookie-based id, but is much more persistent than a typical cookie because it is a network-wide cookie, as opposed to a website’s cookie.

⁵These account for sales made through the online channel, which is significant for the website. There may be more sales occurring through offline channels, which we do not observe.

of an average person visiting that website. Competition from LumberLiquidators is even higher. Moreover, spending on marketing and advertising including retargeting is prevalent among the players in this category. In our investigation, all five of the competitors we considered engaged in retargeting.

B Description of BuildDirect Users from comScore Data

To get a more general description of the users, we acquired data from comScore MediaMetrix. Figure 7 compares demographic characteristics for BuildDirect users with the general population on the internet. It shows that the users are more heavily sampled from the income-bucket \$60k-75k, are more likely to be older and from larger households with children.

Figure 8 plots the website usage across months and shows that the website experiences less usage in the winter months, when home-improvement category experiences a general slowdown.

Table 5 compares average monthly site visitation data with close competitors that sell products in categories similar to BuildDirect. It shows that statistics for BuildDirect are comparable to competitor websites – the number of visitors, number of page views, time spent, and percentage population reached by BuildDirect are of the same order as other competitors in this category. Of course, general retailers such as Amazon are much larger and deal in categories beyond home-improvement, so they get much more traffic in the internet.

C Performance of the Tag Campaign

In this section we provide more details on the tag campaign described in section 2.2 of the paper. Ideally, one would want the tag campaign to tag all individuals eligible for the campaign, so the experiment can estimate an average effect across all possible consumer types. However, in practice, the tag campaign might miss tagging some people, excluding them from the experiment. This could occur if people do not give an opportunity to serve the tagging PSA ad within the time window of the tag campaign. Or if the tag campaign is

not able to deliver its one impression even with the high bid. Excluding some people from the experiment does not invalidate the experiment because the individuals included can be randomized across groups. However, if very few people are tagged, one might worry about the representativeness of the experimented-on population, and the statistical power of the experiment. In our experiment 71.1% of the people who received DoubleClick retargeting impression were tagged, and the rest were not tagged. One might be able to increase this proportion by increasing the time-window of the tag-campaign. In our case we limited the tag window to one day because we wanted to get a tight estimate of when the consumer entered the retargeting campaign, because our objective was to examine the temporal dimension of retargeting. It is difficult to verify in the data whether the date of the tag-impression is the actual date of the beginning of retargeting. Indeed, a purpose of the tag campaign is to inform us about the beginning of the campaign. However, we can examine the difference in days when the first retargeting impression occurred and when the first tag campaign occurred, considering that both of them are noisy measures of campaign start. Figure 9 below displays the distribution, which shows that majority of individuals received their first tag impression and retargeting impression on the same day. Importantly, this distribution is statistically similar regardless of whether the user was assigned to F5 or F15 in week 1; a Kolmogorov-Smirnov test is unable to reject that the differences are same ($p\text{-value} > 0.8$).

D More Randomization Checks

The paper shows that the users are balanced across conditions in terms of observable behavior before they are assigned to an experimental condition. In this section we show more checks based on pre-experiment time period. We check whether the mean consumer activity pre-experiment differed across the various experimental groups. Column (1) in Table 3 reports the p-values from these tests for the product-viewers experimental conditions. The combined bonferroni-adjusted p-value is 0.54, indicating the conditions are similar across these dimensions including the number of visits to the website in the pre-experimental period, the number of carts created, the number of orders placed, the number of free samples ordered, and the number of days a user was active. Looking at individual p-values, we

note that for one measure – the number of samples ordered – the p-value is lower. This is possible by chance, while testing multiple hypotheses. In our case it is likely to be driven by outlying observations in one experimental condition. We test for this in two subsequent checks. First, we consider the incidence i.e., whether free samples are ordered or not (instead of the number of occurrences) in the pre-experimental period, and find that the difference turns out to be insignificant. In a second check, we drop one condition which has outliers, and find that we cannot reject the hypothesis that the conditions have equal mean values. Overall, these randomization checks show that consumers in the different conditions are not systematically different in terms of their baseline behaviors. We repeat these tests for the cart-creators retargeting campaigns. Column (2) of Table 3 reports the corresponding p-values and shows that there is no indication of significant difference between the various experimental conditions in the means of the pre-experiment behaviors of consumers.

Additionally, we use data on ad impressions received by the users during the pre-experimental time to check whether the number of ads seen in the past is same across conditions. Among the users in the product-viewers campaign, the average impressions prior to the experiment is 6.06. The average does not vary significantly across conditions ($p = 0.26$). The corresponding average is 2.38 among users in the cart-creators campaign, and does not vary significantly across conditions ($p = 0.87$). Figure 6 shows the distribution of the number of impressions received by users in the product-viewer campaign during the pre-experimental time period. For ease of presentation, the chart shows data on users who received at least one impression. It shows the distribution separately across groups of users allocated to F0, F5 or F15 in week 1. A visual inspection suggests that the distributions are similar. An F-test indicates that the averages across the three groups are statistically indistinguishable ($p=0.56$).

E Experimental Design Implications

At the core of our analysis, our experiment design gives us multiple conditions to compare, and we do not rely solely on comparisons with the “no advertising throughout” to make inferences about the week by week impact of advertising. This feature of our design enables

us to correctly measure the effects of advertising. The following example illustrates this aspect, which is crucial for our analysis.

Rationale Let’s look at a specific two period example closely. How can we estimate the unbiased marginal effect of week 2 advertising? Given our design, we can pick groups of individuals who got *identical* treatment in week 1, and differ *only* in week 2. For illustration, consider the split of population into four groups shown in Figure 10. In the beginning, individuals are randomly assigned to week 1 and week 2 frequency-caps. So the populations in the four groups are comparable, hence shown in the same color at $t=0$ in the figure. In week 1, groups A, B get Criteo only, and C, D get Criteo and DoubleClick impressions. Therefore, we can compare A,B pooled, with C,D pooled to get the effect of DoubleClick impressions in week 1. By the end of week 1 we may have two different kinds of populations (in green and yellow). However, we can estimate the effect of week 2 advertising for yellow populations and green populations *separately*, comparing A with B, and C with D. In the A,B comparison, we compare people who got assigned to no experimental ads in week 2, with those who got assigned experimental ads in week 2. This way we get the effect of week 2 assignment for people who got no-ad assignment in week 1, and were identically treated until week 2. This is exactly what we do in the bars in the left panel in figure 10 in the paper. The panel on the right-hand-side of the same figure shows the other comparison (C,D), of people who *identically* got the other treatment in week 1.

Data We plot data to support the above logic. Figure 11 shows mean visit probabilities for the four groups of populations at the three different times ($t=0,1,2$). The colors of the bars match the instance in Figure 10. The graph on the top-left shows that the four groups are *a priori* similar in terms of visitation to the website before the experiment. After week 1 treatment, the yellow populations have lower visits in week 1 relative to green, as expected. Importantly, the two yellow bars, and the two green bars are similar. The bottom left graph is the same as figure 10 in the paper, placed here for matching with the conceptual discussion.

Overall, our design is able to get us individuals who are identical up to a point, and receive different treatments starting that point. It would be problematic if we were comparing week

2 outcomes for group A with group D, that is, the behavior of people who got no ads in both weeks with people who got ads in both weeks.

F Additional Analysis

Changes in visits at the intensive margin The paper tests whether retargeting causes more people to return to the website. Tables 6 to 9 expand the analysis to show that the effect exists at the intensive margin as well; proportion of visitors with more pages also increases.

Accidental Ad Clicks We attempt to account for the effects that may be driven by users accidentally clicking on the ads. For this we go back to our activity logs data and remove any session with just one activity. Then we create a new 0/1 dependent variable “Visit_noLoneActivity” for any user x time_period combination, which indicates whether the user had a session with more than one visit (no lone visit). This corresponds to our current dependent variable “Visit” which does not ignore sessions with just one visit. We rerun our main analysis, and the analysis on complementarity with the new dependent variable.⁶ Tables 10 to 11, and Figures 12 and 13 show the new analysis. They show that our substantive findings do not change. We still find retargeting to have an effect, and complementarities to exist. We stress tested this analysis by being even more conservative. We removed all sessions with 2 or less pages browsed. The findings remain unchanged. Overall, we conclude that our findings are robust to removal of sessions with one or two pages browsed, which may be unintended accidental visits.

Controls for Freq-caps in other weeks We note in the paper that five experimental conditions in the product viewers experiment were removed from the analysis because they were not implemented correctly. This leads to an imbalance in the distribution of frequency-caps across experimental conditions. Because of these missing conditions frequency-cap

⁶This test is conservative because a consumer seriously researching a product could also visit just one page in a session. This test would remove such visits.

assignment in one week may not be independent of frequency assignment in another week.⁷ If this was having a significant impact on our estimates, we would expect “controlling for” assignments in other weeks to change our estimates. For example, if week 1 and week 2’s frequency-caps are positively associated, the measured impact of week 2’s frequency-caps on week 2 visits would decrease when we control for week 1’s assignments if week 1 advertising directly affects week 2 visits. To check this we rerun the week-by-week analysis controlling for the frequency-caps in the non-focal week using fixed effects. Table 13 shows the results. The estimates do not change, and the fixed effects are insignificant, suggesting no interference from other week assignments. We also reran the tests in our test for complementarity between weeks 1 and 2 in a regression framework, controlling for frequency-caps in weeks 3 and 4. Here also we found our results to be unchanged. The effect of week 2 advertising is higher when week 1 advertising is switched on in this regression (p-value=0.02).

Examining Ad Impressions All tables in the paper reporting ad effects note the average incremental impressions per day per individual caused by the focal treatment. To compare week-by-week effects “per impression” we normalized the point estimates in Table 3 in the paper, dividing them by the incremental average impressions per day per individual. Table 12 shows these adjusted numbers. We also examine if the incremental impressions due to week 2 advertising depend on week 1 advertising. Table 15 regresses the total number of BuildDirect ad impressions (DoubleClick + Criteo) received by a user in week 2 on the ad assignments in weeks 1 and 2. The coefficients for “Positive F-cap both weeks” is small and statistically insignificant, indicating that the incremental impressions in week 2 (due to ads switched on) is the same irrespective of the frequency-caps in week 1.

Accounting for Ad Impression Changes in Complementarity results The analysis so far shows that week 1 advertising did not affect week 2’s manipulation statistically significantly. Here, we further examine the complementarity result in the paper, accounting for the

⁷Lets take a 2×2 design example (two time-periods 1,2; two levels of advertising No-Ad and Ad), in which the issue becomes apparent. Following our design logic, we would like to have four conditions (No-Ad, No-Ad); (Ad, No-Ad); (No-Ad, Ad); (Ad, Ad); where first and second arguments within parentheses represent the assignments for weeks 1 and 2 respectively. If all four conditions are taken together, week 2 advertising is uncorrelated with week 1. If (No-Ad, Ad) is removed, then Week 2 advertising is positively correlated with week 1 advertising (because week 2 advertising \implies week 1 advertising).

magnitude of the point estimates. Complementarity between week 1 and week 2 advertising relies on comparison of two changes. First change (δ_1) is from *week 2 off* to *week 2 on* when *week 1 off* (left-panel of Figure 10 in the paper). Point estimates in Table 15 show that this change is derived from an average of 8.00 incremental impressions per individual. The second change (δ_2) is from *week 2 off* to *week 2 on* when *week 1 on* (right-panel of Figure 10 in the paper). This change is derived from an average of 8.33 incremental impressions per individual. Therefore, δ_2 occurred with 4.1% ($= \frac{.33}{8.00}$) more impressions. However, we can reject the hypothesis that $1.04 \times \delta_1 = \delta_2$. Furthermore, we can say that the impact of switching on experimental advertising in week 2 has a 50% larger impact when week 1 ads are switched on, with more than 90% confidence. Therefore, it looks like week 1 assignment's impact on the effectiveness of week 2 advertising is not coming through an increase in the number of impressions delivered.

Does Week 1 Assignment Systematically change the Distribution of Week 2 impressions? Having seen that the level of impressions does not change because of week 1 experimental assignment, we further investigate evidence for whether the distribution of impressions changes in a systematic way. One important pre-experimental variable we observe is whether an individual visited the website in the pre-experiment time period. We know that historical visitation behavior is highly predictive of future behavior.⁸ For exposition of the rationale, let H denote people who visited historically, and let L denote the remaining population. We ask the question: is the Doubleclick campaign more likely to target H type people in week 2, when advertising in week 1 is turned on? In other words, does the proportion of impressions delivered to type H individuals in week 2 increase when week 1 frequency cap is high? If Doubleclick is able to learn an individual's type through week 1 advertising, one might see the distribution of week 2 impressions changed because of week 1 advertising. Figure 14 plots the proportion of impressions delivered to people with pre-experimental activity in a given day in week 2, separately for whether those who were assigned a frequency-cap of 0 in week 1 or not. We can see that this proportion remains

⁸Having visited builddirect in the pre-experimental time period makes visiting in any week during the campaign more likely (p-val < 0.001).

between 25-30% over days, and is similar between the two groups of individuals.⁹

G Mechanism: Output Interference

As discussed in the paper, output interference theory predicts that BuildDirect’s advertising decreases the likelihood of the consumer recalling a competitor, which can generate our findings. The objective of this section is to clarify this assertion using a stylized model.

Model setup The setup is similar to the model in Sahni (2016). Consider a scenario with two products A and B . A is the focal advertised product with advertising level a . A consumer chooses one product from the set of products she remembers. This choice could be to buy or to gather information. Let p_A and p_B denote the probability of remembering products A and B . Let π_S be the probability of choosing the focal product A from the set of products S she remembers. S can include either A or B , or both or none. Therefore the probability of choosing A , denoted by $y = p_A(1 - p_B)\pi_A + p_A p_B \pi_{AB}$. The first term corresponds to the user remembering only A and not B ; the second term corresponds to the user remembering both A and B ; other configurations of the set of remembered products do not include A so they have no contribution in the probability that A is chosen.

We assume a increases the likelihood of the user remembering A . Therefore $\frac{\partial p_A}{\partial a} > 0$ and $\frac{\partial p_A^2}{\partial a^2} \leq 0$ following the standard model of positive but diminishing effect of advertising on reminders. If a interferes and decreases the recall of B , then $\frac{\partial p_B}{\partial a} < 0$. As a increases its effect of the decrease of recalling B goes down in absolute value, therefore $\frac{\partial p_B^2}{\partial a^2} \geq 0$.¹⁰ If there is no interference, then this model reduces to the standard informative model with $\frac{\partial p_B}{\partial a} = 0$ and $\frac{\partial^2 p_B}{\partial a^2} = 0$.

⁹To statistically test whether week 2 impressions distribution changes because of week 1 assignment, we regress impressions an individual gets in a day on an indicator of whether he/she had pre-experimental activity, an indicator of week 1 frequency-cap being positive, and an interaction of the two. We find the interaction term to be insignificant (p-value=.53), suggesting that week 1 advertising is not giving more impressions to people with pre-experimental activity.

¹⁰This assumption on the second partial derivative is reasonable, but is not crucial, as can be seen in the later discussion.

Implications of the model Given this model, one can see that

$$\frac{\partial y}{\partial a} = \frac{\partial p_A}{\partial a} ((1 - p_B)\pi_A + \pi_{AB}p_B) + p_A \frac{\partial p_B}{\partial a} (\pi_{AB} - \pi_A) . \quad (1)$$

Note that $\pi_{AB} - \pi_A < 0$ because the likelihood of A being chosen always decreases when B is added to the set of products the user remembers. Given the above setup, and the fact that $\pi_{AB} - \pi_A < 0 \implies \frac{\partial y}{\partial a} > 0$; probability of choosing A increases with its advertising.

Also

$$\frac{\partial^2 y}{\partial a^2} = \underbrace{\frac{\partial^2 p_A}{\partial a^2} ((1 - p_B)\pi_A + \pi_{AB}p_B)}_{\leq 0} + \underbrace{\left(p_A \frac{\partial^2 p_B}{\partial a^2} + 2 \frac{\partial p_A}{\partial a} \frac{\partial p_B}{\partial a} \right) \cdot (\pi_{AB} - \pi_A)}_{?} . \quad (2)$$

The first term is weakly negative because $\frac{\partial^2 p_A}{\partial a^2} \leq 0$. However, the sign of the second term is ambiguous, which can cause $\frac{\partial^2 y}{\partial a^2} > 0$ which implies that advertising can exhibit increasing returns or complementarity. The following discussion further clarifies the implication of interference, that is, $\frac{\partial p_B}{\partial a} < 0$.

Standard model: no interference Under the standard model, p_B does not depend on a . Therefore, (a) response to advertising is concave: $\frac{\partial y}{\partial a} > 0$ and $\frac{\partial^2 y}{\partial a^2} \leq 0$. The latter is because the second term in equation (2) is zero when p_B is independent of a . (b) Additionally, in this case when awareness for A is high, $p_A \rightarrow 1$ then $\frac{\partial p_A}{\partial a} \rightarrow 0$ (because probability cannot increase beyond 1) $\implies \frac{\partial y}{\partial a} \rightarrow 0$. In other words, the effect of a on y vanishes when awareness for A is already high.

Allowing for interference Both constraints (a) and (b) get mitigated when a decreases the likelihood of B being remembered. Firstly, even when awareness for A is high, $p_A \rightarrow 1$ then $\frac{\partial p_A}{\partial a} \rightarrow 0 \nRightarrow \frac{\partial y}{\partial a} \rightarrow 0$ because the second term in equation (1) is positive even when $p_A \rightarrow 1$. In other words, the effect of a on y can be positive even when awareness for A is already high. Secondly, the second derivative can be positive. This is likely to occur when $\frac{\partial^2 p_A}{\partial a^2}$, $\frac{\partial^2 p_B}{\partial a^2}$ are small, and $\frac{\partial p_B}{\partial a}$ and $\frac{\partial p_A}{\partial a}$ have a significant magnitude. Theoretically, this possibility can very well exist.

H Mechanism: Displacing other ads

Retargeted ads can displace other ads that compete for the targeted individual’s attention, and can distract her by showing competing products in the retargeter’s category, or other products that she might get interested in. Therefore, retargeting can make it more likely that the individual returns to the retargeter’s website.

The following example illustrates the mechanism. A consumer, let’s call him Jack, is interested in installing a hardwood floor and goes to BuildDirect.com to search for available options. After browsing a few products Jack exits the website, and browses websites on how he can get hardwood flooring installed. Depending on whether he is retargeted by BuildDirect, he has the following experience. Our description is based on our simulation of the search session, described in Figure 15.

(A) No Retargeting: Jack searches Google for “install hardwood floor” and explores the first organic link which takes him to diynetwork.com, which is a popular website with information on do-it-yourself projects.¹¹ On the top of the page he sees a banner ad from lumberliquidators that mentions a deal on flooring. On the right he sees a banner that invites him to search homeadvisor for local contractors that can help install floors. After visiting this page he explores the next two relevant organic links (not owned by a retailer such as Lowes). He goes to hometips.com where he sees ads by “Cali Bamboo”, and younghouselove.com where he sees a banner by the home depot. Both Cali bamboo and The Home Depot sell hardwood floors.

(B) Retargeting: As in the previous condition, Jack searches Google for “install hardwood floor” and explores the first organic link which takes him to diynetwork.com.¹² On the page he sees two banners by bulddirect, showing the hardwood flooring products he had just browsed. The lumberliquidators and homeadvisor banners are displaced. After visiting this page he explores two organic links. He goes to hometips.com and younghouselove.com where he sees bulddirect banners in both cases.

¹¹We simulated this experience by conducting this search in “incognito” mode on google chrome.

¹²We simulated this experience by conducting this search in “incognito” mode on google chrome. The difference is that we first visited bulddirect.com and searched for hardwood flooring products.

Exposure to competitors’ ads in condition (A) is not surprising, given that sellers of hardwood floors want to show ads to people reading about installing hardwood floors, because this segment of population is likely to buy hardwood floors. Such ads are likely to make Jack aware of options other than builddirect; he might check out lumberliquidators.com or calibamboo.com and never return to builddirect.com or contact a contractor through home-advisor and take the contractor’s suggestion. On the other hand, in condition (B) Jack does not see other ads and is more likely to return to builddirect.com.

In the above example the ads displaced are from close competitors. In other cases, retargeting might displace ads for products in unrelated categories. In those cases the impact on retargeting could be indirect. For example, if Jack had searched for “landscaping” instead of “install hardwood flooring”, he might go to hgtv.com through google and see a banner by Angie’s list (as we saw in our simulation), which can (1) take time away from buying hardwood flooring; (2) remind him of searching for hardwood floors on aggregators like Angie’s list.

I Non-Experimental Estimates

To illustrate the value of experimentation in this context, we conducted the following analysis. What would the advertiser estimate in the absence of an experiment? Specifically, without the experimental control group, what would be an advertiser’s estimate of the effect of the campaign? Is this estimate significantly different from our experimental estimate? To answer this question we picked the treatment group that got high frequency-cap throughout the four weeks, and estimated the effect of setting this frequency cap schedule on product-viewers revisiting the website (our main effect in Table 2). We estimated the model

$$y_i = \beta_1 \text{ads}_i + \beta_2 \text{ads}_i^2 + x_i + \epsilon_i \quad (3)$$

where y_i is an indicator of whether i visits BuildDirect.com in four weeks of the campaign, x_i is a vector of control variables comprising pre-experimental behavioral data we observe, and ads_i is the number of ad impressions i got from BuildDirect’s DoubleClick campaign. We

picked the quadratic functional form because it was most significant; higher order terms were insignificant. We estimated both linear and logistic specifications of the model in equation (3). Then we used the model estimates to predict the change in y between the treatment group and the control group, which received the frequency-cap of 0 throughout. Comparing this non-experimental estimate with the experimentally estimated effect (shown in Table (16)) we find that these models tend to significantly overestimate the effect of retargeting.

J Precluding Users who Exit the Campaign

We repeated the analysis in table 3 of the paper (week-by-week contemporaneous effects), after removing individuals who had created a cart by the beginning of the week. In other words, we estimate the effect of week 2 frequency caps on people who had not created a cart by week 2; these are the people who did not self-select themselves out of the campaign. Table 17 below shows the estimates with cart-creators excluded. Compared to the other analysis in table 3 in the paper, note that the number of observations now decrease across columns (because as people create carts they get excluded from the analysis). In terms of the effects, we find that the estimated effects do not change. None of the estimates are statistically, or economically different from what we have in table 3. Qualitatively, we still see that the percentage of people returning due to the campaign decreases over weeks. Table 18 does a similar analysis for cart creators campaign. Recall that for cart creators we just have two time periods (weeks 1,2 combined, and weeks 3,4 combined). And we just have two levels of frequency-caps. Column (1) of the table shows the effect of week 3,4 advertising on the full sample. Column (2) excludes individuals who made a purchase in weeks 1,2. Again, we don't see any significant change in the estimated effect of advertising.

Figure 1: Example of Retargeting Campaign



Figure 2: A snapshot of BuildDirect homepage

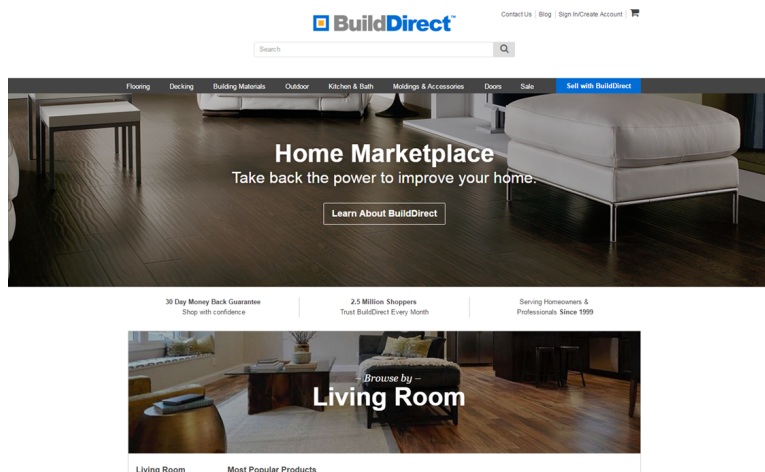


Figure 3: A snapshot of an example search page on builddirect.com

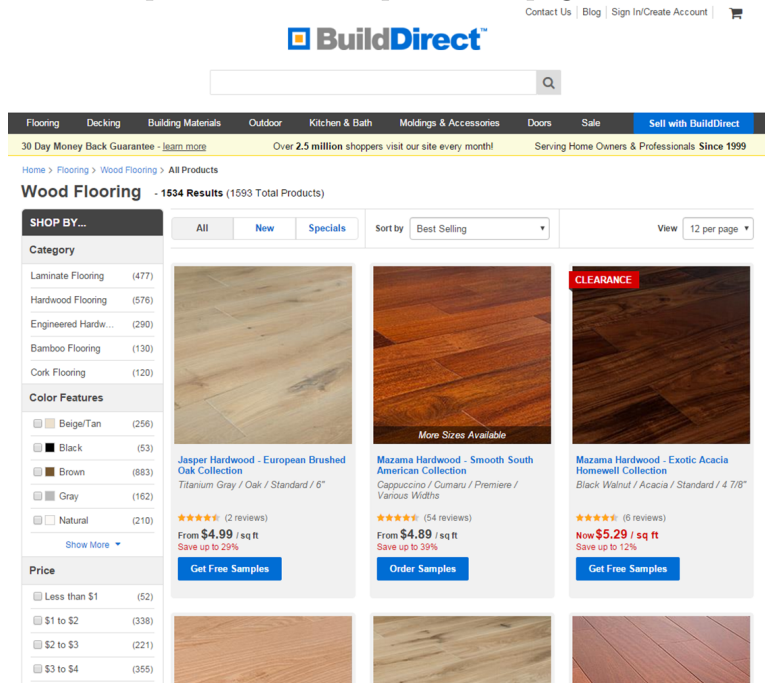


Figure 4: A snapshot of an example product page on builddirect.com

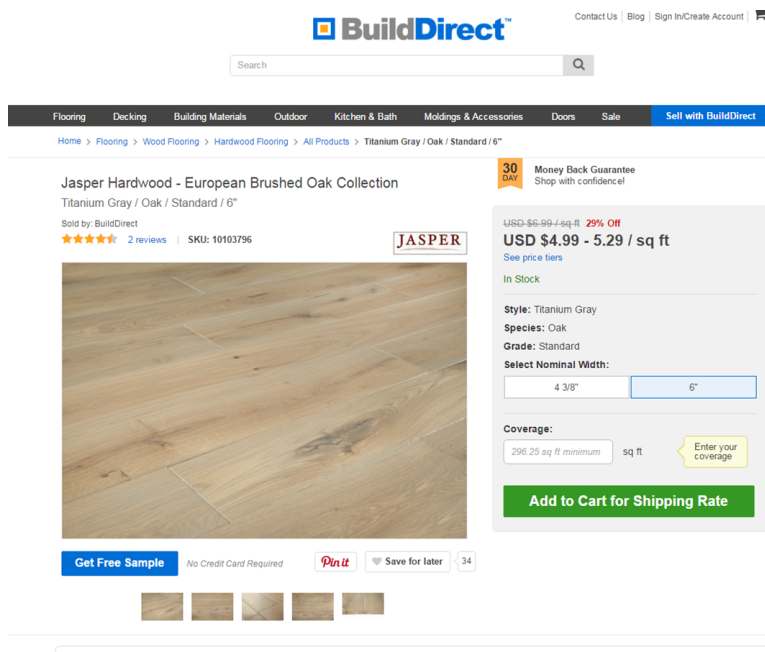


Figure 5: A snapshot of an example page seen after creating a cart

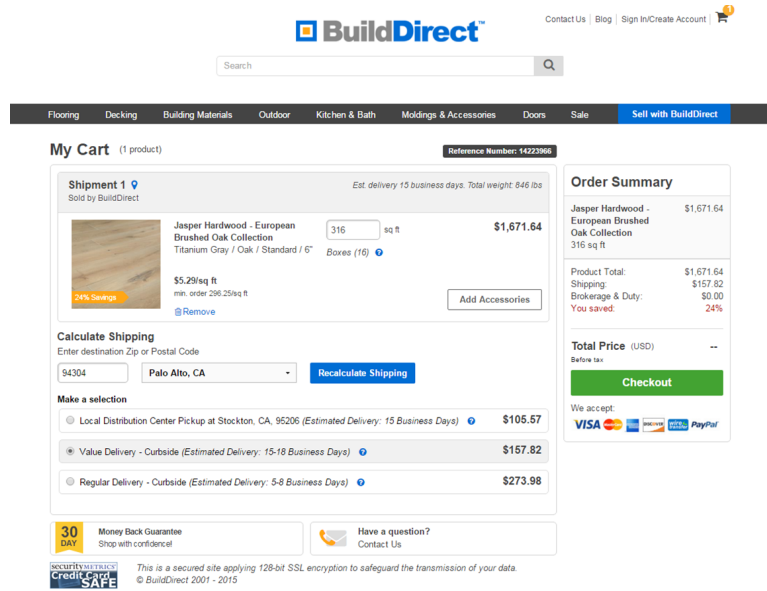
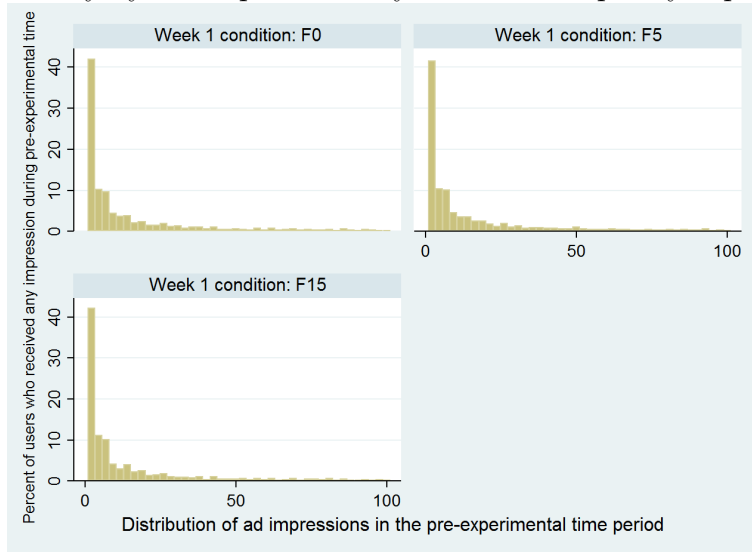


Figure 6: Distribution of ad impressions received in the pre-experimental time period, separately by the experimentally allocated frequency cap for week 1



Notes: The chart shows histograms that display the distribution of the number of ad impressions received by users in the pre-experimental time period. For ease of presentation we use data on users who receive at least one impression during this time period. The distribution is presented separately by the frequency-cap the users were allocated to in the first week of the experiment. A comparison across the figures shows that the distributions are similar across the three groups, which supports that randomization achieved balance across conditions. An F-test indicates that the averages across the three groups are statistically indistinguishable ($p=0.56$).




Figure 7: Demographics of BuildDirect.com users relative to average population on the internet

Notes: Index of 100 represents the general internet population in comScore's data. A number greater than 100 means that the corresponding bucket is over represented on BuildDirect.com, and a number less than 100 means the opposite.




Figure 8: Trend of website usage across months

Figure 9: Distribution of days between the first retargeting impression and the tag impression.

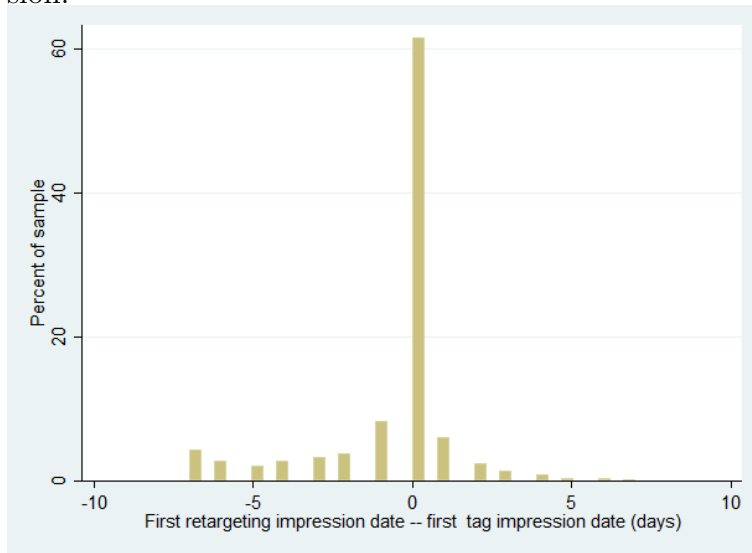


Figure 10: Experiment design implication

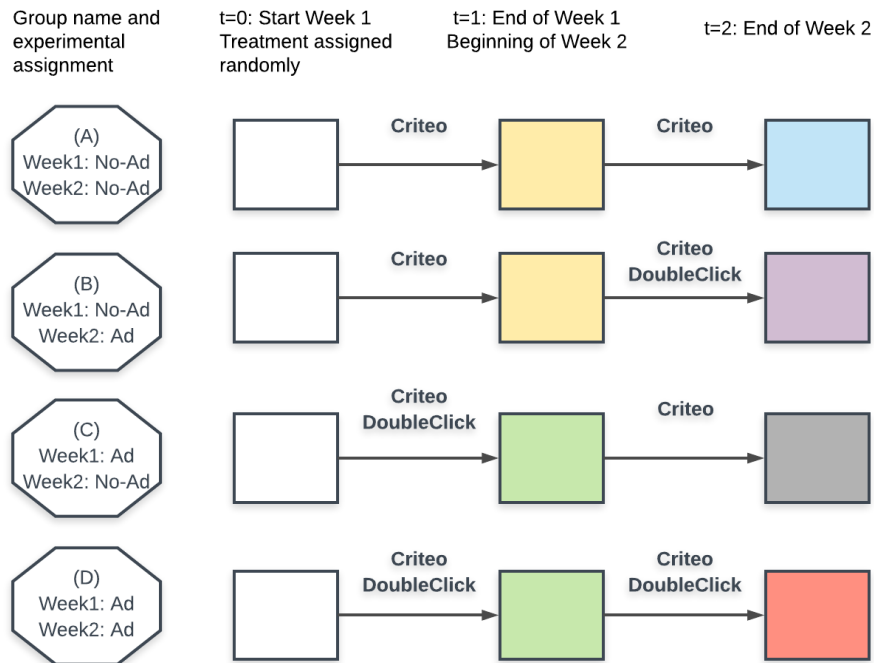


Figure 11: Experiment design implication




Figure 12: Ignoring one-visit sessions: Figure corresponding to Figure 10 in the paper.

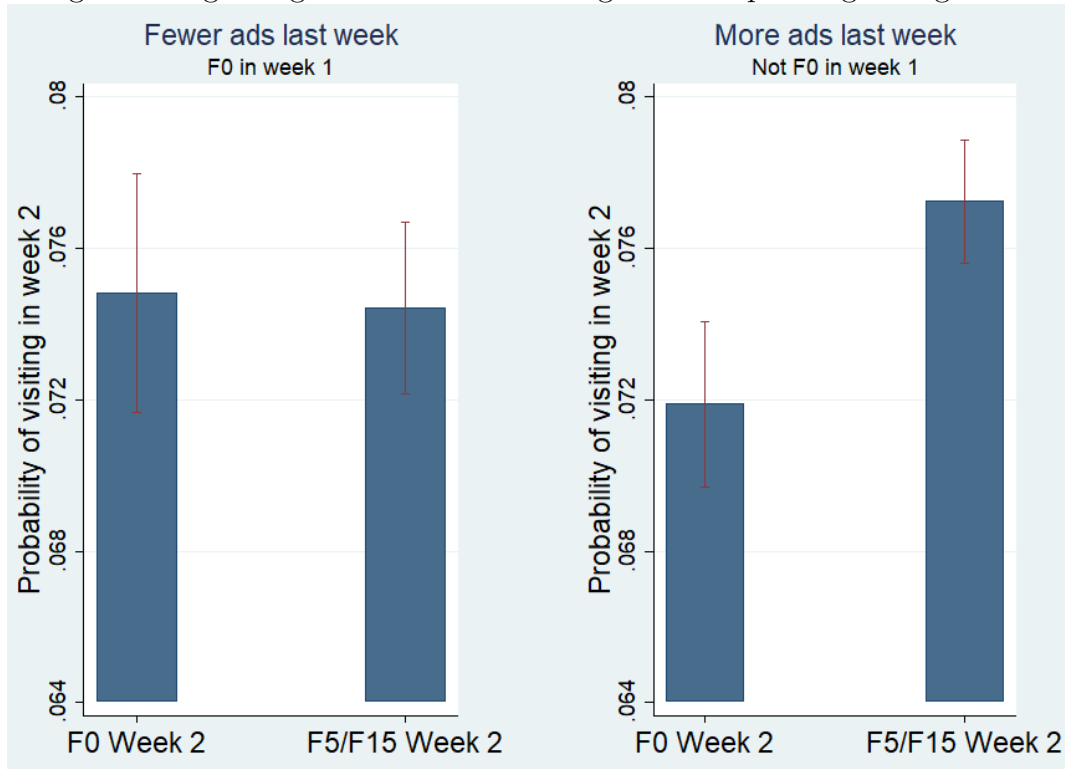


Figure 13: Ignoring one-visit sessions: Figure corresponding to Figure 11 in the paper.

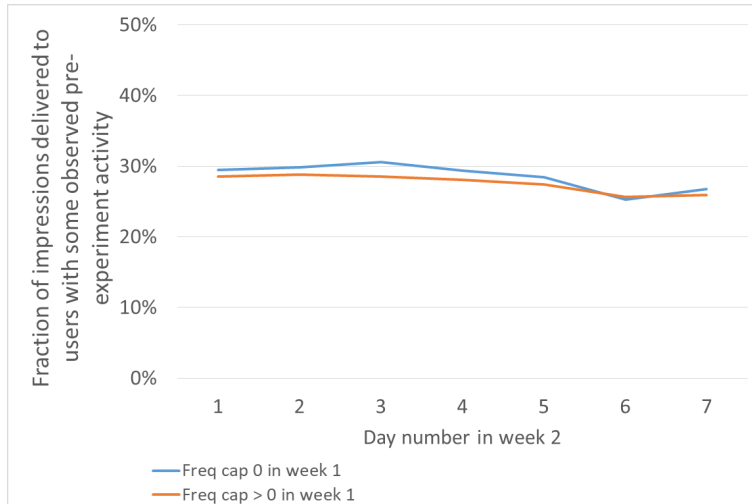




Figure 14: In a given day in week 2, the chart shows the proportion of impressions delivered to people who had any pre-experimental activity.

Notes: The figure shows that the proportion of week 2 impressions delivered to individuals who have any pre-experimental activity remains between 25-30% across days, and does not change by whether the individual is assigned to a positive frequency-cap in week 1. A statistical test shows that the differences are statistically insignificant. Specifically, we regress impressions an individual gets in a day on an indicator of whether he/she had pre-experimental activity, an indicator of week 1 frequency-cap being positive, and an interaction of the two. We find the interaction term to be insignificant (p-value=.53), suggesting that week 1 advertising is not giving more impressions to people with pre-experimental activity.

Simulated user experience

- (1) Search for Hardwood flooring products on BuildDirect.com
- (2) Search google for “install hardwood floor”

When not being retargeted by BuildDirect



When being retargeted by BuildDirect






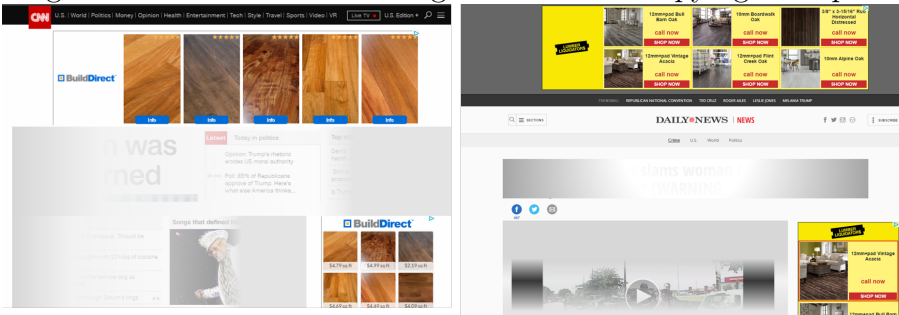


Figure 15: Difference in the user’s advertising exposures by whether he/she is retargeted

Figure 16: Screenshots showing advertisers occupying multiple ad-slots on the same page



Notes: The graph displays two examples of a common situation in which the same advertiser occupies multiple slots on the same page. The screenshot on the left shows a page with two ads by BuildDirect, and the one on the right shows a page with two ads from LumberLiquidator, which is BuildDirect’s competitor

Table 1: Descriptive statistics on users in our data.

Number of users with any activity=234,712

	Mean	Std. dev.
<i>Statistics on browsing behavior</i>		
Number of sessions (days on which the user interacted with the website)	2.55	3.33
Number of days spanning a user’s interaction with the website (last date - first date)	16.34	31.15
Number of days spanning a user’s interaction with the website (last date - first date) conditional on return	35.19	37.77
Number of product pages browsed	24.96	108.40
Number of search pages browsed	18.86	55.11
<i>Statistics on “conversion activity”</i>		
Probability of creating a cart	0.1357	0.3424
Probability of ordering sample	0.0382	0.1917
Probability of ordering a product	0.0039	0.0626
Probability of creating a cart conditional on clicking on a retargeted ad	0.2032	0.4024
<i>Statistics on browsing conditional on conversion</i>		
Number of sessions (days on which the user interacted with the website) for those who created a cart	5.52	6.01
Number of sessions before a cart is created (days on which the user interacted with the website) for those who created a cart	2.12	2.13
Number of days spanning a user’s interaction with the website (last date - first date) conditional on creating a cart	35.26	40.52
Number of days spanning a user’s interaction with the website (last date - first date) conditional on ordering a sample	51.25	41.55
Number of days spanning a user’s interaction with the website (last date - first date) conditional on ordering a product	56.98	41.95
Days between first interaction and when the cart is created (for those who created a cart)	9.44	19.61
Days between first interaction and when a sample is ordered (for those who order a sample)	14.14	22.74
Days between first interaction and when a product is bought (for those who ordered)	23.89	26.84

Competitor	Percentage of individuals who visited BuildDirect, and also visited the competitor	For baseline reference: Percentage of individuals in the population who visit the competitor (unconditional on visiting BuildDirect)
HOMEDEPOT.COM	50.50%	13.6%
LOWES.COM	34.80%	8.6%
WAYFAIR.COM	34.40%	5.6%
LUMBERLIQUIDATORS.COM	17.30%	0.3%
BUILD.COM	13.30%	0.2%

Table 2: Assessing competition. The table shows the percentage of individuals who visited BuildDirect.com and also visited another competitor’s website. Source comScore MediaMetrix, April 2015.

Table 3: Randomization Checks: Test for differences of means across treatment conditions

Dependent measure	(1) Product viewers campaign p(>F)	(2) Cart creators campaign p(>F)
Num of days of activity	0.672	0.562
Num Visits	0.861	0.295
Num Carts Created	0.210	0.747
Num of Orders	0.653	0.460
Num of Samples	0.054	0.489
Days of activity greater than 0	0.446	0.961
Num Visits greater than 0	0.322	0.975
Num Carts Created greater than 0	0.336	0.366
Num of Orders greater than 0	0.382	0.682
Num of Samples greater than 0	0.600	0.825
All DVs for pre-experimental period		

Table 4: Regression: Total Criteo impressions delivered in four weeks on indicators of frequency-caps set across weeks.

DV: Total Criteo impressions across 4 weeks			
	Coef.	Std. Err.	p-value
Week 1 F5	-0.33	0.39	0.39
Week 1 F15	-0.41	0.37	0.28
Week 2 F5	-0.24	0.38	0.53
Week 2 F15	-0.58	0.38	0.12
Week 3 F5	-0.32	0.38	0.40
Week 3 F15	-0.33	0.38	0.39
Week 4 F5	0.08	0.39	0.84
Week 4 F15	-0.02	0.38	0.95
Intercept	41.18**	0.47	<0.01
N	234,595		

Notes: This regression is statistically insignificant; an F-test is unable to reject the hypothesis that all coefficients are zero (p=0.79).

Attribute (per month)	BuildDirect.com	LumberLiquidators.com	EmpireToday.com	Build.com
Unique visitors (thousands)	858	717	190	598
Total Views (millions)	4	7	1	10
Total minutes spent (millions)	6	3	-	7
% reach of the internet	.3%	.3%	.1%	.2%
Average minutes per visit	5.0	2.2	0.9	4.5

Table 5: Comparing BuildDirect’s website browsing activity with close competitors. This table represents data from April 2015, a month when our experiment was in progress.

Table 6: Effect of Retargeting on visits in four weeks after entering the experiment: Product-Viewers Campaign

	(1)		(2)		(3)		(4)	
	DV: visits in four weeks ≥ 1		DV: visits in four weeks ≥ 2		DV: visits in four weeks ≥ 3		DV: visits in four weeks ≥ 4	
	Coeff	Std. err	Coeff	Std. err	Coeff	Std. err	Coeff	Std. err
Indicator for advertising switched on	0.0296***	0.0091	0.0291***	0.0091	0.0232***	0.0088	0.0181**	0.0083
Intercept (Baseline: F0 condition)	0.2023***	0.0073	0.2023***	0.0073	0.1862***	0.0071	0.1625***	0.0067
N	8,999		8,999		8,999		8,999	

Notes: (* $p < 0.1$; ** $p < 0.05$, *** $p < 0.01$) The table presents coefficients and robust standard errors from several OLS regressions across its columns. For the purpose of this analysis, we pool data for three conditions in which advertising frequency-cap remained constant, specifically, F0 throughout, or F5 throughout or F15 throughout. The independent variable in each of the regressions is an indicator of an experimental condition in which retargeting was turned on. The dependent variable for the first column is an indicator of whether the user came back to the website in the four weeks after entering of the experiment. The coefficient for the indicator of advertising being on is positive and significant, suggesting that retargeting brings back people who would not have visited the website in the next four weeks. Columns (2), (3) and (4) investigate whether the users' activity increases beyond just coming back and visiting the website once. The analysis shows that there is a significant shift in distribution of visits beyond 1.

Table 7: Effect of Retargeting on visits in eight weeks after entering the experiment: Product-Viewers Campaign

	(1)		(2)		(3)		(4)	
	DV: visits in eight weeks ≥ 1		DV: visits in eight weeks ≥ 2		DV: visits in eight weeks ≥ 3		DV: visits in eight weeks ≥ 4	
	Coeff	Std. err	Coeff	Std. err	Coeff	Std. err	Coeff	Std. err
Indicator for advertising switched on	0.0303***	0.0097	0.0299***	0.0097	0.0248***	0.0094	0.0224**	0.0089
Intercept (Baseline: F0 condition)	0.2421***	0.0078	0.2421***	0.0078	0.2230***	0.0076	0.1934***	0.0071
N	8,999		8,999		8,999		8,999	

Notes: (* $p < 0.1$; ** $p < 0.05$, *** $p < 0.01$) The table presents coefficients and robust standard errors from several OLS regressions across its columns. For the purpose of this analysis, we pool data for three conditions in which advertising frequency-cap remained constant, specifically, F0 throughout, or F5 throughout or F15 throughout. The independent variable in each of the regressions is an indicator of an experimental condition in which retargeting was turned on. The dependent variable for the first column is an indicator of whether the user came back to the website in the eight weeks after entering the experiment. The coefficient for the indicator of advertising being on is positive and significant, suggesting that retargeting brings back people who would not have visited the website in the next eight weeks. Columns (2), (3) and (4) investigate whether the users' activity increases beyond just coming back and visiting the website once. The analysis shows that there is a significant shift in distribution of visits beyond 1.

Table 8: Effect of Retargeting on visits in four weeks after entering the experiment: Cart-creators Campaign

	(1)		(2)		(3)		(4)	
	DV: visits in four weeks ≥ 1		DV: visits in four weeks ≥ 2		DV: visits in four weeks ≥ 3		DV: visits in four weeks ≥ 4	
	Coeff	Std. err	Coeff	Std. err	Coeff	Std. err	Coeff	Std. err
Indicator for advertising switched on	0.0204**	0.0081	0.0206**	0.0081	0.0169**	0.0080	0.0172**	0.0079
Intercept (Baseline: F0 condition)	0.3751***	0.0058	0.3744***	0.0058	0.3526***	0.0057	0.3223***	0.0056
N	14,351		14,351		14,351		14,351	

Notes: (* $p < 0.1$; ** $p < 0.05$, *** $p < 0.01$) The table presents coefficients and robust standard errors from several OLS regressions across its columns. For the purpose of this analysis, we pool data for three conditions in which advertising frequency-cap remained constant, specifically, F0 throughout, or F15 throughout (recall that the cart-creators campaign did not have a condition with frequency cap of 5). The independent variable in each of the regressions is an indicator of an experimental condition in which retargeting was turned on. The dependent variable for the first column is an indicator of whether the user came back to the website in the four weeks after entering of the experiment. The coefficient for the indicator of advertising being on is positive and significant, suggesting that retargeting brings back people who would not have visited the website in the next four weeks. Columns (2), (3) and (4) investigate whether the users' activity increases beyond just coming back and visiting the website once. The analysis shows that there is a significant shift in distribution of visits beyond 1.

Table 9: Effect of Retargeting on visits in eight weeks after entering the experiment: Cart-creators Campaign

	(1)		(2)		(3)		(4)	
	DV: visits in eight weeks ≥ 1		DV: visits in eight weeks ≥ 2		DV: visits in eight weeks ≥ 3		DV: visits in eight weeks ≥ 4	
	Coeff	Std. err	Coeff	Std. err	Coeff	Std. err	Coeff	Std. err
Indicator for advertising switched on	0.0199**	0.0083	0.0199**	0.0083	0.0166**	0.0082	0.0165**	0.0081
Intercept (Baseline: F0 condition)	0.4207***	0.0059	0.4200***	0.0059	0.3968***	0.0058	0.3664***	0.0057
N	14,351		14,351		14,351		14,351	

Notes: (* $p < 0.1$; ** $p < 0.05$, *** $p < 0.01$) The table presents coefficients and robust standard errors from several OLS regressions across its columns. For the purpose of this analysis, we pool data for three conditions in which advertising frequency-cap remained constant, specifically, F0 throughout, or F15 throughout (recall that the cart-creators campaign did not have a condition with frequency cap of 5). The independent variable in each of the regressions is an indicator of an experimental condition in which retargeting was turned on. The dependent variable for the first column is an indicator of whether the user came back to the website in the eight weeks after entering of the experiment. The coefficient for the indicator of advertising being on is positive and significant, suggesting that retargeting brings back people who would not have visited the website in the next eight weeks. Columns (2), (3) and (4) investigate whether the users' activity increases beyond just coming back and visiting the website once. The analysis shows that there is a significant shift in distribution of visits beyond 1.

Table 10: Ignoring one-visit sessions: Effect of Retargeting on visits

	Product Viewers				Cart Creators			
	(1)		(2)		(3)		(4)	
	DV: visits in four weeks>0 Coeff Std. err		DV: visits in eight weeks>0 Coeff Std. err		DV: visits in four weeks>0 Coeff Std. err		DV: visits in eight weeks>0 Coeff Std. err	
Indicator for advertising switched on	0.0291***	0.0091	0.0317***	0.0097	0.0206**	0.0081	0.0202**	0.0083
Intercept (Baseline: F0 condition)	0.2023***	0.0073	0.239***	0.0078	0.3743***	0.0058	0.4184***	0.0059
N	8,999		8,999		14,351		14,351	

Notes: (* p<0.1; ** p<0.05, *** p<0.01) The table presents coefficients and robust standard errors from OLS regressions across columns. Columns (1) and (2): data is pooled across three conditions, F0 throughout, F5 throughout, F15 throughout for product viewers campaign. Columns (3) and (4): data is pooled across two conditions, F0 throughout, F15 throughout for cart creators campaign. The independent variable in each regression is an indicator of retargeting switched on. The dependent variable (DV) is mentioned with the column number.

Table 11: Ignoring one-visit sessions: Week-by-week contemporaneous effects of advertising on the user visiting the website: Product viewers campaign

	(1)		(2)		(3)		(4)	
	DV: (0/1) visit website in week 1		DV: (0/1) visit website in week 2		DV: (0/1) visit website in week 3		DV: (0/1) visit website in week 4	
	Coeff	Std. err	Coeff	Std. err	Coeff	Std. err	Coeff	Std. err
Indicator for F5	0.0034**	0.0018	0.0026**	0.0013	0.0003	0.0011	0.0011	0.0010
Indicator for F15	0.0084***	0.0017	0.0042***	0.0013	0.0028**	0.0011	0.0022**	0.0010
Intercept (Baseline: F0)	0.1370***	0.0012	0.0728***	0.0009	0.0515***	0.0008	0.0378***	0.0007
N	234,595		234,595		234,595		234,595	

Notes: (* p<0.1; ** p<0.05, *** p<0.01) The table presents coefficients and robust standard errors from several OLS regressions across its columns. For this analysis, we pool data for all the conditions in our product-viewer campaign. In each of the regressions, the dependent measure is an indicator of whether the user visited the website during that week. The explanatory variables are indicators of whether the user is assigned to F5 or F15 during that week. F0 condition serves as the baseline (intercept). Therefore, the coefficients are the change in visit probability relative to F0.

Table 12: Week-by-week effects *per impression*: Product viewers campaign

	(1)	(2)	(3)	(4)
	DV: (0/1) visit website in week 1	DV: (0/1) visit website in week 2	DV: (0/1) visit website in week 3	DV: (0/1) visit website in week 4
Effect of F5 relative to F0	.0041	.0035	.0004	.0019
Effect of F15 relative to F0	.0040	.0027	.0021	.0017

Notes: The table presents point estimates from Table 3 in the paper divided by the average number of impressions delivered per day per individual in the respective experimental condition.

Table 13: Week-by-week contemporaneous effects of Product viewers campaign: with controls

	(1)		(2)		(3)		(4)	
	DV: (0/1) visit website in week 1		DV: (0/1) visit website in week 2		DV: (0/1) visit website in week 3		DV: (0/1) visit website in week 4	
	Coeff	Std. err	Coeff	Std. err	Coeff	Std. err	Coeff	Std. err
Indicator for F5	0.0035**	0.0018	0.0026**	0.0013	0.0003	0.0011	0.0011	0.0010
Indicator for F15	0.0085***	0.0017	0.0042***	0.0013	0.0027**	0.0011	0.0022**	0.0010
Intercept (Baseline: F0)	0.1357***	0.0021	0.0709***	0.0016	0.04977***	0.0014	0.0379***	0.0012
Controls for freq-caps in other weeks using fixed effects	Yes		Yes		Yes		Yes	
N	234,595		234,595		234,595		234,595	
p-value H_0 : fixed effect controls are zero	.97		.11		.48		.25	

Notes: (* $p < 0.1$; ** $p < 0.05$, *** $p < 0.01$) The table presents coefficients and robust standard errors from several OLS regressions across its columns. For this analysis, we pool data for all the conditions in our product-viewer campaign. In each of the regressions, the dependent measure is an indicator of whether the user visited the website during that week. The explanatory variables are two indicator-variables – (1) whether the user was allocated to F5 during that week, and (2) whether the user was allocated to F15 during that week. We also control for frequency-caps in other weeks by adding fixed effects for frequency-caps in other weeks. Compared to the Table 9 in the paper, we note that none of the estimated coefficients changes significantly. Further, the control variables are statistically insignificant in each of the four columns (F-test is unable to reject the hypothesis that the coefficients for control variables are all zero; $p > 0.10$ for each of the four columns).

Table 14: Effect of week 2 experimental ads by quartile (based on exposure to non-experimental ads in week 1)
DV: 0/1 indicator of whether the user visits the website in week 2

	Quartile 4			Quartile 3			Quartile 2			Quartile 1		
	Coef	Std. Err		Coef	Std. Err		Coef	Std. Err		Coef	Std. Err	
F5 in week 2	0.0076**	0.0035		0.0015	0.0028		0.0022	0.0027		0.0018	0.0014	
F15 in week 2	0.0119**	0.0035		0.0030	0.0028		0.0015	0.0027		0.0032**	0.0014	
Intercept (baseline: F0)	0.1366**	0.0024		0.0802**	0.0019		0.055**	0.0019		0.026**	0.0011	
N	58,253			58,338			42,787			75,217		

*Notes:*Quartile 1 sees fewest Criteo ads in week 1, and Quartile 4 sees the most.

Table 15: Change in average impressions (from experimental and non-experimental campaigns combined) received in weeks 1 and 2 with changing experimental assignment

	Impressions in week 2	
	Coef	Std. Err
Positive F-cap in week 1	.48	.54
Positive F-cap in week 2	8.00**	.61
Positive F-cap both weeks	.33	.72
Intercept (baseline: F0 F0)	25.78**	.45
N	234,595	

Notes: For every user we count the total number of BuildDirect ad impressions received in week 2, which is our dependent variable for the regression.

Specification	Controls	Model Estimate (not using the experimental variation)	S.e.	Experimental estimate	S.e.	Bias in the non- experimental estimate (relative difference in estimates)
Linear	No	.0878	.0017	.0325	.0106	170%
Logit	No	.1066	.0020	.0325	.0106	228%
Linear	Yes	.0693	.0014	.0325	.0106	113%
Logit	Yes	.0882	.0040	.0325	.0106	171%

Table 16: Comparison of estimates with/without using the experiment

Notes: The table shows estimated effect of the campaign (with frequency-cap of 15 throughout, which was also BuildDirect’s status quo) from various specifications of the model in equation (3). Comparing these with the experimental estimate we can see that model estimates significantly overestimate the effect.

Table 17: Week-by-week contemporaneous effects precluding cart-creators: Product viewers campaign

	(1)		(2)		(3)		(4)	
	DV: (0/1) visit website in week 1		DV: (0/1) visit website in week 2		DV: (0/1) visit website in week 3		DV: (0/1) visit website in week 4	
	Coeff	Std. err	Coeff	Std. err	Coeff	Std. err	Coeff	Std. err
Indicator for F5	0.0036**	0.0018	0.0028**	0.0013	0.0014	0.0011	0.0011	0.0009
Indicator for F15	0.0085***	0.0017	0.0045***	0.0013	0.0034**	0.0011	0.0023**	0.0009
Intercept (Baseline: F0)	0.1371***	0.0012	0.0677***	0.0009	0.0457***	0.0008	0.0330***	0.0007
N	234,595		230,961		229,528		228,603	

Notes: (* p<0.1; ** p<0.05, *** p<0.01) The table presents coefficients and robust standard errors from several OLS regressions across its columns. For this analysis, we pool data for all the conditions in our product-viewer campaign. We exclude individuals who have created a cart up to the week; because the campaign does not show ads to people who have previously created a cart, we are excluding individuals who are not likely to get impressions during the week. In each of the regressions, the dependent measure is an indicator of whether the user visited the website during that week. The explanatory variables are indicators of whether the user is assigned to F5 or F15 during that week. F0 condition serves as the baseline (intercept). Therefore, the coefficients are the change in visit probability relative to F0.

Table 18: Week-by-week contemporaneous effects precluding buyers: Cart creators campaign

	(1)		(2)	
	DV: (0/1) visit website in weeks 3&4		DV: (0/1) visit website in weeks 3&4	
	Without exclusion		With exclusion	
	Coeff	Std. err	Coeff	Std. err
Advertising switched on in weeks 3,4	0.002	0.0045	0.0012	0.0045
Intercept (Baseline: F0)	0.1426***	0.0032	0.1403***	0.0032
N	23,710		23,442	

Notes: (* p<0.1; ** p<0.05, *** p<0.01) The table presents coefficients and robust standard errors from several OLS regressions across its columns. In column (2) we exclude individuals who made a purchase prior to week 3; because the campaign does not show ads to people who have previously bought, we are excluding individuals who are not likely to get impressions during the week. In each of the regressions, the dependent measure is an indicator of whether the user visited the website during weeks 3, 4. The explanatory variables are indicators of whether advertising was switched on during weeks 3,4. F0 condition serves as the baseline (intercept).