

To be Tough or Soft: Measuring the Impact of Counter-Ad-blocking Strategies on User Engagement

Shuai Zhao

New Jersey Institute of Technology
sz255@njit.edu

Cristian Borcea

New Jersey Institute of Technology
borcea@njit.edu

Achir Kalra

Forbes Media
akalra@forbes.com

Yi Chen

New Jersey Institute of Technology
yi.chen@njit.edu

ABSTRACT

The fast growing ad-blocker usage results in large revenue decrease for ad-supported online websites. Facing this problem, many online publishers choose either to cooperate with ad-blocker software companies to show *acceptable ads* or to build a *wall* that requires users to whitelist the site for content access. However, there is lack of studies on the impact of these two counter-ad-blocking strategies on user behaviors. To address this issue, we conduct a randomized field experiment on the website of Forbes Media, a major US media publisher. The ad-blocker users are divided into a treatment group, which receives the wall strategy, and a control group, which receives the acceptable ads strategy. We utilize the difference-in-differences method to estimate the causal effects. Our study shows that the wall strategy has an overall negative impact on user engagements. However, it has no statistically significant effect on high-engaged users as they would view the pages no matter what strategy is used. It has a big impact on low-engaged users, who have no loyalty to the site. Our study also shows that revisiting behavior decreases over time, but the ratio of session whitelisting increases over time as the remaining users have relatively high loyalty and high engagement. The paper concludes with discussions of managerial insights for publishers when determining counter-ad-blocking strategies.

CCS CONCEPTS

• **General and reference** → **Empirical studies**; • **Information systems** → **Online advertising**; • **Human-centered computing** → **User studies**;

KEYWORDS

Ad Blocking, Randomized Field Experiment, Difference-In-Differences, User Studies, Online Advertising

ACM Reference Format:

Shuai Zhao, Achir Kalra, Cristian Borcea, and Yi Chen. 2020. To be Tough or Soft: Measuring the Impact of Counter-Ad-blocking Strategies on User Engagement. In *Proceedings of The 2020 Web Conference (WWW'20)*. ACM, Taipei, Article xx, 7 pages. https://doi.org/xx.xxx/xxx_x

Permission to make digital or hard copies of part or all 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 third-party components of this work must be honored. For all other uses, contact the owner/author(s).

WWW'20, 2020

© 2020 Copyright held by the owner/author(s).

ACM ISBN 978-x-xxxx-xxxx-x/YY/MM.

https://doi.org/xx.xxx/xxx_x

1 INTRODUCTION

An ad blocker is a tool, most likely a browser plugin, to remove ads while a user is reading online content. The broad usage of ad blockers has a big impact to the ad-supported web publishing system. Web publishers provide content for free, instead they gain revenue from digital advertising, which contributes 333.25 billion US dollars in 2019 [1]. With the increasing usage of ad blockers, it is expected that online publishers will lose revenue of 35 billion US dollars world-wide by 2020, and the loss has a steady increase at 30% per year [2]. Without sufficient revenue, publishers cannot afford to generate high-quality free content, which will ultimately hurt online users' interests.

In response, more and more online publishers (e.g., Wired, Forbes, AdAge, Digiday, Los Angeles Daily News) launched their counter-ad-blocking methods [3]. Rafique et al. found that counter-ad-blocking scripts were used by 16.3% of the 1,000 most popular domains [4]. Currently, there are two popular methods of counter-ad-blocking used by publishers. The first is the *tough* "whitelist-or-leave" strategy, and the second is the *soft* acceptable ads exchange (AAX) strategy. The "whitelist-or-leave" strategy works like a wall (called Wall strategy as well). When an ad blocker is detected, a publisher's website pops up a message requesting the user to turn off or pause the ad blocker, i.e., whitelist the publisher's website. If a user rejects the request, she is forbidden to access the content that she intends to view. The soft AAX strategy shows users acceptable ads, agreed upon with the ad blocking companies, which appear in the page even when an ad blocker is active. Acceptable ads are generally less annoying ads, such as text ads instead of video ads, and also fewer in number.

Despite of the importance of the ad blocking problem, there are few studies on it. Existing work [5, 6] focused on techniques and mechanisms for counter-ad-blocking, but not the effect of different counter-ad-blocking strategies on users' engagement. The work in [7], on the other hand, has studied the effect of such strategies. However, it compared the Wall strategy with the ads-free strategy under a retrospective quasi-experiment setting. It is not realistic for publishers to provide free content with no ads. We have different experiment setting and the goal of our study is to understand in-depth the differences of effect between the Wall strategy and the AAX strategy on user engagement, both of which are actively used by online publishers. Specifically, we want to address the following research questions:

- **RQ1:** What is the overall effect of the Wall strategy on user engagement compared to the AAX strategy? Furthermore, what is the effect if an ad blocker user chooses to whitelist?
- **RQ2:** How does the effect differ for user groups with different characteristics?
- **RQ3:** What is the longer-term effect of the Wall strategy? How would that differ from the short-term effect?

Contributions. To the best of our knowledge, this is the first study to compare the two most commonly used counter-ad-blocking strategies on the web. Our work contributes empirical evidences for understanding the different impact of these two strategies on user engagement. Our study shows that the Wall strategy has an overall negative impact on user engagements. It has no statistically significant effect on highly-engaged users because they would view the pages no matter what strategy is used. On the other hand, it has a big impact on low-engaged users, who have no loyalty to the site, especially in terms of reduced number of page views. Our long-term study finds that revisiting behavior decreases over time, but the ratio of session whitelisting increases over time because the remaining users have relatively high loyalty and high engagement. Although our work uses user behavior data from one publisher, given that the datasets and settings are common to most of publishers, we expect our findings generalizable to most online publishers.

2 EXPERIMENT DESIGN

All users in our study are ad blocker users. The experiment ran for a period of two and half months in 2018, from August 13th to October 22th. On September 13th of 2018, the Wall strategy started. Before that, all users received the AAX strategy. For each incoming user, we randomly assigned her to either control or treatment group and used cookies to track the user over time. It is noted that one inevitable limitation of using cookie is that we will not be able to identify a user if she deletes her cookie. However, identifying users beyond using cookies would violate user privacy regulations (e.g., GDPR¹). There were no other significant changes to the website during the experiment, which avoids confounding or extraneous factors brought by the publisher.

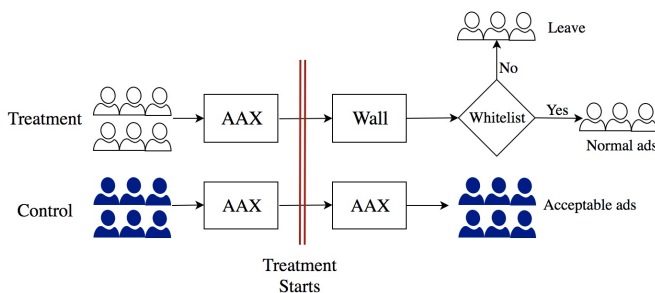


Figure 1: Illustration of our experiment design

Figure 1 illustrates the experiment design. In the post-treatment period, users in the control group were shown *acceptable ads* and

¹General Data Protection Regulation, https://ec.europa.eu/info/law/law-topic/data-protection_en

users in the treatment group were wall-blocked and saw an “Ad-block Detected” message. When facing the Wall strategy, they had to whitelist the web page or the entire web site in order to access the content. Users who did not whitelist left the site.

We utilized JavaScript tracking scripts to detect the existence of ad-blockers. We randomly selected 40,000 ad-blocker users, who were randomly placed into either the control group or the treatment group. The randomization at the user level enables us to track the treatment effect longitudinally, as discussed in Section 6.

3 DATA DESCRIPTION

The dataset contains 40K unique ad-blocker users, equally assigned into either the treatment group or the control group. The data contains a range of user engagement activities and environment measurements such as:

- overall and active browser session time
- numbers of pages in a session (i.e., pageviews)
- hits (i.e., actions) in a session, such as play a video, mouse scroll, or text selection
- date and time
- geographic location
- traffic source (e.g., search engine, social media, or by typing the URL)
- system information, e.g., Operating System (OS), browser, screen resolution

First, we analyze the characteristics of the ad-blocker users, and the statistics is consistent with the random selection of users between the two groups. The majority of users comes from US because the publisher has high influence on US audience (Figure 2). Due to the large variety in user attributes, it is difficult to have the distributions of the two groups exactly the same for every attribute. However, the difference in the country distribution of the two groups is small and similar, which confirms the random assignment on users.




Figure 2: Percentage of users in the study per country

Figure 3 shows the OS and browser distributions for the users in the dataset. The majority of the ad blocker user visits are from PC operating system, such as Windows and MacOS. The reason is that users are keener to utilize ad-blockers on PCs to avoid annoying ads, since web pages viewed on PCs have more ads and these ads can be intrusive (e.g., video ads). Another reason is that it is easier for users to install ad-blockers on PCs than on mobile devices. For browsers, we find that Chrome is mostly used by ad blocker users, since it is highly popular and offers more ad-blocker software options in its plugin-in store compared to other browsers.

Next, we analyze user behavior across the entire dataset to determine patterns and anomalies. Since we use a real-life dataset, it is inevitable to observe outliers in which the behavior data (e.g., pageviews, hits, or session dwell time) is very high. Extremely large values are probably caused by users leaving the browser open

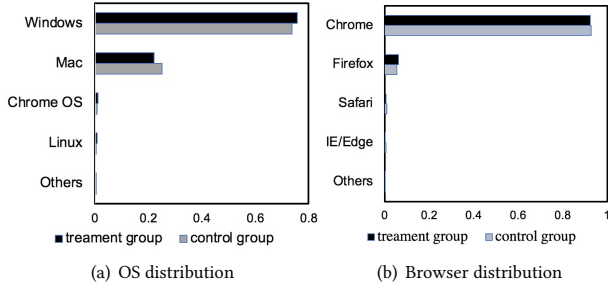


Figure 3: OS and browser distribution for users in the study

and moving away from the computers. We set a session-level filter threshold for each metric based on the observation of the data distribution to remove such outliers.

The user behavior distribution, after removing outliers, is presented in Figure 4. The distribution fitting curve line for each histogram (i.e., pageviews, hits, and dwell time) is estimated by the kernel density estimation approach. The rightmost values in x-axis in each figure is the filter threshold of outliers. The results show that these user behavior features are typically skewed to low values and have a long tail to large values. This is expected because the majority of users tend to have limited interactions with a website in a session. It is also worth mentioning that zero user engagement is recorded if an ad-blocker user choose to leave the website without whitelisting.




Figure 4: User behavior ratio (pageviews, hits, and dwell time). The Y axis is the session count normalized by the total number of sessions

4 MEASURING THE IMPACT OF THE WALL STRATEGY AND WHITELISTING

This section first presents our method for measuring the impact of the wall strategy on user engagement, and further zoom into the whitelist effect, in comparison with the AAX strategy. Then, we analyze the results of applying this method on our data. As suggested by the domain experts from our publisher collaborator, we consider three KPIs (key performance indicators) shown in Table 1 to measure user engagement.

We use the difference-in-differences (DID) estimation methodology, which is a popular method for estimating average treatment effects (ATE) while controlling for unobservables [8]. The key underlying assumption of DID is that differences between treatment

Table 1: KPI Metrics on User Engagement

KPI	Description
#pageviews	the number of pages viewed in a session
#hits	the number of actions in a session, such as play a video, scroll, selection
dwell time	time spent by a user in a session

and control groups would have a common trend in the absence of treatment. It was originally proposed as a “quasi-experimental” method to mitigate the effect of extraneous factors and selection bias. The application of DID in our randomized experiment offers robust checks on whether there are group selection bias and extraneous effects. Let us clarify that there is indeed selection bias when measuring the whitelisting effect because the whitelist behavior cannot be randomized in the experiment (i.e., it is decided by the users). Therefore, DID is suitable to measure both the wall strategy and the whitelisting effect in our experiment.

In DID, let i be an ad-blocker user and $Y_{i,j}$ be the engagement outcomes measured by a KPI metric from Table 1 that are observed in session j . $T_{i,j}$ is a binary variable regarding the randomly-assigned treatment status, where 1 indicates a user receiving the Wall treatment strategy and 0 indicates a user receiving the AAX control strategy. $t_{i,j}$ is another binary variable regarding the time period, where 1 indicates the time period after the treatment group receives the treatment (i.e., **post-treatment**) and 0 indicates the time period before the treatment group receives treatment (i.e., **pre-treatment**). The DID can be modeled as linear regression at individual level as follows.

$$Y_{i,j} = \alpha + \beta T_{i,j} + \gamma t_{i,j} + \delta(T_{i,j} * t_{i,j}) + \lambda C_{i,j} + \epsilon_{i,j} \quad (1)$$

where $\alpha, \beta, \gamma, \delta, \lambda$ are unknown parameters, $C_{i,j}$ is the extraneous factor (i.e., control variables) for user i in session j , and $\epsilon_{i,j}$ is a random unobserved “error” term. Therefore, the treatment effect is calculated as the difference in the differences of two groups as in the following equation.

$$\begin{aligned} \hat{\delta}_{DD} &= E[\overline{Y}_1^T] - E[\overline{Y}_0^T] - (E[\overline{Y}_1^C] - E[\overline{Y}_0^C]) \\ &= \alpha + \beta + \gamma + \delta + \lambda - (\alpha + \beta + \lambda) - (\alpha + \gamma + \lambda - \alpha - \lambda) \\ &= (\gamma + \delta) - \gamma \\ &= \delta \end{aligned}$$

Here, \overline{Y}_0^T and \overline{Y}_1^T are the sample averages of the behavior outcomes for the treatment group before and after treatment, respectively. \overline{Y}_0^C and \overline{Y}_1^C are the corresponding sample averages of the behavior outcomes for the control group. The parameter δ estimates whether the treatment effect is positive or negative, as well as the intensity of the treatment effect.

A linear model could estimate the treatment effect based on equation 2. We add dummy variables in the equation in order to control for the time, the day, and the weekend effect because empirical evidence suggests the user behavior is affected by these factors. hours_evening and hours_night are two dummy control variables

to indicate whether the visit happens in the evening or night time, compared to during the day time by default.

$$y = \beta_0 + \beta_1 * \text{timeperiod} + \beta_2 * \text{group} + \beta_3 * \text{hours_evening} + \beta_4 * \text{hours_night} + \beta_5 * \text{weekend} + \beta_6 * \text{timeperiod} * \text{group} + \epsilon \quad (2)$$

Since a linear model can yield negative predicted values, while our dependent variables should all be non-negative variables, this linear model does not fit well in our study. Inspired by [7], we propose to use a negative binomial (NB) regression for our study. NB regression is based on Poisson regression, which can model non-negative variables. A Poisson regression, however, still possesses one problem for our study. It assumes that the mean and the variance are the same, which may not be satisfied by the real data. In particular, in our study, the distribution of online user behavior features is typically skewed to the low values and have a long tail to the large values (see Figure 4). The variance is substantially larger than the mean, i.e., over-dispersion. To address the over-dispersion problem caused by highly skewed dependent variables, the NB regression adds a new parameter α in the model. The full NB regression is as follows:

$$P(y|X) = \frac{\Gamma(y + \alpha^{-1})}{y! \Gamma(\alpha^{-1})} \left(\frac{\alpha^{-1}}{\alpha^{-1} + \mu}\right)^{\alpha^{-1}} \left(\frac{\mu}{\alpha^{-1} + \mu}\right)^y \quad (3)$$

α is a positive parameter to represent the extent of over-dispersion auto-fitted by the data. It is solved by the maximum likelihood method. The expected value is $E(y) = \mu$, and the variance is $Var(y) = \mu[1 + \alpha\mu]$, which is a larger than $E(y)$.

Table 2: Wall strategy effect on user engagement

	Pageviews	Hits	Dwell time
Constant	0.350*** (0.006)	2.97*** (0.003)	5.793*** (0.002)
hours_evening	-0.024*** (0.004)	-0.004 (0.002)	-0.079*** (0.002)
hours_night	-0.015* (0.006)	-0.015*** (0.003)	-0.098*** (0.003)
group	0.052*** (0.006)	0.014*** (0.003)	0.041*** (0.002)
period	0.044*** (0.005)	0.014*** (0.002)	0.001 (0.002)
weekend	0.032*** (0.005)	-0.032*** (0.002)	-0.018*** (0.002)
group*period	-0.215*** (0.008)	-0.456*** (0.004)	-0.262*** (0.003)

Note. Cluster-robust standard errors in parentheses.
*p < 0.1; **p < 0.05; ***p < 0.01.

Wall Strategy Effect: We start with the session level measurement. The results of the NB regression are in Table 2. The Wall strategy has a negative effect on user engagement according to δ , i.e., the coefficient parameter β_6 in equation 2. The effect includes a statistically significant decrease of $e^{-0.215} - 1 = -19.3\%$ on the number of pageviews, $e^{-0.456} - 1 = -36.6\%$ on the number of hits, and $e^{-0.262} - 1 = -23.0\%$ on the session dwell time, compared to

the AAX strategy. The results are as expected because some users in the treatment group choose not to whitelist, and thus they are denied access to page content, consequently, resulting in less engagement for the treatment population. But when examining the coefficients ahead of *group*, we find it is statistically significant. We think it is due to the large variety of unobserved user attributes instead of problems with our randomization. Also, compared to the true treatment effect, the magnitude of group variable is small and inconsequential.

Zoom Into Users Who Whitelist: We next zoom into the users who choose to whitelist in the treatment group. We compare engagement behaviors in the whitelisted sessions of the treatment group with the control group where users have AAX sessions. As can be seen in Table 3, the whitelist behavior has a statistically significant positive effect on user engagement. The intuition is that whitelist behavior indicates that the ad-blocker user has higher interest in the intended article, and thus is more likely to spend more time and interact more with the website than the users in the control group.

Table 3: Whitelist effect on user engagement

	Pageviews	Hits	Dwell Time
Constant	0.365*** (0.005)	2.967*** (0.003)	5.799*** (0.002)
hours_evening	-0.028*** (0.003)	-0.024*** (0.002)	-0.102*** (0.001)
hours_night	-0.024*** (0.005)	-0.023*** (0.002)	-0.109*** (0.002)
group	0.052*** (0.006)	0.015*** (0.003)	0.041*** (0.002)
period	-0.026*** (0.004)	-0.365*** (0.002)	-0.352*** (0.002)
weekend	0.019*** (0.004)	-0.010*** (0.002)	-0.010*** (0.002)
group*period	0.218*** (0.010)	0.128*** (0.006)	0.158*** (0.005)

Note. Cluster-robust standard errors in parentheses.
*p < 0.1; **p < 0.05; ***p < 0.01.

5 ANALYZING THE WALL EFFECT ON USER GROUPS

This section measures the impact of the Wall strategy on users with different characteristics. User loyalty is a major characteristic impacting a user's behavior, and it is represented by her engagement with the website. Therefore, we propose to cluster users based on their engagement level, observed in




Figure 5: Choosing the number of clusters

the pre-treatment period. In other words, user characteristics are identified before the treatment starts.

The clustering features include the total number of sessions, the numbers of pageviews and hits, and the dwell time. The K-Means method is used for clustering, and the Euclidean distance is used to measure the similarity between users. Figure 5 shows the sum of Euclidean distances for each user to its nearest centroid (y-axis), with varying numbers of clusters (x-axis). The shape of the fitting curve suggests the number of cluster $K = 3$ is a good choice because it is at the elbow of the curve.

The user cluster results are described in Table 4. The user engagement increases from group 1 to group 3, with group 1 consisting of low-engaged users, group 2 consisting of medium-engaged users, and group 3 consisting of high-engaged users. It shows that the majority of users are low engaged. We use the principal component analysis method to reduce the engagement features into two dimensions and visualize the clusters in Figure 6. The figure shows the three groups of users are clustered in different areas in space. The low-engaged users are crowded together in one small area due to few sessions and limited engagement. The high-engaged users spread into a much larger area due to more frequent visits and higher variant engagement, and medium-engaged users are in the middle.

In order to measure the cluster-level impact of the Wall strategy, we use the coarsened exact matching method to match individual users in the treatment and control groups on a one-to-one basis to make sure that the samples are balanced and the users are similar to each other. The reason is that user engagement tends to exhibit the “regression-towards-the-mean” (RTM) phenomenon [9]. In order to avoid the RTM influence and sample bias, we design a matching procedure, as illustrated in Algorithm 1. We utilize the same engagement features as in the clustering. Euclidean distance is selected to measure the user similarity. If there is no similar user in the control group (i.e., exceeds the threshold), we will discard the corresponding user in the treatment group. Overall, 99% of users in the treatment groups are matched to users in the control group.

The DID method is utilized to measure the Wall strategy effect per cluster and the results are shown in Table 5. We notice that the coefficients of *group* are close to 0, and they are not statistically significant. This validates the effectiveness of our user matching procedure. It avoids the selection bias per cluster, and it indeed matches similar users in the treatment and control groups.

As shown in Table 5, the Wall strategy does not have statistically significant effect on high-engaged users. This is expected because these users are loyal, and they would view the pages no matter what strategy is used. For medium-engaged users, the Wall strategy hurts the most on their interactions with the pages (hits and dwell time), but less on pageviews. Medium-engaged users still need




Figure 6: Visualization of user clustering results by PCA; Axes are latent dimensions

Algorithm 1: User Matching

```

1 Map store=new Map(user_id, user_id);
2 Function Matcher
3   forall user i in one cluster of treatment group do
4     forall user j in control group C do
5       calculate the distance between i and j, find the
        closest user  $j^*$  for i;
6     end
7     if  $distance(i, j^*) < threshold$  then
8       add (i,  $j^*$ ) into store; delete  $j^*$  from C;
9     end
10  end
11  return store;
12 end

```

Table 4: Descriptive Analysis of User Clusters

Group	% Users	% Visits	Average session #pageviews	Average session #hits	Session dwell time
1	85.76	55.76	1.37	17.3	218.6
2	12.88	32.96	1.53	20.4	414.7
3	1.36	11.28	1.85	21.6	528.6

to access the pages, but their activities are largely weakened by the intrusiveness of the Wall strategy or the annoying ads after whitelisting (i.e., the medium-engaged users in the control AAX strategy see less annoying ads).

For low-engaged users, the Wall strategy has a large negative effect on pageviews, but relatively low impact on hits and dwell time. This is because low-engaged users have no loyalty to the website and their whitelist decisions are driven by their interest on the intended page. Thus, a significant amount of low-engaged users will refuse to whitelist and leave, resulting in a large decrease on pageviews. On the other hand, their original base of hits and dwell time are the lowest, and they cannot decrease much after the treatment. Therefore, the effect of the Wall strategy on hits and dwell time for this user group is smaller than the one for the medium-engaged user group. From a publisher’s point of view, the number of pageviews is more important than hits and dwell time because of the popular cost-per-view business charging model to advertisers. Also, since the majority of users are low-engaged users, the revenue of the publisher is expected to suffer a lot when using the Wall strategy.

6 LONG-TERM STUDY

Next, we study the effect of the Wall strategy on user engagement over time. We separate our post-treatment




Figure 7: Short-term and long-term post-treatment measurements

Table 5: Effect of The Wall Strategy per User Cluster

	cluster 1			cluster 2			cluster 3		
	Pageviews	Hits	Dwell time	Pageviews	Hits	Dwell Time	Pageviews	Hits	Dwell time
constant	0.365*** (0.011)	2.898*** (0.008)	5.490*** (0.008)	0.525*** (0.038)	3.202*** (0.031)	6.220*** (0.030)	0.671*** (0.085)	3.218*** (0.071)	6.413*** (0.069)
period	-0.030** (0.015)	-0.294*** (0.012)	-0.196*** (0.012)	-0.160*** (0.055)	-0.462*** (0.044)	-0.535*** (0.043)	-0.231* (0.123)	-0.403*** (0.100)	-0.439*** (0.098)
group	0.022 (0.015)	0.001 (0.001)	0.006 (0.012)	0.028 (0.054)	-0.005 (0.044)	0.007 (0.043)	0.046 (0.120)	-0.008 (0.102)	0.022 (0.098)
group*period	-0.258*** (0.022)	-0.162*** (0.017)	-0.034** (0.016)	-0.114* (0.078)	-0.226*** (0.062)	-0.114*** (0.044)	-0.020 (0.173)	-0.021 (0.042)	-0.079 (0.138)

Note. Cluster-robust standard errors in parentheses. *p < 0.1; **p < 0.05; ***p < 0.01.

(20 days interval in each sub-period), as illustrated in Figure 7. We refer to the first 20 days as short-term and to the next 20 days as long-term.

We consider three aspects to measure the long-term effect of the Wall strategy on user engagement: the frequency of revisiting, whitelist ratio, and user engagement per session. First, we compare the number of visits per user as well as the session-level whitelist ratio for short-term and long-term periods in the treatment group.




Figure 8: Comparison between short-term and long-term visit behavior

As shown in Figure 8(a), the average number of visits per user within a subperiod decreases from 1.96 to 1.33, which indicates that fewer revisits happen over time. To better examine the effect of the Wall strategy on revisiting, inspired by [10], we utilize the Kaplan-Meier estimator to fit the survival curves of revisits. We consider the duration gap between the first initial Wall treatment and the next visit. The results are shown in Figure 8(b), in which X axis is the timeline in terms of days, and the Y axis is the percentage of no revisits. The figure shows the black dashed line (the Wall strategy) is above the the blue solid line (the AAX strategy), indicating that the Wall strategy postpones the next revisit of the same user in the treatment group. Quantitatively, we find that the Wall strategy causes a 20.5% increase of the visit




Figure 9: Engagement in a session

duration gap. The reason is probably that the ad-blocker users feel disturbed when facing the Wall strategy, and they are less willing to come back.

We also observe that the whitelisted-session ratio increases from 48.3% to 65.5%. The reason is that, with the Wall strategy, loyal users are likelier to whitelist gradually over time. On the other hand, the ad-blocker users who refuse to whitelist previously would probably not come back again.

Finally, we measure the user engagement behavior for each session over time, where we consider only the whitelist sessions. Similar to the method presented in Section 4, we use the DID method to control the extraneous variables. As shown in Figure 9, there is a slight increase in the engagement behavior in a whitelist session over time. This indicates that users get accustomed to the Wall strategy and the annoying ads over time. It also shows the Wall strategy effect is stronger in the short-term, but its negative effect is reduced gradually over time.

7 DISCUSSIONS

In this paper, we conduct a randomized field experiment on two counter-ad-blocking strategies, benefiting from collaboration with Forbes Media, a major US media company. Our analysis shows that the Wall strategy has indeed a filtering effect on high-engaged users. They have strong loyalty to websites, and are more likely to whitelist. Therefore, we do not recommend the Wall strategy to publishers unless they have a large portion of loyal users.

If a publisher indeed wants to adopt the Wall strategy, the problem is how to convert casual users to high-engaged users, since casual users are more likely to leave forever when facing the Wall strategy. Our suggestion is to allow new users to bypass the Wall in order to strengthen their attachment to the website. The Wall can then be shown later, after noticing a significant increase in their engagement. Nevertheless, future research can expand this work by designing dynamic wall blocking strategies using machine learning methods to optimize user conversion to high-engaged users.

ACKNOWLEDGEMENT

This work is partially supported by NSF under grant No. DGE 1565478, and by the Leir Foundation. Any opinions, findings, and conclusions expressed in this material are those of the authors and do not necessarily reflect the views of the funding agencies.

REFERENCES

- [1] Jasmine Enberg. Digital ad spending 2019. <https://www.emarketer.com/content/global-digital-ad-spending-2019>, 2019.
- [2] Jessica Davies. Uh-oh: Ad blocking forecast to cost \$35 billion by 2020. <https://digiday.com/uk/uh-oh-ad-blocking-forecast-cost-35-billion-2020/>, 2019.
- [3] Shuai Zhao, Chong Wang, Achir Kalra, Leon Vaks, Cristian Borcea, and Yi Chen. Ad blocking and counter-ad blocking: Analysis of online ad blocker usage. 2017.
- [4] M Zubair Rafique, Tom Van Goethem, Wouter Joosen, Christophe Huygens, and Nick Nikiforakis. It's free for a reason: Exploring the ecosystem of free live streaming services. In *Proceedings of the 23rd Network and Distributed System Security Symposium (NDSS 2016)*, pages 1–15. Internet Society, 2016.
- [5] Enric Pujol, Oliver Hohlfeld, and Anja Feldmann. Annoyed users: Ads and ad-block usage in the wild. In *Proceedings of the 2015 Internet Measurement Conference*, pages 93–106. ACM, 2015.
- [6] Umar Iqbal, Zubair Shafiq, and Zhiyun Qian. The ad wars: retrospective measurement and analysis of anti-adblock filter lists. In *Proceedings of the 2017 Internet Measurement Conference*, pages 171–183. ACM, 2017.
- [7] Atanu R Sinha, Meghanath Macha, Pranav Maneriker, Sopan Khosla, Avani Samdariya, and Navjot Singh. Anti-ad blocking strategy: Measuring its true impact. In *Proceedings of the ADKDD'17*. ACM, 2017.
- [8] Brett Danaher, Michael D Smith, Rahul Telang, and Siwen Chen. The effect of graduated response anti-piracy laws on music sales: evidence from an event study in france. *The Journal of Industrial Economics*, 62(3):541–553, 2014.
- [9] Adrian G Barnett, Jolieke C Van Der Pols, and Annette J Dobson. Regression to the mean: what it is and how to deal with it. *International journal of epidemiology*, 34(1):215–220, 2004.
- [10] Mike Stoolmiller and James Snyder. Modeling heterogeneity in social interaction processes using multilevel survival analysis. *Psychological Methods*, 11(2):164, 2006.