

A Field Experiment to Study the Effect of Ad-Blocking and Anti-Tracking on Consumer Behavior

Cristobal Cheyre
Cornell University

Zijun Ding
Carnegie Mellon University

Li Jiang
The George Washington University

Florian Schnaub
University of Michigan

Jeremy Thomas
University of Michigan

Alessandro Acquisti
Carnegie Mellon University



for Network & Data Sciences





Economics of Digital Services (EODS)

EODS is an initiative of the University of Pennsylvania's Center for Technology, Innovation & Competition (CTIC) and Warren Center for Network & Data Sciences. Its aim is to generate independent research on the ways digital platforms access user data in support of their business models and the implications for antitrust policy and law. The initiative was funded by a major grant from the John S. and James L. Knight Foundation to support scholarly inquiry and novel approaches in the evolving digital age.

To learn more about the initiative, visit www.pennEODS.org.

Center for Technology, Innovation & Competition (CTIC)

CTIC is an interdisciplinary academic center at the University of Pennsylvania Carey Law School that bridges law and technology for academia and students. Focusing on intellectual property, competition, Internet and privacy, emerging technologies, health care, and entertainment and media law and policy, CTIC conducts foundational research that shapes the way legislators, regulatory authorities, and scholars develop policy and legal frameworks. CTIC also produces programming that explores the full range of scholarly perspectives, engages with industry experts, and educates the next generation of technology scholars, lawyers, and policymakers. www.pennCTIC.org

The Warren Center for Network & Data Sciences

The University of Pennsylvania's Warren Center fosters research and innovation in interconnected social, economic, and technological systems. Collaborating with Penn affiliates, it focuses on the role of data and algorithms to understand networked systems and how they can improve lives. The center also produces events that connect researchers, students, and entrepreneurs across the spectrum of network science.

www.warrencenter.upenn.edu

A Field Experiment to Study The Effect of Ad-Blocking and Anti-Tracking on Consumer Behavior

*Cristobal Cheyre¹, Zijun Ding², Li Jiang³,
Florian Schaub⁴, Jeremy Thomas⁴, Alessandro Acquisti^{2,5}*

Abstract

We present the design and the development of the technical infrastructure of a large-scale field experiment on the economic impact of ad-blocking and anti-tracking technologies on consumers' behavior and economic outcomes. The online advertising industry has often heralded the economic benefits of (targeted) online advertising. Its claims are juxtaposed by the privacy concerns associated with the vast number of ad-tech companies tracking and analyzing consumers' online behavior – often without consumers' awareness. We have developed a technical infrastructure to conduct a long-term field experiment to analyze the economic impact of ad-blocking and anti-tracking technologies, focusing on consumers' online behaviors (such as browsing and shopping), and their ultimate purchasing outcomes (as measured by amounts of money spent online, product prices paid, time spent on product searching, and purchase satisfaction). In this paper, we describe the rationale and motivations behind our study; the experimental design and the instrumentation infrastructure we developed for the experiments; the results of three pilots of the study; and additional settings in which our infrastructure could be deployed.

¹ Cornell University

² Carnegie Mellon University

³ The George Washington University

⁴ University of Michigan

⁵ An earlier version of this paper was presented at the Privacy Law Scholar Conference (2019). The authors thank Adhishree Abhyankar, Vikash Anand, Akshyan Anandakrishnan, Gabriella Ardiles, George Chang, Gaurav Deshpande, Ronak Dedhia, Prina Doshi, Gurinder Gill, Paritosh Gupta, Jas Kamlesh Haria, Yucheng He, Helen Huang, Akshay Hundia, Jaison Jose, Lynn Kim, Preeti Kumari, Madhukar Mohta, Sarthak Munshi, Keshav Pandey, Abhishek Parikh, Phani Krishna Pasumarthi, Yucheng Li, Cathryn Lin, Akanksha Rawat, Vinit Shah, Kyle Yang, Yaman Yu, Richie Varghese, Divya Virmani, Zhengda Wu, and Anru Xu for excellent research assistantship, and Sarah Pearman and Logan Warberg for comments. The authors also gratefully acknowledge support from the Alfred P. Sloan Foundation, Carnegie Mellon's CyLab Security and Privacy Institute, and the University of Pennsylvania Center for Technology, Innovation and Competition and the Warren Center for Network & Data Sciences. For a complete list of Acquisti's additional grants, please visit <https://www.heinz.cmu.edu/~acquisti/cv.htm>.

1. Introduction

In the last few years, online advertising has grown to capture over half of all advertising expenditures in the US. This growth has been fueled by a shift of consumers' attention from traditional to online media, and by the apparent cost and efficiency advantages of online advertising. Online advertising has become a primary monetization model for online services and mobile apps. A key innovation that has fueled the rise of online advertising is the ability to determine in real time what ad to show to individual consumers, by targeting advertising messages to the context of webpages and search queries (contextual ads), and individual user behaviors (behaviorally targeted ads). The latter – targeting based on the visitor's past behaviors and inferred characteristics and interests – is possible thanks to online tracking and behavioral profiling. According to the advertising industry, targeting is an economic “win-win” for all stakeholders within the advertising ecosystem (publishers, advertisers, consumers, and of course data intermediaries, such as ad networks), as it reduces less efficient untargeted advertising and provides a better experience for Internet visitors by reducing the exposure to irrelevant ads (AudienceScience and DM2PRO 2010; AdExchanger 2011).

The remarkable growth of online advertising and the sophistication of targeting technologies have not gone unnoticed by consumers. The large number of online ads users receive daily has led to appreciation of their usefulness, but also to frequent critiques against online advertising for being intrusive and invasive (Ur et al. 2012), and as a potential enabler of market manipulation (Calo 2014). Moreover, due to repeated scandals associated with the handling of personal data by online (advertising) firms, concerns have kept growing over the risks users face for having their data continuously mined and analyzed by companies they may not even know in ways they are often not aware of. Even when not individually targeted, the mere over-abundance or the potentially inappropriate content of online ads can cause annoyance among users (Turow et al. 2009; An 2018). This has led a significant number of Internet users to adopt privacy-enhancing technologies that limit the ability of firms to track them (e.g., anti-trackers), and/or limit the amount of online advertising they receive (e.g., ad-blockers): In a 2020 survey, 40% of US respondents report using some type of ad-blocking software (AudienceProject 2020).

The growing popularity of ad-blockers and anti-trackers among consumers hasn't been well received by online advertising companies and online publishers. Recent research has attempted to

estimate online publishers' revenue losses due to ad-blockers and concluded that "ad-blocking poses a substantial threat to the ad-supported web" (Shiller, Waldfoegel, and Ryan 2017). Regulators have responded to increasing concerns regarding online data collection by proposing, and in many cases enacting, more stringent privacy regulations, including the European Data Protection Regulation (GDPR) and the California Consumer Privacy Act (CCPA). Technology firms have also reacted in different ways. Some firms, such as Apple, have incorporated privacy protection in their products to limit tracking. Other firms, such as Google, instead of limiting data collection are attempting to increase anonymity with technologies such as the Google Topics API.

While firms in the online publishing and online advertising ecosystem have been very vocal in arguing how limiting online tracking and advertising may harm their revenues, surprisingly little effort has been devoted to understanding how online advertising in general, and the ubiquitous mining and analyzing of users' personal data in particular, may affect users' welfare – as well as how curtailing of tracking and/or advertising ultimately affects online users' online purchasing behavior and economic outcomes.

To fill this gap, we present the design of the first large-scale field experiment on the economic impact of ad-blocking and anti-tracking technologies on *consumers'* online behaviors, purchases, and satisfaction. Understanding the impact of these technologies on consumer outcomes is crucial to assess the benefits and costs associated with the provision of online advertising, and especially with the collection of personal information for the purpose of serving behaviorally targeted advertising, and thus to understand if the purported benefits of online advertising compensate for the invasion of consumers' privacy as claimed by the online advertising industry.

Our effort and focus are distinctly different from most of the (large, and growing) body of work in economics, marketing, and information systems on the advertising economy. Most of the existing literature on online advertising has focused on the effects of online advertising on advertisers, advertising networks, and online publishers. Typically, research on online advertising has measured the effectiveness of specific campaigns by certain advertisers. Researchers have sought to understand the relationship between targeting, click-through rates, and conversion rates (Manchanda, Ansari, and Gupta 1999; Farahat and Bailey 2012; Rutz, Bucklin, and Sonnier 2012; Bleier and Eisenbeiss 2015; Sahni 2015). Another topic of inquiry has been how to determine the impact and posterior effect on sales of viewed but not clicked ads (e.g., Ghose and Todri-

Adamopoulos 2016). As for current research on consumers and online ads, works in that area have been mostly limited to determining how consumers react and interact with advertising. For example, some researchers have studied whether users find targeted advertising intrusive and “creepy,” or whether over-exposure to advertising leads to “banner blindness” (Benway and Lane 1998). Finally, literature on privacy-enhancing technologies has primarily focused on determining how the use of these technologies can affect publishers’ revenues, and thus the availability and quality of content. Our research objective differs from prior literature in two important ways. Firstly, we focus on consumers and we seek to understand whether being shielded from online advertising in general, and from behaviorally targeted advertising in particular, affects consumers’ welfare. In particular, we analyze how the presence of anti-tracking or ad-blocking technologies influences users’ search and purchase behavior, and ultimately purchase outcomes (such as amounts spent, prices paid, and satisfaction). Secondly, our methodology is based on a carefully designed longitudinal field study, which has greater ecological validity than a lab experiment would, thus enabling us to identify the real-world impact of ad-blocking and anti-tracking on consumer welfare.

Using a study software specially developed for this project (based on a previous project conducted at Carnegie Mellon University, the Security Behavior Observatory, Forget et al. 2014), we will track consumer product searching and purchasing behavior longitudinally for several months. Our study software consists of a browser extension that records participants’ browsing behaviors (related to product searching and shopping), advertising exposure, and online purchases, and an email client that detects promotional emails and online purchase confirmations received by the participant. Moreover, the software allows us to remotely exert ad-blocking and anti-tracking manipulations. In this way, we assign participants into three experimental conditions: control, ad-blocking, and anti-tracking. In the control condition, we exert no intervention in the software; participants are exposed to all, including behaviorally targeted, advertising as it naturally occurs online. In the anti-tracking condition, we opt participants out of targeted advertising using the online advertising industry’s self-regulatory approach for consumer control over targeted advertising (DAA 2009; FTC 2009); participants are exposed to advertising, but those ads are not targeted to them. In the ad-blocking condition, online ads are blocked by our software, leveraging a prominent ad-blocking tool (AdBlockPlus, configured to block online ads at its maximum capacity). In this condition, participants’ exposure to ads is minimal or none (limited by the state-

of-the-art effectiveness of ad-blocking technologies). We also encouraged users to implement the manipulations on their smartphones and tablets (and we are able to detect if they did so). Our unique setting allows us to answer a number of research question that uncover the actual impact of ad-blocking and anti-tracking, such as:

- Do consumers spend more or less money to purchase products when they are not exposed to ads (vs. exposed to ads)?
- Do consumers spend more or less time searching for product information if they are not exposed to ads?
- Are consumers more or less satisfied with their online purchases across experimental conditions?
- Does it matter (for expenditures, searching time, and product satisfaction) whether the ads received are behaviorally targeted or non-targeted?
- Are participants in the ad-blocking condition more satisfied with their overall browsing experience, and do they spend more time-consuming online content?

In this paper, we first explain the motivation for our study, and its intended contribution to the literature. We next introduce our study design and software, explaining its functionality and main features, including its security and confidentiality safeguards. We then present our experiences in three pilot studies. We conclude with a description of the main study that we will complete using this infrastructure, and some suggestions of additional research questions that could be addresses with the help of our software.

2 Related Literature

The academic literature on advertising and its impact on consumer behavior and welfare is vast. That literature includes theoretical and empirical studies as well as lab and field experiments. Given the rich and detailed type of data we capture through our study, our efforts therefore relate to multiple streams of this literature, which we discuss below.

2.1 The economics of advertising

The literature on the effects of advertising on consumer behavior is relevant for our examination of the impact of ad-blockers and anti-trackers on consumers' welfare. The interest of researchers

in investigating the impact of advertising well predates the Internet age. A number of contributions in economics and marketing have explored different views on the role of advertising and its effect on consumer choice and consumer welfare. Some authors have argued that the role of advertising is chiefly informative (Nelson 1974). According to this view, advertising serves the purpose of communicating the existence of a product, or information such as its attributes, price, and where to buy it. The informative view supports the idea that advertising does provide value to consumers, as it reduces search cost. The main critique against the informative view stems from the fact that the information in the ad is provided by an interested party, and thus it may be tainted in favor of the advertiser. A contrasting view proposes that the purpose of advertising is chiefly persuasive (O'Shaughnessy and O'Shaughnessy 2004). Under this view, ads seek to manipulate consumers' preferences to shift a product or service's demand curve to the advantage of the advertised brand. Under this scenario, advertising can lead to a larger and less elastic demand for highly advertised products, which would in turn cause higher prices and lower quality. The third and more recent view of advertising proposes that its role is complementary (Hansen and Christensen 2003). According to this view, advertising affects the demand curve, but not by distorting preferences. Rather, it does so by becoming an attribute of the advertised product itself and building brand equity. This view fits well with luxury brands, which are heavily advertised to create an image for their consumers. The goal of those efforts is that consumers, by acquiring the product, are also buying into the prestige associated with the brand.

Over the years, theoretical and empirical studies have found support for each of these views in different settings. This likely means that, depending on the product, market, and strategy, advertising can be used to achieve different goals, and have different purposes (informative, persuasive, and complementary), even at the same time. In our study, we capture not just purchase outcomes, but also browsing and search behaviors (such as participants' time spent searching for products, or whether they were or were not informed about a purchased product by online ads). Thus, through our experimental conditions, we hope to be able to capture the effects of blocking different types of online advertising on consumer behavior.

2.1 Online advertising

A recent stream of literature has focused specifically on online advertising, and particular attributes and challenges arising in online contexts, such as whether ads are actually viewable and do, in fact,

get noticed by the intended consumers. One of the key characteristics of online advertising is that showing an additional impression carries an almost null marginal cost for the server hosting the ads and the site where the ads are shown. Since the early days of the Internet, this particular cost structure has naturally led to many ads being shown on webpages. This was not always well received by consumers. Several studies have documented what has been referred to as “banner blindness” (the tendency of Internet visitors to ignore banner-like information on a website; Benway and Lane 1998); and numerous studies using different methods have determined that in fact readers do not pay attention to parts of webpages that commonly display advertising (Lapa 2007; Herve et al. 2011). A related but separate issue that plagues online advertising is the viewability of ad impressions. A study by Google (2014) determined that over half of all served impressions were not viewable on users’ screens, because they were placed outside of the browser window frame. In our study, we are able to capture which ads appear within the active and viewable windows of users’ screens.

With all the viewability and attention issues that are common in online advertising, an area of research that has received substantial attention is the measurement of advertising effectiveness. One of the advantages of online advertising is that, in principle, it makes it possible to record how users interact with the ad. The most common and accepted measure of effectiveness has been the use of click-through rate and consecutive conversions into purchases (Hollis 2005). However, as advertising has evolved, these measures have proven ineffective and a number of new methodologies are being tried in order to establish, in a causal framework, the effectiveness of ad impressions (Ghose and Todri-Adamopoulos 2016). In our study, we also capture which ads users click on. Unlike prior work (which, again, focused on the effectiveness of specific campaigns), we look at the effectiveness of advertising in general, and compare user behaviors under full and reduced exposure to ads.

In addition to investigating participants’ reactions to the presence or absence of online ads (whether they are behavioral, contextual, or untargeted), another motivation for our work relates to the increasing use of personal data for behaviorally targeted advertising. In an effort to increase the effectiveness of digital advertising expenditures, advertising technology companies have developed sophisticated methods to track users and to offer advertisers the ability to target their campaigns based on users’ characteristics and behaviors. A number of studies have shown that the

use of behavioral targeting increases the likelihood of users clicking on ads. For example, Farahat and Bailey (2012) find that targeting increases click-through rate (CTR) by 67%. Yang and Ghose (2010) find that paid search results increase CTR by a factor of 2.4 and conversions by a factor of 3.5. On the flip side, when privacy policy changes impede behavioral targeting, consumer self-stated purchase intent falls (Goldfarb and Tucker 2010). Advertisers' willingness to pay for targeted advertising also reveals that such advertising must be valuable, as advertisers pay on average 2.7 times more for targeted impressions (Beales 2010). Despite these impressive figures, some authors argue that targeted advertising may not be as effective as it seems. For example, Frick and Telang (2017) use a randomized field experiment to reveal that users who would have been shown a re-targeted ad have a higher baseline probability of purchasing the advertised product, and thus are more likely to purchase the product even if they are not shown the ad. Similarly, Lambrecht and Tucker (2013) suggest that over-targeting can backfire. Whereas the studies on digital advertising presented above focus primarily on the merchants' perspective, we focus on consumer purchases and welfare. By contrasting user behavior and purchase outcomes under three experimental conditions (control with targeted ads; treatment with non-targeted ads; and second treatment without ads), we intend to isolate and capture the specific role of behaviorally targeted ads on consumer behavior. In this sense, our study is also related to the economic literature on privacy (Acquisti, Taylor, and Wagman 2016), which focuses on the economic consequences of both the sharing and the protection of personal information.

The studies presented above focus primarily on the merchants' perspective. Our research also relates to studies that examine how users react to digital advertising. The attitudes of users towards advertising, including targeted ads, are complex. Some users believe that online ads are invasive, annoying, often contain inappropriate content, and slow down the loading of webpages (An 2018). Moreover, online ads raise privacy concerns and related psychological discomfort (Turow et al. 2009; Morales 2010). Some users recognize that targeted advertising may confer some benefits to them (McDonald and Cranor 2010; Ur et al. 2012). However, most users find that the use of personal data to target advertising is invasive (Ur et al. 2012) and leads to "creepy" advertising (Moore et al. 2015). These attitudes, coupled with the thousands of online ads users receive in a month, have led to wide adoption of ad-blocking and anti-tracking software.

2.1 Ad-blocking and anti-tracking

Finally, our study is related to the emerging literature on ad-blocking and anti-tracking. The studies that have looked at the impact of ad-blocking and anti-tracking technologies have taken the perspective of online publishers and the advertising industry. The ad industry has suggested that \$15.8 billion was lost due to the use of ad-blockers in 2017 (Sullivan 2017). The Interactive Advertising Bureau (IAB) believes that ad-blockers are a threat to the industry rather than a benefit (IAB 2016). If reducing ads or targeting imposes a cost, it may be worthwhile to figure out who bears the cost – whether consumers, advertisers, or publishers. Whereas research suggests that ad-blocking reduces site visits and reduces publishers’ revenues (Shiller, Waldfoegel, and Ryan 2017), the effects of ad-blocking and anti-tracking software on consumers’ actual purchasing and search behaviors and satisfaction have not been studied comprehensively. Miroglio et al. (2018) touched on the effect of ad-blocking on user engagement in browsing using observational Firefox browser usage data. Using propensity score matching, they found that ad-blocking users spend more time browsing relative to a control group without ad-blockers. In other words, ad-blocking benefits consumers by enhancing consumer engagement rather than inducing a cost. Yet, without random assignment, it is difficult to draw a causal relationship between ad-blocking usage and engagement. Furthermore, engagement is only one part of the consumer experience and consumer welfare. A related study by Budak et al. (2016) uses observational data to investigate the extent to which display advertising leads to online retailers’ customer acquisition. The authors analyze web browsing histories from 13.6 million users over a 12-month period, and find that the vast majority of shopping sessions start from channels other than display ads: web searches, search ads, email marketing, or direct navigation.

To date, no experimental research has comprehensively studied the actual effects of the presence or absence of online ads, and their being or not being targeted, on metrics such as consumer browsing, search, and purchase behavior, and consumer satisfaction. Only recently some authors have started looking at how exposing or shielding consumers from advertising influences their browsing and/or purchasing behavior (Todri 2020; Moshary 2021), but those efforts have focused on experiments on a single platform or used observational data. We contribute by exploiting a large-scale field experiment that (1) captures detailed individual user online behavior, (2) focuses on the actual economic impact of ad-blockers and anti-trackers on end-users (consumers), (3)

examines consumer behavior across all merchants they visit, and (4) maximizes the contexts in which participants are subject to the experimental manipulation. In so doing, our paper offers a real-world estimation of the economic impact of privacy-enhancing technologies on consumers.

3. Experimental Design

3.1 Procedures

Informed consent. Informed consent and strict safeguards for handling participants' data are pivotal in a study of this type. The study was approved by the Institutional Review Board (IRB) of Carnegie Mellon University. Before participants can be enrolled in our study, we inform them about the purpose and procedures. We inform them that the study will include and require the installation on their desktop/laptop computers of software designed at CMU for collection of data about their Internet activities and shopping-related emails, and that the software may include ad-blocking or anti-tracking features. We also inform them that, depending on their experimental condition, they may be requested to install on their mobile devices third-party software with ad-blocking or anti-tracking functionalities. Participants are also informed what data will be collected: (1) web browser activities and configuration (content settings, browsing history, cookies, extensions, etc.), and details related to the use of online shopping and search engines; and (2) shopping-related emails received from online vendors (e.g., from Amazon, Target, Ebay, etc.). Additionally, we inform participants about what data we will avoid collecting, for example, bank card/credit card details, content of personal or business emails, documents on file-sharing websites, financial sites, social media sites, and content of government sites.

Screening out and entry survey. Upon recruitment, we ask participants to answer an online entry survey in the popular survey platform Qualtrics. We collect basic information about them, screen out ineligible participants, and redirect to the study software installation instructions and downloading page. Specifically, we collect participants' age, metadata about their computers (operating system, browser), as well as information about their tablets and smartphones, device usage, prior use of ad-blocking or anti-tracking technologies, frequency and preferences of online shopping, and general advertising attitude. We also include a few attention check questions to ensure the quality of data and to filter out bots.

Participants need to meet the system requirements and they cannot already be using ad-blocking or anti-tracking software at the time of recruiting. We deliberately exclude individuals who are using ad-blocking or anti-tracking to be able to use randomized assignments. Assigning someone who typically blocks ads to a condition in which they would be exposed to ads may affect their perception and satisfaction of their browsing experience in unpredictable ways and could negatively impact adherence to study protocol (e.g., not using an additional ad-blocker). For people who may have used ad-blockers before and uninstalled such software, we ask about past usage, and we control for it in the analysis. This way, we do not unduly decrease sample population and we can still check for diverging behaviors (or lack thereof) between the whole sample and the more restricted one (those who had never used ad-blocking or anti-tracking software).

Because our software is designed to work in the Windows system and for Chrome browsers, our pilot includes only participants who meet these system requirements. We are aware that, in doing so, we do not include MacOS users or people who mainly use Safari or other browsers. We discuss this limitation in the general discussion section.

Experimental conditions and software installation. We randomly assign participants to three between-subject conditions: the control, ad-blocking, and anti-tracking conditions. In the control condition, we exert no intervention in the software; participants are exposed to all, including behaviorally targeted, advertising as it naturally occurs online. In the anti-tracking condition, we opt participants out of targeted advertising using the online advertising industry's self-regulatory approach for consumer control over targeted advertising (DAA 2009; FTC 2009); participants are exposed to advertising, but those ads are not targeted to them. In the ad-blocking condition, online ads are blocked by our software, leveraging a prominent ad-blocking tool (AdBlockPlus, configured to block online ads at its maximum capacity) so that participants' exposure to ads is minimal or none (limited by the state-of-the-art effectiveness of ad-blocking technologies).

Participants install our software on their personal computers, which enables us to exert anti-tracking and ad-blocking interventions, and further to capture participants' real-time web browsing and purchasing behaviors for a period of three months. We also encourage users in ad-blocking and anti-tracking conditions to install an ad-blocker or opt out of tracking on their smartphones and tablets, and provide instructions on how to do this. After the installation is complete, we ask

participants to visit a page from their mobile devices that allows us to determine if they installed the ad-blocking or anti-tracking software in their devices.

Periodic surveys. The eligible participants, who successfully install the study software, are periodically (once a month) asked to answer a short survey about their browsing and online shopping experience, and satisfaction with the products purchased during the last month.

Exit survey. At the end of the experimental period, we ask participants to answer an exit survey about their demographic information, overall browsing and online shopping experience, perceived exposure to ads, and privacy attitudes. We also provide instructions on how to install the study software. (In the pilot studies, as the experimental period is shorter, we combine the questions for monthly and exit survey and present them once, at the end of the pilot.)

3.2 Sample size

The target sample size for our planned experiment is 1,200 participants - that is, 400 participants per condition, to be recruited from Amazon Mechanical Turk as well as from other channels, including inviting participants through online advertising as some recent online experiments (Allcott et al. 2020). The target size was decided based on the following rationale. First, 400 participants per condition is a good sample size to draw causal inferences even for a small effect size; second, as we estimated the numbers of high-quality workers (approval rate > 99%) on MTurk, and the proportion of non-users of ad-blocking and anti-tracking software on MTurk, 400 participants per condition may be an eligible number we can reach. In Section 7, we discuss the actual sample size of the three pilots we ran. They included over two dozen subjects and were recruited from Amazon Mechanical Turk.

3.3 Duration of the experiment

The planned experiment is a longitudinal study lasting 3 months. This length of time was selected on the basis of a couple of factors. First, advertising has a cumulative effect, i.e. one exposure may lead to purchases weeks later. On the flip side, we may not observe an immediate effect after installation of ad-blockers, because there is a residual effect of all previously seen advertising. Furthermore, repeated exposures may have an additive effect because repeated exposures raise awareness of the brand, etc. Therefore, the study needs to be long enough to capture the cumulative

effect. Second, 3 months is a reasonable duration to record purchases. According to a past lab experiment (Frik, Haviland, and Acquisti 2019), 57% of people purchased once or several times a month, so there is a high likelihood that we will capture multiple online purchasing instances during the 3-month period.

4. Infrastructure

Our study leverages a data collection infrastructure based on a client-server architecture. The client components include two distinct software applications that we provide to participants to install on their devices. The server components provide web and database services to maintain secure and reliable collection and storage of the collected dataset. The infrastructure described in this section reflects the evolution of our project throughout the first three pilots, with each pilot highlighting successes (our ability to collect the intended data) and issues (such as inability to collect intended data, degradation of performance in participants' computers, and so forth). The overall goal of the infrastructure is to collect the required data while minimizing the collection of undesirable data (e.g., sensitive, identifiable, or unnecessary data). We achieve this goal by taking advantage of data-driven classification and redaction systems combined with curated whitelist and blacklist filtering in our client components.

The sensitive nature of the dataset collected as part of the study requires that we follow security best practices to safeguard each participant's information. Towards this goal, we first designed a secure server environment that is maintained by our university's IT department and limits access to the research team. Within this environment, we deployed three server components: a web server for communicating with client applications, a storage server for recording collected data, and a database server for providing the dataset to the research team for analysis. We further integrated end-to-end encryption of the collected data using a system designed as part of the Security Behavior Observatory project (Forget et al. 2014). This system uses encryption to ensure that data collected from the client applications is kept confidential from the time of collection within the client application through storage on our server. Each participant's data is encrypted by the client application with a unique key. The client application uses public key encryption to share this key with our web server which then stores the key securely and separately from the collected data. In the client application, the collected data is continuously created, encrypted with the client's key, and sent to our web server. Communications with the web server use Hypertext Transfer Protocol

Secure (HTTPS) to provide additional security properties (server authentication) for the data in transit. Once the web server sends the collected data to our storage server, the data is decrypted, parsed, and inserted into an analysis database that uses role-based permissions and database-level encryption to further protect the participant data.

4.1 Browser Extension

We include a browser extension that serves the dual function of both collecting relevant data from the participant's browsing behavior and, in several of the experimental conditions, altering the browsing experience through blocking ad content or ad-tracking mechanisms. Through this browser extension, we extract details such as browsing history, browser context and settings, web searches, and various data points relevant to the content observed by the participant (e.g., ads displayed to the participant, search results across relevant e-commerce sites, etc.). The browsing history allows us to re-create the full browsing timeline of the participant and identify important events such as web searches performed, advertisements clicked, and time spent on any webpage. The browser context and settings allow us to fully understand the configuration of the browser and provide details as to the participant's experience while using the browser. This includes information related to other third-party extensions, web page permissions (e.g., "Allow website X to access your location?"), and use of tabs and multiple windows.

The majority of the data we collect is the full web content (HTML, CSS, javascript, images, etc.) gathered on targeted websites. For the purposes of our study, we target the full collection of content from both shopping websites and web searches. The collection of full content on shopping sites allows us to extract the products viewed, compared, searched for, added to shopping carts, and a variety of other behaviors. To allow for this more complex and diverse analysis, we collect the full content of these websites to our servers. We use a classification system to identify, in real time, all of the websites visited by the participants that contain shopping or e-commerce content. This classification system uses a number of heuristics extracted from both the website metadata (e.g., domain, meta tags, etc.) and the full textual content on the page. We trained this model on a set of labeled websites built primarily from publicly available domain categorizations (e.g., Alexa rankings). Additionally, we perform similar data collection on web searches (e.g., Google, Bing, etc.). To identify these searches we simply include a curated list of known, popular websites that provide search functionality in our data collection rules. To avoid cases of over-collecting

information, we include a blacklist functionality in our system. This blacklist covers a set of websites that we think do not contain shopping content and will not be needed for our analysis. This covers websites such as banking, social media, and web applications such as Google Docs.

To block advertising content we leverage the code of a popular, open-source ad-blocking extension (Adblock Plus⁶). We modified it to allow us to both control and measure the ad-blocking functionality. In our blocking experimental condition, we enable the full ad-blocking functionality of the extension. The extension itself is configured to use community-maintained filter lists that identify the advertising content based on known advertising domains and characteristics of the ad content within the page. In our other experimental conditions, we disable this functionality. In any case where the ad-blocking functionality disrupts the experience of the participant, our software provides a user interface with controls to disable the software on a specific website. We further instrument our software to record any cases of this behavior from our participants. However, in all three conditions we use the extensions advertisement identification mechanism to collect and measure the advertising content sent to the participant's browser. This provides us with details of the advertising content observed by the participant (in the control and anti-tracking conditions) or blocked by the extension (in the ad-blocking condition). Included in this advertising content is the HTML for the identified advertisement content, images, HTTP requests, and relevant metadata such as domains, HTTP headers, and cookies. Furthermore, advertising content often includes tracking components (e.g., pixels, web beacons, etc.). Our system collects these when identified. This system is limited to collecting advertising content that can be identified by the ad-blocking extension and by stringent performance requirements to avoid any degradation in the participant's browsing experience.

4.2 E-Mail Extension

We developed an email extension that, similarly to the Browser extension, allows us to collect information from the participant's email. To accommodate this requirement, we include an email client application, Mozilla Thunderbird, that has our data collection extension integrated with the software provided to our participants. This allows us to take full advantage of both the stability and functionality of an existing email client application. Additionally, it allows the participant to

⁶ The source code for Adblock Plus (<https://adblockplus.org/>) is open source which allowed us to modify it for our study purposes.

reliably and securely connect their email accounts without us learning their email account credentials. With this component, we target the collection of purchase confirmation emails. To meet this requirement, we developed a data-driven classification system to identify the targeted emails from all others. We trained this system on a corpus of purchase emails provided and labeled by the research team. Additionally, we included a variety of public email datasets to provide a large sample of negative data points.

To reduce the amount of personal information collected as part of this study, we focused on developing components that allow for targeted collection (e.g., only collecting web content on shopping sites with a high success rate), and we developed a system for redacting sensitive information that integrates into all of the client software components. This redaction system uses a set of weighted heuristics to identify sensitive information (e.g., participant's name, physical address, credit card numbers, etc.) within the client applications to strip sensitive information before any data is sent to our servers.

4.3. Implementation of experimental conditions: Mobile ad-blocking apps or tracking opt-out

Depending on the experimental conditions, participants are also asked to install ad-blocking apps or to opt out of tracking in their smartphones and tablets. This practice helps us mitigate the limitation and concern about the effect of mobile advertising on online shopping. We selected popular and highly effective ad-blocking apps in the market as the mobile ad-blocker app in our experiment. Even so, we do not claim that our intervention is a complete prevention of ad exposure because participants may still see ads offline, in apps, on TV, in magazines, or hear radio ads. Because of the limitation of our infrastructure, we are not able to capture mobile data.

4.4 Data types

4.4.1 Browser data

From the browser, we target data in three sets: data collected on all websites visited by the participant, data collected on all shopping websites visited by the participant, and data collected on all search websites visited by the participant. The method used to classify websites into these categories is explained in Section 4.2. Below we describe the type of data collected for each of these.

- All Website Visits
 - Web Navigation Details
 - URL (domain, parameters)
 - HTTP Headers (cookies)
 - Transition Type: Clicked link, typed, redirect, bookmark
 - Advertisement Content on Page
 - Web requests to known ad domains
 - Visual Ad Content (HTML, images)
 - Non-visual Ad Content
 - Tracking elements (pixels, beacons, cookies, etc.)
- All Shopping Website Visits
 - Full Page Content (HTML, CSS, images, scripts)
 - Information Extracted:
 - Products: Viewed, compared, searched, added to cart, purchased
 - Timeline: Time spent on each stage of product research and purchasing
- All Search Website Visits
 - Full Page Content (HTML, CSS, images, scripts)
 - Information Extracted:
 - Search: Queries, results, sponsored results, clicks
 - Timeline: Time spent reviewing results, paging through results
- Browser Context and Configuration
 - Settings and permissions for both the browser and for specific websites
 - Participant's use of multiple windows, tabs, third-party extensions

4.4.2 Email data

From the participant's email, we identify and collect purchase confirmation emails and promotional emails. The method used to classify emails into each of these categories is explained in Section 4.3. For each of these email types, we collect the full content of the email as well as email metadata.

- Purchase Confirmation Email
 - Email Metadata: Timestamp, sender, recipient

- Information Extracted:
 - Products: Vendor, purchase price, quantity
- Promotional Email
 - Email Metadata: Timestamp, sender, recipient
 - Information Extracted:
 - Advertisement information

4.4.3 Survey data

Survey data includes responses to entry, periodic, and exit surveys, described earlier.

5. Limitations of the Study Design

Our study design has a few limitations. First, we do not monitor offline spending. Therefore, we do not capture the effects of blocking online ads or online tracking on offline purchasing behavior. This is indeed an interesting sub-question, and we did consider using personal financial management tools (such as Mint by Intuit) to track offline spending and control for it in the actual experiment. However, we believe that this is a second-order effect for which it would be hard to disentangle and eliminate confounding factors in the current experimental setup. Additionally, collecting data on offline purchases adds a layer of privacy and security concerns, which future research design should address carefully.

Second, in the current study, we do not track mobile browsing behavior (e.g., on smartphones and tablets). However, we are likely to detect online purchases completed on other devices if their confirmation emails are captured in our email data. Moreover, we attempt to reduce exposure to ads and tracking on mobile devices in ad-blocking and anti-tracking conditions.

Third, our study may be subject to selection bias in several ways, due to exclusion criteria. We recruit people who are currently not using ad-blocking and anti-tracking technologies. The reasons for not using these technologies may vary. For example, the participants may have a low intrinsic motivation to use these technologies, be not aware of such technologies, not perceive the benefits of using them, not feel annoyed by online ads, be concerned about the usability of such tools and potential inconveniences or technical issues with online browsing, or have (or have not) other concerns. Moreover, as our study collects an extensive amount of information for a prolonged period of time, people who agree to participate in it may be less privacy-concerned than the general

population. However, we believe that none of these reasons interfere with our main dependent variables (prices paid, product searching time, and product satisfaction). Additionally, we control for participants' attitudes to privacy and advertising elicited in the exit survey.

Currently, we do not include participants who use operating systems other than Windows and browsers other than Chrome. Additionally, we exclude infrequent online shoppers (who made an online purchase longer than 12 months prior to the experiment). It is possible that MacOS users or frequent shoppers have different socio-demographic profiles, and shopping preferences and habits, than Windows users or frequent online shoppers. However, in our experiment, we include the potential proxy variables (e.g., demographic data on age, education, income, computer usage) to control for such differences.

6. Research Questions

The rich and detailed dataset we capture through the study allows us to examine different dimensions of the influence of online targeted ads on consumers' behavior and welfare. In this section we highlight a non-exhaustive list of research questions that we plan to tackle in the experiment to be completed with our infrastructure.

(1) How does consumer aggregate spending vary as a function of exposure to online ads?

It is not clear or obvious whether consumers end up buying more, the same, or fewer products online when they are shielded from online ads by an ad-blocker (vs. exposed to ads), and whether they end up, on average, increasing or decreasing their overall online spending. In broad terms, exploring this question can shed light on the role of advertising, and in particular whether advertising ultimately has a modifying effect on preferences, and thus consumer demand. For instance: does online advertising alter consumer preferences in a manner that affects aggregate consumer demand (that is, consumers end up buying and spending more, or less)? Or does, in fact, online advertising merely affect the allocation of consumer spending (that is, consumers, on average, keep consuming the same amount of products, but the choices of merchants they acquire products from are affected by the presence of ads)? If the latter, do ads end up steering consumers towards merchants and products with certain price or quality characteristics, rather than others (that is: do consumers end up being more likely to purchase more or less expensive products, and higher or lower quality products)? Do ads influence consumer purchases of utilitarian and hedonic

products differently? Previous theoretical work does not provide unequivocal support for one specific possible direction of results. As such, it is in fact possible to build reasonable hypotheses for opposite scenarios and outcomes. If participants exposed to ads spend more money than those in the ad-blocking condition, for instance, it would indicate a role of advertising that is consistent with the persuasive or the complementary roles of advertising.

(2) Do consumers spend more or less time searching for product information if ads are targeted (vs. non-targeted vs. no ads)?

If online advertising and behavioral targeting conform to the informative view of advertising, we should expect that consumers who receive ads spend less time searching for products. Moreover, consumers who receive behaviorally targeted ads should spend even less time searching for products, as advertising should support their discovery process.

(3) Are consumers more or less satisfied with their online purchases across conditions?

Whether advertising influences consumers to spend more online (or not) does not fully characterize consumer welfare. It may well be the case that exposure to advertising leads consumers to buy more products because they discover purchases that leave them with more consumer surplus than they would receive otherwise, leading to higher satisfaction. On the flip side, exposure to advertising may lead to more impulsive purchases that the consumer later regrets, leading to lower satisfaction. Thus, it is important to evaluate consumers' degree of satisfaction with their purchases across the different experimental conditions.

(4) Does it matter (for expenditures, searching time, and product satisfaction), whether the ads received are behaviorally targeted or non-targeted?

One of the arguments the advertising industry has advanced to support tracking of consumer behavior to target ads is that this practice benefits consumers. With targeted advertising, consumers only receive ads for products they are likely to be interested in. This argument fits well with the informative view of advertising. In this context, it is interesting to explore whether behavioral targeting is instead being used in a persuasive way to lead consumers to over-consume and whether opting out of ad targeting can reduce over-consumption.

(5) Are participants in the ad-blocking condition more satisfied with their overall browsing experience, and do they spend more time consuming online content?

The impact of advertising on consumers is not limited to its influence on their purchase behavior. As advertising is often perceived as a nuisance, it is worthwhile exploring whether limiting the amount of advertising that consumers receive online affects their satisfaction with the browsing experience. We can do this in two different ways. The first is simply to ask participants to rate their overall browsing experience. The second method relies on an indirect measure. If participants in the ad-block condition are more satisfied with their browsing experience, they should spend more time online than participants in the other conditions consuming online publications (news sites, blogs, videos, etc.).

(6) Which group of participants buy products from a more diverse selection of online vendors, those that receive ads, or those that are shielded from ads?

For some online sellers it may be difficult to make themselves known to potential customers. Notably Amazon holds an unrivaled position and captures over half of all online retail sales, almost 10 times more than its closest rival (Statista, 2021). One of the often touted advantages of online advertising is its efficiency, as with targeting techniques advertisers no longer have to broadcast messages to reach potential consumers. In the past, it would have been impossible for a small merchant to advertise in a major national newspaper. Nowadays, they can, as they can purchase a few impressions to reach people in specific geographic areas or with specific interests. It is not clear then which type of retailers will benefit the most from advertising. Will participants in the ad-blocking condition concentrate their purchases in the major retailers to a greater extent than other participants?

7. Pilot Studies

To determine how users would react to the study software and to detect development issues, we have conducted three study pilots. The first pilot was conducted in late 2018, the second pilot in May 2019, and the third pilot in June 2021. The pilots have been successful in demonstrating the feasibility of recruiting participants, and our ability to collect and process the data necessary to answer the research questions. In each pilot we recruited 30 participants that used the software for a two-week period. In our latest pilot we were able to extract the following data:

1. The URLs of all webpages visited by participants. We are able to classify them and determined that ~50% of pages visited correspond to shopping websites.
2. The purchases completed by participants. This includes 78 purchases from 20 different stores. 68% of purchases were concentrated in just one large online retailer, with the second largest retailer in our sample only attracting 10% of purchases.
3. When available, we can extract the reasons stated for showing ads from the “why this ad” link included in some ads. This allowed to verify our experimental conditions were functioning as intended.
4. Additionally, we collected survey data about perception regarding prevalence and relevance of online advertising, satisfaction with online purchases, subjective online browsing experiences, and question related with participants experiences while using our software.

8. Other Potential Uses for the Experimental Infrastructure

Our experimental infrastructure allows collecting rich data in a secure and privacy preserving manner, which enables the exploration of several research questions that have been difficult to assess in the past. In this manuscript we have explained the development and functionalities of our software and described the first study that will be conducted with it, and that motivated its development: The effect of ad-blocking and anti-tracking on online browsing and purchase behavior. However, we also plan future, separate investigations that would focus on additional questions related to advertising and consumer behavior. Additionally, we are open to collaborations with researchers interested in deploying our infrastructure to explore novel ideas. Below we provide some examples of research questions we intend to tackle in the future with the help of our software:

(1) How do participants in the ad-blocking condition react to websites that restrict access to visitors that employ ad-blockers?

Given the widespread adoption of ad-blocking software, some publishers have started using anti ad-blocking technologies. When a user who has installed an ad-blocking software tries to access these publishers, they are asked to deactivate the ad-blocker or whitelist the site, and are sometimes denied access to the content until they do so. Other publishers have adopted a less restrictive

approach, as they fear that denying access to content may lead users to choose to leave the site, and in the long term, lead to decreased popularity and relevance of the publisher. Our experimental setting provides an excellent opportunity to evaluate how users respond to such requests; for example, do they whitelist the sites, or abandon them?

(2) Is the attitude towards online advertising different for participants in the anti-tracking condition vs. participants in the control condition?

Studies on consumers' attitudes towards advertising have shown a complicated relation. Some consumers recognize that targeted advertising may produce some benefits, but most also find the use of personal data for targeting purposes invasive (Goldfarb and Tucker 2010). Eliciting preferences towards advertising by surveying consumers is not ideal, as what respondents state as their beliefs often doesn't match real-life behaviors. If personal experiences shape preferences, evaluating how attitudes towards advertising differ among our three conditions can provide a less biased elicitation of preferences towards behavioral targeted advertising.

(3) Measuring the degree of ad targeting.

One of the difficulties when evaluating behaviorally targeted advertising is that advertisers may choose to use users' data in different ways. On some occasions, even when advertisers have many data points about a user, they may choose not to display a targeted ad. Thus, the presence of requests to advertising networks cannot be interpreted as a signal that a user is receiving targeted ads. In a recent contribution, Carrascosa et al. (2015) devise a methodology to evaluate the degree of targeting of ad impressions. If we implement this methodology to our data on advertising, we can effectively estimate to what extent increasing degrees of targeting are actually associated with increased effectiveness of advertising.

(4) Evaluating the effectiveness of anti-tracking tools.

The ability to measure the degree of targeting of ad impressions could also enable us to evaluate the effectiveness of different anti-tracking tools and mechanisms. One of the main criticisms the "do not track" mechanism has faced is that servers may choose not to honor the request. Conversely, industry programs, such as AdChoices, have the advantage of the commitment of participants to honor the requests, but the disadvantage of limited coverage. The lack of

effectiveness of the anti-tracking measure that rely on requests has led some people to adopt active strategies to neutralize tracking attempts. For example, there are many tools that block third-party cookies or third-party scripts. Using our infrastructure we could set up a study where participants are assigned to different anti-tracking conditions and evaluate the degree of targetedness of the advertising they receive.

(5) Impact of online advertising on product exploration and purchase decisions.

Measuring the effect of advertising on purchase decisions is difficult. The methodology we propose in this study is based on comparing overall purchases between groups that are shielded from or exposed to online advertising. A difficulty that may arise with this method is that our sample size, although large, may be insufficient to obtain a significant effect. An additional way we can try to estimate the effect of advertising is to elicit, upon enrollment in the study, if participants are considering making a major purchase in the coming weeks/months. This will allow us to identify participants who are in the early discovery process of a major purchase. Then, we can use this information to determine the rate of actual purchases among study participants in different conditions who were contemplating major purchases, as well as the characteristics of the purchased product. The logic behind this idea is twofold: 1) As online purchases may be low probability events, restricting our analysis to a subsample of participants who are considering a purchase may help us identify an effect; 2) consumers who are considering a large purchase are likely to have done some research about the product they are interested in buying. One of the most common types of behaviorally targeted advertising is retargeting, which refers to the practice of advertising products the consumer has recently visited. Evaluating the effectiveness of retargeting is challenging, as being exposed to retargeted ads is not a random event. Being able to explore whether consumers who are considering a purchase and who are shielded from online advertising have lower purchase rates than similar users who are exposed to advertising is a unique opportunity to test the effectiveness of behaviorally targeted advertising.

9. Conclusions

We describe the development and functionalities of a technical infrastructure to conduct a field experiment on the effects of privacy enabling technologies on consumers. This tool will facilitate the collection of holistic data that will enable us to examine the economic impact for consumers

(e.g., prices paid for products, and the total amount spent on online purchases) of ad-blocking and anti-tracking software as a function of randomly assigned experimental conditions. It will also allow us to investigate whether the effects of our experimental treatments vary depending on consumers' individual characteristics (e.g., previous experience with the brand and vendor, position in the purchase funnel, purchase involvement, advertising attitude, etc.). Furthermore, we can examine how the usage of anti-tracking or ad-blocking software influences consumer search behaviors (e.g., time spent on information or product searching), and satisfaction; as previous research highlighted the psychological and cognitive effects online advertising can have on consumers, such as information overload (Van Zandt 2004), reduced ability to use information for decision-making (Jacoby 1977), and increased time required to reach a decision as well as reduced level of satisfaction with the choice (Jacoby 1984; Hiltz and Turoff 1985). Therefore, our field experiment is designed to make multiple novel contributions over prior work. First, our study is the first to examine the effects of ad-blocking and anti-tracking technology comprehensively: instead of focusing solely on user engagement like Miroglio et al. (2018), our study more broadly investigates the impact on search effectiveness and costs, the economic impact from the consumer perspective, and satisfaction over a number of dimensions. Second, conducting a longitudinal field experiment allows us to draw causal relationships between ad-blocking / anti-tracking software and consumer purchase behaviors. Finally, due to randomization, we overcome the methodological limitations of natural observational experiments.

Bibliography

- Acquisti, Alessandro, Curtis Taylor, and Liad Wagman. 2016. "The Economics of Privacy." *Journal of Economic Literature* 54 (2): 442–92. <https://doi.org/10.1257/jel.54.2.442>.
- AdExchanger. 2011. "If A Consumer Asked You, 'Why Is Tracking Good?', What Would You Say?" October 28, 2011. <https://adexchanger.com/online-advertising/why-is-tracking-good/>.
- Allcott, Hunt, Luca Braghieri, Sarah Eichmeyer, and Matthew Gentzkow. 2020. "The Welfare Effects of Social Media." *American Economic Review* 110 (3): 629–76. <https://doi.org/10.1257/aer.20190658>.
- An, Mimi. 2018. "Why People Block Ads (And What It Means for Marketers and Advertisers)." December 18, 2018. <https://blog.hubspot.com/news-trends/why-people-block-ads-and-what-it-means-for-marketers-and-advertisers>.
- AudienceProject. 2020. "Attitude towards Advertising and Use of Ad Blocking."
- AudienceScience and DM2PRO. 2010. "Audience Targeting: State of the Industry Study."
- Beales, Howard. 2010. "The Value of Behavioral Targeting." https://www.networkadvertising.org/pdfs/Beales_NAI_Study.pdf.
- Benway, Jan Panero, and David M. Lane. 1998. "Banner Blindness: Web Searchers Often Miss 'Obvious'

- Links.” *Itg Newsletter* 1 (3): 1–22.
- Bleier, Alexander, and Maik Eisenbeiss. 2015. “The Importance of Trust for Personalized Online Advertising.” *Journal of Retailing* 91 (3): 390–409. <https://doi.org/10.1016/j.jretai.2015.04.001>.
- Budak, Ceren, Sharad Goel, Justin Rao, and Georgios Zervas. 2016. “Understanding Emerging Threats to Online Advertising.” In , 561–78. ACM. <https://doi.org/10.1145/2940716.2940787>.
- Calo, Ryan. 2014. “Digital Market Manipulation.” *George Washington Law Review* 82 (January): 995.
- Carrascosa, Juan Miguel, Jakub Mikians, Ruben Cuevas, Vijay Erramilli, and Nikolaos Laoutaris. 2015. “I Always Feel Like Somebody’s Watching Me: Measuring Online Behavioural Advertising.” In *Proceedings of the 11th ACM Conference on Emerging Networking Experiments and Technologies*, 13:1-13:13. CoNEXT ’15. New York, NY, USA: ACM. <https://doi.org/10.1145/2716281.2836098>.
- DAA. 2009. “Self-Regulatory Principles for Online Behavioral Advertising.” http://digitaladvertisingalliance.org/sites/aboutads/files/DAA_files/seven-principles-07-01-09.pdf.
- Farahat, Ayman, and Michael C. Bailey. 2012. “How Effective Is Targeted Advertising?” In *Proceedings of the 21st International Conference on World Wide Web*, 111–20. WWW ’12. New York, NY, USA: ACM. <https://doi.org/10.1145/2187836.2187852>.
- Forget, Alain, Saranga Komanduri, Alessandro Acquisti, Nicolas Christin, Lorrie Faith Cranor, and Rahul Telang. 2014. “Building the Security Behavior Observatory: An Infrastructure for Long-Term Monitoring of Client Machines.” In *Proceedings of the 2014 Symposium and Bootcamp on the Science of Security*, 24:1-24:2. HotSoS ’14. New York, NY, USA: ACM. <https://doi.org/10.1145/2600176.2600200>.
- Frick, Thomas, and Rahul Telang. 2017. “Pay for What You Get - Incentive Misalignments in Programmatic Advertising.”
- Frik, Alisa, Amelia Haviland, and Alessandro Acquisti. 2019. “The Impact of Ad-Blocking on Consumer Behavior. A Lab Experiment.” In .
- FTC. 2009. “FTC Staff Report: Self-Regulatory Principles for Online Behavioral Advertising.” <https://www.ftc.gov/sites/default/files/documents/reports/federal-trade-commission-staff-report-self-regulatory-principles-online-behavioral-advertising/p085400behavadreport.pdf>.
- Ghose, Anindya, and Vilma Todri-Adamopoulos. 2016. “Toward a Digital Attribution Model: Measuring the Impact of Display Advertising on Online Consumer Behavior.” *MIS Quarterly* 40 (4): 889–910. <https://doi.org/10.25300/MISQ/2016/40.4.05>.
- Goldfarb, Avi, and Catherine E. Tucker. 2010. “Privacy Regulation and Online Advertising.” *Management Science* 57 (1): 57–71. <https://doi.org/10.1287/mnsc.1100.1246>.
- Google. 2014. “5 Factors of Display Viewability.” Think with Google. December 2014. <https://www.thinkwithgoogle.com/marketing-resources/data-measurement/5-factors-of-viewability/>.
- Hansen, Flemming, and Lars Bech Christensen. 2003. *Branding and Advertising*. Copenhagen Business School Press DK.
- Hervet, Guillaume, Katherine Guérard, Sébastien Tremblay, and Mohamed Saber Chtourou. 2011. “Is Banner Blindness Genuine? Eye Tracking Internet Text Advertising.” *Applied Cognitive Psychology* 25 (5): 708–16. <https://doi.org/10.1002/acp.1742>.
- Hiltz, Starr R., and Murray Turoff. 1985. “Structuring Computer-Mediated Communication Systems to Avoid Information Overload.” *Communications of the ACM* 28 (7): 680–89.
- Hollis, Nigel. 2005. “Ten Years of Learning on How Online Advertising Builds Brands.” *Journal of Advertising Research* 45 (June): 255–68. <https://doi.org/10.1017/S0021849905050270>.
- IAB. 2016. “Ad Blocking: Who Blocks Ads, Why and How to Win Them Back.”
- Jacoby, Jacob. 1977. “Information Load and Decision Quality: Some Contested Issues.” *Journal of Marketing Research* 14 (4): 569–73. <https://doi.org/10.1177/002224377701400414>.
- . 1984. “Perspectives on Information Overload.” *Journal of Consumer Research* 10 (4): 432–35.
- Lambrecht, Anja, and Catherine Tucker. 2013. “When Does Retargeting Work? Information Specificity

- in Online Advertising.” *Journal of Marketing Research* 50 (5): 561–76.
- Lapa, Chad. 2007. “Using Eye Tracking to Understand Banner Blindness and Improve Website Design.” Rochester, NY: The Rochester Institute of Technology. RIT Digital Media Library. <https://scholarworks.rit.edu/cgi/viewcontent.cgi?referer=http://scholarworks.rit.edu/&httpsredir=1&article=1696&context=theses>.
- Manchanda, Puneet, Asim Ansari, and Sunil Gupta. 1999. “The ‘Shopping Basket’: A Model for Multicategory Purchase Incidence Decisions.” *Marketing Science* 18 (2): 95–114. <https://doi.org/10.1287/mksc.18.2.95>.
- McDonald, Aleecia, and Lorrie Faith Cranor. 2010. “Beliefs and Behaviors: Internet Users’ Understanding of Behavioral Advertising.” SSRN Scholarly Paper ID 1989092. Rochester, NY: Social Science Research Network. <https://papers.ssrn.com/abstract=1989092>.
- Miroglio, Ben, David Zeber, Jofish Kaye, and Rebecca Weiss. 2018. “The Effect of Ad Blocking on User Engagement with the Web.” In *Proceedings of the 2018 World Wide Web Conference*, 813–21. WWW ’18. Republic and Canton of Geneva, Switzerland: International World Wide Web Conferences Steering Committee. <https://doi.org/10.1145/3178876.3186162>.
- Moore, Robert S., Melissa L. Moore, Kevin J. Shanahan, and Britney Mack. 2015. “Creepy Marketing: Three Dimensions of Perceived Excessive Online Privacy Violation.” *Marketing Management* 25 (1): 42–53.
- Morales, Lymari. 2010. “U.S. Internet Users Ready to Limit Online Tracking for Ads.” Gallup.Com. December 21, 2010. <https://news.gallup.com/poll/145337/Internet-Users-Ready-Limit-Online-Tracking-Ads.aspx>.
- Moshary, Sarah. 2021. “Sponsored Search in Equilibrium: Evidence from Two Experiments.” SSRN Scholarly Paper ID 3903602. Rochester, NY: Social Science Research Network. <https://doi.org/10.2139/ssrn.3903602>.
- Nelson, Phillip. 1974. “Advertising as Information.” *Journal of Political Economy* 82 (4): 729–54. <https://doi.org/10.1086/260231>.
- O’Shaughnessy, John, and Nicholas J. O’Shaughnessy. 2004. *Persuasion in Advertising*. Routledge.
- Rutz, Oliver J., Randolph E. Bucklin, and Garrett P. Sonnier. 2012. “A Latent Instrumental Variables Approach to Modeling Keyword Conversion in Paid Search Advertising.” *Journal of Marketing Research* 49 (3): 306–19. <https://doi.org/10.1509/jmr.10.0354>.
- Sahni, Navdeep S. 2015. “Effect of Temporal Spacing between Advertising Exposures: Evidence from Online Field Experiments.” *Quantitative Marketing and Economics* 13 (3): 203–47. <https://doi.org/10.1007/s11129-015-9159-9>.
- Shiller, Ben, Joel Waldfogel, and Johnny Ryan. 2017. “Will Ad Blocking Break the Internet?” Working Paper 23058. National Bureau of Economic Research. <https://doi.org/10.3386/w23058>.
- Sullivan, Laurie. 2017. “U.S. Publishers: \$15.8B Annual Revenue Lost To Ad Blocking.” October 16, 2017. <https://www.mediapost.com/publications/article/308814/us-publishers-158b-annual-revenue-lost-to-ad.html>.
- Todri, Vilma. 2020. “The Impact of Ad-Blockers on Online Consumer Behavior.” SSRN Scholarly Paper ID 3795713. Rochester, NY: Social Science Research Network. <https://papers.ssrn.com/abstract=3795713>.
- Turow, Joseph, Jennifer King, Chris Jay Hoofnagle, Amy Bleakley, and Michael Hennessy. 2009. “Americans Reject Tailored Advertising and Three Activities That Enable It.” SSRN Scholarly Paper ID 1478214. Rochester, NY: Social Science Research Network. <https://papers.ssrn.com/abstract=1478214>.
- Ur, Blase, Pedro Giovanni Leon, Lorrie Faith Cranor, Richard Shay, and Yang Wang. 2012. “Smart, Useful, Scary, Creepy: Perceptions of Online Behavioral Advertising.” In *Proceedings of the Eighth Symposium on Usable Privacy and Security*, 4:1-4:15. SOUPS ’12. New York, NY, USA: ACM. <https://doi.org/10.1145/2335356.2335362>.
- Van Zandt, Timothy. 2004. “Information Overload in a Network of Targeted Communication.” *The RAND Journal of Economics* 35 (3): 542–60. <https://doi.org/10.2307/1593707>.

Yang, Sha, and Anindya Ghose. 2010. "Analyzing the Relationship Between Organic and Sponsored Search Advertising: Positive, Negative, or Zero Interdependence?" *Marketing Science* 29 (4): 602–23. <https://doi.org/10.1287/mksc.1090.0552>.

Appendix 1

Screenshot of online installation guide (None of the steps completed yet):

Ad-Blocker Experiment

Installation steps

Please feel free to email us to seek assistance at **cmu-app-study@lists.andrew.cmu.edu**

Please follow the steps in order. Click on the step to get more details.
Note: It is important that your Chrome Browser is open and running during the whole installation process.

1. Download the software installer
2. Beginning Installation
3. Launch Thunderbird
4. Add email account
5. Close Thunderbird
6. Install the Chrome extension
7. Install on Mobile Devices

[Click here to download the installer](#)

A Thunderbird app shortcut will be placed on the desktop once the installation of the software is complete. You may remove the icon from the desktop, but do not uninstall the app. Thunderbird will continue to run in the background. Users are not required to use the Thunderbird app as their primary email client. In fact, you do not need to open Thunderbird while in the study, nor do you need to use it to send or read emails. It simply works in the background to perform research data collection.

Screenshot of installation guide (step 1 completed):

Ad-Blocker Experiment

Installation steps

Please feel free to email us to seek assistance at **cmu-app-study@lists.andrew.cmu.edu**

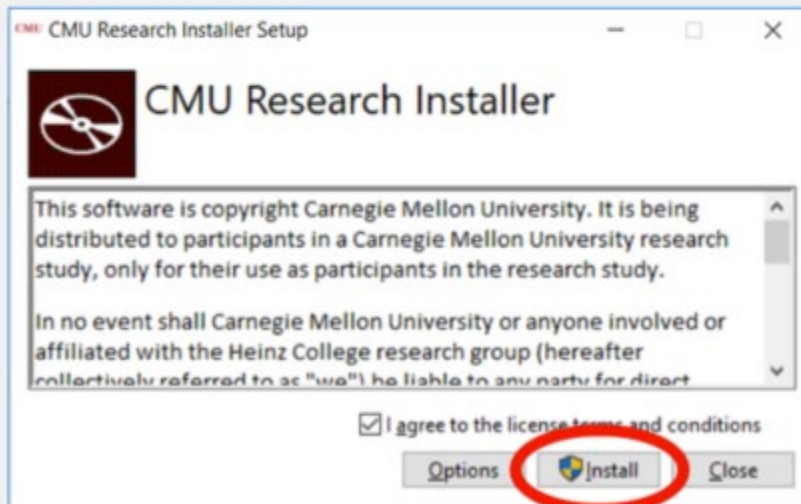
Please follow the steps in order. Click on the step to get more details.

Note: It is important that your Chrome Browser is open and running during the whole installation process.

1. Download the software installer

2. Beginning Installation

Click the download button that appears in the bottom left corner of your browser to begin installation. Follow the on-screen prompts as seen in the image below. Read the license terms and conditions and check the box to agree to the license terms and conditions. Complete the installation by clicking the "Install" button.



A prompt may appear asking you, "Do you want to allow this app from an unknown publisher to make changes to your device?" Please click on Yes to proceed with the installation.