

An Empirical Study of How People Perceive Online Behavioral Advertising

Aleecia M. McDonald and Lorrie Faith Cranor

November 10, 2009

CMU-CyLab-09-015

CyLab
Carnegie Mellon University
Pittsburgh, PA 15213

An Empirical Study of How People Perceive Online Behavioral Advertising

Aleecia M. McDonald¹ and Lorrie Faith Cranor²
Carnegie Mellon University
November 10, 2009

Abstract

We performed a series of in-depth qualitative interviews with 14 subjects who answered advertisements to participate in a university study about Internet advertising. Subjects were not informed this study had to do with behavioral advertising privacy, but raised privacy concerns on their own unprompted. We asked, “what are the best and worst things about Internet advertising?” and “what do you think about Internet advertising?” Participants held a wide range of views ranging from enthusiasm about ads that inform them of new products and discounts they would not otherwise know about, to resignation that ads are “a fact of life,” to resentment of ads that they find “insulting.” Many participants raised privacy issues in the first few minutes of discussion without any prompting about privacy. We discovered that many participants have a poor understanding of how Internet advertising works, do not understand the use of first-party cookies, let alone third-party cookies, did not realize that behavioral advertising already takes place, believe that their actions online are completely anonymous unless they are logged into a website, and believe that there are legal protections that prohibit companies from sharing information they collect online. We found that participants have substantial confusion about the results of the actions they take within their browsers, do not understand the technology they work with now, and clear cookies as much out of a notion of hygiene as for privacy. When we asked participants to read the NAI opt-out cookie description, only one understood the text. One participant expressed concern the NAI opt-out program was actually a scam to gather additional personal information. No participants had heard of opt-out cookies or flash cookies. We also found divergent views on what constitutes advertising. Industry self-regulation guidelines assume consumers can distinguish third-party widgets from first-party content, and further assume that consumers understand data flows to third-party advertisers. Instead, we find some people are not even aware of when they are being advertised to, let alone aware of what data is collected or how it is used.

1. Introduction

Behavioral advertising, also known as targeted advertising, is the practice of collecting data about an individual’s online activities for use in selecting which advertisement to display. Third party cookies are one of several of the mechanisms to enable behavioral advertising: a central advertising network with ads across thousands of websites can set and read cookies, noting every

¹ Aleecia M. McDonald is a PhD candidate in the Engineering & Public Policy Department of Carnegie Mellon University. <http://www.aleecia.com>

² Lorrie Faith Cranor is an associate professor in the School of Computer Science and in the Engineering & Public Policy Department of Carnegie Mellon University. <http://lorrie.cranor.org/>

time a given user visits any of the sites in the network.³ By correlating which sites an individual visits, advertisers can build profiles of likely characteristics and interests, and display advertisements to people most likely to purchase a given product or service. Targeted ads command a premium but also offer the potential for more cost-effective advertisements. While each advertisement costs slightly more, the specific ads go to fewer people than they would in a non-targeted campaign, and the hope is that a higher percentage of ad views will result in sales.

Behavioral advertising has received a lot of attention in the past few years. Questions about consumer's online privacy, how easily seemingly anonymous information can be re-identified,⁴ and the legality of some behavioral advertising business practices⁵ are at issue. The advertising industry⁶ and their allies⁷ favor the continuation of an "industry self-regulation" approach. The Federal Trade Commission has held numerous workshops and released guidelines for self-regulation,⁸ and there are several legislative proposals at the Federal⁹ and State¹⁰ level. In 2008, TRUSTe commissioned a report on behavioral advertising, finding 57% of respondents are "not comfortable" with browsing history-based behavioral advertising, "even when that information

³ Kristol, D., "HTTP Cookies: Standards, privacy, and politics," *ACM Transactions on Internet Technology* (TOIT) Volume 1 , Issue 2 (November 2001.) Pages 151 – 198. Available from: <http://doi.acm.org/10.1145/502152.502153>

⁴ Ohm, P. Broken Promises of Privacy: Responding to the Surprising Failure of Anonymization, 57 *UCLA L. Rev.* ____ (forthcoming 2010). Available from: <http://ssrn.com/abstract=1450006>

⁵ In particular: did NebuAd or their business partners violate wiretap and other laws? To date, the only settled case law for NebuAd pertains to jurisdiction. See Davis, W., *Online Media Daily*, "Judge Dismisses Case Against ISPs That Worked With Closed NebuAd," (October 12, 2009). Available from: http://www.mediapost.com/publications/?fa=Articles.showArticle&art_aid=115259

⁶ AAAA, ANA, BBB, DMA, and IAB. "Self-Regulatory Program for Online Behavioral Advertising," (2009). Available from: <http://www.iab.net/media/file/ven-principles-07-01-09.pdf>

⁷ Szoka, B. M. and Thierer, A. D., Targeted Online Advertising: What's the Harm And Where Are We Heading? (February 13, 2009). Progress & Freedom Foundation Progress on Point Paper, Vol. 16, No. 2, February 2009. Available at SSRN: <http://ssrn.com/abstract=1348246>

⁸ Federal Trade Commission Staff Report, "Self-Regulatory Principles For Online Behavioral Advertising: Tracking, Targeting, and Technology" (February 2009). Available from: <http://www.ftc.gov/os/2009/02/P085400behavadreport.pdf>

⁹ Boortz, A. R. "New Federal Privacy Bill in the Works: Behavioral Advertising "Beneficial," But Must Be Done "Appropriately"" *AdLaw By Request* (August 12, 2009). Available from <http://www.adlawbyrequest.com/2009/08/articles/legislation/new-federal-privacy-bill-in-the-works-behavioral-advertising-beneficial-but-must-be-done-appropriately/>

¹⁰ Arias, M. L. "Internet Law – Behavioral Advertising in the United States," *Internet Business Law Services* (June 30, 2009). Available from: http://www.ibls.com/internet_law_news_portal_view.aspx?s=latestnews&id=2237

cannot be tied to their names or any other personal information.”¹¹ Several academic scholars have also investigated this area. Anton et al. studied privacy concerns in 2002 and again in 2008, and found that “individuals have become more concerned about personalization with regard to customized browsing experiences, monitored purchasing patterns, and targeted marketing and research” in 2008.¹² Gomez et al. estimated that Google Analytics tracks at least 329,330 unique domains, and found confusion in privacy policies containing “conflicting statements that third-party sharing is not allowed but third-party tracking and affiliate sharing are.”¹³ Most recently, Turow et al. conducted a representative sample of Americans and found 66% do not want behavioral advertising, with three quarters or more rejecting common behavioral advertising practices.¹⁴ While the Turow work is valuable because it quantifies the percentage of Americans holding particular views, the standardized phone interview format meant they were unable to discover why people hold those views. In this paper, we overcome that limitation to investigate people’s mental models of online advertising.

2. Methods

We performed a series of in-depth qualitative interviews with 14 subjects who answered advertisements to participate in a university study about Internet advertising. Subjects were not informed this study had to do with behavioral advertising privacy, but raised privacy concerns on their own, unprompted. We followed a modified mental models protocol of semi-structured interviews, using standard preliminary questions for all participants while also following up individually to gather participants’ understanding of and reaction to behavioral advertising in particular.

Our study ran from September 28th through October 1, 2009 in Pittsburgh, PA. We recruited participants with a notice on the Center for Behavioral Decision Research website, which is run by Carnegie Mellon to notify the Pittsburgh community of research opportunities. Participants were compensated \$10 for an hour of their time. Of our 14 subjects, 8 were male and 6 female. Half were age 21–29 and half were age 30–59. Participants had diverse professional backgrounds including health, architecture, photography, marketing, and information technology.

¹¹ TRUSTe, “2008 Study: Consumer Attitudes About Behavioral Targeting,” (March 28, 2008). Available from: http://danskprivacynet.files.wordpress.com/2009/02/truste2008_tns_bt_study_summary1.pdf

¹² Antón, A. I., Earp, J. B., and Young, J. D. “How Internet Users’ Privacy Concerns Have Evolved Since 2002,” North Carolina State University Computer Science Technical Report # TR-2009-16 Submitted to *IEEE Security & Privacy* (July 29, 2009) Available from: http://theprivacyplace.org/blog/wp-content/uploads/2009/07/tr_2009_16.pdf

¹³ Gomez, J., Pinnick, T., and Soltani, A. “KnowPrivacy,” UC Berkeley School of Information Report 2009-037, (October 10, 2009). Available from <http://www.escholarship.org/uc/item/9ss1m46b>

¹⁴ Turow, J., King, J., Hoofnagle, C., Bleakley, A., Hennessy, M. “Americans Reject Tailored Advertising and Three Activities that Enable It,” (September 29, 2009). Available at SSRN: <http://ssrn.com/abstract=1478214>

Because our sample size is small, we do not yet know how well our results generalize to the full United States population. What we can offer are insights to *why* people hold the views they hold, and their motivations behind the actions they take online. We were able to follow up on participants' comments and engage them in dialog to elicit their views, rather than just ask fixed questions. We are still analyzing our rich qualitative data set. We describe preliminary unpublished findings below and will update the CUPS website with final publications as they become available (<http://cups.cs.cmu.edu>). We also expect to conduct a follow-up survey to determine the prevalence of the views held by our interview participants in the larger population.

3. Consumer Expectations

Much of the current self-regulation approach to online privacy is grounded in the Fair Information Principle of notice. Notice, by its nature, requires communication. As Morgan et al. wrote, "An effective communication must focus on the things that people need to know but do not already. This seemingly simple norm is violated remarkably often in risk communication."¹⁵ To follow this guidance we must find out what people already know about online privacy, what they do not, and what information they require to make decisions based on privacy policies. We need to investigate people's pre-existing mental models to see what beliefs they hold about online privacy risks, remedies, and mitigation. *Mental models* are the beliefs people hold about how a system works, interacts, or behaves. Incorrect views may form a view of the world that leads to poor evaluation of options and ultimately to bad decisions. For example, if people hold the mental model that any company with a privacy policy is bound by law not to release data to third parties, and if that is the only threat that worries them, why would people bother to read the policy? The existence of a link to a privacy policy would seem sufficient in and of itself.¹⁶ Our research contributes to understanding consumer expectations.

3.1. *Limited Knowledge of Types of Internet Advertising*

We began all interviews by asking the open-ended question "What is Internet advertising?" The answer given most immediately was "pop ups," with all but four participants mentioning pop ups. This is an intriguing response since modern browsers block pop ups by default, and indeed, participants discussed their interactions with pop up blocking. However, participants call many things "pop ups," including interstitial and hover ads. For one participant the association is so strong that she talks about all ads "popping up" on her screen, even while clearly giving examples of banner ads. For her, all ads are pop ups. Banner ads are tied with pop ups for the most prevalent response when we asked participants, "What is Internet advertising?" Banner ads were not usually mentioned first (as pop ups were) and were rarely mentioned by name. However, participants were quite capable of describing banner ads even without the vocabulary to name them. Over a third of respondents mentioned spam as a form of Internet advertising. We found it surprising that a few participants mentioned Google AdSense by name. While Google's brand is well known, we had not expected AdSense to reach beyond the advertising industry. Instead, several participants had either used AdSense to try to monetize their own blogs or knew friends who had used AdSense.

¹⁵ Morgan, M. G., Fischhoff, B., Bostrom, A., and Atman, C. J. *Risk Communication: A Mental Models Approach*. (Cambridge: Cambridge University Press, 2002).

¹⁶ Research shows that people do, in fact, believe the words "privacy policy" mean they are protected by law. See Hoofnagle, C. and King, J. "What Californians Understand About Privacy Online," (September 3, 2008) <<http://ssrn.com/abstract=1262130>> Accessed 11 September 2008.

Some participants gave characteristics of ads, rather than examples of ads. Less than half mentioned video and audio ads, usually while expressing displeasure at ads they find distracting. Participants also mentioned difficulty closing ads, and in particular complained that pop ups do not necessarily have a close button in the same place (here, again, we see confusion between true pop up ads and similar forms of advertising like interstitials.) The following concepts were mentioned by one participant each: viruses, hijacked links within articles, a constant stream of pop ups, and behavioral advertising (not mentioned my name, but described by the participant as a way to “exploit a person's history”). The other thirteen respondents did not mention or allude to behavioral advertising at all when asked to define Internet advertising. Overall, the picture that emerges includes only a general familiarity with advertising, and some user frustration with specific advertising methods and modalities.

3.1.1. Mixed Identification of Internet Advertising

Contextual search advertisements are well understood. All participants said Google is their search engine of choice. When asked if Google has ads, all participants answered correctly. Participants knew there are ads down the right hand side, that “sponsored” links frequently appear at the top of results pages, and that these links are also advertisements. They were all able to recall these details of Google’s advertisements with no prompting beyond asking if there are ads and where they are located.

We asked how advertising on Google works. All participants understood that advertisers pay Google to run ads. Participants were less clear on the mechanics of payment. Some expected Google charges for all ads displayed, and some thought Google only charges for ads when people click on them. No one thought anything that was impossible or has not occurred at one time. All told, this is a surprisingly sophisticated understanding of Google’s contextual advertising during search tasks.

In contrast, when we gave participants a printout of a webpage from the *New York Times* and asked them to identify the advertisements, answers varied widely. On the low end, participants looked at the graphics only, and discounted anything that came from the *Times* itself (e.g. home delivery and subscriptions.) At the other extreme, one participant counted every single item on the page as an advertisement, including hyperlinks in the article to other *Times* articles — and even the article itself. She reasoned the article text was likely a press release and therefore also an advertisement. Some of the differences in answers stemmed from participants skipping over parts of the page, discounting anything other than an image as a possible advertisement. Even while asking specifically about ads, a few people suffered from “ad blindness” and simply did not notice smaller ads that were in unexpected places (e.g. flush against the masthead instead of the right-hand column.) But much of the difference was definitional. While they did not phrase it this way, some participants saw advertisement as strictly a third party endeavor. Anything from the *Times* itself was therefore not an ad. Some participants also discounted all text as a potential source of advertisement.

Clearly participants do understand that text can be advertising, or they would not have been able to answer correctly about Google search ads. Why do some people then discount text as a source of advertisement on the *Times*? We have two hypotheses. First, it could be that Google is uncommonly good at communicating with their users. Ads are always in the same place, the “sponsored” label and yellow background are understood, and the right side is the place people expect to find ads. Second, it could be that people’s pre-existing mental models of print media come into play with the *Times*. People have learned with experience that ads in printed newspapers and magazines are usually graphics. To look for text ads on the *Times* people must first unlearn what they already knew, where Google was a blank slate with no direct offline

analog. Or it may be a combination of factors that people react to in different ways, which might account for why participants reacted uniformly to Google but with great variance to *Times* advertisements.

3.1.2. Inability to Distinguish Widgets

Regardless of the cause, what the *Times* results mean is that even absent any confusion over technology, participants had different mental models of advertising. We find participants have a wide range of expectations on the simple question of what is or is not an advertisement. Industry guidelines assume people can distinguish third party widgets from first party content and assume that people understand that data flows differently to third party advertisers. Therefore they treat third party widget providers as first party data collectors, subject to fewer guidelines¹⁷:

In addition, in certain situations where it is clear that the consumer is interacting with a portion of a Web site that is not an advertisement and is being operated by a different entity than the owner of the Web site, the different entity would not be a Third Party for purposes of the Principles, because the consumer would reasonably understand the nature of the direct interaction with that entity. The situation where this occurs most frequently today is where an entity through a “widget” or “video player” enables content on a Web site and it is clear that such content is not an advertisement and that portion of the Web site is provided by the other entity and not the First Party Web site. The other entity (e.g., the “widget” or “video player”) is directly interacting with the consumer and, from the consumer’s perspective, acting as a First Party. Thus, it is unnecessary to apply to these activities the Principles governing data collection and use by Third Parties with which the consumer is not directly interacting.

Instead, we find some people are not even aware of when they are being advertised to, never mind being aware of what data is collected or how it is used. It appears that self-regulatory guidelines may assume an unrealistic level of media literacy on the part of Internet users.

3.2. Misperceptions of First Party Cookies

We asked several questions regarding cookies. All participants had heard of cookies before. However, there was widespread confusion about what cookies are or how they are used. When asked, “What is a cookie?” nearly a third of participants replied immediately that they were not sure. Slightly more than a third of participants gave an answer that was at least partially correct without also saying something factually incorrect. Only one person articulated that a cookie can contain a unique identifier. We asked follow up questions of “are there ways cookies can help you?” and “are there ways cookies do not help you?”

More than half of participants confused cookies with browser history, including one participant who believed the backward and forward arrows in a web browser depend on cookies. Participants did not understand that browser history is stored independently of cookies. One participant told us cookies contain a “history of websites” visited and that if he deletes cookies, then “hyperlinks in different colors goes away, that’s what it does. It clears the navigation history.” He related how

¹⁷ AAAA, ANA, BBB, DMA, and IAB. “Self-Regulatory Program for Online Behavioral Advertising,” (2009) page 24. Available from: <http://www.iab.net/media/file/ven-principles-07-01-09.pdf>

when he was a child living at home with his mother, he lost his computer privileges because she could see where he had been based on the color of web links, which he blamed on cookies. Cookies mean “someone else can follow your previous path, and can see what you’ve read before, but that means they can get into your [computer].” More exploration revealed that in his view, cookies were only an issue on computers where he shared a single account with multiple people as he had in his mother’s home. At work, where he signed into his computer account with his own password, he believed cookies could not provide details of his browsing history because he was the only one with access to the account. Notice the confusion around password-protected accounts and privacy protections: several participants had confusion in similar areas. From follow up questions we learned that participants clear cookies and browser history at the same time, so they do not distinguish the effects. Browser user interfaces may contribute greatly to confusion as we describe later.

While participants generally did not understand what cookies store locally rather than what cookies provide a key to in a remote database, perhaps it is more important that they understand the effects of cookies rather than their mechanism. Over a third of participants said that cookies can be related to saving passwords, though during follow-up questions they did not know if their web browsers were storing the passwords or if cookies were involved. Similarly, three participants answered that cookies allow them to remain logged in to websites without retyping a password. Three participants believed cookies store their preferences for websites, including details like preferred colors and placement of site elements.

Most people believed something that was not correct about cookies. A small number of people mistakenly believed that cookies store far more than they do, such as believing cookies record all actions they take online. A couple of people thought cookies store personally identifiable or sensitive data like social security numbers, credit card numbers, and IP addresses. Three participants believe cookies are a form of malware (virus, spyware, or spam.) Several people described warnings for self-signed certificates and mistakenly believed that those warnings pertained to accepting or rejecting cookies. A few believe cookies make websites display more quickly. Again, if we consider this in the context of what cookies enable, some of these answers are more reasonable, but these answers show people do not understand what risks and benefits cookies pose.

Only three of our fourteen participants said that cookies are related to personalized advertisement. Notice three very different perspectives ranging from outright rejection to seeing benefits but finding harms outweigh them to support that is conditioned on the mistaken view that current practices are illegal.

- One participant said she has no choices about cookies, because if you “say no then you don’t get to go to the site. That’s not much of an option.” She could not think of any way cookies help her. For ways cookies do not help her, she said sites use cookies to personalize, and that “could mean more personalized advertising. It makes me feel like they expect me to be gullible.”
- A second said cookies are things “that programs use to gather information about sites [visited], functionality, and demographics for an ad.” Cookies “factor in” when advertisers “decide if it’s worth the charge they pay to advertise to that person at that time.” He said that “if asked for information [people] would say no,” and believes he has “no choices” about cookies. He said that cookies are good when “a set pattern of behaviors, sites, topics, or hobbies” can give “information on products and services that are more interesting,” but “some [cookies] are used negatively to exploit a person’s history,” and “cookies open pools of information one might prefer to stay private.” He

was concerned that IP addresses identify “a processor or individual computer,” and that the type of information collected for advertising could also be used to guess people’s passwords. He was skeptical about how well behavioral advertising works, saying “maybe you’re buying a gift” but would continue to see ads about that purchase in the future. “Patterns may be a coincidence,” and advertisers “may put you in more of a box than you are in.” Drawing an analogy to shopping offline, he said “you may be shopping in a public place but there is a privacy issue” with companies “knowing where you spend money and time.” Even with a computer collecting and storing the data, there still must be a “person manipulating and interpreting that,” and that invites “bias” because “some manipulate facts to serve a goal.”

- A third participant said advertisers use cookies to “find out as much as [advertisers] can without asking for names,” to gain an “idea of what sort of person” you are. He mentioned ISPs trying to “find ways to catalog this wealth of information,” to pair ads to an audience. He described this practice as a “smart thing” and “reasonable.” He then volunteered that he believes ISPs are constrained by law not to share information. When we asked what the law entails, he answered he was not sure and perhaps constraints were not from law but that there would be a “public uproar” and a “bad image” for any company sharing even anonymous customer data. He made the analogy to phone service where recording conversations can be illegal, and said there are “certain cultural norms and expectations” to privacy.

3.3. *Unclear on Clearing Cookies*

Nine of our 14 participants self-reported that they clear cookies. Only one of those nine said they clear cookies on their computer for privacy. Another three clear cookies on shared machines out of privacy concerns. People told us they clear cookies for the following reasons:

- To delete history
- To avoid malware (viruses, spyware)
- To reduce clutter
- To save space
- Out of habit
- For “hygiene”

Participants have a vague notion that too many cookies are bad, do not know why, and are not sure why they should delete cookies.

3.4. *Ignorance of Cookie Variants*

We asked participants if they had every heard of several technologies and if so, to define them. For anything they had not heard of we asked them to guess what the phrase might mean. We asked about:

- Session cookies
- Third party cookies
- Flash cookies

A few people had heard of session cookies or third party cookies. Those few who had heard of them were able to give mostly accurate answers. No one had heard of flash cookies before, with participants guessing things like they are cookies that “appear in a flash and are gone.”

4. Risks, Concerns, and Benefits

We asked participants for their views of “what are the best and worst things about Internet advertising?” and, “what do you think about Internet advertising?” Overall, participants held a wide range of views. Two participants responded enthusiastically not just to the idea of advertising-supported content, but to the ads themselves, which inform them of new products and discounts they would not otherwise know about. Two people were against online advertising, finding the content “insulting” and an attempt to reach “the vulnerable.” The remaining ten participants were neutral to resigned. Ads are simply “a fact of life,” multiple participants said.

4.1. *Perception 1: Internet Advertising is Necessary*

Participants named several benefits from Internet advertising such as:

- “Necessary” for the Internet to function and to enable free content
- “Good if you can control [them]”
- “Great” or “beneficial” because ads are a source of information
- “Can totally ignore” ads, unlike television or billboards
- “Short and sweet” ads
- Ads tend to be “related to [the] page”
- Ads are “more of what I want” and “not random”

With the notable exception of being able to ignore ads, this list is very similar to the benefits touted by advertisers themselves. Participants generally feel advertisements are annoying, but also see advertisements as an essential element of online life. They understand advertisements as the payment for otherwise “free” online content. A minority of participants volunteered a preference for relevant ads. However, this does not mean they understand or like data collection for behavioral advertising. When participants ask for more relevant advertisements, they almost always express a preference for contextual, not behavioral, targeting. “Relevance” means ads related to the website they are visiting, rather than related to them individually.

4.2. *Perception 2: Internet Advertising is Annoying*

The single most frequent response volunteered was that Internet advertising is “annoying,” a word used by nearly half of all participants. Participants mentioned harms from Internet advertising such as:

- “Annoying”
- “Insulting”
- “Distracting”
- “Crude graphics”
- “Clogs up Internet access” / “Slower”
- Unrelated / Off topic / Awkward mismatches
- “Opens pools of information one might prefer to stay private”
- “Not regulated”

Participants complained about being distracted by ads while trying to work or perform other primary tasks, which made pop ups and streams of ads particularly unpopular. Participants mentioned flashy colors, over-reliance on primary colors, movement, and sound as distracting elements.

4.3. Perception 3: Internet Advertising is Concerning

One straight male participant complained that he kept “getting male companion [advertisements].” He explained that this “mismatch is awkward sometimes” because it “makes you feel targeted as someone you’re not.” A second participant explicitly raised behavioral advertising and “threats to privacy.” A third participant discussed the “two way communication” of the web, and volunteered that a “privacy issue comes up” due to “creepier” advertisements “based on personal messages and keywords.” A fourth participant called for a complete “reboot” of the Internet. A fifth participant worried about “obscene” and “inappropriate” ads, particularly as she is considering starting a family. She worried about how to keep children safe online. A sixth participant raised lack of regulation. She mentioned “horror stories” of friends who signed up to get free iPods but had to submit their friends’ names first, and then never even received the promised iPods. She was most disturbed about an ad for a prescription drug to grow longer eyelashes, which was advertised just like mascara but without discussion of potential medical side as other media require. She said with TV it “seems more obvious what you can trust” but for Internet advertising a “well-designed website can be a scam.” She concluded that regulation for online advertisements is necessary and “all of it needs some kind of change.”

Four things were striking about these opening conversations. First, discussion of “relevant” ads ran the gamut from support to deep concerns about privacy. As the interviews continued the diversity of opinions became even more marked and we learned how little people understand of current practices. Second, participants were largely pragmatic about advertising. Even when they had scathing remarks about bad experiences, on the whole they understand and accept the model that advertising supports content. Their frustrations are generally not due to the existence of advertising, but rather to the specific practices. Third, participants expressed real anger and frustration about advertising tactics they see, even when they do not understand the data being amassed about their online activities that they do not see. Finally, all of the issues raised above were volunteered, not prompted, after very open-ended questions at the start of the interviews. Participants’ concerns about advertising practices, content, lack of regulation, behavioral targeting, and privacy surfaced in the first few minutes of discussion. These issues are central to how participants perceive online advertising.

5. Mechanisms of Consumer Privacy Protection

From what we have observed to date, it appears behavioral advertising violates consumer expectations and is understood as a source of privacy harm. While we do not attempt a full analysis of possible policy responses here, we note several things. First and foremost, consumers cannot protect themselves from risks they do not understand. One younger participant said in frustration that she did not learn about how to protect her online privacy in school, she was just taught typing. We believe there is a serious need not just for improved notice of practices, but for the education requisite to understand disclosures. Most non-regulatory approaches require consumers to understand tradeoffs and know enough to take whatever actions will enable their privacy preferences. At the current moment that seems unrealistic, but the outlook could improve in the future. Below, we offer some preliminary findings about the industry self-regulation mechanism of NAI opt out cookies, and some observations about web browsers’ roles in cookie management.

5.1. Consumers Do Not Understand Opt Out Cookies

None of our fourteen participants had heard of cookie-based methods for opting out of tracking cookies, including TACO¹⁸ and NAI opt out cookies.¹⁹ At the end of the protocol, we showed four participants a text description of NAI opt out cookies from the NAI opt out website.²⁰

All four participants understood they would continue to see at least some online advertisements. However, there is substantial confusion about what the NAI opt out does. The text does not explain that companies may choose to continue all data collection and profiling, and that in some cases the only thing that changes is the type of ads displayed.²¹ One participant understood this but the other three did not.

- The first participant believed the NAI opt out “sets your computer or ethernet so information doesn’t get sent.” She still expected to see ads, but now the ads would be “random.” She said it might “sound old fashioned” but in a choice between “convenience and privacy, I’m going to pick privacy.” She was afraid that by clicking to opt out “all these people get your information” and therefore “this could be a phishing expedition.”
- A second participant began his comments by saying “Where do I click? I want this!” He believed the NAI opt out to be an “opt out tool so users opt out of being tracked.” He thought “the ads are still there, they just get no data.”
- A third participant believed the purpose of the NAI opt out text was “reducing the amount of online advertising you receive.” He understood data collection was also involved, but not how, just “some sort of control over what companies use that information.” He would choose to opt out of some companies, “ones I thought the information they would seek would be too personal to share with a group.”
- Our final participant did understand the NAI text better. At first he said by way of example that it means if you use GMail the opt out cookie means “stop reading my email and tailoring ads.” However, he later clarified “What you search is Google property, it’s theirs. They’re going to profile you but not show you that they are.”

¹⁸ Targeted Advertising Cookie Opt-Out (TACO) is a plugin for the Firefox browser that stores persistent opt out cookies, available from: <https://addons.mozilla.org/en-US/firefox/addon/11073>

¹⁹ The Network Advertising Initiative (NAI) offers non-persistent opt out cookies for all browsers, available from: http://www.networkadvertising.org/managing/opt_out.asp

²⁰ Our study used printed materials so we did not test the NAI video, which may communicate more clearly. The degree to which the video’s clarity is important hinges on how visitors engage the site. The NAI may be able to provide information about what percentage of website visitors watch the video to completion, but four calls to NAI asking to speak about their opt out cookies went unreturned.

²¹ Anderson, Stacy. “House Subcommittees Hold Joint Hearing On Behavioral Advertising,” *Security, Privacy and the Law* (July 2009.) Available from: <http://www.securityprivacyandthelaw.com/2009/07/articles/recent-legislation-1/house-subcommittees-hold-joint-hearing-on-behavioral-advertising/> Original testimony available from: <http://www.youtube.com/watch?v=-Wk1p2qdbmw>

5.2. Web Browsers May Promote Consumer Confusion

While we did not study web browser interactions specifically, participants explained ways they use their web browsers to interact with cookies. In the section “Misperceptions of First Party Cookies,” we documented consumer confusion between cookies and browser history. One component of this confusion is temporal: participants reported they delete cookies and clear history at the same time, which leads them to misattribute properties of browser history to cookies. The reason participants clear cookies and history together likely stems from the way they are swirled together in the user interfaces of web browsers. For example, as shown in Figure 1, Firefox presents choices about cookies, history, and bookmarks on the same tab. There is no visual hint that these three topics are distinct. To the contrary, cookies are in the middle of options for history, which serves to convey history and cookies are related. Moreover, Firefox does not expose any cookie options unless users know to change a setting from “Remember history” to “Use custom settings for history.” Anyone looking through preference tabs for cookies will not find them in the default configuration.

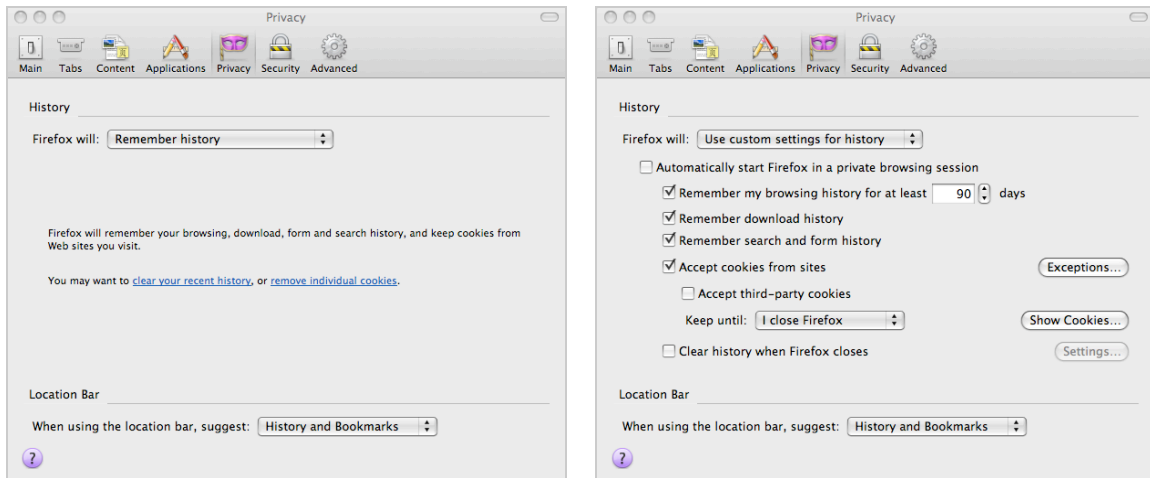


Figure 1: Firefox’s Macintosh user interface mixes cookies, history, and bookmarks

Mixing cookies, history, and bookmarks is not the only area where web browsers interfaces contribute to lack of understanding of Internet privacy issues. As another example, web browsers give no notice of or access to Flash cookies. Even technologically sophisticated users are unfamiliar with Flash cookies and how they can “respawn” deleted cookies.²² As another example, Internet Explorer implements P3P support, but information about P3P is buried in the user interface, so much so that a study of online trust markers found none of the participants were familiar with the P3P icon.²³ The Internet Explorer P3P implementation works well in that it does not require user intervention. Based on default settings, users do not accept any third party cookie that does not have an associated P3P policy with an opt out. In this way browsers can provide an enforcement mechanism that may be stronger and faster to take effect than any regulations.

²² Soltani, A., Canty, S., Mayo, Q., Thomas, F., and Hoofnagle, C. “Flash Cookies and Privacy,” (August 10, 2009). Available at SSRN: <http://ssrn.com/abstract=1446862>

²³ Jenson, C., Potts, C., and Jenson, C. “Privacy practices of Internet users: Self-reports versus observed behavior,” *International Journal of Human-Computer Studies*, Volume 63, Issues 1-2, (July 2005) Pages 203-227.

However, as the early history of cookies themselves and the current example of Flash cookies and P3P amply demonstrate, just because browsers *can* provide user control does not mean they *will*. Cookies were introduced fifteen years ago, yet we observed most people do not understand even first party cookies. Browsers can be an important part of user empowerment but as the lack of cookie knowledge illustrates, informed privacy decision making is not something the free market is solving.

6. Observations

Netscape introduced cookies fifteen years ago, yet today approximately two thirds of our respondents were unable to explain what cookies do without volunteering incorrect information. Half of participants confuse cookies with browser history, and that confusion may be promoted by web browsers' user interfaces. Participants had no understanding of flash cookies or that flash cookies can respawn deleted cookies across domains.

None of our participants were familiar with NAI opt out cookies. Participants who incorrectly believed NAI opt outs mean they are no longer subject to profiling were very enthusiastic supporters. Based on NAI's text, participants had a difficult time understanding what the NAI opt out cookies do.

Consumers have a very clear understanding of when and where Google search displays advertisements. However, consumers do not understand which parts of the *New York Times* website are advertisements. They lack the knowledge to distinguish widgets from first party content. Consequently, it is overly optimistic to believe consumers know their data flows to widget providers as a first party.

One of the questions posed by the advertising industry is "where's the harm" in behavioral advertising, with a suggestion that a formal benefit cost analysis should occur before regulation. This question seems to ignore privacy loss as a distinct harm. In contrast, our participants spoke frequently about their privacy concerns. One technically savvy participant even described withdrawing from online life as a result of privacy concerns. We refer readers to our prior research where we estimated the high value of people's time if they were to actually read privacy policies. We found "it appears the balance between the costs borne by Internet users versus the benefits of targeted ads for industry is out of kilter," and "suggest that any such cost-benefit analysis should include the value of time for reading privacy policies."²⁴

²⁴ McDonald, A. M. and Cranor, L. F. "The Cost of Reading Privacy Policies," *I/S: A Journal of Law and Policy for the Information Society* (2008). Forthcoming from <http://www.is-journal.org/> and available from <http://cups.cs.cmu.edu>