

# User Perceptions of Online Advertising

## Yale ISP Conference, March 25-26, 2011

Aleecia M. McDonald

### 1 Introduction

This paper presents empirical data on American adult Internet users' knowledge about and perceptions of Internet advertising. This paper contains a subset of results from several research studies including in-depth interviews, an online survey of participants' views of online advertising and their ability to make privacy decisions, and preliminary results from a pilot study of user expectations for Do Not Track. This paper condenses over 100 pages of research down to NN pages. Interested readers may consult the original papers for related work, research methods, and additional results.

### 2 Online advertising

In the Fall of 2009, we performed a series of in-depth qualitative interviews with 14 subjects who answered advertisements to participate in a university study about Internet advertising.<sup>1</sup> 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.

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. Banner ads are tied with pop ups for the most prevalent response. 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.

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. 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 by name, but described). 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.

Four things were striking about these opening conversations. First, discussion of "relevant" ads ran the gamut from support to deep concerns about privacy. 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 specific practices. Third, participants expressed anger and frustration about advertising tactics they see, even when they do not understand the data being amassed about their

---

<sup>1</sup>This section is taken from a technical report co-authored with Lorrie Faith Cranor, [2].

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 voiced concerns about advertising practices, behavioral targeting, and privacy in the first few minutes of discussion. Privacy is central to how participants perceive online advertising.

### 3 Behavioral advertising

When we described current advertising practices in our lab study, participants told us they did not believe such things happened. One participant said behavioral advertising sounded like something her “paranoid” friend would dream up, but not something that would ever occur in real life. We asked our online participants about two pervasive current practices described as hypotheticals.<sup>2</sup> First we asked about behavioral ads with the following description:

*Imagine you visit the New York Times website. One of the ads is for Continental airlines. That ad does not come to you directly from the airline. Instead, there is an ad company that determines what ad to show to you, personally, based on the history of prior websites you have visited. Your friends might see different ads if they visited the New York Times.*

We asked about ads based on content in hosted email, which describes systems in use like Gmail:

*Imagine you are online and your email provider displays ads to you. The ads are based on what you write in email you send, as well as email you receive.*

Table 1: Perceived likelihood of practices occurring

Response	Behavioral Ads	Email Ads
This happens a lot right now	51%	25%
This happens a little right now	35%	14%
This does not happen now but could happen in the future	11%	28%
This will never happen because it is not allowed by law	1%	16%
This will never happen because there would be consumer backlash against companies that engaged in this practice	1%	13%
Other	1%	5%

As shown in Table 1, participants seem to have a high degree of understanding that behavioral advertising happens, with only 13% of respondents casting doubt that current practices occur. Yet only 40% believe advertising based on email content is happening today, and 29% believe this common practice will never occur.

Recall 41% of our participants reported that they check gmail accounts. We found statistically significant differences between gmail users and non-gmail users for the email scenario ( $\chi^2=20.1$ , d.f.=5,  $p<.001$ ). Gmail users were far more aware that this practice occurs today, with 51% of gmail users saying it happens either

<sup>2</sup>This online study is described in a paper co-authored with Lorrie Faith Cranor presented at the 2010 Research Conference on Communication, Information and Internet Policy (TPRC) [4]. The TPRC paper substantially expands upon a paper co-authored with Lorrie Faith Cranor presented at the 2010 Workshop on Privacy in the Electronic Society (WPES) [3]. The 314 participants in the study were recruited from Mechanical Turk.

a lot or a little now, in contrast to 30% of non-gmail users. It is encouraging to see gmail users are more likely to understand the practices gmail follows, but surprising that half of gmail users do not understand how gmail works. This suggests a lack of informed consent for gmail’s business model and a potential for surprise. Gmail users were half as likely to think ads based on email would never happen due to backlash (8% v. 16%) but equally likely to think ads based on email are barred by law (15% v. 16%).

For both scenarios we asked, “How would you feel about this practice?” (Participants were able to select more than one answer.) As shown in Table 2, the most popular answer is that 46% of participants find behavioral advertising “creepy,” but a small group of 18% welcome targeted advertisements. Responses on how people feel about advertising based on email are markedly more negative, with 62% saying email should be private and that they find ads based on email creepy. Only 4% of respondents saw email-based advertising as a benefit, and only 9% supported the trade off of data and advertising for free services. This matches what we heard in interviews: people understand ads support free content, but do not believe data are part of the deal.

Table 2: Attitudes toward current practices

Response	Behavioral Ads	Email Ads
No one should use data from email because it is private like postal mail	N/A	62%
It’s creepy to have advertisements based on my emails	N/A	62%
It’s creepy to have advertisements based on sites I’ve visited	46%	N/A
Wouldn’t even notice the advertisements, just ignore them	38%	18%
No one should use data from Internet history	30%	28%
Glad to have relevant advertisements about things I am interested in instead of random advertisements	18%	4%
It’s ok as long as the email service is free	N/A	9%
Other	3%	5%

We again contrasted our gmail users to non-gmail users for the email scenario. We did not find statistically significant differences between gmail users and non-gmail users for the email scenario ( $\chi^2=9.96$ , d.f.=5,  $p=.076$ ). This means gmail users are as likely as non-gmail users to find the practices predominately creepy, and believe their email should be private like postal mail. Those who choose to use gmail are not doing so out of lack of concern for privacy in comparison to non-gmail users.

## 4 Opt-out cookies

The Network Advertising Initiative (NAI) offers non-persistent opt out cookies for all browsers.<sup>3</sup> The NAI opt out cookies are an industry self regulation approach to providing privacy choices to users. During interviews<sup>4</sup> we learned that not only did our participants fail to understand the NAI opt-out page, several of them thought it was a scam. In our online study we learned that is not a widely held view, but neither is the correct explanation for the page’s function. We showed a screenshot (see Figure 1) and asked “Based on the image above, if you visited this web site, what would you think it is?”

- 34% answered “A website that lets you tell companies not to collect data about you.” There are some companies for which this is the case. However, some NAI members like Yahoo! continue to collect data exactly as before; they just do not tailor ads to reflect that data.

<sup>3</sup>See [http://www.networkadvertising.org/managing/opt\\_out.asp](http://www.networkadvertising.org/managing/opt_out.asp).

<sup>4</sup>See [2].

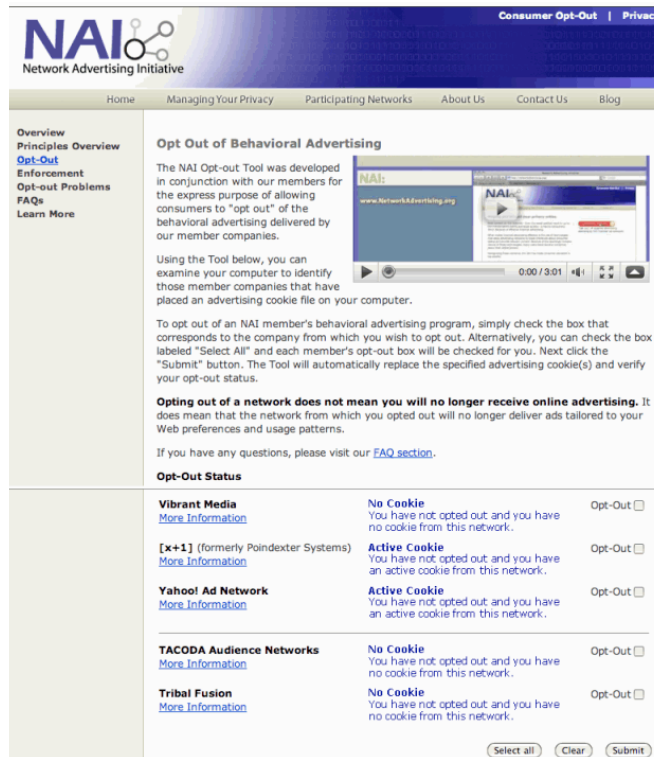


Figure 1: Screenshot of the NAI Opt Out page

- 25% answered "A website that lets you tell companies you do not want to see ads from them, but you will still see as many ads overall." This is incorrect because companies continue to serve ads, just not targeted ads. The ad source is unchanged.
- 18% answered "A website that lets you see fewer online ads." This is both wrong and prominently disclaimed in the NAI text.
- 11% answered "A website that allows companies to profile you, but not show you ads based on those profiles." **Correct answer.**
- 6% answered "A scam website to collect your private information."
- 5% answered "A scam website to find out which websites you have visited."

These results paint a bleak picture of users' abilities to make sense of opt-out cookies. Our largest group of respondents misunderstood the NAI text and believed their information would not be collected if they opted out. NAI visitors may think they are selecting which ads they see, rather than targeted v. random ads from the same sources, and make choices that do not reflect their actual preferences. People think the site is a scam at the same rate they understand what it is for. NAI opt-out cookies may not currently be working well as instruments of self-regulation.

## 5 Do Not Track

This section is an early look at preliminary results from a pilot study completed on March 11, 2011. The 41 participants in the study were recruited from Mechanical Turk. Please do not cite to this early work — a revised survey is in progress now, with questions refined based on ambiguities in responses. The overall picture that emerges should be similar to the pilot results, but the finer details will change.

One of the central challenges to implementing a Do Not Track mechanism comes in defining what, exactly, tracking means. Google, Microsoft, and Mozilla have three very different approaches to implementing the browser side of Do Not Track, but even running code does not clear up the matter.

- Google built their Chrome solution around NAI opt-out cookies. As mentioned above, different NAI members do different things when they read an opt-out cookie, ranging from collecting exactly the same data but not showing ads targeted based on that data to not collecting data at all.
- The Mozilla Firefox browser can optionally send a header expressing a user’s preference not to be tracked. It is likely that different parties honoring the Do-not-track header will do so in different ways, as happened with opt-out cookies.
- Microsoft’s approach for Internet Explorer allows groups to publish blacklists for companies engaged in tracking. Each group that publishes a blacklist has to decide what they consider to be tracking, and which companies belong on their blacklist.

As just one example of how complicated defining “tracking” has become, Figure 2 contains a list of data uses that the Center for Democracy and Technology consider to be tracking, or not [1]. To understand this chart, users would need to understand at least the difference between first- and third-party websites, what behavioral advertising is, the types of data collected for behavioral advertising, the difference between identifiable and non-identifiable data, reporting, and analytics.

<b>Tracking</b>	<b>Not tracking</b>
Third-party online behavioral advertising	Third-party ad and content delivery
Third-party behavioral data collection for first party uses	Third-party reporting
Third-party behavioral data collection for other uses	Third-party analytics
Behavioral data collected by first parties and transferred to third parties in identifiable form	Third-party contextual advertising
	First-party data collection and use
	Federated identity transaction data
	Data collection required by law and for legitimate fraud prevention purposes

Figure 2: The Center for Democracy and Technology list of examples of data used for tracking and not tracking, illustrating their definition of tracking

Preliminary results suggest that users do not share nearly so nuanced view of tracking. We opened the study by showing a screenshot of a button simply labeled “Do not track,” and asked participants to

imagine this was a new button in their current web browser. We asked them what they would expect it to do. Answers ranging from very general answers, like “Do not track my Internet activities,” and “I would be anonymous,” to specifics like “Turn off cookies,” and “It would make it impossible for anyone to track my I.P address,” to a few skeptical answers like “the opposite - will track” and “send a virus to my computer.” We proceeded through the rest of the study without ever defining tracking or Do Not Track for participants, relying on their own expectations.

We asked participants “What data do you think websites could collect before you clicked a Do Not Track button? (Please check all that apply)” followed by asking the same question for after clicking a Do Not Track button. Unfortunately, the “check all” format did not work perfectly in practice in our pilot, so we have restructured this question for our final study. Even with that flaw, the results in Table 3 provide early insight into user expectations for Do Not Track.

Table 3: Percentage of participants who selected that a website can collect a type of data before, and after, clicking Do Not Track

<b>Data collection type</b>	<b>Before DNT</b>	<b>After DNT</b>	<b>After-Before Difference</b>
No data at all	10%	61%	-51
Anything so long as the information was anonymized and combined with other people’s so you could not be specifically identified	88%	5%	83
The names of websites you have visited on the Internet	78%	7%	71
Anything so long as it did not identify you by name	71%	5%	66
Which web pages you saw on the site you are currently visiting	76%	10%	66
Which ads you have seen	76%	12%	64
Which web browser you use (Internet Explorer, Firefox, etc.)	71%	24%	47
Your IP address (the address for your computer, which may remain the same for years at a time, or may change each time you connect to the Internet, depending on your set up)	59%	29%	30
Your search terms (for example, what you type into Google, Yahoo!, or Bing)	0%	12%	-12
Which ads you have clicked on	0%	12%	-12
Any data they collect now	N/A	7%	N/A

To highlight a few of the more interesting results:

- 61% of respondents expected that if they clicked a Do Not Track button, the site they visited would collect no data at all. None of the current proposals for Do Not Track contemplate limiting data collection to nothing for first party use, yet that is what many users expect.
- Respondents did not expect Do Not Track to work by collecting the same information, but aggregating it with other user’s data, with only 5% selecting that as a possibility for Do Not Track. Similarly, participants did not expect Do Not Track to work by collecting the same information, but anonymize it, with only 7% selecting that as a possibility for Do Not Track.
- Only 7% of respondents expected websites could collect the same data before and after users click Do Not Track.

We expect to see different results for search terms and ads clicked on, which currently show an increase for the percentage of people who expect websites could collect such data after users click a Do Not Track button. We believe this is an error in study design, with participants checking choices like “Anything...” ending in aggregate data or anonymized data and thinking they did not need to check all additional choices.

We expect the results to hold in general terms: users expect Do Not Track to work in ways that differ from current Do Not Track implementations.

## 6 Discussion

Consumers cannot protect themselves from risks they do not understand. We find a gap between the knowledge users currently have and the knowledge they would need to possess in order to make effective decisions about their online privacy. This has implications for public policy, commerce, and technologists. Most non-regulatory approaches require consumers to understand tradeoffs and to 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.

In general, users do not appear to want targeted advertisement at this time, and do not find value in it. However, a small but vocal subset of users are genuinely eager for relevant ads. They are matched by a subset of users vehemently against the practices that enable targeted ads. In the middle, the majority attempt to ignore ads and see no benefit to giving data to advertisers. Ideally, users could choose for themselves but at present they lack the knowledge to be able to make informed decisions, and lack effective mechanisms to enact those choices.

We found people generally unwilling to pay for privacy, not because they do not value it, but because they believe it is wrong to pay. Paying to keep data private was termed “extortion” by some participants. 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. 40% of participants in our online study agree or strongly agree they would watch what they do online more carefully if advertisers were collecting data, which suggests advertising may cause a chilling effect — and suggests participants do not realize advertisers already are collecting data. In our lab study, one technically-savvy participant described withdrawing from online life as a result of privacy concerns.

The NAI is as a major player in behavioral advertising but their opt-out cookie page is very confusing, with only 11% understanding what it is for. With their leadership role in self-regulation, the NAI may not be supporting Internet users’ ability to avail themselves of self-help options. Because Google’s Chrome browser uses NAI opt-out cookies in their implementation of Do Not Track, Chrome may also confuse users.

Preliminary work on Do Not Track suggests users expect data collection to halt. 61% of participants expect no data collection at all, and users do not expect Do Not Track to work by de-identifying or aggregating data. In contrast to the nuance involved in many definitions of tracking, it appears users expect a Do Not Track feature to provide a straight-forward halt to data collection online.

With lack of understanding of and a lack of interest in tailored content, unless industry moves rapidly towards an effective self-regulatory solution, regulation may be needed. One possible path for regulation is to require opt-in for all forms of advertising other than contextual. However, opt-in systems are not a panacea: they can be designed so users click them away without understanding them, and once users opt-in it may be difficult to reverse the choice. If industry elected to, they could use self-regulation mechanisms to improve decision making through education, improved technology and tools, and more privacy-protective policies far more quickly than regulators could act. These tasks will be challenging no matter which parties take the initiative.

## 7 Acknowledgments

Please see individual papers for details on funding. Most of the studies summarized in this paper benefited from the wisdom of my co-author and PhD advisor, Lorrie Faith Cranor. However, this summary was completed without her assistance, and I do not speak for Professor Cranor. Similarly, while I am currently

consulting on Mozilla Firefox's Do-not-track feature, I do not speak for Mozilla. Thanks to Jon M. Peha and to the Institute for Informed Technology Policy (IITP) for support.

## References

- [1] CENTER FOR DEMOCRACY & TECHNOLOGY. What does "Do Not Track" mean? A scoping proposal by the Center for Democracy & Technology, January 2011. <http://cdt.org/files/pdfs/CDT-DNT-Report.pdf>.
- [2] MCDONALD, A., AND CRANOR, L. F. An empirical study of how people perceive online behavioral advertising. Tech. Rep. CyLab Technical Report 09-015, Carnegie Mellon, November 2009. [http://www.cylab.cmu.edu/research/techreports/tr\\_cylab09015.html](http://www.cylab.cmu.edu/research/techreports/tr_cylab09015.html).
- [3] MCDONALD, A. M., AND CRANOR, L. F. Americans' attitudes about internet behavioral advertising practices. In *Proceedings of the 9th Workshop on Privacy in the Electronic Society (WPES)* (October 4 2010).
- [4] MCDONALD, A. M., AND CRANOR, L. F. Beliefs and behaviors: Internet users' understanding of behavioral advertising. In *38th Research Conference on Communication, Information and Internet Policy (Telecommunications Policy Research Conference)* (October 2 2010).