Americans' Attitudes About Internet Behavioral Advertising Practices

Aleecia M. McDonald
am40@andrew.cmu.edu

Lorrie Faith Cranor
lorrie@cs.cmu.edu

Carnegie Mellon University
Pittsburgh, PA

ABSTRACT

AUTHORS' PRE-PRESS VERSION
Please cite to the published paper from WPES

This paper presents empirical data on American Internet users' knowledge about and perceptions of Internet advertising techniques. We present the results of in-depth interviews and an online survey focusing on participants' views of online advertising and their ability to make decisions about privacy tradeoffs. We find users hold misconceptions about the purpose of cookies and the effects of clearing them. Only 11% of respondents understood the text description of NAI opt-out cookies, which are a self-help mechanism that enables user choice. 86% believe ads are tailored to websites they have visited in the past, but only 39% believe there are currently ads based on email content, and only 9% think it is ok to see ads based on email content as long as their email service is free. About 20% of participants want the benefits of targeted advertising, but 64% find the idea invasive, and we see signs of a possible chilling effect with 40% self-reporting they would change their online behavior if advertisers were collecting data. We find a gap between people's willingness to pay to protect their privacy and their willingness to accept discounts in exchange for private information. 69% believe privacy is a right and 61% think it is "extortion" to pay to keep their data private. Only 11% say they would pay to avoid ads. We find participants are comfortable with the idea that advertising supports free online content, but they do not believe their data are part of that exchange.

WPES'10, October 4, 2010, Chicago, Illinois, USA.

Categories and Subject Descriptors

K.4.1 [**Computers and Society**]: Public Policy Issues—*privacy*

General Terms

[Human factors]

1. INTRODUCTION

Real-time mass media was born with national radio networks in the 1920s. As mass media gave rise to mass advertising, advertisers' campaigns became national. However, typically only a subset of people are interested in any given product or service advertised. As the old advertisers' lament has it, "We know we're wasting half our ad dollars, we just don't know which half." Online advertising can be targeted to users most likely to be interested in a particular product or service. Customers may benefit from ads targeted to their personal interests, reducing irrelevant ads and the time it takes to find products.

Behavioral advertising, which is one form of targeted advertising, is the practice of collecting data about an individual's online activities for use in selecting which advertisement to display. Behavioral advertising creates profiles for Internet users based on a variety of different data types and inferences drawn from those data. Third-party cookies are one of several mechanisms used to enable behavioral advertising: a central advertising network with ads across thousands of websites can set and read cookies, noting every time a given user visits any of the sites in the network. By correlating which sites an individual visits, ads clicked, inferences about age range and sex, and approximate physical location based on the computer's IP address, advertisers build profiles of that individual's characteristics and likely interests. Profiles indicate if a given user is a good target for certain ads, with interest categories like "cars" or "Hawaiian travel." Google and Yahoo! both use behavioral advertising and made their interest categories public at the end of 2009.

In this paper we review related work in section 2 and describe our methods in section 3. We present our findings regarding using cookie management as a self-help mechanism, participants' views of tailored advertising, and their willingness to pay for privacy in sections 4, 5, and 6 respectively. We conclude in section 7.

2. BACKGROUND AND RELATED WORK

Targeted advertising has received a lot of scrutiny in the past few years. There are questions about consumer's online privacy, how easily seemingly anonymous information can be re-identified [10], and the legality of some behavioral advertising business practices. The advertising industry favors an "industry self-regulation" approach. The Federal Trade Commission has held workshops and released guidelines for self-regulation [7, 6]. State and Federal legislatures are considering new regulations around Internet privacy, including proposals from Representatives Boucher, Stearns, Rush, and Senator Kerry.

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 cannot be tied to their names or any other personal information" [14]. In 2009, TRUSTe found that even if it "cannot be tied to my name or other personal information," only 28% of Internet users would feel comfortable with advertisers using web browsing history, and 35% believe their privacy has been invaded in the past year due to information on the Internet [15]. Anton, et. al., performed some of the earliest work on behavioral advertising in 2002, with a follow up study in 2009 [3]. They found the types of privacy concerns remained stable, but the level of concern has increased around information used for behavioral advertising. Turow et. al. conducted a nationally representative phone survey in 2009. They found 66% of adults do not want tailored advertising, which increased to as high as 86% when participants were informed of three common techniques used in advertising [17]. In 2003, Turow found that when offered a choice between paying for their favorite website with cash or with their personal information, over half of respondents said they would rather stop using the site all together [16].

Economics literature suggests that the most someone is willing to pay (WTP) to buy something should be equal to the minimum they are willing to accept (WTA) in payment for it: there should be a point of indifference between the good and cash. A difference between WTP and WTA may be indicative of an *endowment effect*, a phrase coined by Richard Thaler to describe when people place more value on an object that they own. The canonical example is that if two groups are asked to put a value on a coffee mug, people answering without owning the mug will generally suggest a lower price than people who first receive the mug as their own property. The endowment effect does not always occur with abstract items. For example, giving people a token that they can redeem for a mug does not have the same effect as giving them the actual mug [5]. Prior work shows a gap between WTP and WTA for revealing private data (for example, number of sexual partners) in an offline experiment [8]. Acquisti et. al. found substantial differences between WTP and WTA with gift cards and inexpensive tangible goods, including "subjects who started from positions of greater privacy protection were five times more likely than other subjects to forego money to preserve that protection" [1]. We examine the Acquisti hypothesis in an online context. If there is also a gap between WTP and WTA online, then the way privacy choices are framed may affect the decisions people make about online privacy.

3. RESEARCH METHODS

We followed a two-part approach. First we performed a laboratory study to identify a range of views through qualitative interviews. Then we conducted an online survey to test and validate our qualitative results.

In the first study 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 primed for privacy. We followed a modified mental models protocol of semi-structured interviews, using standard preliminary questions for all participants, then following up to explore participants' understanding. Our study ran from September 28th through October 1, 2009 in Pittsburgh, PA. We recruited participants with a notice on a website that lists research opportunities. Participants were compensated $10 for an hour of their time.

In the second study, we recruited 314 participants from the Mechancal Turk[1] website at the end of April, 2010. We paid participants $2 for a 20-30 minute study. We removed two outliers from our dataset; they had unusually short response times and response patterns that suggested they had not read the questions. We saw a drop-out rate of 37%. We deliberately started the study with short-answer questions to encourage people not to take the survey unless they were willing to invest some time. We coded free-form responses to tabulate categories of responses, or "unclear" when we were unsure what participants had in mind.

3.1 Demographics

Of the 14 subjects we interviewed, 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.

For the online study, 41% of our participants were male. Our population was notably skewed younger than the adult American Internet population. 25% were 18-24, 39% were 25-34, 17% were 35-44, 12% were 45-54, and 8% were over 55. 74% were white, 9% American Indian or Alaskan Native, 6% Asian, 4% Black or African American, and 2% Latina/Latino or Hispanic.[2] Our online survey participants have been using the Internet for an average of 13 years, with 15% online for over 15 years and 2% online less than five years. 85% use Windows, 11% use Macintosh, 4% all other operating systems or unsure. This matches Quantcast's estimate of 87% Windows, 11% Macintosh [11]. The most popular browser was Firefox at 48%, followed by Internet Explorer (34%), Chrome (10%), Safari (5%), Opera (2%) and 1% answering they are unsure or other. Our sample is skewed toward Firefox users, with an estimated 25% of the market share, and Chrome (6%), at the expense of Internet Explorer (60%) users [4]. Most use at least one remotely-hosted and professionally-managed email service: Yahoo Mail (50%), Gmail (50%), Hotmail (23%), and AOL mail (16%). Only 9% of participants reported they do not check at least one email account of this type.

3.2 Transferability

Early in the online study, before we asked questions that might affect participants' views, we asked the same three questions Turow et al. asked in their study designed to be

[1]Mechanical Turk is crowd-source web portal run by Amazon. See `www.mturk.com` for details.

[2]Both our survey and the Census allow more than one selection for race which is why results sum to more than 100%.

representative of the US population [17]. As our sample is not a statistically representative sample of United States Internet users, we contrasted to the Turow work to understand the transferability of results to other contexts [9]. We found similar results for two of their three questions, as shown in Table 1.

Table 1: Percentage of respondents who want tailored content

Do you want websites you visit to show you...	Turow et al.'s	Our results
ads that are tailored to your interests?	32%	45%
discounts that are tailored to your interests?	47%	80%
news that is tailored to your interests?	40%	41%

Where the representative Turow sample is comprised of 35% of people aged 18 – 34, our sample is 69% in that age range. However, despite age-linked differences in responses, our younger sample does not explain why we saw a substantially higher percentage interested in tailored discounts. We had approximately 20% more interest in tailored discounts in all of our age categories as compared to the Turow work. One possible explanation: we recruited participants willing to spend 20 minutes to answer our survey for $2. Our participants may be unusually sensitive to financial incentives. For tailored ads and news, our findings mirrored the Turow paper: most respondents are not interested in tailored advertisements or news.

4. PERCEPTIONS ABOUT COOKIES

All participants in the interviews had heard of cookies before but we observed widespread confusion. 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.

4.1 Misperceptions of First Party Cookies

While interview participants generally did not understand what cookies are, perhaps it is more important that they understand the effects of cookies. We asked follow up questions of "are there ways cookies can help you?" and "are there ways cookies do not help you?" Over a third of participants said that cookies can be related to saving passwords. Similarly, three participants answered that cookies allow them to remain logged in to websites without retyping a password, though during follow-up questions they did not actually know if cookies were involved (as opposed to Apple's Keychain Access., etc.). Three participants believed cookies store their preferences for websites, including details like preferred colors and placement of site elements.

Only three participants said that cookies are related to personalized advertisement. They expressed three very different perspectives. 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. 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." 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." 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." 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 described this practice as a "smart thing" and "reasonable." He then volunteered that he believes companies are constrained by law not to share information. He later said 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. Notice the analogies to off-line settings as participants form their views of how privacy works online. Legal protection of privacy in telephone conversations and postal mail are often assumed to carry over to Internet communications as well.

4.2 Managing Cookies

There are three ways people manage cookies: by not letting them save to their hard drive in the first place, by deleting them automatically, or deleting them "by hand." We asked about all three methods in our online study.

Several major web browsers offer a "private browsing" feature that allows users to toggle to a private mode that never saves cookies, history, and cache data. When finished, users exit private browsing and have access to their normal set of cookies, history, and cache data. Only 23% reported they ever use private browsing, 50% do not use private browsing, and 27% are not sure if they use private browsing.

17% use software that deletes cookies for them, 23% are not sure, and 60% answered no. Those who answered yes predominately use either anti-malware software or CC Cleaner, though sometimes they had trouble naming the specific product they use (e.g., "malware by anti-malware.") Some may delete cookies via anti-malware programs without understanding they are doing so. One participant answered "TACO, NoScript, & Firefox," which is a sophisticated approach.

9% said they never clear cookies, 9% believe they clear cookies themselves annually or less than once a year, 16% a few times a year, 10% monthly, 17% a few times a month, 16% a few times a week, 12% daily, and 8% clear cookies every time they close their browser. This is self-reported data, but about 70% believe they clear cookies at least annually.

4.3 Unclear on Clearing Cookies

Why do people clear cookies? Interestingly, they are not

always sure themselves. Participants in our lab study had a vague notion that too many cookies are bad. They are not sure under which conditions they should delete or retain cookies. For all that they do not understand about how cookies work, they do understand some of the benefits of cookies, such as not needing to log in again.

For the online study, we asked an open-ended question about why they deleted or saved cookies and coded the responses. Participants wrote answers that reflect an underlying lack of knowledge like "Someone recommended it to me once and I have done it ever since," or "I'm not very sure what [cookies] are. I have cleared them before because it was suggested to me that I do." Family is sometimes mentioned as the source of advice, including "Mom told me to," "My daughter told me to," and "My husband doesn't want them." Similarly for why people do not clear cookies frequently, participants gave answers like "I don't really know" or "No particular reason." We coded these vague responses along with a variety of other non-reason or unclear answers as "Other," which comprised 8% of all responses. In total, our 314 participants gave 390 reasons to delete or not delete cookies. Of 80 reasons not to delete cookies:

- 31% were some form of apathy, either that cookies do not bother them or they do not care about cookies.
- 27% have software that deletes cookies automatically.
- 20% were not sure what cookies are, or why they would delete them.
- 19% were unsure how to delete cookies.
- 3% (two people) wrote that they do not care about being tracked online.

Of 278 reasons given to delete cookies:

- 33% were based on the idea that "many cookies slow down my computer." This seems unlikely in practice.[3]
- 30% had to do with privacy and security. About a fifth of the privacy and security reasons mentioned deleting history; history is commonly confused with cookies. The remaining four-fifths of privacy and security reasons generally reflected some understanding of how cookies work, for example, "I wouldn't want someone being able to get on my computer and remain logged into my accounts. Also, I don't want a website tracking me through them."
- 28% had to do with freeing up hard drive space, reducing clutter, or a notion of hygiene and cleanliness. Answers included "[I] like having a clean slate on the computer all the time," "[to] clear up clutter," and "to make space on my computer." Few modern computers will run into space problems due to cookies.[4]
- 8% mention viruses, spam, or malware. Some tracking cookies are classified as spyware by Norton Anti-virus and other anti-malware programs.

User confusion is high. Some do not know how to delete cookies and might wish to do so, which limits self-help mechanisms in privacy decision making. Some participants reported what seems to be over-clearing of cookies: they delete cookies to avoid issues that cookies do not cause. Cookie deletion creates uncertainty in measuring the number of people — and unique people — who have seen a given online ad, or have visited a given website. Over- and under-counting ad impressions causes economic harms to members of the advertising community, with hundreds of thousands of dollars disputed in large ad campaigns. When users delete their cookies for reasons that do not match their actual preferences, it causes harm without the gains users expect.

4.4 Cookies and Browser History

More than half of our interview participants confused cookies with browser history. Participants did not understand that browser history is stored independently of cookies, which may make it difficult for people to enact their privacy preferences. One participant in our lab study told us cookies contain a "history of websites" visited and when he deletes cookies, "hyperlinks in different colors goes [sic] away, that's what it does. It clears the navigation history." When he was a child he lost his computer privileges because his mother 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..." In his view, cookies were only an issue on computers where he shared a single account with multiple people. 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 and believe they cannot be tracked unless they log in to a website. 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 in which clearing cookies, clearing history, and clearing cache data settings are intermingled may contribute to user confusion.

In the online study, we asked "Sometimes you hear about web browser history. Are cookies and history the same?" 35% of participants incorrectly answered yes. Those who answered no generally had a good working understanding of the difference between cookies and history, with free-form responses like "History is a list of your previous browsing, and cookies are files that registered each site visited." Of those who correctly answered no, 79% were able to give at least partially correct answers explaining how cookies and history differ, 12% gave clearly incorrect answers, and 8% gave answers that were too unclear to tell.

4.5 Consumers Do Not Understand NAI Opt-Out Cookies

None of our interview participants had heard of cookie-based methods to opt-out of tracking cookies, including TACO[5]

[3]For DSL users, a webpage with a 3000 byte cookie takes approximately 80 milliseconds longer to load [13] so users are not wrong to associate cookies with delay. However, just deleting all cookies without blocking them does not improve time to load the page: websites would simply download new cookies to replace the deleted cookies. Participants may be confusing cookies with cached images.

[4]Browsers typically set a maximum size of 4k per cookie and a maximum number of cookies per domain to avoid denial of service attacks from malicious servers filling hard drives, and hard drives today are typically measured in gigabytes.

[5]Targeted Advertising Cookie Opt-Out (TACO) is a plugin for the Firefox browser that stores persistent opt out cook-

and NAI opt-out cookies.[6] At the end of the protocol, we showed four participants a text description of NAI opt-out cookies from the NAI opt-out website, with a list of a few NAI member companies shortened to fit on a single page (see Figure 1.)[7]

All four participants understood they would continue to see at least some online advertisements. However, there was substantial confusion about what the NAI opt-out does. The text does not disclose 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 [2]. 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 opt-out meant "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 thought it would "reduce 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 companies where "the information they would seek would be too personal to share with a group." Our final participant understood the NAI text. At first he said if you use Gmail, the opt-out cookie means "stop reading my email and tailoring ads." 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."

During interviews 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 the same screenshot 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.
- 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. TAILORED CONTENT & PRIVACY CONCERNS

Advertisers claim that consumers are clamoring for more interesting and relevant advertisements, while privacy advocates claim citizens' privacy rights are being trampled. We found support for both views: there are sizable groups of people with each of those views. In the middle, we found a large group of people who are disinterested in better ads since their goal is to ignore ads. They see no benefit to targeted advertising, so they do not see reason to share data with advertisers. They accept the idea that ads support free content, but do not expect data to be part of the exchange.

5.1 Mixed Understanding of Current Practices

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. 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.

As shown in Table 2, participants seem to have a high degree of understanding that behavioral advertising happens,

ies, available from: `https://addons.mozilla.org/en-US/firefox/addon/11073`

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

[7]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 NAI site.

Figure 1: Screenshot of the NAI Opt Out page

Table 2: 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%

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.

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 3, 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 3: Feelings 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%

5.2 Reasons to Accept or Reject Tailored Advertising

Based on discussions in the laboratory study, we compiled a list of reasons participants gave for being for or against behavioral advertising. We presented online participants with a seven point Likert scale from Strongly Agree (7) to Strongly Disagree (1), summarized in Table 4.

Table 4: Mean Likert scores to accept or reject behavioral advertising (Strongly Agree = 7, Strongly Disagree = 1.)

Description	Mean	Agree	Disagree
Someone keeping track of my activities online is invasive	5.7	64%	4%
Behavioral targeting works poorly and I get ads that are not relevant to me, even when they are supposed to be	4.8	34%	7%
I would watch what I do online more carefully if I knew advertisers were collecting data	4.7	40%	15%
I ignore ads, so there is no benefit to me if ads are targeted to my interests	4.7	36%	11%
I ignore ads, so I do not care if ads are targeted to my interests or if ads are random	4.4	31%	16%
I ignore ads, so there is no harm to me if ads are targeted to my interests	4.2	24%	17%
I want the benefits of relevant advertising	4.1	21%	21%
I would stop using any site that uses behavioral advertising	3.6	15%	29%
I am protected by law against advertisers collecting data about me	3.6	16%	34%
I do not care if advertisers collect data about my search terms	2.9	10%	51%
I do not care if advertisers collect data about which websites I visit	2.8	12%	53%

Privacy concerns are top priorities. Nearly two-thirds of our participants agreed or strongly agreed that "someone keeping track of my activities online is invasive," with only 4% disagreeing or strongly disagreeing. This phrase comes directly from a participant we interviewed in the lab study, and reflects the way she thought about behavioral advertising. It is phrased in a way that would likely garner maximum response by mentioning an unnamed, but presumably human, "someone" and using the possessive "my." We suggest the way to understand this result is that if behavioral advertising is framed this way, most Americans will respond poorly to it.

40% agreed or strongly agreed they would be more careful online if they knew advertisers were collecting data. The wording of this question limits data use to advertisers, which may reduce concern. It also explorers the notion of a chilling effect. Respondents at least believe they would self-censor if they knew advertisers were collecting data. While self-reported data is not always indicative of actual behavior, it appears people are considering leaving FaceBook in response to publicity about data flows to advertisers [12]. Advertiser's practices have the potential to reduce Internet adoption and use, and may already be doing so.

Despite claims that users do not care about privacy, half of participants disagreed or strongly disagreed that they do not care if advertisers collect search terms, or if advertisers collect data about websites visited, both of which occur regularly for behavioral advertising and analytics data. Only around a tenth of respondents agreed that they do not care. However, only 15% self-report that they would stop using sites with behavioral advertising.

In our laboratory study we heard two conflicting attitudes from people who ignored ads. Several people told us that because they ignore ads, they get no benefit from targeted advertising and would therefore rather not have any data collected about them. Other people told us that because they ignore ads, they do not care if ads are targeted or random and they do not care if data is collected. We also wondered if there might be people who just do not care at all. In the online study we found the strongest agreement with the statement "I ignore ads, so there is no benefit to me if ads are targeted to my interests" (36% agree or strongly agree,) the weakest agreement on "no harm to me" for targeted ads (24% agree or strongly agree,) with the most strictly apathetic option of not caring if there are targeted or random ads in the middle (31% agree or strongly agree.)[8] This suggests that of those who ignore ads, they are likely to prefer data not be collected about them, since they do not see any benefit. However, just because someone claims to ignore ads does not mean that is always the case. Advertisers may still gain benefit from targeting these users. But an argument that targeted ads are a benefit will likely fall flat with the people who are not interested in any ads, let alone better ads. Interestingly, when we put that question to participants directly, we saw an even split. 21% agree or strongly agree that they want the benefits of relevant advertising while 21% disagree or strongly disagree, with a neutral Likert mean of 4.1. What emerges is neither a strong clamoring for nor a backlash against behavioral advertising, but rather several distinct groups with quite different preferences.

5.3 Privacy and Security Among Top Priorities for Buying Online

98% of our participants indicated they make purchases online. More than half said they never make purchases based on Internet ads or email advertising, as summarized in Table 5. This is self-reported data; people may make buying

[8]We found statistically significant differences in means between "no benefit" and "no harm" as well as "do not care" and "no benefit" ($p < .05$, df=312, paired two-tailed t-Test, $\alpha = .05$). We did not find significance between "no harm" and "do not care" ($p = .060$).

decisions based on ads without being aware they are doing so.

Table 5: Respondents who buy online

How often	Buy online	Buy based on Internet ads	Buy based on email ads
Never	2%	52%	54%
A few times / month	42%	7%	6%
A few times / year	52%	38%	38%

We asked participants how sellers could entice participants to purchase more products online, and listed 13 possible approaches with responses on a four point Likert scale of "Matters a lot," "Matters," "Matters a little," and "Does not matter." We created our 13 categories based on responses to a pilot test with an open-ended question. See Table 6 for results.

Table 6: How sellers can entice more online purchases (Matters a lot = 4, Does not matter = 1)

Description	Mean	Matters a lot	Does not matter
Free shipping	3.7	75%	1%
Will not share your data with advertising partners	3.6	70%	3%
No spam policy	3.6	70%	3%
Improved fraud protection for credit card transactions	3.6	68%	3%
No hassle return policy	3.6	67%	2%
Clear information about products	3.6	66%	2%
Web discounts	3.5	57%	1%
Easy-to-use website	3.4	55%	2%
Online coupons	3.2	46%	4%
Local pickup	2.4	18%	26%
Will only retain data about your purchases for three months	2.3	14%	24%
Products recommended based on your past purchases	2.3	10%	23%
Products recommended based on your friends' past purchases	1.8	7%	47%

The most popular item was free shipping.[9] The next three most popular were clustered around privacy and security: not sharing data with advertisers, a policy against spam, and fraud protection. In contrast, the remaining privacy and security item on data retention scored near the bottom. This may be a function of the specific description, or due to lack of understanding of how data retention limits reduce privacy and security risks, but suggests data retention is not currently a major concern for users.

Return policies and clear information about products scored higher than discounts, all of which scored better than an easy-to-use website or online coupons. No clear story emerges about usability vs. financial incentives. Recommending additional products did not interest our respondents, regardless of whether recommendations came from their own purchasing history or their friends. From the discussions we had during our lab-based study, many people find it "creepy" to get suggestions based on friends' purchasing history. However, we are surprised to see their own purchasing history score nearly as low, when well-known companies like Amazon have successful services in production. This may suggest users do not think about the mechanics behind such recommendations, or just that they think themselves more immune to advertisements than they are in actual practice.

6. PAYMENT FOR PRIVACY

We have observed that some people who are highly concerned with privacy are strongly disinclined to spend money to preserve privacy. This can seem counterintuitive, especially since in many domains the amount someone is willing to pay for something indicates how highly it is valued. Instead, some people who believe privacy is a right respond negatively to the idea of paying to protect their privacy.

6.1 Gap Between Willingness to Pay and Willingness to Accept

We split our participants into two groups. First we asked them to name their favorite online news source, and answer how frequently they visit it to make our next questions more salient. Then one group answered the question "Would you pay an additional $1 per month to your Internet service provider (ISP) to avoid having your favorite news site collect your data for behavioral advertisements?" The second group answered a similar question of "Would you accept a discount of $1 per month off your Internet service provider (ISP) bill to allow your favorite news site to collect your data for behavioral advertisements?" In theory, there should be no difference between the price someone is willing to pay (WTP) to protect privacy and their willingness to accept (WTA) payment for revealing information.

We did find a gap between WTP and WTA. Only 11% of respondents were willing to pay $1 per month to keep their favorite news site from collecting data, while 31% of respondents were willing to accept a $1 per month discount to disclose the information. Thus, 11% said they were willing to pay $1 extra to gain privacy while 69% said they were unwilling to accept a $1 discount to give up privacy. In the privacy sphere this could have two very interesting effects. First, people who think they have already lost the ability to control private information — that privacy is not something they are endowed with — may value privacy less as a result. Those who believe they have control over information may value privacy more as a result. Second, the difference between opt-in and opt-out rates for online privacy may not just be due to the well-documented tendency for people to keep defaults unchanged. If a service collects

[9]The word "free" often gets a strong response. It would be interesting to see if this result is robust when phrased as "discounted shipping."

data by default and users must opt-out of data collection, that suggests users are not endowed with privacy, and they may respond to that cue by valuing their privacy less.

6.2 Reasons to Pay or Refuse to Pay for Privacy

We followed up by asking questions to better understand why people would decide to pay or accept $1, based on reasons we heard from our lab study participants. We asked "Some websites may offer you a choice of paying for content or receiving content for free in exchange for letting them send you targeted advertising. How strongly do you agree or disagree with the following statements?" with a seven point Likert scale from Strongly Agree (7) to Strongly Disagree (1). See Table 7 for details.

Table 7: Reasons to pay for privacy or accept a discount

Description	Mean Likert	Agree	Disagree
Privacy is a right and it is wrong to be asked to pay to keep companies from invading my privacy	5.9	69%	3%
Companies asking me to pay for them not to collect data is extortion	5.6	61%	5%
It is not worth paying extra to avoid targeted ads	5.5	59%	5%
Advertisers will collect data whether I pay or not, so there is no point paying	5.4	55%	4%
I hate ads and would pay to avoid them	3.3	11%	36%

Only 3% of respondents either disagreed or strongly disagreed that privacy is a right and it is wrong to be asked to pay for privacy online, even in exchange for free content. The top two ranking replies suggest that one reason people will not pay for privacy is because they feel they should not have to: that privacy should be theirs by right. Yet when phrased as an economic proposition, that it is "not worth paying extra," participants also predominately agree. One might expect that participants who highly value privacy would disagree, and would think it is worth paying for privacy even if they also believe they should not have to do so, but only 5% did. Distrust of the advertising industry, or perhaps of actors on the Internet as a whole, is another reason people may not be willing to pay for online privacy with just over a majority agreeing or strongly agreeing that data will be collected even if they pay companies not to collect data. Finally, we can rule out dislike of advertising as a major factor in online privacy decision making, with only 11% willing to pay to avoid ads because they "hate" them. Most participants are accustomed to advertising. Mass media advertising has been part of life since before they were born. It is the data collection that is new, and, to many, a troubling aspect of online advertising.

7. CONCLUSIONS AND DISCUSSION

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. 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. 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 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 in giving data to advertisers. Ideally, users could choose for themselves but at present they lack the knowledge to be able to make informed decisions.

Most users understand that cookies store data on their computers, enable tailored ads, and allow tracking across sites. They are unclear on important details like whether cookies may be combined with other data, what data is stored in cookies, if blocking cookies preserves geolocational privacy, and they are particularly unclear about laws and law enforcement. Web browsers may contribute to users' confusion. Browsers may also be an avenue to help with user understanding and decision making in the future. Thus far, browser makers have been largely absent from behavioral advertising issues. However, Microsoft has been involved with behavioral advertising for years, and their adoption of P3P in Internet Explorer changed the third-party cookie landscape. Yet the *Wall Street Journal* reported that Microsoft re-designed Internet Explorer 8 specifically to enable third party tracking for business reasons [18]. It may be naive to expect browser makers to support user privacy at their own expense, and most major browser makers are involved with Internet advertising. The NAI is a leader in behavioral advertising self-regulation but their opt-out cookie page is very confusing, with only 11% understanding what it is for. The NAI may not be supporting Internet users' ability to avail themselves of self-help options.

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. We also found a gap between willingness to pay to protect data and willingness to accept a discount in exchange for releasing the same data. People may ascribe more value to what they possess. People may value their privacy less when presented with an opt-out for data collection, which suggests data belongs to the company collecting it, rather than an opt-in choice for data collection, which suggests data belongs to the individual.

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 interview 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. In our lab study, one technically-savvy participant even described withdrawing from online life as a result of privacy concerns.

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.

8. ACKNOWLEDGMENTS

Thanks to Faisal N. Jawdat and to Greg Norcie for coding open-ended questions. Thanks to anonymous reviewers, participants at the Privacy Law Scholars Conference (PLSC 2010) and to CUPS laboratory members for invaluable feedback. This research was supported in part by CyLab at Carnegie Mellon under grant DAAD19-02-1-0389 from the Army Research Office and by Microsoft Research.

9. REFERENCES

[1] A. Acquisti, L. John, and G. Loewenstein. What is privacy worth? Technical report, Heinz College, Carnegie Mellon University, 2009.

[2] S. Anderson. House subcommittees hold joint hearing on behavioral advertising. *Security, Privacy and the Law*, July 2009. 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=-Wklp2qdbmw. Accessed 9 November 2009.

[3] A. I. Anton, J. B. Earp, and J. D. Young. How Internet users' privacy concerns have evolved since 2002. Technical Report Computer Science Technical Report TR-2009-16, North Carolina State, 2009. http://theprivacyplace.org/blog/wp-content/uploads/2009/07/tr_2009_16.pdf Accessed 3 March 2010.

[4] S. Axon. Chrome gains, IE slumps in browser wars. The Social Media Guide, May 2010. http://mashable.com/2010/05/04/chrome-firefox-ie-stats/ Accessed 15 May 2010.

[5] The endowment effect: It's mine, I tell you. *The Economist*, June 2008. http://www.economist.com/science/displaystory.cfm?story_id=11579107 Accessed 2 September 2008.

[6] Federal Trade Commission. FTC staff revises online behavioral advertising principles, February 2009. http://www.ftc.gov/opa/2009/02/behavad.shtm Accessed 15 May 2009.

[7] Federal Trade Commission. Self-regulatory principles for online behavioral advertising: Tracking, targeting, and technology. Staff Report, February 2009. http://www.ftc.gov/os/2009/02/P085400behavadreport.pdf Accessed 9 November 2009.

[8] J. Grossklags and A. Acquisti. When 25 cents is too much: An experiment on willingness-to-sell and willingness-to-protect personal information. In *Workshop on the Economics of Information Security (WEIS)*, 2007.

[9] S. R. Harper and G. D. Kuh. Myths and misconceptions about using qualitative methods in assessment. In S. R. Harper and S. D. Museus, editors, *Using qualitative methods in institutional assessment. New Directions for Institutional Research*, number 136, pages 5–14. Jossey-Bass, San Francisco, 2007.

[10] P. Ohm. Broken promises of privacy: Responding to the surprising failure of anonymization. *UCLA Law Review*, (forthcoming), 2010. http://papers.ssrn.com/sol3/papers.cfm?abstract_id=1450006 Accessed 10 November 2009.

[11] L. Prescott. OS X share up 29% in past year, slowly chipping away at Microsoft. VentureBeat, February 2010. http://venturebeat.com/2010/02/26/os-x-share-up-29-in-past-year-slowly-chipping-away-at-microsoft/ Accessed 15 May 2010.

[12] J. D. Sutte. Some quitting facebook as privacy concerns escalate. *CNN Tech*, May 2010. http://www.cnn.com/2010/TECH/05/13/facebook.delete.privacy/index.html?iref=allsearch Accessed 15 May 2010.

[13] T. Theurer. Performance research, part 3: When the cookie crumbles. *Yahoo! User Interface Blog*, March 2007. http://yuiblog.com/blog/2007/03/01/performance-research-part-3/ Accessed 10 May 2010.

[14] TRUSTe. 2008 study: Consumer attitudes about behavioral targeting, March 2008. http://danskprivacynet.files/wordpress.com/2009/02/truste2008_tns_bt_study_summary1.pdf Accessed 9 November 2009.

[15] TRUSTe and TNS. 2009 study: Consumer attitudes about behavioral targeting, March 2009.

[16] J. Turow. Americans & Online Privacy: The System is Broken. Annenberg Public Policy Center Report, 2003.

[17] J. Turow, J. King, C. J. Hoofnagle, A. Bleakley, and M. Hennessy. Americans reject tailored advertising and three activities that enable it. Technical report, Annenberg School for Communications, University of Pennsylvania, September 2009. http://repository.upenn.edu/asc_papers/137/ Accessed 4 March 2010.

[18] N. Wingfield. Microsoft quashed effort to boost online privacy. *Wall Street Journal*, August 2 2010.