Folk Models of Online Behavioral Advertising

Yaxing Yao
Syracuse University

Davide Lo Re
Sapienza University of Rome

Yang Wang
Syracuse University

ABSTRACT

Online Behavioral Advertising (OBA) is pervasive on the Internet. While there is a line of empirical research that studies Internet users' attitudes and privacy preferences of OBA, little is known about their actual understandings of how OBA works. This is an important question to answer because people often draw on their understanding to make decisions. Through a qualitative study conducted in an iterative manner, we identify four "folk models" held by our participants about how OBA works and show how these models are either incomplete or inaccurate in representing common OBA practices. We also discuss how privacy tools can be designed to consider these folk models. In addition, most of our participants felt that the information being tracked was more important than the web trackers themselves. This suggests the potential for an information-based blocking scheme rather than a tracker-based blocking scheme used by most existing ad-blocking tools.

Author Keywords

Mental model, web tracking, Online Behavioral Advertising (OBA), Privacy-Enhancing Technologies (PETs)

ACM Classification Keywords

H.5.m. Information Interfaces and Presentation (e.g. HCI): Miscellaneous

INTRODUCTION

Online Behavioral Advertising (OBA), or targeted advertising, is prevalent on today's Internet [36]. OBA is "the practice of tracking an individual's online activities in order to deliver advertising tailored to the individual's interests" [11]. A common practice of OBA is that first-party sites (i.e., sites that a user visits voluntarily) rely on third-party entities (e.g., ad networks) to track a user's browsing activities across websites and to provision ads targeted at the user [26]. OBA can benefit both advertising companies (e.g., increasing click-through rates and prices of ads [4]) and Internet users (e.g., providing ads that better match their potential interests [27, 39]). However, since OBA involves online tracking and profiling of users, it has raised significant privacy issues [38, 27, 39].

Prior studies have found various user attitudes and perceptions of OBA (e.g., [38, 27, 39, 23, 33]). For instance, Ur et al. note that people find OBA "creepy and scary" because of its online tracking practices, but sometimes people also find OBA "smart and useful" [39]. As such, individual users seem to have varying acceptance of OBA depending on the context [10, 40, 28]. However, most of these studies either (1) did not study people's understandings of how OBA works (e.g., [38, 22]) or (2) investigated people's perceptions of OBA after the researchers explained OBA (e.g., [23, 33, 10, 40, 28]), therefore it is not clear to what extent ordinary Internet users actually understand how OBA works now what their understandings are.

Drawing from the literature on mental models, we examine people's understandings of how OBA works. Psychologist Kenneth Craik pioneered the concept of mental models, describing "the mind constructs 'small-scale models' of reality that it uses to anticipate events, to reason, and to underlie explanation" [12]. Since then, the notion of mental models has been further developed. For instance, Phil Johnson-Laird, an influential scholar of mental models defines them as "psychological representations of real, hypothetical, or imaginary situations" [17]. Mental models have also been studied extensively to understand how people comprehend various things such as language and music [15]. In addition, "mental models affect people's reasoning" [17] and people draw from their mental models to make various decisions [17, 18]. For instance, people's mental models of how thermostats work influence the ways in which they control these devices [42].

The mental model approach has also been applied in the domain of privacy and security (e.g., [1, 3, 9, 41, 8, 24, 29, 13]), but has not been systematically used in the context of OBA. Rick Wash conducted an interview study to examine people's mental models of home computer security [41]. He notes, "to understand the rationale for people's behavior, it's important to understand the decision model that people use" [41]. Drawing from prior literature (e.g., [35, 2]), he uses the term *folk models* to denote mental models that can be incorrect representations of reality but are used by people in practice [41].

Our work was in part inspired by Wash's study [41]. We aim to uncover people's folk models of OBA, regardless of whether these models accurately represent the reality of OBA. We note that mental models can encompass more than a picture of how things work [15], but here we use folk models to denote people's understanding of how OBA works. There is little work that touches on this question, and our study aims to fill the gap. We believe that understanding people's folk models of OBA is important because these models can influence people's behavior or decisions regarding OBA, for instance,

how they control or manage OBA. Furthermore, privacy tools for OBA can be more effective when they incorporate people's folk models, for example, by helping people recognize privacy risks (e.g., third-party tracking) and adopt countermeasures (e.g., blocking third-party trackers).

We inductively developed four folk models of OBA held by our participants through a qualitative study conducted in an iterative manner. In addition to a pilot study with eight people, we conducted two rounds of semi-structured interviews with another 21 Internet users from different U.S. states and cities. These models differ in terms of the following: who tracks Internet users' information; where the tracked information is stored; and how targeted ads are selected or provisioned.

Similar to Wash's study [41], our qualitative research does not support claims that can be generalized to all Internet users, but it instead aims to uncover folk models that people have about OBA and that can inform future privacy-enhancing designs for OBA. In the sense of theoretical sampling [25], the discovered folk models are held by real people but the study says little about how common or statistically representative these models are in the general population.

To guide future privacy tools for OBA, we also asked participants' opinions about what tools or features they desire in order to help them protect their privacy in the context of OBA. While most OBA tools focus on trackers, most of our interviewees felt that the information being tracked is more important than trackers. This result suggests the potential for an information-based blocking scheme rather than a tracker-based blocking scheme used by most existing ad-blocking tools such as Ghostery.

This paper makes two main contributions. First, we uncover different folk models of OBA that ordinary Internet users have. These models have implications for privacy designs and public policies of OBA. Second, we identify people's desired features of privacy-enhancing tools for OBA. These features should be incorporated into future privacy tools.

RELATED WORK

Our work was mainly inspired by prior research on people's attitudes and perceptions of online tracking and OBA, as well as people's mental models of privacy and security.

People's Attitudes and Perceptions of OBA

There is a line of empirical research that examines people's attitudes towards OBA mostly via surveys. Several surveys have shown people's objection of online tracking and OBA. For instance, Turow et al. polled 1000 Internet users in the U.S. and found that 87% of them did not want advertisers to track them online [38]. Similarly, McDonald and Cranor found that 64% of their survey respondents considered targeted ads to be "invasive" [27]. Another survey found that one major reason why the respondents disliked OBA was because of online tracking and subsequent analyses of that tracked data [32].

However, Ur et al.'s interview study has painted a more nuanced picture. They found that many of their interviewees considered OBA "creepy and scary" because of its online tracking practices, but sometimes people also found targeted ads "smart and useful" [39]. This study also suggested people's acceptance of OBA may vary depending on the context.

A number of subsequent survey studies focused on people's context-based preferences of OBA [10, 40, 28]. Leon et al. found that the data retention period and scope of data use significantly affected their respondents' willingness to share data for OBA [23]. Chanchary and Chiasso found that people's OBA preferences differ by the first-party sites they visit [10]. Melicher et al. combined their participants' browsing histories and interview data in identifying additional situational factors such as the types of information being tracked and the frequency of visiting first-party sites that can affect people's attitudes towards online tracking [28]. Wang et al. surveyed both American and Chinese Internet users and found that both user groups had different OBA preferences based on the type of first-party sites, despite the fact that the former had more privacy concerns over OBA than the latter [40].

While these prior studies offer invaluable insights into people's perceptions of OBA, most of these studies (e.g., [23, 33, 10, 40, 28]) provided a detailed explanation of OBA before examining people's preferences of it. In contrast, four prior studies asked people's perceptions of OBA before explaining OBA [38, 27, 39, 22] and two of them did not ask about people's understandings of how OBA works [38, 22]. The other two studies touched on this question but did not yield mental models that represent people's understandings of OBA [27, 39]. Ur et al.'s study focused on people's attitudes towards OBA rather than their understandings of OBA [39]. The remaining study investigated people's beliefs about OBA [27] but differed significantly from our study.

More specifically, McDonald and Cranor provided their survey respondents four diagrams depicting different configurations of first- and third-party cookies in OBA and asked the respondents to select the configuration which was not possible [27]. Unlike their approach, we sought to discover people's folk models of OBA without providing any a priori models or pictures to constrain or influence their thinking. We have discovered folk models (e.g., browser-based models) that differ from the models they provided in their study. We will present our folk models in the results section.

In addition, few studies have touched on people's understanding of online tracking and OBA. Rader conducted an online experiment and found that most participants were aware that sites like Google or Facebook can collect information about their users' activities on them (e.g., what pages they visit or what links they click) [33]. This is a case of first-party tracking. Ur et al. asked their interviewees the ways in which ads are tailored to them. The two most common methods mentioned were based on users' browsing histories and web searches [39]. Another survey study found that people have various understandings of the type of data (e.g., personal information or location) web trackers can track online [10]. Some of these perceptions were incorrect, e.g., people thought online tracking was malware and online tracking directly involved local browsing history [28]. Our work differs

from these studies in that we focus on people's folk models of OBA rather than exploring them in passing.

Overall, the extant literature does not provide a clear picture of the folk models people have about how OBA works. Our study aims to fill this gap.

People's Mental Models of Privacy and Security

The mental model approach has been employed by a number of researchers to investigate people's understandings of the Internet [37, 19]. Thatcher and Grey's work utilized drawing as a means of understanding people's mental models [37]. Their work revealed several typical understandings of how the Internet works, such as considering the Internet as a central database, or as a modular structure network [37]. Our study adopted a similar drawing task to solicit people's understandings of OBA.

Kang et al. observed that people's mental models of how the Internet works can be very different, and these models were partially influenced by people's technical knowledge [19]. The researchers suggested that users with more technical knowledge tend to have a more sophisticated mental model, but the level of technical knowledge barely affects users' security and privacy practices [19].

Researchers have also used the mental model approach to investigate users' perceptions related to their privacy and security. Camp proposed five possible mental models that can be used to explain people's understandings of computer risks, including models of physical safety, medical infections, criminal behavior, warfare activities, and market failures [9]. Asgharpour et al. conducted a card sorting study and found that computer security experts and non-experts have different mental models of computer security [3]. For instance, experts associated passwords with a criminal model whereas non-experts thought of a physical safety model [3]. Wash's work on people's mental models of threats towards their home computers suggested eight folk models, including four virus-centered models and four hacker-centered models [41]. Bravo-Lillo et al. used a mental model approach to understand computer users' psychological processes and reactions toward computer warnings [8]. They were able to identify different perceptions of novice and advanced users and to obtain insights in improving computer warnings [8]. Most recently, Naiakshina et al. studied people's mental models of the security of mobile messaging tools and found that people overestimated the capabilities of attackers [29].

The above studies shed light on people's mental models of the Internet and privacy and security risks. However, people's mental models of OBA still remain unclear. Our study aims to address this gap by inductively analyzing people's understandings of how OBA works.

Our primary research question is what folk models people employ in practice about OBA, for instance, regarding the information flow in OBA. This was in part inspired by Helen Nissenbaum's theory of contextual integrity which presents a framework to determine privacy violations based on the norms and appropriateness of information flow in a particular context [30]. A secondary research question is what privacy-enhancing features or tools people desire for OBA. Answers to both questions will inform future privacy designs for OBA.

METHOD

We designed and conducted a qualitative study in an iterative manner to understand people's folk models of OBA. This study was approved by the IRB. We started with a pilot study to test the interview script and explore people's understandings of OBA. We then conducted a first-round of interviews to develop initial folk models, followed by a second-round of interviews to further verify the models.

Pilot Study

Drawing from prior research examining people's attitudes and perceptions of OBA [38, 27, 39, 23, 33, 10, 40, 28], we developed a list of interview questions that investigate people's understandings, attitudes, and experiences of OBA. To assess the quality of these questions, we pilot tested this interview protocol with eight family members and friends during January and February, 2016. The pilot results suggest that they understood the questions albeit most of them did not understand how OBA works. For instance, most of them did not know that third-party entities (e.g., ad networks) are likely involved in OBA. These pilot study participants' understandings of OBA were covered by the four folk models developed in the subsequent two rounds of interviews. For instance, many of them held the connected-first-party model.

The pilot results also suggest that they varied in their opinions of OBA after we explained the concept and that they differed in their interests in learning more about OBA and/or using tools to control OBA. In order to further identify their understanding of OBA (i.e., mental models) and their preferences of OBA, we added a drawing task and a card sorting task.

First-Round Interviews

We revised the interview protocol based on the feedback from the pilot study. Next, we describe the updated protocol.

Questions about Internet usage

We began our interviews with questions about interviewees' demographics such as age, gender, and occupation. We then asked about their background in using computers and the Internet, e.g., "What do you usually do when you browse the web? What devices do you use to browse the web?" We also asked about their usage of web browsers, e.g., "Do you know that you can change your browser settings? Do you know what a browser extension/add-on is? Do you save any of your account information in your browser? What kind of information do you save?"

We then asked them to sort 18 cards, each containing an information item (e.g., name or home address), based on their comfortableness with saving the data into their browser. This card sorting task was designed to assess their perceived sensitivity of different information. Most interviewees put the information items into two or three clusters based on their perceived sensitivity, for instance, social security numbers as highly sensitive and religion as mildly or moderately sensitive. Since these card sorting results mostly corroborate the

findings reported in the prior literature (e.g., [23, 20, 40]), we removed this task from the second-round of interviews.

Mental models of OBA

Next, we asked about interviewees' attitudes toward and interactions with online ads, e.g., "Do you notice that there are ads on websites? Do you generally click on ads?"

Similar to the use of hypothetical scenarios in Wash's mental model study [41], we presented a hypothetical ad scenario in which a user first looks for shoes in Amazon.com and a few hours later he or she visits Facebook and sees other shoe ads there. This scenario was designed to represent common OBA practices that interviewees can easily understand since Amazon and Facebook are popular sites that people visit. We then asked them to draw what they think happened in this scenario on a piece of paper and to explain their drawing. This drawing with think-aloud task explored interviewees' own understandings of how OBA works before we offered our definition and explanation of OBA. These drawings visualized the interviewees' folk models of OBA (i.e., their own theories of how OBA works).

We followed up with additional questions about their knowledge and understanding of OBA and web trackers, e.g., "Have your heard of targeted ads? Do you know how targeted ads work? Have you ever heard of web trackers? What do you think web trackers are, who they are and what they do?"

Then we offered the same explanation of web trackers to each interviewee. Specifically, we explained that the sites they visit voluntarily are first-party entities, and that web trackers are typically third-party entities which track user information and can provide ads targeted to the user based on the collected user data (e.g., browsing activities, page visits). We then answered any questions that interviewees had about web trackers. We also asked them "What do you think trackers are collecting when they are tracking you? What's more important to you, the trackers or the data is being tracked?"

Privacy-enhancing tools for OBA

Finally, to help inform future privacy design for OBA, we asked interviewees questions about their desired features in helping them deal with web trackers, e.g., "If there was a magic tool that can do anything, what types of features would you like this tool to have pertaining to web trackers?"

We asked these questions after explaining OBA with the rationale that if interviewees did not have a correct understanding of OBA, they may miss features that they would need or want. For instance, similar to what we found in the pilot study, many participants in this round of interviews were not aware that web trackers are often third-party entities. These participants requested the privacy tools to provide more information about OBA, including the third-party trackers involved. If we asked these tool-related questions before explaining OBA, these participants would not know the existence of third-party trackers and thus are unlikely to ask for corresponding tool support. However, asking these tool-related questions before explaining OBA might discover that people having different folk models desire different privacy features. Therefore, we asked these tool questions both before and after explaining OBA in the second-round interviews.

Second-Round Interviews

We analyzed the first-round interviews and developed four folk models that our participants had about how OBA works. Similar to the iterative methodology used in Wash's mental model study [41], we conducted a second-round of interviews with new participants to check the validity of these models by seeking "negative" examples [31] that are not covered by these models.

There were two major updates of the interview protocol in this round. First, we removed the card sorting task. Second, we asked the questions related to privacy tools both before and after explaining OBA. In other words, we updated the sequence of study components: questions about Internet usage, questions about mental models (with the same hypothetical scenario), questions about privacy tools, our explanation of OBA, and the questions about privacy tools (second time).

Participant Recruitment

We recruited prospective participants from a university campus, shopping malls, public libraries, and online communities (e.g., Craigslist). We also used snowball sampling, i.e., asking participants to refer our study to their contacts [5]. We deliberately selected participants in order to create a diversified sample in which participants have various demographic characteristics and occupational backgrounds.

From March to May 2016, we recruited and conducted our 1st-round of interviews with 14 participants from an urban area in the Eastern US. These interviews were face-to-face. From July to August 2016, we recruited and conducted our second-round of interviews with seven additional participants from another urban area in the Eastern US and two urban areas in the Western US. These interviews were conducted online using services such as Skype. Participants showed and explained their drawings in the interviews and sent their drawings to the researchers afterwards. Each interview took about one to two hours and was compensated $10.

It is worth noting that our sample is not statistically representative of the general Internet user population, but it is diverse in terms of participants' age, geographic locations and occupations. Similar to Wash's study [41], we do not believe our sample is particularly special. There are probably other people similar to our participants in the general population. In addition, we did not observe any significantly new findings, particularly regarding people's understandings of how OBA works, from our second-round of interviews. This suggests theoretical saturation [16] and thus we did not conduct any more interviews.

Data Analysis

We audio recorded all interviews upon participants' permission, and then transcribed the audio recordings. We then conducted a *thematic analysis* [7], a common approach for analyzing qualitative data.

First, we read through all the interview transcriptions multiple times to immerse ourselves in the data. Second, two

Table 1. Participants used three factors in reasoning about OBA and constructing their folk models.

Folk model	Who tracks info	Where info stored	How ads selected or provisioned
Browser-pull	Browser	Browser	Browser pulls ads
1st-party-pull	Browser	Browser	1st-party sites pulls ads
Connected 1st-party	1st-party	1st-party	1st-party sites share data directly and pull ads
3rd-party	1st-party	3rd-party	1st-party shares data with 3rd-party, 3rd-party pulls ads

co-authors coded one interview together at the sentence level to develop a code book.

Then, the two coders coded the same subset of interviews independently using the code book. When they encountered concepts not covered by the existing code book, they added new codes accordingly. Once finished, the two coders compared, discussed and converged the codes into an updated code book of 210 unique codes, such as, "Internet experience," "attitudes toward OBA," and "PETs features." We wrote the codes on post-it notes and created an affinity diagram to group these codes into nine themes: background, misconception, advertisement, specific information concerns, privacy-enhancing technologies, mental models, privacy and security practices, privacy expectations, and web trackers.

Finally, we read the associated interview quotes to ensure the coherence within each theme. Based on our review, we adjusted the inappropriately grouped codes and the affinity diagram accordingly. Both rounds of interviews were captured in this diagram.

RESULTS

In this section, we report the results from the 21 interviews, focusing on our participants' folk models of how OBA works and their preferences of privacy tools for OBA.

Participants

The ages of the 21 participants ranged from 19 to 67, with an average of 34. Six participants were female and 15 were male. They were from a wide range of locations, including large and small cities in the states of New York, Pennsylvania, California and Washington. Various occupations such as university staff, college students, software engineers, business professionals, retired workers, a mechanical engineer and a waitress were represented among the participants.

All of our interviewees use computers and the Internet on a daily basis. Two of them use the Internet less than 2 hours a day, the rest of them use the Internet more than 7 hours a day. The primary purposes of using the Internet include checking emails, using social media, doing research for their jobs, contacting friends and families, and reading news. In addition, 19 of our interviewees had heard about targeted ads. Some of them voluntarily talked about their experiences of targeted ads. Four interviewees said that they have heard of web trackers, but only one understood what a web tracker is.

Folk Models of OBA

We provided our interviewees a detailed scenario to understand their thoughts about how OBA works and how information flows. The interview results suggested that our participants' understandings of how OBA works mainly differed by three factors: who tracks users' information; where the information is stored; and how ads are selected or provisioned. Based on these three factors, we identified four major models. Table 1 summarizes the factors that our participants used to reason about OBA and construct their folk models. Table 2 summarizes participants' folk models as well as their attitudes toward web trackers and OBA.

Browser-Pull Model

Five interviewees held this model. They believed that all tracking is done by the browser, which would pull from advertisers relevant ads that target user data/profiles stored locally by the browser. In this model, the web browser plays the primary role in OBA. For instance, P5 thought that the web browser monitors and detects his browsing patterns and pulls ads based on those patterns. He also believed that all tracked information is saved in his local computer.

Table 2. Participants' folk models and attitudes of trackers and OBA.

ID	Folk model	Accept trackers	Accept OBA
P1	3rd-party	Yes	Yes
P2	Connected 1st-party	No	Yes
P3	3rd-party	Yes	Yes
P4	3rd-party	No	No
P5	Browser-pull	No	Yes
P6	Connected 1st-party	No	No
P7	Connected 1st-party	No	No
P8	3rd-party	No	No
P9	Browser-pull	No	No
P10	1st-party-pull	Yes	Yes
P11	Browser-pull	Yes	Yes
P12	1st-party-pull	Yes	Yes
P13	Browser-pull	Yes	Yes
P14	Connected 1st-party	No	No
P15	Browser-pull	No	No
P16	1st-party-pull	Yes	Yes
P17	3rd-party	Yes	Yes
P18	Connected 1st-party	No	Yes
P19	3rd-party	Yes	Yes
P20	3rd-party	Yes	Yes
P21	1st-party-pull	Yes	Yes

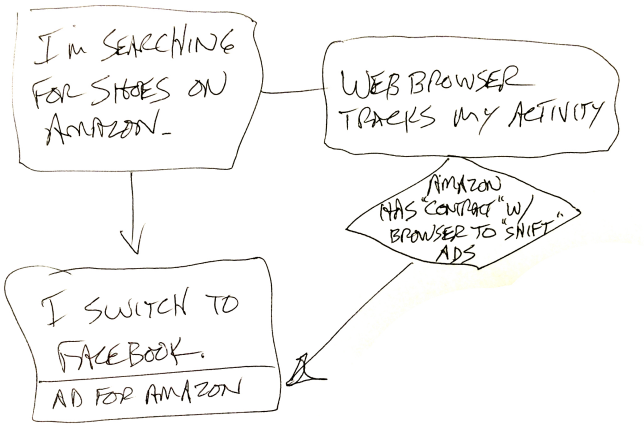

Figure 1. Browser-pull model: an example from P9. When a user searches for a pair of shoes on Amazon, the web browser will save the search information. The web browser has contracted with Amazon. When the user visits Facebook, the browser will pull the saved information and display ads for Amazon on the user's Facebook page.

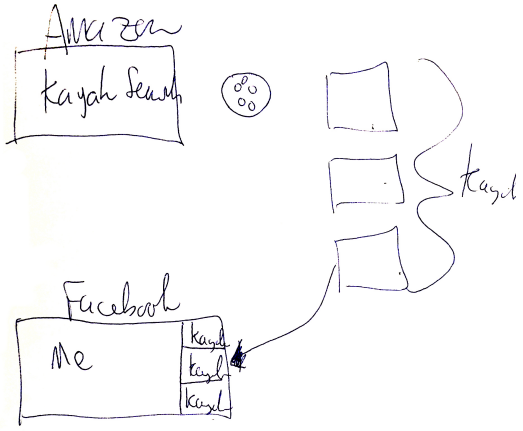

Figure 2. First-party-pull model: an example from P10. When a user searches for a kayak on Amazon, the browser will save the search information in a browser cookie. The browser will find other sites that also sell kayaks. Later, the user visits Facebook, which will pull these sites from the browser and display them on the user's Facebook page.

"The system is set up to notice your patterns and to pull information that seems relevant to you...I'm just thinking [the information] is [transmitted to] my computer." (P5)

P9 had a similar view as illustrated in his drawing (see Figure 1) in which the browser tracks his online activities and has contracted with Amazon to ship their ads. He explained,

"I'm searching on Amazon and looking for shoes, web browser tracks my activity, and, you know, I'm just thinking that Amazon and ads are contracted with web browser, and browser just ships ads. There's when I'm on Facebook, the ads just pops up." (P9)

P15 also held this model but also felt that he can control the browser's tracking through the browser settings. He said,

"I think it is all based on your Internet options what you allow. I think it is the browser that allows this...No matter whatever browser I'm on...I can go to the Internet options and mess around the way it looks into my information." (P15)

The essence of this model is that the web browser is key – the browser tracks users' activities, saves their information on the local computer, and selects and displays the relevant ads. Because the browser is on users' computers, some participants holding this model (e.g., P15) also had the perceived agency to limit or control OBA through the browser settings.

First-Party-Pull Model

Four interviewees held the first-party-pull model. Similar to the browser-pull model, participants of this model also believed that all tracking is done by the browser. However, unlike the browser-pull model, people of this model thought that first-party sites (e.g., Amazon or Facebook) rather than the browser pull relevant ads based on the user's data/profile stored in the browser. In this model, both the web browser and first-party websites play active roles in OBA.

For instance, P10 explained the use of cookies and the retrieval of targeted ads by the first-party sites, as shown in his drawing (see Figure 2). In this example, when he searches for kayaks on Amazon, the browser will save the search information in a browser cookie. Then, the browser will find other sites that also sell kayaks. When he visits Facebook later, Facebook will pull these sites that sell kayaks from the browser and display them on his Facebook page. Here again, the first-party site (Facebook) pulls the relevant ads. In addition, P10 believed that first-party websites can only access the cookies from the last website that the user visited. P12 shared a similar model but described his theory in a more technically sophisticated way, highlighting the use of the HTML meta tag on first-party sites (e.g., eBay or Facebook). He believed that these websites are designed in a way (with similar meta tag structures) so that they can directly access all of the user's browsing/searching history and cookies in order to select targeted ads.

"So this is the eBay webpage, and in your meta-tag you're gonna have embedded information that not only pulls up the information from your cookie and consent your account to automatically login...but it also contains advertising tracking data...And then if you log into, for instance, Facebook, if they have a similar meta-tag structure they can access the search data from this tracking cookie, so that this controls the same search criteria." (P12)

P16 is a web developer with technical knowledge of the Internet. His drawing (see Figure 3) illustrated that the browser stores the user's Amazon activities in its local cache; then Facebook pulls that user information from the cache, bids ads with that user information, and finally displays the targeted ads on the user's Facebook page. He was our only participant who mentioned ad bidding, which suggests that he had more knowledge about the online ad ecosystem than other participants. However, he was not aware of third-party tracking

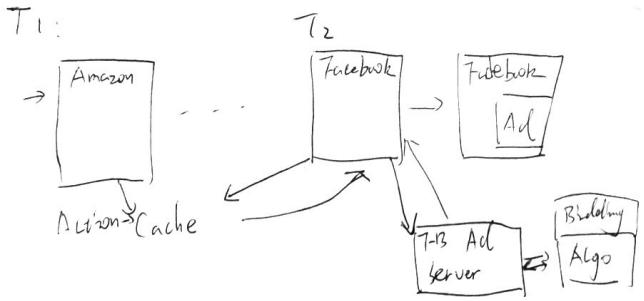

Figure 3. First-party-pull model: an example from P16. He searches shoes on Amazon, then the browser stores the action in the cache; later, he visits Facebook, which then goes to its advertising server, and bids ads with the cached user information using a bidding algorithm. Facebook then displays the ads on his Facebook page.

in OBA. After he described his mental model, we explained OBA to him and asked about his feelings of OBA again. He was surprised to learn that often third-party entities (e.g., ad networks) track users' online activities for OBA purposes and his concerns about his information and privacy remained the same. He said,

"I never thought it is third party. It matters to me since I'm not sure how they can use my information legally and the purposes. But I think I feel the same, because before I know this, I'm still worried about my information and privacy." (P16)

P16's example is rather telling because we would normally assume that technically savvy users would know how OBA works. But, that was not the case. This is somewhat surprising because even web developers like him who seemed to have knowledge about ad bidding (a rather advanced understanding of online ad systems) did not know about common online tracking done by third-party entities.

Other less technically savvy participants also held this model albeit with less details. For instance, P21 thought that the browser records his activities on Amazon and then Facebook pulls his information from the browser. However, unlike P10 and P16, he did not know the technical specifics of how Facebook can actually access his information stored by the browser. He commented,

"Chrome gets all my transaction from Amazon, and for some reason, Facebook can access this information." (P21)

While these participants provided different levels of technical details in their explanations, the underlying theory is the same. In this model, the browser still plays an important role in tracking users' browsing activities and storing this information locally. But, first-party sites (e.g., Facebook) rather than the browser select relevant ads based on the user profile stored in the browser. This means, unlike the browser-pull model, the browser cannot single-handedly deliver the targeted ads. Instead, first-party sites select ads that they think are relevant to users.

Connected-First-Party Model

Five interviewees believed that first-party websites (e.g., Amazon and Facebook) are directly connected and collaborate with each other. In this model, users' data is tracked and stored by each first-party site that they visit. Different first-party websites are connected, exchanging the user data that each of them tracks and saves. As such, first-party websites form a kind of a collaborative network and play the main role for delivering targeted ads.

For instance, P18 believed that first-party sites have shared resources between them so they can share their user information stored in their databases.

"It definitely goes into the database on Amazon, and then it will probably, I guess there is some kind of shared resource between them, so it will basically go into the database on Facebook, then shows on my page." (P18)

P7 shared the idea and also explicitly mentioned a partnership between Amazon and Facebook, which makes the information sharing possible.

"Amazon and Facebook have some type of partnership and so Amazon gives them certain information and are able to locate certain people for specific products and advertise certain products to certain people for whatever they're looking for."

P6 went further and suggested that Amazon sells her data to Facebook, as shown in her drawing (see Figure 4). She said,

"Amazon has decided to work in conjunction with Facebook, this is my personal belief and transferred all of my data about my shoe experience and gave it to Facebook, to say 'Hey, she likes shoes.' And now Facebook is going to get bulk dollars from other shoe companies because now I know about other places beside or other shoe types beside the one I just bought...I think amazon sells that to Facebook." (P6)

P6 emphasized the economic or business model in her understanding. Her drawing shows that Facebook gives money to Amazon (with one dollar sign) for her data that Amazon shares and Facebook receives more money from other shoe companies (with two dollar signs) to serve their ads on her

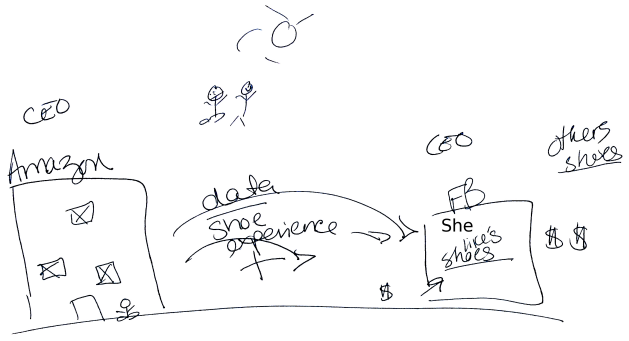

Figure 4. Connected-first-party model: an example from P6. When she searches shoes on Amazon, Amazon saves the search information. Amazon sells the data to Facebook (indicated by a dollar sign at the bottom left of the Facebook box). Facebook then gets more money from *"other shoe"* companies (two dollar signs) to show the shoe ads on her Facebook. The CEOs of Amazon and Facebook under the sun would not share her data if it is not for money.

Facebook page. As such, P6 believed that money drives the connection between Amazon and Facebook. She disapproved an alternative explanation in a witted fashion and articulated money as the driving force behind this connection.

"I don't see why Amazon would do this because I don't see like the CEO of Amazon and the CEO of Facebook hanging out under the sun as best friends smiling...so there's got to be a reason...the biggest lubricant I ever come across is money, or at least some kind of gain of some sorts." (P6)

The key of this model is that first-party sites are directly connected and they share user data with each other in order to select targeted ads. According to this model, the connected first-party websites enable OBA, regardless of the reasons for their connections (e.g., a partnership or user data purchases).

Third-Party Model

Seven interviewees held this model. In this model, people believed that first-party sites track and collect user data then contribute the data to a third-party entity, and then the third-party entity leverages the user data it has (presumably from different first-party sites) to select relevant ads for users. As such, various first-party websites and third-party entities are involved in OBA, according to this model.

Some participants believed that there are third-party entities involved but they knew almost nothing else about these third-party entities, for instance, who they are or whom they belong to. For example, P4 drew a big bubble that she called an "Internet space" that stores and provides user data to different sites such as Facebook (see Figure 5). But, she cannot tell what this Internet space is or who controls it.

"I don't know, it must be like some Internet thing, Internet space I don't know, and somehow it just goes to like Facebook and whatever else there is out there." (P4)

P19 drew a more detailed graph, illustrating the existence of some database that all companies such as Amazon and Google share (see Figure 6). But, he knew nothing else about this database.

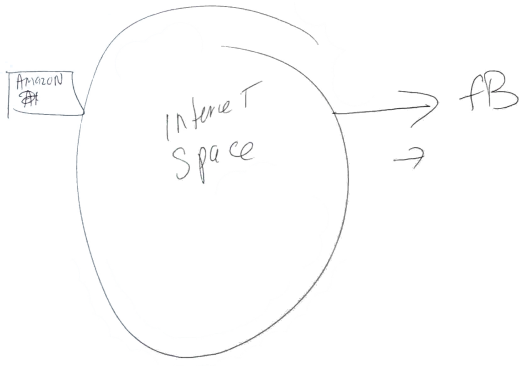

Figure 5. Third-party model: an example from P4. When a user searches shoes on Amazon, Amazon collects the user's data and then transmits the data to an Internet space. This Internet space sends ads to Facebook, which will display the shoe ads on the user's Facebook.

"I don't really know. I guess there should be some sort of database in the middle, then not only Amazon and Facebook, but also other companies, have access to it, keep injecting new information to it. It's more of a shared space, or common space for all companies who are involved in this ecosystem."

P19 also questioned how Amazon and Facebook match the same user. He hypothesized the use of cookies, which include a user's IP address. He also doubted first-party sites (Amazon and Facebook in our scenario) are directly connected. This is an important difference from the connected-first-party model in which first-party sites are directly connected.

P17 also believed there is a central database and he thought it is dominated by Google.

"There must be some central database or data center...I think it's like Google. Google has something like this, like many big companies have this kind of data center. But it is dominated by Google. In the example you mentioned, there is no Google involved, so I guess it is third party." (P17)

In this case, he suggested that large companies like Google represent the third-party entities. This understanding was fairly accurate since Google indeed represents a major web tracker and serves targeted ads across the Internet [36].

Like P6 who had a connected-first-party model, P1 also focused her understanding on the economic aspects. However, P1 believed that those third-party entities rather than the connected first-party sites make the ads ecosystem work.

"These guys [third parties] have an agreement with Amazon, they are like, 'Oh, I'm just going to take information from this guy'. Facebook gets money by displaying the ads sent by these guys [third parties]...this branch [third parties] allows that to happen. So in a way it is a neutral third party." (P1)

In her view, the third-party entities connect Amazon and Facebook, collect and store users' data, and then send targeted ads to Facebook.

Regardless of whether these participants knew who the third-party entities are or represent, they shared the key understand-

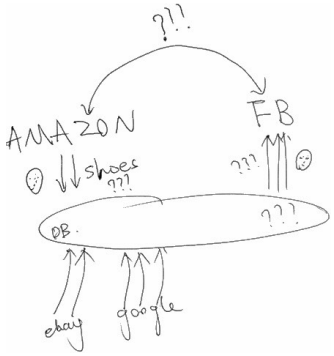

Figure 6. Third-party model: an example from P19. All companies share a common database. When a user searches shoes on Amazon, Amazon sends that information to the shared database. Other companies such as eBay and Google also contribute user information to this database. When the user visits Facebook, the site obtains user data from this database to select relevant ads.

ing that these third-party entities rather than the first-party sites that users visit voluntarily make the OBA work. Both user tracking and selection of targeted ads are done by these third-party entities. According to this model, first-party sites are not connected directly but are bridged through third-party entities.

Misconceptions and Speculations of OBA

During our interviews, we also observed participants' recurring misconceptions and speculations about OBA. We use the word "misconceptions" to denote our participants' inaccurate understandings of web trackers and what trackers collect, mainly from a technical standpoint. Typical misconceptions our participants had include: trackers are hackers, and trackers are viruses. Wash's home computer security study has uncovered several hacker-based and virus-based mental models [41], however, his participants did not report considering web trackers as hackers or viruses. Furthermore, we use the word "speculations" to represent our participants' views that are technically possible but their applications for OBA are not clear. Common speculations our participants made include: trackers access local files on a user's computers, and trackers resides locally on users' computers.

Misconception: trackers are hackers

Some interviewees identified trackers as hackers, people with malicious intentions. For instance, P2 believed that web trackers can hack into his online accounts.

"They say it's a secure site and you got to login, but of course I login with the same password that I always use...I'm sure that those web trackers can hack in there too." (P2)

P2 seemed to confuse web trackers with hackers that aim to break into people's accounts and steal their personal data.

Misconception: trackers are viruses

Considering trackers as computer viruses was another common misconception among our participants. For instance, P4 expressed her belief that her anti-virus software will protect her from trackers.

"Thank God for Norton because sometimes it comes up oh so and so just attacked you or something, so I don't even pay attention because I figure that will save me." (P4)

These participants seemed to misconstrue web trackers as computer viruses designed to attack their computers.

Speculation: trackers access local files

Some participants thought that trackers can access files stored on their local computers. For instance, when asked whether he would be interested in a tool that can block trackers, P5 expressed his lack of interest in such a tool because there is very little on his computer that he worries about. He mentioned, *"Even things that are around my desktop, besides my resume and cover letters and that's about it."* His explanation reflected his overestimation of the capabilities of trackers in which they can access (arbitrary) files stored locally on his computer. P5 also believed that trackers can log his typing, saying *"everything you type in can technically be downloaded."* While tracking users' typing is technically possible, we are not aware of any reports of this kind of tracker behavior in practice.

Speculation: trackers reside locally on user computers

Some participants indicated that trackers can not only be something in the browser but also reside locally on their computers. For instance, P5 said *"I think it's in the web browser. I also think there's something on your computer."* But he could not elaborate what he meant by "something" on his computer.

Privacy-Enhancing Tools for OBA

To inform future design of privacy tools for OBA, we asked our interviewees questions about tools or features that can help protect their privacy in the context of OBA.

Trackers vs. the information being tracked

Existing ad blockers such as Ghostery are structured by trackers. When a user visits a website, the ad blocker shows a list of trackers on the site that the user can selectively block. However, these tools do not show what type of information each tracker tracks. In addition, prior research has shown that ordinary Internet users do not recognize the names of most trackers (e.g., BlueKay) with few exceptions being household names such as Google [22]. Furthermore, our card sorting results support the prior literature (e.g., [23, 20, 40]) that people perceive different levels of sensitivity for different information items (e.g., home address is perceived more sensitive than educational level). Given these observations, we wondered whether the information being tracked is more recognizable and thus more useful to users than the trackers. Therefore, we asked our participants *"what is more important to you, the tracker or the information being tracked?"*

All but one interviewee answered that the information being tracked is more important. For instance, P1 cared more about the information being tracked because this information can be used to make assumptions about her.

"I would say what is being tracked. I guess they use the information to build out their profile, I guess it is a little strange using the information they collect to make assumptions about me, what type of person or Internet user." (P1)

P7 provided a different justification, arguing that the tracked information can be used to identify individuals.

"I mean the biggest thing is the information. I mean trackers are replaceable, but information is not because that's a specific set of info per person." (P7)

P8 was the only participant that did not perceive the information being tracked to be more important than the trackers because he valued and wanted to know both.

"What information is being collected for sure, but I also want know who is collecting it. I want to say both, because, you know, I would want to know who that person, or the entity is, how they are gonna use that information." (P8)

Desired privacy features for OBA

When asked about their expectations of a magic tool that can help protect their privacy regarding OBA, our participants suggested many features.

Block tracking. A commonly desired feature is to block tracking. For instance, P17 would like to automatically block trackers based on his preferences. P16 desired a feature that allows him to select the type(s) of information that he wants the trackers to track or not to track.

Interestingly, when we asked participants' experiences with online ads, some participants reported using ad blockers to block ads but they did not relate these ad blockers to web trackers. This might be because they are called ad blockers rather than tracker blockers.

Transparency. Several interviewees were also interested in knowing more about trackers and their behaviors. For instance, P1 commented,

"it is a scary technology, but maybe if I have a better understanding of how it connects with companies or something like that. Maybe I can see like the scope of web trackers? Like how many people it affects, how many places my information is going." (P1)

P1 hoped to know detailed information about the scope and effect of tracking. In addition, P19 was interested in knowing what data is being tracked by whom and for what purposes.

In our second-round interviews, we asked this privacy tool-related question both before and after we explained OBA. P18 held a connected-first-party model and requested additional privacy tool support after our OBA explanation, which made him realize the existence of third-party trackers. He then suggested the tool to provide detailed information about third-party trackers and their behaviors.

Effortless to use. In addition to concrete features, many interviewees emphasized the tool should be effortless to use. P10, for instance, expressed that he would only use such a tool if it only needs a one-time setup for all websites.

"This is per website or do I do it one time and it does it for every website? That was my first thing cause I don't want to have to do it per website." (P10)

This is understandable because privacy protection is often not people's primary or direct task. Therefore, they would not want to divert from their main task to spend too much time in using a privacy tool. For example, automatic blocking of tracking as suggested by P17 would satisfy this criterion.

DISCUSSION

Drawing from the literature on mental models and particularly Rick Wash's work on folk models of home computer security [41], we examine Internet users' understandings of how OBA works through a qualitative study including a pilot study and two rounds of semi-structured interviews.

We discover four folk models of how OBA works. The *browser-pull* model assumes that all tracking is done by the browser, which would pull from advertisers relevant ads that tailor to the user data/profile the browser stores locally. In this case, the browser is the "middleman" between the first-party site and advertisers. The *first-party-pull* model presumes that all tracking is still done by the browser, but first-party sites pull relevant ads based on the user data/profile stored in the browser (e.g., cookies). In this case, first-party sites decide which ads to show. The *connected-first-party* model posits that different first-party sites directly share and even sell user data that they collect and one first-party site can use another first-party site's user data to pull relevant ads directly from advertisers. In this model, first-party sites directly interact with each other and with the advertisers. Lastly, the *third-party* model assumes that first-party sites first track and collect user data then contribute the data to a third-party entity, then this third-party entity uses the user data it has (presumably from different first-party sites) to select relevant ads. This model is closer to common OBA practices than other models but it is still not detailed enough, e.g., some participants hardly knew anything about the third-party entities.

As discussed in the related work section, our work is one of the first studies that investigate people's mental models of OBA. The body of literature on mental models of privacy and security rarely touches on the topic of online tracking or OBA, for instance, Wash's study focuses on home computer security [41]. The extant research on people's privacy perceptions of OBA does not focus on people's understanding and mental models of OBA. The notable exception is the work of McDonald and Cranor in which they provided their survey respondents four diagrams of OBA, focusing on who have access to users' cookies [27]. They then asked their respondents which diagram is unlikely to happen [27]. In comparison, our folk models emerged from our interviews rather than pre-defined by us. Our folk models differ from their cookie-centered models [27] because ours are based on three factors: who tracks user information, where the tracked information is stored, and how the targeted ads are selected or provisioned.

Why Folk Models of OBA Matter

The folk models uncovered by our study are novel, but why do they matter? There are several reasons why they matter.

User education

All four folk models are either inaccurate or incomplete. Similar to Camp's suggestions that risk communication should be designed based on non-expert mental models [9], we believe that it would be useful to customize user education of OBA based on the folk models.

In our second round of interviews, some of our interviewees changed their attitudes towards OBA because of our explanation of OBA. For instance, some participants of the connected-first-party model were surprised to learn that their information can be tracked or even sold by third-party entities. Therefore, their attitudes towards OBA were changed from neutral to negative. Knowing a user's current folk model can tailor the education to reduce the knowledge asymmetry between the user and the OBA practices.

Previous studies have suggested that technically savvy users have more accurate or sophisticated mental models than their less technically savvy counterparts (e.g., [19]). However, we did not observe a clear relationship between technical knowledge and folk models. Somewhat surprisingly, our arguably most technically savvy participants P16 and P18, two web

developers, held the 1st-party-pull model and the connected-first-party model, respectively. Both of them were not aware of third-party trackers. This is important because even technically savvy users can have inaccurate or incomplete models and need user education to gain a more accurate picture of OBA.

Attitudes towards OBA

Capturing people's folk models can help understand people's attitudes towards OBA. We observed some associations between the two.

Interviewees of the browser-pull model had different attitudes toward tracking and OBA. These participants believed that the browser tracks and stores their data. Interviewees who were aware of different browser settings (e.g., clear browser history and cookies) tended to be positive about OBA because of their perceived ability to control tracking by setting the browser options. In contrast, those who did not know about browser settings tended to be critical of OBA.

Interviewees of the first-party-pull model generally accept OBA because they only expected first-party sites to access their information in order to select relevant ads. They had little concern because they generally trusted the sites that they visit voluntarily. However, they were unaware of the existence and impact of third-party tracking.

For interviewees holding the connected-first-party model, they were generally not against online tracking because they thought their data is only shared between first-party sites that they trust. However, they did not appreciate the idea of first-party sites selling their information between each other. They understood that this is one of the main business models of the Internet, but they still disliked it.

Participants of the third-party model all included third-party entities in their explanations. However, their descriptions of third-party entities varied significantly, ranging from a clear idea of a specific organization to a vague notion of an *"Internet space."* Their attitudes toward OBA also varied. We did not observe any significant patterns in this group.

User behavior

The literature of mental models suggest that these models can influence people's reasoning and decision making (e.g., [17, 18]). We also encountered some examples of certain folk models affecting people's behavior in our study. For example, some participants of the browser-pull model rely on browser settings to control online tracking because they believed that web tracking and OBA are carried out by the browser. Another example is that P18 of the connected-first-party model requested a transparency feature that provides detailed information about third-party trackers only after we explained OBA. This suggests that people of different folk models may need different privacy features (particularly educational features) that tailor to their (lack of) understanding.

Implications for Design and Policy

Our results have a number of implications for privacy designs and public policies of OBA.

First, as mentioned before, future privacy tools for OBA could highlight different information to cater to people with different folk models. For example, for people having the connected-first-party model, the tools can emphasize that third-party entities can be tracking and sharing their online activities. In addition, governmental policies or industry best practices could require or encourage privacy policies of web tracker companies to include simple but visual representations of how they work in the OBA ecosystem, similar to the way that our interviewees drew their folk models.

Wash argued that technologies should be designed to work with people's mental models even if these models are incorrect because it is more difficult to educate users about the correct mental model [41]. We agree with this viewpoint to some extent. For instance, while the browser-pull model does not capture the common OBA practice, researchers have proposed privacy-preserving, client-based OBA systems, resembling the browser-pull model [6]. However, we still believe there are benefits to educate people about OBA practices that are common on the Internet. For instance, people holding the browser-pull model might think they can control or stop OBA by just setting their browser options. Therefore, that folk model could discourage them from adopting more effective privacy tools such as ad blockers that can block third-party trackers.

Second, popular tools such as Ghostery and AdBlock are capable of blocking third-party trackers. These tools list the trackers on a site and allow people to block them selectively. However, most of our interviewees felt that the information being tracked is more important than the trackers themselves. This is a significant finding because it suggests that a completely different blocking scheme, one based on the type of information being tracked, might be perceived more useful by Internet users than the status quo, a tracker-based blocking scheme. In other words, the tools can be structured by the information being tracked rather than by a list of trackers. In addition, these tools can allow users to block tracking of certain types of information. Alternatively, future tools can support both schemes.

Emerging technologies, such as OpenWPM [14], Sunlight [21] and ReCon [34] are promising in identifying or inferring what information is being tracked by a tracker and the purpose of tracking to some extent. They pave the way for information-based blocking tools. On the policy front, we advocate that web trackers and ad networks should clearly explain what information they collect and why they collect them in their privacy policies and preferably in a machine-readable format. This could enable future privacy tools that automatically analyze and compare the behaviors of different trackers and the OBA practices of different sites.

Limitations and Future work

We outline our study limitations and directions for future work. First, we did not have a particularly large sample. But our study was conducted in an iterative manner including a pilot study with eight people and two rounds of interviews with a total of 21 participants. The results from the pilot study and the actual interview study were consistent. In fact, we did

not learn any significantly new things from our second-round interviews, suggesting theoretical saturation. Our sample is also diversified in that our participants came from various age groups and geographical areas, representing different occupations. Therefore, we are confident our results are valid.

Second, our qualitative study aims to examine people's folk models of OBA in depth rather than assess how statistically representative these models are in the generic population. In future work, we plan to conduct a large-scale survey to further examine how common these models are.

Third, when we asked our interviewees to draw their mental models of OBA and web tracking, we only used one hypothetical scenario. This may prevent us from discovering additional models. Future work can include multiple scenarios and ideally ones that people have experienced themselves.

Fourth, we asked participants to do the card sorting task before the drawing task in our first-round interviews. The card sorting task asked about participants' comfortableness with saving their data into their browser. This might prime people to think more about browsers. However, we believe the priming is minimum because we removed the card sorting task in our second-round interviews and there were participants having the browser-pull model and the first-party-pull model. In both models, the browser is responsible for tracking users.

Fifth, our interviews are self-reported data and thus do not include participants' actual behavioral data. To further examine the impact of these folk models on people's behavior, future work can consider collecting and analyzing user behavior data, for instance, through experiments and/or log analyses.

Finally, our study did not test a concrete privacy tool. However, we have learned a great deal about people's desired privacy features for OBA. We plan to implement some of these features such as information-based blocking.

Conclusion

Online Behavior Advertising is pervasive on the Internet. We interviewed 21 people from the US to investigate their understandings of how OBA works. We identified four folk models held by our interviewees. These models are either inaccurate or incomplete in representing common OBA practices. User education tailoring to people's folk models of OBA is likely to be more effective. In addition, most of our interviewees felt that the information being tracked is more important than the trackers. Future privacy tools should consider these folk models and user preferences of OBA.

REFERENCES

1. Alessandro Acquisti and Jens Grossklags. 2005. Privacy and Rationality in Individual Decision Making. *IEEE Security & Privacy* 3, 1 (2005), 26–33.
2. Anne Adams and Martina Angela Sasse. 1999. Users Are Not the Enemy. *Commun. ACM* 42, 12 (1999), 40–46.
3. Farzaneh Asgharpour, Debin Liu, and L. Jean Camp. 2007. Mental Models of Security Risks. In *Financial Cryptography and Data Security*, Sven Dietrich and Rachna Dhamija (Eds.). Number 4886 in Lecture Notes in Computer Science. Springer Berlin Heidelberg, 367–377.
4. Howard Beales. 2010. The value of behavioral targeting. *Network Advertising Initiative* (2010).
5. Patrick Biernacki and Dan Waldorf. 1981. Snowball Sampling: Problems and Techniques of Chain Referral Sampling. *Sociological Methods & Research* 10, 2 (Nov. 1981), 141–163.
6. Mikhail Bilenko, Matthew Richardson, and Janice Tsai. 2011. Targeted, Not Tracked: Client-Side Solutions for Privacy-Friendly Behavioral Advertising. Rochester, NY. http://papers.ssrn.com/abstract=1995127
7. Richard E. Boyatzis. 1998. *Transforming Qualitative Information: Thematic Analysis and Code Development*. SAGE.
8. Cristian Bravo-Lillo, Lorrie Faith Cranor, Julie Downs, and Saranga Komanduri. 2011. Bridging the Gap in Computer Security Warnings: A Mental Model Approach. *IEEE Security and Privacy* 9, 2 (March 2011), 18–26.
9. L. J. Camp. 2009. Mental models of privacy and security. *IEEE Technology and Society Magazine* 28, 3 (2009), 37–46.
10. Farah Chanchary and Sonia Chiasson. 2015. User Perceptions of Sharing, Advertising, and Tracking. In *Proceedings of the Eleventh Symposium On Usable Privacy and Security (SOUPS 2015)*. 53–67.
11. Federal Trade Commission and others. 2009. FTC staff report: Self-regulatory principles for online behavioral advertising, 2009. *Federal Trade Commission, Washington, DC* (2009).
12. K. J. W. Craik. 1967. *The Nature of Explanation*. Cambridge University Press.
13. Janna Lynn Dupree, Richard Devries, Daniel M. Berry, and Edward Lank. 2016. Privacy Personas: Clustering Users via Attitudes and Behaviors Toward Security Practices. In *Proceedings of the 2016 CHI Conference on Human Factors in Computing Systems (CHI '16)*. ACM, New York, NY, USA, 5228–5239.
14. Steven Englehardt, Chris Eubank, Peter Zimmerman, Dillon Reisman, and Arvind Narayanan. 2015. OpenWPM: An automated platform for web privacy measurement. (2015).
15. Alan Garnham and Jane Oakhill. 1996. *Mental Models In Cognitive Science: Essays In Honour Of Phil Johnson-Laird*. Psychology Press.
16. Barney G. Glaser and Anselm L. Strauss. 2006. *The discovery of grounded theory: strategies for qualitative research*. Transaction Publishers.
17. Philip Johnson-Laird, Vittorio Girotto, and Paolo Legrenzi. 1998. Mental models: a gentle guide for outsiders. *Sistemi Intelligenti* 9, 68 (1998), 33.

18. Natalie A. Jones, Helen Ross, Timothy Lynam, Pascal Perez, and Anne Leitch. 2011. Mental models: an interdisciplinary synthesis of theory and methods. *Ecology and Society* 16 (March 2011), 1–13.
19. Ruogu Kang, Laura Dabbish, Nathaniel Fruchter, and Sara Kiesler. 2015. My Data Just Goes Everywhere: User Mental Models of the Internet and Implications for Privacy and Security. In *Eleventh Symposium On Usable Privacy and Security (SOUPS 2015)*. 39–52.
20. Bart P. Knijnenburg, Alfred Kobsa, and Hongxia Jin. 2013. Dimensionality of information disclosure behavior. *International Journal of Human-Computer Studies* 71, 12 (2013), 1144–1162.
21. Mathias Lecuyer, Riley Spahn, Yannis Spiliopolous, Augustin Chaintreau, Roxana Geambasu, and Daniel Hsu. 2015. Sunlight: Fine-grained Targeting Detection at Scale with Statistical Confidence. In *Proceedings of the 22nd ACM SIGSAC Conference on Computer and Communications Security*. ACM, 554–566.
22. Pedro Leon, Blase Ur, Richard Shay, Yang Wang, Rebecca Balebako, and Lorrie Cranor. 2012. Why Johnny can't opt out: a usability evaluation of tools to limit online behavioral advertising. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems*. ACM, 589–598.
23. Pedro Giovanni Leon, Blase Ur, Yang Wang, Manya Sleeper, Rebecca Balebako, Richard Shay, Lujo Bauer, Mihai Christodorescu, and Lorrie Faith Cranor. 2013. What matters to users?: factors that affect users' willingness to share information with online advertisers. In *Proceedings of the Symposium on Usable Privacy and Security*. ACM, 7–26.
24. Jialiu Lin, Shahriyar Amini, Jason I. Hong, Norman Sadeh, Janne Lindqvist, and Joy Zhang. 2012. Expectation and Purpose: Understanding Users' Mental Models of Mobile App Privacy Through Crowdsourcing. In *Proceedings of the 2012 ACM Conference on Ubiquitous Computing (UbiComp '12)*. ACM, New York, NY, USA, 501–510.
25. John Lofland, David A. Snow, Leon Anderson, and Lyn H. Lofland. 2005. *Analyzing Social Settings: A Guide to Qualitative Observation and Analysis* (4 edition ed.). Cengage Learning, Belmont, CA.
26. Jonathan R. Mayer and John C. Mitchell. 2012. Third-Party Web Tracking: Policy and Technology. In *Proceedings of the 2012 IEEE Symposium on Security and Privacy (SP '12)*. IEEE Computer Society, Washington, DC, USA, 413–427.
27. Aleecia McDonald and Lorrie Faith Cranor. 2010. Beliefs and behaviors: Internet users' understanding of behavioral advertising. TPRC.
28. William Melicher, Mahmood Sharif, Joshua Tan, Lujo Bauer, Mihai Christodorescu, and Pedro Giovanni Leon. 2016. (Do Not) Track Me Sometimes: Users Contextual Preferences for Web Tracking. *Proceedings on Privacy Enhancing Technologies* 2016, 2 (2016), 135–154.
29. Alena Naiakshina, Anastasia Danilova, Sergej Dechand, Kat Krol, M. Angela Sasse, and Matthew Smith. 2016. Poster: Mental Models-User understanding of messaging and encryption. In *Proceedings of European Symposium on Security and Privacy*. http://www.ieee-security.org/TC/EuroSP2016/posters/number18.pdf
30. Helen Nissenbaum. 2004. Privacy as contextual integrity. *Wash. L. Rev.* 79 (2004), 119.
31. Anthony J. Onwuegbuzie and Nancy L. Leech. 2006. Validity and Qualitative Research: An Oxymoron? *Quality & Quantity* 41, 2 (May 2006), 233–249.
32. Kristin Purcell, Joanna Brenner, and Lee Rainie. 2012. Search engine use 2012. (2012).
33. Emilee Rader. 2014. Awareness of Behavioral Tracking and Information Privacy Concern in Facebook and Google. In *Symposium on Usable Privacy and Security (SOUPS)*. 51–67.
34. Jingjing Ren, Ashwin Rao, Martina Lindorfer, Arnaud Legout, and David Choffnes. 2016. ReCon: Revealing and Controlling Privacy Leaks in Mobile Network Traffic. In *Proceedings of The 14th ACM International Conference on Mobile Systems, Applications, and Services (MobiSys)*.
35. Roy G. D'Andrade. 1995. *The Development of Cognitive Anthropology*. Cambridge University Press.
36. Steven Englehardt and Arvind Narayanan. 2016. *Online tracking: A 1-million-site measurement and analysis*. Technical Report. Princeton University. http://randomwalker.info/publications/OpenWPM_1_million_site_tracking_measurement.pdf
37. Andrew Thatcher and Mike Greyling. 1998. Mental models of the Internet. *International journal of industrial ergonomics* 22, 4 (1998), 299–305.
38. Joseph Turow, Jennifer King, Chris Jay Hoofnagle, Amy Bleakley, and Michael Hennessy. 2009. Americans reject tailored advertising and three activities that enable it. *Available at SSRN 1478214* (2009).
39. Blase Ur, Pedro Giovanni Leon, Lorrie Faith Cranor, Richard Shay, and Yang Wang. 2012. Smart, useful, scary, creepy: perceptions of online behavioral advertising. In *Proceedings of the Eighth Symposium on Usable Privacy and Security (SOUPS'12)*. ACM, 4–19.
40. Yang Wang, Huichuan Xia, and Yun Huang. 2016. Examining American and Chinese Internet Users Contextual Privacy Preferences of Behavioral Advertising. In *Proceedings of ACM Conference on Computer-Supported Cooperative Work and Social Computing (CSCW 2016)*.
41. Rick Wash. 2010. Folk models of home computer security. In *Proceedings of the Sixth Symposium on Usable Privacy and Security*. ACM, 11.
42. Willett Kempton. 1986. Two theories of home heat control. *Cognitive Science* 10 (1986), 75–90.