

Examining American and Chinese Internet Users' Contextual Privacy Preferences of Behavioral Advertising

Yang Wang, Huichuan Xia, Yun Huang

SALT (Social Computing Systems) Lab, School of Information Studies, Syracuse University
{ywang, hxia, yhuang}@syr.edu

ABSTRACT

Online Behavioral Advertising (OBA), which involves tracking people's online behaviors, raises serious privacy concerns. We present results from a scenario-based online survey study on American and Chinese Internet users' privacy preferences of OBA. Since privacy is context-dependent, we investigated the effects of country (US vs. China), online activities (e.g., online shopping vs. online banking), and platform (desktop/laptop vs. mobile app) on people's willingness to share their information for OBA. We found that American respondents were significantly less willing to share their data and had more specific concerns than their Chinese counterparts. We situate these differences in the broader historical, legal, and social scenes of these countries. We also found that respondents' OBA preferences varied significantly across different online activities, suggesting the potential of context-aware privacy tools for OBA. However, we did not find a significant effect of platform on people's OBA preferences. Lastly, we discuss design implications for privacy tools.

Author Keywords

Online Behavioral Advertising; OBA; Privacy; Context; China; USA; Cross-Country Study.

ACM Classification Keywords

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

INTRODUCTION

Advertising networks are increasingly using Online Behavioral Advertising (OBA) to provide ads tailored to individual Internet users. The US Federal Trade Commission (FTC) defines OBA as "the practice of tracking an individual's online activities in order to deliver advertising tailored to the individual's interests." [10] By profiling an individual's online activities and characteristics, OBA could build models to infer users' preferences and display tailored ads accordingly [16]. As indicated by prior research [5], OBA could benefit advertising companies greatly by increasing the click-through rates and the price of the ads [5]. OBA could also benefit Internet

users by tailoring ads to their potential interests, and to a certain degree, by supporting websites to provide free services to these users [31, 52].

However, OBA has also raised privacy concerns due to its pervasive tracking. Turow et al.'s 2009 survey of 1,000 Internet users in United States found that 68% of them "definitely would not" and 19% "probably would not" allow advertisers to track them online, even anonymously [51]. Similarly, the 2011 TRUSTe and Harris Interactive online survey found that up to 85% of the respondents had a negative attitude towards online tracking and OBA [50]. While most users do not like being tracked by advertising companies, users' attitudes towards privacy regarding OBA vary [52]. For instance, one study indicates that even though many users consider OBA to be creepy and dislike being tracked online, some of them still find targeted ads relevant and favorable [31]. Despite people's concerns, online tracking and OBA are clearly happening [1, 16, 30].

There is a growing stream of empirical research on people's attitudes towards OBA, and most of them focused on the US (e.g., [26, 31, 38, 51, 52]). Little is known about how people in other countries feel about OBA, such as in China. Compared to the US ranked the second in the number of Internet users, accounting for about 10% of the world's Internet population, China has the world's largest Internet population, more than 20% of the world's total. According to a 2014 report published by the China Internet Network Information Center (CNNIC), online advertising in China has become more prevalent and targeted. In addition, 70% of Chinese mobile phone users would prefer to view ads on mobile apps in order to use them for free, and 32.4% are tolerant of mobile advertising so long as it does not influence their normal user experience¹. These results motivate us to examine and compare American and Chinese Internet users' perceptions of OBA in a more systematic manner.

In addition, whether something violates user privacy is highly context-dependent [4, 33]. This is particularly important to the domain of OBA, as some Internet users prefer to receive relevant ads despite their privacy concerns [31, 52]. Thus, if users only have privacy concerns for OBA in certain contexts, then privacy tools may be built to selectively block OBA based on contexts. This has the potential to allow both the ad industry and Internet users to reap the benefits of OBA, while respecting people's privacy. Despite the importance of

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than ACM must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from Permissions@acm.org.

CSCW '16, February 27-March 02, 2016, San Francisco, CA, USA
© 2016 ACM. ISBN 978-1-4503-3592-8/16/02...\$15.00

DOI: <http://dx.doi.org/10.1145/2818048.2819941>

¹<http://www.cnnic.cn/hlwfzyj/>

context, few studies have systematically examined the impact of context on user privacy preferences of OBA.

To examine how context may affect user preferences of OBA, we conducted a scenario-based survey study. We focus on two aspects of context: *activity* and *platform*. For *activity*, we created five types of online activities including online shopping, online banking, online health information seeking, online dating, and online social networking. For *platform*, we included two platforms - desktop/laptops, and mobile apps - on which online activities and OBA can occur. We investigated survey respondents' willingness to share their information for OBA when they are conducting these five types of online activities on one of the two platforms.

There are two main reasons why we chose to focus on these three variables in our study: *country* (US/CN), *platform* (mobile/desktop), and *activity* (five different online activities). First, they are all important aspects of the context in which people experience OBA. While they have been suggested in the literature as playing important roles in people's privacy/security decision-making (e.g. [6, 7, 13]), to the best of our knowledge, they have not been systematically investigated in the context of OBA. Second, similar to prior privacy studies on OBA (e.g., [26]), our primary goal is to assess the relative effect of each variable after controlling for other variables. The interactions between these three variables are also of interest. For instance, the interaction between country and platform is interesting because China's mobile service providers (e.g., China UniCom) have publicly announced tracking the web browsing history of their mobile phone users². However, we did not find any public information about whether Chinese service providers track their users browsing behavior on the desktop/laptop. Therefore, Chinese users might have different expectations about being tracked by the service providers between their desktop/laptop computers and mobile devices. The interaction between country and website is also interesting because the most popular websites for the same category (e.g., online shopping) in China and the US are comparable but different websites.

We found that American respondents were significantly less willing to share their information for OBA and had more specific concerns than their Chinese counterparts. We also found respondents' preferences of OBA varied significantly across different online activities. However, we did not find a significant effect of platform on respondents' preferences of OBA. These results can inform the design of privacy tools for OBA.

In summary, this paper makes three main contributions. First, we found that the type of online activity has a significant impact on users' willingness to share personal information for OBA purposes, contributing to the empirical evidence of the contextual nature of privacy preferences and highlighting the potential of context-aware privacy tools for OBA. Second, we identified Chinese Internet users' perceptions of OBA, which, to our knowledge, is non-existent in the literature. Third, we systematically compared American and Chinese Internet

users' perceptions of OBA, highlighting the nuanced differences against the broader historical, legal and social backdrop in these two countries.

RELATED WORK

Conceptualization of Privacy

In this paper we focus on information privacy [44]. One definition that fits with the OBA context is "the claim of an individual to determine what information about himself or herself should be known to others" [51]. Smith et al. develop a validated scale for measuring main dimensions of users' general privacy concerns about organizational practices: (1) collection, or the concern about the extensive collection of personal identifiable information; (2) unauthorized secondary use, which refers to information being collected for one purpose but used for a different purpose; (3) improper access, or the concern of personal information being available for unauthorized view or use; and (4) errors, or the concern for deliberate or accidental errors in personal information [45].

A growing body of literature advocates the study of privacy in specific context. Helen Nissenbaum's theory of contextual integrity eloquently points out that human behaviors, e.g., an event or transaction that occurs, are situated in some context: geographical space and certain constituted norms in a political, or cultural environment [33]. In this paper, we study privacy in the context of OBA.

Privacy in OBA

OBA is usually employed by profiling users based on tracking their online activities, such as the websites they visit over time [31]. Because of the pervasive and invasive online tracking [1], OBA has raised privacy concerns [2, 31, 51, 52]. For instance, users generally understand the need of web ads for receiving free online services, however, users do not think tracking them is part of the deal [31].

Various factors could influence users' privacy preferences toward OBA. For instance, one study found that data retention (i.e., the length of time that this data is retained) and scope of use (i.e., how such data would be used) have substantial influence on users' willingness to share their information with OBA [26]. Users also have varying preferences for sharing different types of personal information. For instance, prior research finds that users are relatively more willing to share information about their computers than their personally identifiable information with OBA [31].

Government agencies, Internet standard groups, and browser vendors are also paying close attention to OBA and its related privacy issues. In 2012, the U.S. Federal Trade Commission and the White House released several reports to discuss privacy issues and possible consequences of OBA [11]. W3C, the group that establishes Internet standards, started a working group on tracking protection in 2011. The working group has also been working on the specifications of a "Do Not Track" (DNT) mechanism that allows people to signal their intent of opting out of online tracking in web browsers that websites and ad networks can choose to honor. Browsers such as Microsoft Internet Explorer have implemented mechanisms to support DNT. However, how much the industry is

²<https://www.techinasia.com/china-unicom-tracking-mobile-browsing-history-check-database/>

honoring the DNT signals is unclear. California requires online service providers to disclose whether and how they respond to DNT signals in their privacy policies [8], however websites that include such information in their policies are found very limited.

On the regulation side, the US currently does not have legislation that regulates OBA. Industry's self-regulation was also found to be insufficient. For instance, researchers found various usability issues in tools designed for enabling Internet users to opt out of online tracking [25]. Various privacy-enhancing OBA schemes have been proposed (e.g., [3, 49]), but none have been adopted in real-world practice. In other words, OBA could still put Internet users' privacy at risk.

Cross-Country Study of Privacy and Advertising

Prior literature shows that people in different countries can have very different privacy perceptions of the same thing [21, 53]. For example, American students are found demanding for more privacy protection in their residency than their Turkish counterparts [21]. Another study on the use of social networking sites (SNS) found that American users are more concerned about online privacy in general than their Chinese and Indian counterparts, but Chinese users are more concerned about their online identities [53]. Some scholars argue that differences in countries associated with privacy perceptions are rooted in history and conventions. For example, Smith et al. suggest that Europeans tend to consider their privacy to be a fundamental human right, whereas people in the US regard privacy more as an individual matter and a sort of contractual negotiation with some organization [43].

In terms of users' perceptions of advertising, prior studies have also suggested differences in countries. For instance, one study found that Japanese users are more irritated by mobile advertising (via SMS) than their Austrian counterparts, however, both groups consider entertainment and credibility as key factors in assessing the value of ads [42]. However, another study found that Chinese users' attitudes towards mobile ads (via SMS) are more affected by the attitudes of other Chinese users than intrinsic features of the ads themselves. In particular, younger Chinese women are more easily affected by their peers because of their pursuit of fashion [17].

While OBA is encountered by Internet users around the world, most of the empirical research of OBA focuses on users in the US (e.g., [31, 51, 52]). One exception is a study of Internet users from Romania and US, which found Romanian users have more positive attitudes toward targeted advertising and they are more prone to click the ads than American users [54]. In our study, we examine and compare American and Chinese Internet users' perceptions of OBA. This comparison is interesting because they are the two biggest Internet-using countries. Since prior research has found that Americans and Chinese have different privacy differences in domains such as online commerce [56], SNS [53], and location sharing apps [27], we suspect that there will be differences for OBA too. Therefore, our first research question is:

RQ1 (Country Context): *Do American and Chinese Internet users have different privacy preferences of OBA?*

Contextual Privacy Preferences

There is a strong theoretical rationale and empirical evidence that context affects people's privacy decision-making. Nissenbaum's theory of contextual integrity [33] is particularly relevant. She proposes context-based rules and expectations for protecting users' privacy against what specific information may be collected, with whom the websites may share, and under what conditions such information sharing may occur [34]. Prior studies have found that users' privacy preferences vary across different contexts such as location sharing [4, 22], mobile apps [28], and ubiquitous computing systems [39].

In the domain of OBA, few studies have systematically examined the impact of contexts on users' privacy preferences of OBA. While there has been research on factors that could affect Internet users' attitudes towards OBA [26], these studies focus on different "features" of the site (e.g., length of data retention) rather than different usage "contexts" in which OBA could occur (e.g., online shopping). In our study, we examine users' willingness to share their information with OBA in specific online activities (e.g., finding medical information for a health issue online). This matters because if people's preferences vary across different online activities, then privacy-sensitive mechanisms can be built to collect certain user data in certain circumstances for OBA.

RQ2 (Activity Context): *Do the Internet users' privacy preferences of OBA vary across different online activities?*

Another aspect of context is the platform on which OBA occurs. Prior research has found that people can be more risk-seeking on their mobile devices than on desktop computers (e.g., searching more sensitive topics on mobile devices) [36]. We are curious whether the platform in which OBA occurs would affect people's preferences of OBA. We are particularly interested in OBA in mobile apps because researchers have identified privacy and security risks of ads appearing within the apps [15]. Furthermore, people may already have some privacy preferences (e.g., the exact resources or data on the device that the app can access) during the installation of mobile apps, but these preferences do not specify nor regulate whether these resources or data can be used for OBA.

RQ3 (Platform Context): *Do the Internet users' privacy preferences of OBA differ between the two platforms: desktop/laptop and mobile apps?*

METHODOLOGY

We conducted an online survey study hosted by Qualtrics in the summer of 2014. This study was designed using a mixed-model approach where *country* (China or the US) and *platform* (desktop/laptop or mobile) are between-subjects variables and *online activities* (e.g., online shopping) is a within-subjects variable. Each respondent viewed five hypothetical web usage activities, which are described in detail below. This study was approved by our Institutional Review Board.

Our study design adopted many aspects of a previous OBA study by Leon et al. including the explanation of OBA, the roleplaying aspect of the study, and some specific questions (e.g., what personal information people are willing to share

for OBA) [26]. While Leon et al. studied the impact of varying stated data-collection practices (e.g., data retention policies) on users' attitudes towards OBA [26], we instead use a similar methodology to explore the impacts of country, online activity, and platform on people's OBA privacy preferences.

Survey Content

We designed the survey in English, which was then translated into another version in Simplified Chinese (CN) by two native Chinese speakers. The survey has two branches corresponding to the two platforms: desktop/laptop and mobile apps. We randomly assigned each respondent to one of the two branches. Each branch has a total of 24 questions.

Our respondents started the survey with a few questions about general web advertising. For instance, the first question, which was open-ended, asks: "In a sentence or two, please tell us what you think about website advertising." The survey then introduced respondents to the concept of OBA with a typical example: "Imagine that you are looking to buy a car and decide to visit the cars.com website. Cars.com has contracted with XYZ Advertising Company which collects information about your interactions with the cars.com website in order predict your preferences and to show you ads that are most likely to be of interest to you. These ads are known targeted ads, or online behavioral advertising (OBA). For example, if you search for 'Audi A8' or read an article about Audi cars on the cars.com website, XYZ Advertising Company could show you ads for BMW cars and other luxury cars." This introduction of OBA was adopted from [26].

Respondents were then introduced to the two platforms: "Targeted ads can appear on websites or on mobile apps. For the following questions, we ask you to focus on using [your desktop or laptop PC / apps on your smart phone]" (for the desktop branch and the mobile branch, respectively). We did this because we want to ensure that our respondents can clearly distinguish the two platforms and then focus on the platform that they were assigned to. Respondents were then asked about their generic preferences about OBA. For instance, one question was "Would you be willing to allow OBA to use and store the information about your:" with answers representing various personal information such as age, gender, and information about their online interactions (e.g., searched items).

Next, respondents were presented with the five types of online activities: online shopping, health, dating, banking, and social networking. These activities represented a wide range of typical online activities and we portrayed each of them taking place on a representative website or its mobile app in China and the US, respectively. The five concrete activities provided in the study were: (1) buying books for oneself from a website or an app (*Amazon.com* in the US, *JD.com* in CN); (2) finding medication for a health condition (flaky scalp) on a website or an app (*CVS.com* in the US, *HaoDF.com* in CN); (3) Dating with some person online on a website or an app (*Match.com* in the US, *JiaYuan.com* in CN); (4) Transferring money online on a banking website or on its app (*Chase.com* in the US, *cmbChina.com* in CN); and (5) Joining an interest group on a SNS website or an app (*Facebook.com* in the US,

Renren.com in CN). In order for us to make reasonable comparisons, the websites/apps for China and the US are similar in terms of their popularity and functions.

These descriptions of online activities differ by activity and platform, while other factors are kept constant (e.g., data retention time). As an example, we show the description of the shopping activity on the Amazon mobile app below:

Imagine that you want to buy some books on Amazon for yourself. Amazon has contracted with an advertising company, which collects information about your interactions with Amazon in order to predict your preferences and show you ads that most likely to be of interest to you. For example, if you search "camera" or read a custom review about a camera lens, the advertising company could show you ads for Canon, or another camera brand.

*Now in particular, imagine that you decide to buy the books via the **Amazon app on your smart phone.***

- 1. The advertising company, which has contracted with Amazon would collect your browsing data within the app.*
- 2. The advertising company would retain and use collected information about you for one month.*
- 3. Amazon can share your information with its affiliated partners (which are other companies which have business partnership with the website).*

Would you be willing to allow the advertising company to collect and share your browsing data as follows within the Amazon app?

Similar to the study by Leon et al. [26], the respondents were asked to answer a series of questions on a five-point Likert-scale about their willingness to share ten types of personal information for each online activity. These ten types of personal information were name, age, gender, income bracket, marital status, home address, email, search items entered, information inferred from web browsing (e.g., "genres of books I like") and information on pages, as well as the time spent on them. The possible responses range from "I would not be willing at all" to "I would be completely willing." These types of personal information were also adopted from [26]. Prior studies have also shown that some of this information (e.g., gender, age) can be inferred from users' browsing data if not collected directly from them [19]. The sequence of online activities were randomized. For each online activity, the ten types of personal information were also randomized.

After these activity-based questions, the survey asked respondents additional questions about their opinions of the five websites/apps, for instance, a five-point Likert-scale question on how much they trust the corresponding service provider. We asked these questions because they could play a role in users' willingness to share information with these services. Before ending the survey with a set of demographic questions, we asked a second open-ended question: "Do you have any further comments on OBA, which might appear on [websites / mobile apps on your smart phone]?"

We originally designed the survey to ask each participant about ten situations (five online activities x two platforms). We conducted a pilot test of the initial survey design and received the feedback that the survey had too many questions and many seemed redundant. Thus, we decided to have respondents focus on a platform to reduce the number of survey questions and the need to switch between too many situations.

Participants

For the US version of the survey, we recruited respondents from Amazon Mechanical Turk (MTurk) who are from the US and have a 95% or higher approval rating and reported to be 18 or older. We also implemented a mechanism on MTurk to ensure that respondents who have taken one branch of our survey were excluded from taking the same one again or the other branch. [35]. For the Chinese version of the survey, we recruited respondents from ZhuBaJie.com (ZBJ), the main crowdsourcing site (equivalent of MTurk) in China. However, ZBJ does not have the API support for us to implement the automatic mechanism that prevents respondents from answering the survey repeatedly. For this purpose, we manually examined the responses on Qualtrics to filter out repeated responses by using the ZBJ user IDs. The American and Chinese respondents took a comparable amount of time to finish the survey (US: median 11 minutes; CN: median 14 minutes). MTurk and ZBJ Respondents who completed the survey were compensated \$2 and 15 RMB (about \$2), respectively.

To help determine whether our respondents took the survey seriously, we used a combination of measures such as unusually short completion time and nonsense responses to the open-ended questions. After the filtering, we had a total of 379 valid responses: 190 from the US and 189 from China. For each country, responses were roughly evenly split between the mobile app and desktop branches (US-desktop: 95, US-mobile: 95; CN-desktop: 97, CN-mobile: 92).

In terms of demographics, there were similarities and differences between the US and Chinese samples. First, the US sample had more male respondents (54%) than female respondents (46%) while the Chinese sample had an equal split in gender (male 50%, female 50%). Second, the US sample was significantly older than the Chinese sample (US: M=35, SD=11.2; CN: M=27, SD=5.7; Mann-Whitney U Test, $p < 0.001$). In terms of education level and IT background, the US and Chinese samples were comparable. For both countries, about 93% of respondents had at least some college education. The percentages of American and Chinese samples with IT background were 38.9% and 36.5%, respectively. For each country, the participant demographics were similar in the two platform branches.

Method of Analysis

Factor Analysis

The respondents were asked about their willingness to share ten types of personal information in the five specific online activities. To reduce these types to a smaller set of latent factors, we conducted a parallel analysis (PA), followed by an exploratory factor analysis (EFA) with promax rotation to orthogonal oblique structure. PA is a Monte Carlo simulation

method to determine the number of latent factors [24]. EFA is useful in exploring the underlying latent factor structure by combining highly correlated variables into latent factors.

The PA analysis suggested three latent factors for both the US and Chinese samples across different online activities. The EFA suggested the same underlying structure of the three factors as shown in Table 1. To assess the convergent validity of the latent factors (i.e., whether each item in a factor is measuring the same underlying concept), we calculated Cronbach's α value for each factor. The results indicated that the Cronbach's α values of the three factors are all above .85. Alpha values of 0.7 or higher are considered acceptable [48].

Repeated Measures ANCOVA

We performed a series of repeated measures ANCOVA (analysis of covariance), each for the index variable of a particular latent factor (browsing, demographic, or contact information) to assess the effects of country (US vs. China), online activities (e.g., online shopping vs. online dating), and platform (mobile vs. desktop) on respondents' willingness to share that type of information for OBA purposes. The independent variable is a factor score, which is the average of the constituent items of that factor. The country and platform are between-subject independent variables while the online activity is a within-subject independent variable (i.e., repeated measure). Respondents age, gender, IT background as well as their levels of trust with the website or app in a particular online activity were treated as between-subject covariates. We also included interaction terms of the three independent variables (country, platform, and online activity).

Content Analysis

We analyzed answers to the two open-ended questions on website advertising and OBA. Our coding scheme drew from prior literature on OBA (i.e., directed content analysis) but was also partly created from additional themes emerging from

	Cronbach's α	Factor loading
Browsing Information	0.92	
Info inferred from my browsing		0.87
Info on pages I visited		0.81
Search items I entered		0.91
Demographic Information	0.90	
Age		0.96
Gender		0.85
Income bracket		0.59
Marital status		0.86
Contact Information	0.87	
Name		0.80
Email		0.87
Home address		0.85

Table 1. Factor structure and loadings of the ten information items for American respondents in the scenario of online shopping (Amazon) on the desktop/laptop. The factor structure is consistent across American and Chinese samples across different online activities and platforms. Cronbach's α value is calculated for each factor.

the data (i.e., conventional content analysis) [18]. In particular, prior studies have identified different perceptions of OBA such as useful, creepy, annoying and intrusive [2, 52]. Prior literature has also identified specific concerns about OBA such as privacy violation, unauthorized data collection, and information being used by unauthorized 3rd parties [52]. We included these in our initial coding book.

Besides, after reviewing the answers, we found new themes not covered by the code book. Examples of these new codes include “overwhelming,” “slowing down the system,” and “user option and control of OBA.” We added these new codes into the code book, which has a final count of 13 codes. Two researchers coded the data independently using the same coding scheme. The Cohen’s κ value [9] for our inter-coder reliability is .72, exceeding the acceptable threshold of .7 [23].

RESULTS

Overall, we found that country and online activity have significant effects on people’s preferences of OBA, but we did not find a significant effect of platform.

Country: US vs. China

We start by comparing the perceptions of web ads and OBA of the American and Chinese respondents. We include English translation of the Chinese answers in this paper.

Attitudes Towards Web Advertising

The American and Chinese samples had several differences in their answers to the first open-ended question about website advertising (i.e., before the introduction of OBA). First, the Chinese sample was significantly more accepting of web advertising than their American counterparts, as shown in Figure 1. We categorized the overall perception of each response to website advertising as positive, neutral, or negative.

The positive reactions include feeling web advertising as acceptable, useful or necessary while the negative reactions include being seen as annoying, intrusive, distracting, overwhelming, slowing down the website, untrustworthy or even unsafe. For instance, one Chinese respondent said “*There are a lot of web ads. The good and bad are jumbled together. The authenticity is low.*” Some American respondents were not only aware of the risk but also crafted their own strategies to test out the ads. For example, one US respondent said “*I think it is unsafe, and it can be amusing and or informative. If I see attractive advertising on the peripheral, I may check Google the name in the ad, perhaps even google the word ‘scam’ after it, and see what I see. If I think it is safe, I will go to their website with care, always on the ready to close the window if I sense danger. So If I do check out the ad, it will NOT be by*

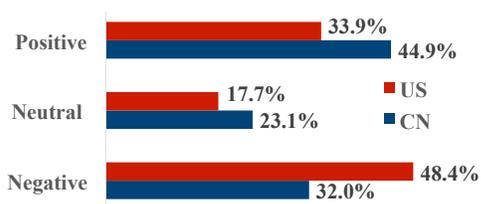


Figure 1. American and Chinese respondents’ sentiments of web ads.

clicking on the ad but rather by investigation and then maybe going to the site directly not via their ad.”

Second, 53% of American respondents mentioned targeted ads without prompting whereas only one Chinese respondent did so. For instance, one US respondent answered “*I like to opt out of targeted ads, but it’s a hassle and the ones done by cookies come back when I clear my cookies. I don’t mind ads in general because I know it keeps the websites free, but targeted ads are annoying because it’s things I already bought, decided not to buy, or was a gift.*” The Chinese respondent said “*After searching things in Baidu, no matter what websites I open, the same products I searched before would jump out.*” This difference between the American and Chinese samples suggests that the American respondents are more aware of OBA than the Chinese respondents. We will situate these differences in the broader contexts of these two countries (e.g., legislation, industry self-regulations, and media portrayals of OBA) in the discussion section.

Generic OBA Privacy Preferences

After the introduction of OBA and the two platforms, our respondents were then asked about their willingness to allow OBA to store and use various types of their personal information. We call these generic OBA privacy preferences because no specific scenario was given in the question. Figure 2 shows that the American respondents were less willing to share each type of information than their Chinese counterparts. However, the largest percentage of American and Chinese respondents were willing to share information about their hobbies and gender while the smallest for their home address.

More specifically, there are two notable differences. First, for credit score/report, only 5% of the American respondents were willing to share compared with 34% of the Chinese respondents. Since there is no established credit score system in China, we actually translated it as credit report, which exists in China but is not as common as credit score in the US. In addition, while 30% of the Chinese sample would be willing to share email, only 6% of the US sample would. The 2014 CNNIC report of Internet Development in China shows that 42.5% of Chinese Internet population uses email compared to 89.3% uses instant messaging. This result may suggest that

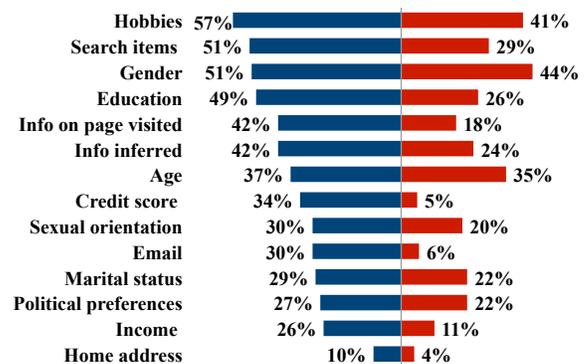


Figure 2. The percentages of Chinese (CN, left) and American (US, right) respondents who would be “willing” or “completely willing” to share each type of information for generic OBA purposes.

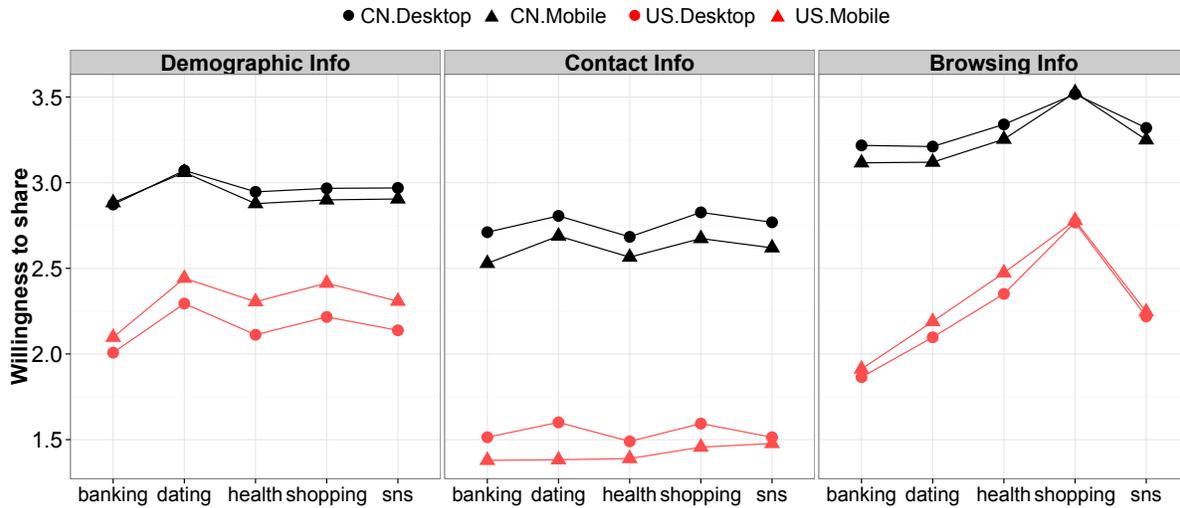


Figure 3. Willingness to share the three types of information with OBA in the five specific online activities under four conditions (CN/US x Desktop/Mobile). American respondents were significantly more concerned about sharing information with OBA than Chinese respondents (the two lines for CN are always higher than the two lines for US). Different online activities also vary significantly in terms of respondents' willingness to share information with OBA (the lines are zig-zagged across different activities, with the browsing info being the most notable case). Platforms (desktop vs. mobile) do not have much difference (the two lines for each country do not have much space).

Dep. Vars	Indep. Vars	Mean sq	F value	Pr(> F)	
Contact Info	Name	Country (Cty)	427.6	136.5	2.00E-16 ***
	Email	Platform (Pla)	5.3	1.698	0.193
	Address	Activity (Act)	0.8018	6.063	7.71E-05 ***
		Age	140.6	44.89	7.70E-11 ***
		Gender	20.2	6.437	0.0116 *
		IT background	103.5	33.04	1.88E-08 ***
		Trust level	179.5	57.31	2.95E-13 ***
Demo. info	Age	Country	107.3	27.61	2.51E-07 ***
	Gender	Platform	3.62	0.931	0.335
	Income	Activity	3.194	16.46	3.33E-13 ***
	Marital status	Cty x Act	0.584	3.011	0.0173 *
		Age	122.67	31.56	3.79E-08 ***
		Gender	29.67	7.63	0.006 **
		IT background	10	2.573	0.1096
	Trust level	242.26	62.33	3.28E-14 ***	
Browsing info	Search items	Country	232	75.7	2.00E-16 ***
	Info inferred	Platform	1	0.322	0.571
	Page visited	Activity	15.982	45.32	2.00E-16 ***
		Cty x Act	3.618	10.26	3.39E-08 ***
		Age	143.5	46.84	3.17E-11 **
		Gender	13.7	4.46	0.0354 *
		IT background	64.2	20.95	6.44E-06 ***
		Trust level	404.7	132.1	2.00E-16 ***

*** p<0.001, ** p<0.01, * p<0.05

Table 2. The repeated measures ANCOVA results of the three types of personal information: contact, demographic, and browsing information. The interaction terms (Cty x Act, Cty x Pla, Act x Pla, Cty x Act x Pla) were included in the analysis but only shown here when they are significant. The adjusted R² for the ANCOVA models of (contact, demographic, and browsing information) are 0.36, 0.19, and 0.31, respectively.

because Chinese Internet users do not use email that much and therefore they disregard email address as critical or sensitive.

Scenario-Based OBA Privacy Preferences

Having examined our respondents generic privacy preferences for OBA, we now focus on their contextual preferences under various specific settings. To assess the effects of country (US vs. China), online activities (e.g., online shopping vs. online dating), and platform (mobile vs. desktop) on respondents' willingness to share that type of information for OBA purposes, we conducted a series of repeated measures ANCOVA. As shown in Table 2, after controlling for the covariates (age, gender, IT background, and trust level with the site), the country and online activity had significant effects on respondents' willingness to share information for OBA. However, we did not find a significant effect of the platform. We also ran the analysis without the interaction terms and analyzed the American and Chinese samples separately. The results of the main effects were consistent with that in Table 2.

Similar to the findings of the first open-ended question on web advertising and the generic OBA privacy preferences, Chinese respondents were significantly more willing to share all three types of information (browsing, contact and demographic information) for OBA than their American counterparts (p<0.001) when controlling for online activity and platform as well as the covariates as shown in Table 2. Figure 3 also clearly shows that the lines of the Chinese sample are always much higher than that of the American sample. Particularly, Chinese respondents were generally willing to share browsing information while American respondents were not. This may be because our Chinese respondents did not know the implications of sharing this browsing information (e.g., companies could infer their demographics, create profiles for them and generate targeted ads [19]).

Attitudes Towards OBA

Right before the last part of the survey on demographics, we asked our respondents whether they have any further comments about OBA in an open-ended question. A large number of respondents did not provide any substantial answers.

They provided simple answers such as “No more,” “N/A,” or “Thank you.” To capture each respondent’s free response of OBA, we combined each respondent’s answers to both open-ended questions (on web ads and OBA) and coded the answers using our code book.

Overall, our coding results show that 25.1% of Chinese respondents felt OBA is helpful, compared with only 10% of American respondents holding that view. 14.8% of Chinese respondents and 12.5% of American respondents felt neutral about OBA. However, for negative perceptions, the most frequently reported perception for Chinese respondents (27.2%) is “privacy violation” in a general sense without specifying how their privacy might be intruded by OBA. For example, one Chinese respondent said “*I suspect this will violate personal privacy.*” In comparison, the American respondents had more specific concerns such as being tracked, which is the most common negative perception of the US sample (10%). For instance, one American respondent said “*I feel that this is a serious invasion of privacy. While I understand the need for advertising on free websites, I don’t like being cyber-stalked in this manner.*” This difference between American and Chinese respondents also implies that the American respondents may be more aware of OBA and its specific privacy implications than the Chinese respondents.

Online Activity

Next, we examine the effect of online activities on respondents’ preferences towards OBA. As shown in Table 2, our repeated measures ANCOVA shows a statistically significant impact of online activity on respondents’ willingness to share each of the three types of information (browsing, contact and demographic information) for OBA ($p < 0.001$) when controlling for country, platform, and the covariates. Figure 3 also shows the zig-zagged lines of both US and Chinese samples, suggesting the differences among these online activities.

We then performed post-hoc pairwise comparisons between different online activities using the Tukey’s HSD (honest significant difference) test with Bonferroni correction. For demographic information, respondents were significantly more willing to share with the dating site than the banking and health sites ($p < 0.001$). Similarly, for contact information, respondents were significantly more willing to share with the dating site than the banking and health sites ($p < 0.05$). These differences are intuitive because online dating sites usually ask for demographic and contact information to categorize and match users. For browsing information (e.g., search items), respondents were significantly more willing to share with the shopping site than the banking and dating sites ($p < 0.001$). This may be because users have experienced and understood the fact that shopping sites usually log user behavior (e.g. products users have purchased or viewed before) to make personalized product recommendations.

We also found a significant interaction effect between country and online activity for browsing and demographic information. As shown in Figure 3, the American respondents were more willing to share their demographic information with the shopping site than the social networking site whereas the Chinese respondents were somewhat the opposite. Demographic

information would seem more reasonable to provide to social networking sites (e.g., in the user profile) than online shopping sites. However, perhaps the American respondents felt less willing to share information with the SNS site (Facebook) because of the privacy-related discussion about Facebook in the popular media [40].

In the open-ended question on OBA, our respondents touched on general perceptions towards OBA in different online activities. Respondents were generally aware of and accepted OBA in the online shopping context. One American respondent mentioned “*I have noticed an influx in ads that tailor to my shopping behaviors over the last year. For example I was recently in the market for a backpack and now I FREQUENTLY get ads on various websites offering me deals on backpacks.*” One Chinese respondent said “*On one hand, OBA could facilitate my online shopping, but on the other hand I am also concerned about my privacy violation by it.*”

Respondents were generally against OBA in the online banking context. One American respondent said “*I would not want to give my information to a second party when visiting a website. ESPECIALLY something like banking. If that second party’s security was compromised, my banking information could be at risk.*” Similarly, a Chinese respondent answered “*I could accept OBA, but it must tell me what kind of risks I may take, particularly regarding online banking.*”

Platform: Desktop/Laptop vs. Mobile Apps

Lastly, we examine the effect of platform on people’s perceptions of OBA. As shown in Table 2, we did not find a statistically significant effect of platform on respondents’ willingness to share various types of information with OBA in specific online activities. Visually, we see that the two lines of each country (i.e., one for desktop/laptop, and the other for mobile app) in Figure 3 were relatively close, suggesting lack of significant difference.

In the open-ended question about OBA, respondents in the mobile branch of the survey brought up many points related to mobile apps. For instance, one American respondent was concerned about one such app being able to monitor what the user does in other apps: “*My biggest concern is how much data they’re collecting about what I do outside of the particular app in question. Amazon knowing what I do in the Amazon app is less concerning to me than Amazon knowing all the browsing and app habits on my phone in general.*” One Chinese respondent was also picky about where OBA ads should appear in a mobile app: “*it’s acceptable if it appears when the mobile app was first opened. It’d be very annoying if it appears in other places [of the app].*” These specific requests should be considered for OBA in mobile apps.

DISCUSSION

In summary, we found country and online activity having significant effects on our respondents’ willingness to share information for OBA but not for platform.

Contextual Privacy Preferences of OBA

Country: US vs. China. There is an existing body of research that suggests that Americans are more concerned about online

privacy than Chinese in a number of domains (e.g., online commerce [56], SNS [53], and location sharing apps [27]). Our study extends this literature by examining OBA - an important yet under-studied domain - and uncovers significant yet nuanced differences between Internet users from these two countries.

In particular, we found that American respondents were significantly less willing to share their information for OBA purposes than their Chinese counterparts. A large percentage of the US sample voluntarily expressed their opinions about OBA in response to the first open ended question about website advertising in the survey before any introduction of OBA, while only one Chinese respondent did so. In addition, when expressing their opinions about OBA, most Chinese respondents discussed privacy concerns in a general sense whereas the American respondents expressed more specific misgivings, such as being tracked. These results suggest that the American respondents were more aware of OBA and its specific privacy implications than their Chinese counterparts.

In the US, privacy has been long been recognized as an essential human right. Warren and Brandeis' seminal article "The right to privacy" in 1890 has arguably ushered in an era of public recognition of and discussion about privacy and subsequent legislation in the US [55]. The American government created the Fair Information Practice Principles (FIPPs) in the 1970's, which underly many privacy laws worldwide. While the US does not have an omnibus or cross-domain privacy law, there are many domain-specific privacy laws such as RFPA (The Right to Financial Privacy Act) and HIPPA (Health Insurance Portability and Accountability Act).

Online privacy has become one of the priorities for the US government. For instance, in Dec. 2010, the Federal Trade Commission (FTC) released a report entitled "Protecting Consumer Privacy in an Era of Rapid Change: A Proposed Framework for Businesses and Policymakers," which describes a new regulatory framework including the "Do Not Track" mechanism allowing people to opt out of OBA. In Feb. 2015, the Obama administration released a draft of the Consumer Privacy Bill of Rights Act which allows individual consumers to sue service providers for privacy violations.

The mass media frequently covers online privacy topics including those regarding OBA³. There are also many US-based civil liberty organizations such as American Civil Liberties Union (ACLU) and Electronic Privacy Information Center (EPIC) that regularly blow the whistle on industry and government practices that may infringe people's privacy. The 2013 Snowden revelation of government surveillance also has limited and short-lived effects on the American public's interests in privacy [37]. In addition, Ad industry organizations such as Digital Advertising Alliance (DAA) have launched public campaigns about OBA and their self-regulations (e.g., the OBA icon and opt-out page).

³Examples: <http://www.nytimes.com/2009/03/16/technology/internet/16privacy.html>; <http://www.wsj.com/public/page/what-they-know-digital-privacy.html>; <http://www.cbsnews.com/videos/60-minutes-probes-data-brokers-and-online-tracking/>

In sum, the topic of online privacy seems to be highly visible in contemporary American society, which provides many opportunities for ordinary citizens to gain knowledge about OBA, along with the privacy implications related to it as well as how to deal with OBA (e.g., opt out of OBA).

In contrast, privacy is a relatively recent concept in China. Historian Peter Zarrow wrote "[In the late Qing dynasty (c. 1890-1912)] No single word existed that was equivalent to the English privacy in the sense of personal, closed off from the public, inner life, family life, private (individual) rights and related concepts." [32] He also noted that the discussion on privacy mostly revolved around a "realm of *si* [personal, self, selfish, private] in terms of its relationship to the realm of *gong* [public, public space, open, communal]." Under the traditional Chinese culture of collectivism, the concept of *si* even has negative connotations of suspicious and selfish. The privatization of state-owned industry in the late 1980's and early 1990's may have triggered a sea change in legitimizing *si* (personal) and shaping the contemporary conceptualization of privacy in China. Hu noted that the Chinese word *yin si* for "privacy" did not appear in the two most popular and authoritative Chinese dictionaries in 1993 [20]. This is more than a century after Warren and Brandeis' classic article on privacy.

On the regulatory front, the first national standard on personal information protection entitled "Information Security Technology Guidelines for Personal Information Protection Within Public and Commercial Services Information Systems" in Mainland China was drafted in 2008 and then issued in 2013 by the Ministry of Industry and Information Technology. While this standard provides guidelines regarding collection, use, and sharing of personal information in public and commercial systems, it does not have any legal force.

In March 2014, the Chinese Advertising Association Interactive Network Branch issued the first Chinese industry self-regulatory OBA guidelines entitled "Chinese Internet OBA User Information Protection Framework Standard."⁴ Many Chinese Internet Juggernauts such as Baidu are part of this self-regulatory group. They adopted an OBA icon (the letter *i* inside of a circle) very similar to the DAA's OBA icon (the letter *i* inside of a triangle, officially called "AdChoices")⁵. However, we are not aware of any public campaign about this Chinese OBA icon nor any opt-out options provided by this Chinese industry self-regulatory group. We are also not aware of any civil liberty organization in China that fights for ordinary citizen's privacy, let alone privacy in OBA.

The topic of online privacy occasionally appears in the Chinese news media, including some articles on OBA. However, these articles often are about privacy studies or opinions from the western countries⁶. This may explain why our Chinese respondents lack knowledge of OBA and its privacy implications because user education is almost absent. This lack of knowledge could also play a role in Chinese respondents' more willingness to share information with OBA as well as

⁴<http://www.iac-i.org/privacy/index.html>

⁵<http://www.youradchoices.com/>

⁶Examples, <http://www.kexuehome.com/articles/201504211120.html>; <http://3g.forbeschina.com/review/201407/0034073.shtml>

lack of specific privacy concerns when some of them are indeed concerned about OBA. One interesting question is that if China had a Snowden-like incident, would it drastically make Chinese people more privacy concerned? It is widely known that the Chinese government is actively monitoring Internet usage and filtering information that is deemed inappropriate [47]. While this is purely speculative, we suspect that Snowden-like incidents are unlikely to happen in China or have similar impact as that in the US.

We recognize that these two countries have huge populations, and that the concepts and the practices of privacy may not be uniform within and unique to these countries. Nevertheless, considering American's and China's historical, legal, and cultural factors, while speculative, may help us situate these differences in the broader social contexts of these two countries.

One factor that may play a role is the acceptable "power distance" in these two countries. According to Geert Hofstede's cultural dimensions theory, the power distance dimension measures "the extent to which the less powerful members of institutions and organisations within a country expect and accept that power is distributed unequally." [14] China has a high rating of 80 for "power distance," while the US has a low rating of 40 [14]. These scores suggest that Chinese society is much more acceptable of power inequalities than American society. In the context of OBA, the ad industry has the upper hand over Internet users. China's high score of power distance may suggest that Internet users tend to accept that the ad industry has more power in controlling OBA than themselves, and thus they are less concerned about OBA because they accept that they do not have much say or control in OBA.

The differences in these historical, legal, and social contexts may have influenced the divergent opinions of our Chinese and American respondents on OBA. Our study not only sheds lights into Chinese Internet users' perceptions of OBA, but also highlights the nuanced differences between the American and Chinese Internet users.

However, it is important to note that since our study did not investigate these cultural or social factors, we cannot make any conclusive claims about the impact of these factors on people's privacy attitudes towards OBA. Nonetheless, this is a promising area for future research.

Online activity. Several empirical studies have found that the majority of American people do not like OBA [31, 51]. They felt OBA is sometimes smart and useful in providing the ads that interest them; at the same time scary and creepy because of the amount of information tracked and collected in order to create the targeted ads [52]. While privacy is a very situation-dependent concept [46], few prior studies have explicitly considered specific online activities when soliciting people's perceptions and preferences about OBA. One notable exception is the study by Leon et al. [26], which inspires many aspects of our own study. Their study investigated different factors (e.g., time of data retention) that may impact people's willingness to share information for OBA where they used a concrete online activity for seeking medical advice [26]. Our study makes a unique contribution by exploring beyond a single on-

line activity, as well as creating and testing a set of common yet different online activities.

We found that the type of online activity indeed has a significant effect on our respondents' willingness to share their information for OBA. For instance, for demographic information, respondents were significantly more willing to share with the dating site than with the banking and health sites. This result is intuitive as online dating sites usually ask demographic information to match users, so people are likely to understand and accept these practices. Results like this also align well with the concept of contextual integrity [33, 34] because these data collection/usage practices are reasonably understood and expected under the particular context.

Platform. Prior research has found that people behave differently on their mobile devices versus desktop e.g., searching more sensitive topics on mobile devices [36]. In addition, people tend to have different security experiences and considerations between the desktop and the mobile environment [7]. These results suggest that platform may play a role in people's privacy-related decision making. However, according to our ANCOVA results, we did not find a significant effect of platform (desktop/laptop vs. mobile apps) on respondents' willingness to share information for OBA. Statistically speaking, this means that our exploration of this factor is inconclusive.

We introduced the two platforms to our respondents. However, based on the pilot testing of the original survey, we decided to ask each participant to think about OBA only on one platform and they were not given an opportunity to explicitly compare these two platforms. This might contribute to the non-significant result on platform. Future studies could consider asking participants to explicitly compare their preferences on different platforms.

Design Implications

What do these results mean for designing privacy-friendly online experiences, particularly related to OBA? First, our results show that contextual factors such as country and the type of online activity matter. Understanding the contextual nature of privacy preferences can inform tools that selectively block online tracking or collection of certain data based on context. Intelligent privacy tools that recognize and consider these contextual factors as well as learn and apply people's context-dependent preferences are more likely to be effective in mitigating peoples concerns about OBA. For instance, when users are browsing products on a reputable e-commerce site, the protection mechanism could by default allow the sharing of browsing information (e.g., items searched) but prohibit the collection and inference of users' contact information.

Second, we value the perspective that Dourish advocates — context is not a static representation that can be separated from the activity at hand but rather is actively produced and enacted in the course of the activity [12]. Taking this perspective, we advocate that privacy tools should show the relevant entities and resources that are involved in OBA to help people produce and make sense of the context around their online activities (e.g., buy a birthday gift online). Tools such

as Lightbeam⁷ embody some of this perspective by providing dynamic visualization of third party trackers as people browse a website, but it does not really show what kind of information is being collected. It would be promising to combine this sort of dynamic visualization with powerful web measurement tools such as OpenWPM⁸ that can detect the kind of information being collected.

Third, our respondents, particularly the Chinese respondents, have little or insufficient knowledge about OBA and its specific privacy implications, which is worrisome. User education is critical. Government agencies, ad industry self-regulatory groups, privacy tool developers, website designers, and browser developers all play a role in educating Internet users about OBA. Educational tools could be designed to improve people's awareness of OBA. Similar to how games are designed to teach people about security [41] or a company's privacy practices⁹, games could be designed to teach people about OBA, its privacy implications, and available user controls. For instance, one could envision a game like Monopoly where users selectively share (or "invest") their personal information for OBA to maximize their benefits (e.g., coupons, free access to premium content) while safe-guarding their privacy (e.g., against third-parties).

Lastly, many popular privacy tools such as AdblockPlus, Ghostery, and Privacy Badger primarily work on desktop/laptop at the moment. Extending these tools to work on mobile devices and in mobile apps would be useful. In addition, translating and customizing these tools for non-English speaking (e.g., Chinese) users would have a much broader impact.

Study Limitations

Our study has a number of limitations. First, while we recruited our American and Chinese respondents from the main crowdsourcing site in each country, this methodology is subject to any bias that may reside in these crowdsourcing sites. Therefore, we cannot make claims about whether our two samples are representative samples of the Internet users in these two countries and consequently, whether our two samples are completely comparable. While the Chinese sample is significantly younger than the US sample (which we have controlled for in our ANCOVA analyses), the two samples are comparable in terms of education level and IT background. We also note that recruiting perfectly comparable and representative samples across different countries will be extremely challenging, if not impossible.

Second, we only included one popular website for each type of online activity in China and the US, respectively. While we did control for respondents' perceived trust with each website as a co-variate in the analysis, our results might still not be generalizable for these types of online activities. Future studies could test multiple websites for each type of online activity. In addition, we only explored five types of online activities. While our pilot testing of the survey did not receive

any complaint about having unrealistic online activities, future work can look into experience sampling and diary study that could better capture people's perceptions of OBA in a wide range of their own online activities. However, recruiting a large number of participants for such studies could be a challenge.

Third, only after conducting the study, we realized that the descriptions of the five online activities have one minor difference - the first sentence of the health scenario specified a concrete topic (i.e., "flaky scalp"), while the other four scenarios did not. For instance, the shopping scenario said "Imagine that you want to buy some books on Amazon for yourself," but did not provide a particular type of book. While this is a difference of specificity in the first sentence, we believe that this difference has little impact if any on our results. The willingness to share information for the health scenario is almost always in the middle among the five scenarios. Besides, the four other scenarios do not differ in the level of specificity. Even if we exclude the health scenario and just compare the other four scenarios, the results of online activity still hold. There are significant differences in willingness to share among the four scenarios where shopping is the highest while banking is the lowest.

Fourth, our respondents' privacy preferences are self-reported and may divert from their actual behavior in the real world. However, this type of self-reported data could still provide valuable insights into people's perceptions of OBA. Many empirical studies of privacy have used this kind of self-reported, scenario-based method (e.g., [26, 29]).

Lastly, the mobile branch/version of our survey was not conducted on mobile phones which may affect some respondents' imagination of OBA on smart phones. While there are ways to recruit respondents and have them answer the survey on their phones, we were concerned that the small screen and keyboard of a mobile phone would make the survey more difficult to answer and thus respondents may pay less attention.

CONCLUSION

Online behavioral advertising (OBA) is prevalent on the Internet. While it could provide ads that people find interesting and useful, its underlying tracking and profiling of users is unsettling. Prior research has studied Internet users' perceptions of OBA, most of them focused on American users and did not study contextual factors (e.g., online activity and platform) systematically. In this study, we surveyed both American and Chinese Internet users about their perceptions of OBA as well as their willingness to share personal information for OBA in different specific situations (e.g., searching for information about medication for a health issue in a health mobile app). We found that Chinese respondents were much more willing to share their data for OBA than their American counterparts. Given OBA is a global phenomenon, further research is warranted to examine how people in different countries perceive OBA and make decisions. We also found that respondents' willingness to share information for OBA vary significantly amongst different online activities. We advocate designing privacy tools for OBA that can support people to seamlessly

⁷<https://www.mozilla.org/en-US/lightbeam/>

⁸<https://github.com/citp/OpenWPM/>

⁹Zynga's PrivacyVille: <https://zynga.com/privacy/privacyville>

make sense of the prevailing context and make privacy decisions congruent with the prevailing context.

ACKNOWLEDGEMENT

We thank our respondents for sharing their insights. We are also grateful to Pedro Leon, Blase Ur, Jeffery Stanton, Jason Dedrick and anonymous reviewers for their thoughtful comments on earlier versions of this paper. This work was supported in part by NSF Grant CNS-1464347.

REFERENCES

1. Gunes Acar, Christian Eubank, Steven Englehardt, Marc Juarez, Arvind Narayanan, and Claudia Diaz. 2014. The Web Never Forgets: Persistent Tracking Mechanisms in the Wild. In *Proceedings of the 2014 ACM SIGSAC Conference on Computer and Communications Security (CCS '14)*. ACM, New York, NY, USA, 674–689. DOI : <http://dx.doi.org/10.1145/2660267.2660347>
2. Lalit Agarwal, Nisheeth Shrivastava, Sharad Jaiswal, and Saurabh Panjwani. 2013. Do not embarrass: re-examining user concerns for online tracking and advertising. In *Proceedings of the Ninth Symposium on Usable Privacy and Security*. ACM, 8–24.
3. Michael Backes, Aniket Kate, Matteo Maffei, and Kim Pecina. 2012. Obliviad: Provably secure and practical online behavioral advertising. In *Security and Privacy (SP), 2012 IEEE Symposium on*. IEEE, 257–271.
4. Louise Barkhuus. 2012. The mismeasurement of privacy: using contextual integrity to reconsider privacy in HCI. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems*. ACM, 367–376.
5. Howard Beales. 2010. The value of behavioral targeting. *Network Advertising Initiative* (2010).
6. Steven Bellman, Eric J Johnson, Stephen J Kobrin, and Gerald L Lohse. 2004. International differences in information privacy concerns: A global survey of consumers. *The Information Society* 20, 5 (2004), 313–324.
7. Reinhardt A Botha, Steven M Furnell, and Nathan L Clarke. 2009. From desktop to mobile: Examining the security experience. *Computers & Security* 28, 3 (2009), 130–137.
8. California Legislature. 2013. AB-370 Consumers: internet privacy. (2013). https://leginfo.legislature.ca.gov/faces/billNavClient.xhtml?bill_id=201320140AB370
9. Jacob Cohen. 1960. A Coefficient of Agreement for Nominal Scales. *Educational and Psychological Measurement* 20, 1 (April 1960), 37–46.
10. Federal Trade Commission and others. 2009. FTC staff report: Self-regulatory principles for online behavioral advertising, 2009. *Federal Trade Commission, Washington, DC* (2009).
11. Federal Trade Commission and others. 2012. Protecting Consumer Privacy in an Era of Rapid Change: Recommendations for Businesses and Policymakers. (2012). <https://www.ftc.gov/sites/default/files/documents/reports/>
12. Paul Dourish. 2004. What We Talk About when We Talk About Context. *Personal Ubiquitous Comput.* 8, 1 (Feb. 2004), 19–30.
13. Xianyi Gao, Yulong Yang, Huiqing Fu, Janne Lindqvist, and Yang Wang. 2014. Private Browsing: An Inquiry on Usability and Privacy Protection. In *Proceedings of the 13th Workshop on Privacy in the Electronic Society (WPES '14)*. ACM, New York, NY, USA, 97–106. DOI : <http://dx.doi.org/10.1145/2665943.2665953>
14. Geert H. Hofstede. 1984. *Culture's consequences: international differences in work-related values*. Sage Publications.
15. Michael C. Grace, Wu Zhou, Xuxian Jiang, and Ahmad-Reza Sadeghi. 2012. Unsafe Exposure Analysis of Mobile In-app Advertisements. In *Proceedings of the Fifth ACM Conference on Security and Privacy in Wireless and Mobile Networks (WISEC '12)*. ACM, New York, NY, USA, 101–112.
16. Saikat Guha, Bin Cheng, and Paul Francis. 2010. Challenges in measuring online advertising systems. In *Proceedings of the 10th ACM SIGCOMM conference on Internet measurement*. ACM, 81–87.
17. Dehua He and Yaobin Lu. 2007. Consumers perceptions and acceptances towards mobile advertising: An empirical study in China. In *Wireless Communications, Networking and Mobile Computing, 2007. WiCom 2007. International Conference on*. IEEE, 3775–3778.
18. Hsiu-Fang Hsieh and Sarah E. Shannon. 2005. Three Approaches to Qualitative Content Analysis. *Qualitative Health Research* 15, 9 (Nov. 2005), 1277–1288.
19. Jian Hu, Hua-Jun Zeng, Hua Li, Cheng Niu, and Zheng Chen. 2007. Demographic prediction based on user's browsing behavior. In *Proceedings of the 16th international conference on World Wide Web*. ACM, 151–160.
20. S. Hu. 2000. *Individual Privacy in Modernizing China*. Ph.D. Dissertation. University of Colorado, Boulder.
21. Naz Kaya and Margaret J Weber. 2003. Cross-cultural differences in the perception of crowding and privacy regulation: American and Turkish students. *Journal of environmental psychology* 23, 3 (2003), 301–309.
22. Patrick Gage Kelley, Michael Benisch, Lorrie Faith Cranor, and Norman Sadeh. 2011. When are users comfortable sharing locations with advertisers?. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems*. ACM, 2449–2452.
23. Klaus Krippendorff. 2004. Reliability in Content Analysis: Some Common Misconceptions and Recommendations. *Human Communication Research* (July 2004), 411–433.

24. Rubén Daniel Ledesma and Pedro Valero-Mora. 2007. Determining the number of factors to retain in EFA: an easy to use computer program for carrying out parallel analysis. *Practical Assessment, Research & Evaluation* 12, 2 (2007), 1–11.
25. 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.
26. 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 Ninth Symposium on Usable Privacy and Security (SOUPS'13)*. ACM, 7–26.
27. Jialiu Lin, Michael Benisch, Norman Sadeh, Jianwei Niu, Jason Hong, Banghui Lu, and Shaohui Guo. 2013. A Comparative Study of Location-sharing Privacy Preferences in the United States and China. *Personal Ubiquitous Comput.* 17, 4 (April 2013), 697–711.
28. Jialiu Lin, Bin Liu, Norman Sadeh, and Jason I. Hong. 2014. Modeling Users Mobile App Privacy Preferences: Restoring Usability in a Sea of Permission Settings. In *Proceedings of Symposium On Usable Privacy and Security (SOUPS'14)*. 199–212.
<https://www.usenix.org/conference/soups2014/proceedings/presentation/lin>
29. Kirsten E. Martin. 2012. Diminished or Just Different? A Factorial Vignette Study of Privacy as a Social Contract. *Journal of Business Ethics* 111, 4 (March 2012), 519–539.
30. J.R. Mayer and J.C. Mitchell. 2012. Third-Party Web Tracking: Policy and Technology. In *2012 IEEE Symposium on Security and Privacy (SP)*. 413–427.
DOI: <http://dx.doi.org/10.1109/SP.2012.47>
31. Aleecia McDonald and Lorrie Faith Cranor. 2010. Beliefs and behaviors: Internet users' understanding of behavioral advertising. TPRC.
32. Bonnie S. McDougall and Anders Hansson (Eds.). 2002. *Chinese Concepts of Privacy*. Vol. 55. Brill Academic Pub, Leiden Netherlands ; Boston.
33. Helen Nissenbaum. 2004. Privacy as contextual integrity. *Washington law review* 79, 1 (2004), 119–158.
34. Helen Nissenbaum. 2011. A contextual approach to privacy online. *Daedalus* 140, 4 (2011), 32–48.
35. Eyal Peer, Gabriele Paolacci, Jesse Chandler, and Pam Mueller. 2012. Screening participants from previous studies on Amazon Mechanical Turk and Qualtrics. (2012).
experimentalturk.files.wordpress.com/2012/02/screening-amt-workers-on-qualtrics-5-2.pdf
36. Dan Pelleg, Denis Savenkov, and Eugene Agichtein. 2013. Touch Screens for Touchy Issues: Analysis of Accessing Sensitive Information from Mobile Devices. In *Proceedings of the Seventh International AAAI Conference on Weblogs and Social Media (ICWSM)*. 496–505.
37. Sren Preibusch. 2015. Privacy behaviors after Snowden. *Commun. ACM* 58, 5 (April 2015), 48–55. DOI :
<http://dx.doi.org/10.1145/2663341>
38. 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.
39. Florian Schaub, Bastian Konings, and Michael Weber. 2015. Context-Adaptive Privacy: Leveraging Context Awareness to Support Privacy Decision Making. *Pervasive Computing, IEEE* 14, 1 (2015), 34–43.
40. Somini Sengupta. 2013. Protecting Your Privacy on the New Facebook. *The New York Times* (Feb. 2013).
<http://www.nytimes.com/2013/02/07/technology/personaltech/protecting-your-privacy-on-the-new-facebook.html>
41. Steve Sheng, Bryant Magnien, Ponnurangam Kumaraguru, Alessandro Acquisti, Lorrie Faith Cranor, Jason Hong, and Elizabeth Nunge. 2007. Anti-Phishing Phil: The Design and Evaluation of a Game That Teaches People Not to Fall for Phish. In *Proceedings of the 3rd Symposium on Usable Privacy and Security (SOUPS '07)*. ACM, New York, NY, USA, 88–99.
42. Rudolf R Sinkovics, Noemi Pezderka, Parissa Haghirian, and others. 2012. Determinants of consumer perceptions toward mobile advertising: a comparison between Japan and Austria. *Journal of Interactive Marketing* 26, 1 (2012), 21–32.
43. Edith G Smit, Guda Van Noort, and Hilde AM Voorveld. 2014. Understanding online behavioural advertising: User knowledge, privacy concerns and online coping behaviour in europe. *Computers in Human Behavior* 32 (2014), 15–22.
44. H Jeff Smith, Tamara Dinev, and Heng Xu. 2011. Information privacy research: an interdisciplinary review. *MIS quarterly* 35, 4 (2011), 989–1016.
45. H Jeff Smith, Sandra J Milberg, and Sandra J Burke. 1996. Information privacy: measuring individuals' concerns about organizational practices. *MIS quarterly* (1996), 167–196.
46. Daniel J Solove. 2006. A taxonomy of privacy. *University of Pennsylvania law review* (2006), 477–564.
47. Zixue Tai. 2010. Casting the Ubiquitous Net of Information Control: Internet Surveillance in China from Golden Shield to Green Dam. *International Journal of Advanced Pervasive and Ubiquitous Computing* 2, 1 (2010), 53–70.

48. Mohsen Tavakol and Reg Dennick. 2011. Making sense of Cronbach's alpha. *International Journal of Medical Education* 2 (June 2011), 53–55. DOI : <http://dx.doi.org/10.5116/ijme.4dfb.8dfd>
49. Vincent Toubiana, Arvind Narayanan, Dan Boneh, Helen Nissenbaum, and Solon Barocas. 2010. Adnostic: Privacy Preserving Targeted Advertising. In *Network and Distributed System Security Symposium (NDSS'10)*.
50. TRUSTe. 2011. Privacy and online behavioral advertising. (2011). <http://www.truste.com/adprivacy/>
51. Joseph Turow, Jennifer King, Chris Jay Hoofnagle, Amy Bleakley, and Michael Hennessy. 2009. Americans reject tailored advertising and three activities that enable it. *Technical report, Annenberg School for Communications, University of Pennsylvania* (2009). http://repository.upenn.edu/asc_papers/137/
52. 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.
53. Yang Wang, Gregory Norcie, and Lorrie Faith Cranor. 2011. Who is concerned about what? A study of American, Chinese and Indian users privacy concerns on social network sites. In *Trust and trustworthy computing*. Springer, 146–153.
54. Ying Wang and Shaojing Sun. 2010. Examining the role of beliefs and attitudes in online advertising: A comparison between the USA and Romania. *International Marketing Review* 27, 1 (2010), 87–107.
55. Samuel D. Warren and Louis D. Brandeis. 1890. The Right to Privacy. *Harvard Law Review* 4, 5 (1890), 193–220.
56. Yue Jeff Zhang, Jim Q. Chen, and Kuang-Wei Wen. 2002. Characteristics of Internet Users and Their Privacy Concerns. *Journal of Internet Commerce* 1, 2 (March 2002), 1–16.