Unexpected and Unpredictable: Factors That Make Personalized Advertisements Creepy

Eelco Herder and Boping Zhang
eelcoherder@acm.org
Radboud University
Institute for Computing and Information Sciences
Nijmegen, The Netherlands

ABSTRACT

Personalized advertisements are the price we have to pay for free social media platforms. Various studies have been carried out on user acceptance of such advertisements in general and most countries have adopted laws and regulations with respect to privacy and data protection. However, not all advertisements evoke the same responses: some ads are considered more annoying, intrusive or creepy than others. In this paper, we present the results of an observational study on user responses to actual Facebook advertisements. The results show that mismatches in terms of context, unexpected data collection or inference, overly generic explanations and repetition are common causes of anxiety and distrust.

CCS CONCEPTS

• **Information systems → Personalization**; **Display advertising**; • **Security and privacy → Usability in security and privacy**; • **Human-centered computing → Laboratory experiments**.

KEYWORDS

online advertisements, user acceptance, empirical study

ACM Reference Format:
Eelco Herder and Boping Zhang. 2019. Unexpected and Unpredictable: Factors That Make Personalized Advertisements Creepy. In *23rd International Workshop on Personalization and Recommendation on the Web and Beyond (ABIS'19), September 17, 2019, Hof, Germany*. ACM, New York, NY, USA, 6 pages. https://doi.org/10.1145/3345002.3349285

1 INTRODUCTION

In commercial contexts, recommendations are often not only provided for the user's benefits, but also to serve a platform's goal, which is typically to make money. For example, web stores typically recommend products in the hope that users will buy these products and video platforms recommend videos that aim to keep the user entertained – and exposed to personalized advertisements, which can be seen as recommendations as well, but with the difference that the main beneficiaries are the advertising platform and the advertiser [4].

Many users encounter Facebook's personalized advertisements on a daily basis. As discussed in more detail in the related work section, the ad revenue enables Facebook to provide its platform for free to billions of users. However, there is another price that users have to pay: as stated in a popular Forbes article[1], "*if you're not paying for it, you become the product*". Several privacy concerns have been voiced regarding Facebook as a platform and Facebook ads in particular [15]. Furthermore, Facebook has been shown to have the power to influence a user's mood, attitude or political orientation [8]. Despite all these concerns and issues around Facebook and its ads, many users still continue to use the platform and therewith – implicitly – to accept the personalized, targeted advertisements.

In this paper, we aim to pinpoint factors that influence user acceptance of online advertisements: what are the reasons for users to consider particular ads creepy, and if an advertisement is considered not-creepy, does this mean that it is appreciated by the user? To what extent do explanations and transparency contribute to user acceptance, and what are good and bad practices from a user's point of view?

The paper is structured as follows. In the next section, we discuss related work on creepiness, privacy matters, personalized advertisements and explanations. In section 3, we explain the study setup, the participant pool and the coding process. The study results are presented in section 4, in which we provide insight in user responses to creepy and non-creepy ads, ad explanations and ad settings. After a discussion on the design implications in section 5, we end the paper with a summary and conclusions.

2 RELATED WORK

In this section, we discuss background literature on creepiness and technology anxiety in general, followed by a global overview on the relation between personalization and privacy. We then focus on online (personalized) advertisements and general user attitudes towards them.

2.1 Creepiness and Technology Anxiety

An empirical study on creepiness by McAndrew [10] revealed several factors that make us feel uneasy. The study starts with the observation that it is a non-pleasant sensation as a response to a *social mismatch* that puts us on our guard against a potentially

ABIS'19, September 17, 2019, Hof, Germany

ACM ISBN 978-1-4503-6896-4/19/09…$15.00
https://doi.org/10.1145/3345002.3349285

[1] https://www.forbes.com/sites/marketshare/2012/03/05/if-youre-not-paying-for-it-you-become-the-product/

threatening person or situation. The study results reveal contributing factors that include unusual patterns of nonverbal behavior, odd emotional responses, or physical characteristics that are outside of the norm. What all these factors have in common is the concept of *unpredictability*.

The introduction of new technologies has consistently met *anxiety* and fear: for instance, at the time of introduction, the train and the telegraph were described as 'creepy', mainly because social values were yet to line up[2]. In a similar vein, recent research showed that generally, users tend to find personalized advertising creepy, but at the same time only had limited understanding of how they work [17].

In sum, the key to creepiness appears to be the uncertainty about possible threats and the uneasiness due to a lack of social norms.

2.2 Personalization and Privacy

Nissenbaum [12] formulated three principles that dominate the discussion surrounding (online) privacy. The first principle concerns the protection of individuals against intrusive agents, most notably the government, the second principle concerns restriction of access to intimate, sensitive or confidential information, and the third principle concerns intrusion into spaces that are deemed private or personal.

Kobsa [7] observed that personalized interaction and user modeling have significant privacy implications, due to the personal data needed to be collected for this purpose. He argued that highly *sensitive data categories* (including purchases, income and political party affiliation) should never be requested without mitigating factors, most notably the (perceived) user benefits of personalization as well as knowledge of and control over the use of personal information.

Paramythis et al [13] proposed the concept of *layered evaluation* of adaptation, with several evaluation criteria regarding the quality, appropriateness and necessity with respect to the collection of input data, the interpretation of the data, the way it is represented in a user model, how it is used for deciding upon adaptation decisions (for example a recommendation or an advertisement) and the way it is being presented to the user.

Knijnenburg et al [5] argue that users constantly face a *trade-off* between the benefits of personalization and potential risks, such as providers sharing data with third parties, exposure of sensitive information, targeted advertising, and discrimination. In the article, several technical solutions are discussed, but also note that there are limitations: most users hardly make use of privacy controls, and platform providers have a wide range of strategies for justifying the use of personal data.

In this paper, we focus on user perception of Facebook advertisements, particularly on the factors that make users consider the advertisements creepy. Facebook has a long history of privacy incidents [15], including the release of an ad platform that targets ads based on a user's profile and behavior within the platform, overly sharing user activity and photos maintain overview and control.

Burke [4] stated that recommender systems often serve the needs of multiple parties and proposed several types of *fairness-aware recommender systems* that balance these different needs. As an example, recommendations (or advertisements) targeted at a user should be of the user interest, but also satisfy the corresponding advertisers, and the needs of the platform.

[2]https://ieet.org/index.php/IEET2/more/selinger201208271

2.3 Personalized Advertisements

The concept of targeted ads is not new. Already in the 1930s, sociologist Paul Lazarsfeld on which occasions laundry services were used and, based on the results, suggested that marketeers for laundry services should target housewives after announcements of births, deaths, and weddings [16].

Baek et al [1] define personalized advertisements as "a form of *customized* promotional messages that are delivered to each consumer through paid media based on personal information". This personal information includes demographics, psychographics, lifestyle, interests or (online and physical) behavioral information.

A particular personalization strategy that can be observed, among others on Facebook, is to include personally identifiable information in the content [9], as a cue to remind the user that the ad was specifically designed for him or her [2]. However, when personalization takes relevance *too far*, this may lead to the uneasy sense that marketeers are capitalizing on information or online activities that the user perceives as *private*.

2.4 Attitude to Personalized Ads

Ur et al [21] found that users justified the existence of ads, because it pays for free online services. Nevertheless, most participants stated that ads should not be too obtrusive. Users also often expressed substantial *privacy concerns* [20]. The most prominent factor that invokes anxiety are *unexpected data collection practices* [11]. Users do not expect or accept extensive or aggressive data collection practices, and at the same time they do not know exactly how they work [21].

Barnard [3] showed that user attitudes toward tailored online media were increasingly negative when the tailoring involved information that was *more personal*. Specifically, behavioral information, such as the websites visited by a user, is regarded as more personal than demographic information, such as age or school. Attitudes towards tailored political ads were especially negative [19].

2.5 Explanations and Transparency of Advertisements

Personalized ads essentially aim to recommend more relevant ads to users in order to improve ad performance, creating value for both advertisers and users, and therefore can be considered as a type of recommender systems. An effective strategy for increasing user acceptance of personalized ads are *explanations*, which increase transparency, scrutability, trust, persuasiveness, effectiveness, efficiency and satisfaction [18].

For this purpose, Facebook's ad system provides three types of explanations: ad-specific explanations ("Why am I seeing this ad"), general explanations ("About Facebook Ads") and ad settings. In a recent study, Kim et al [6] came to the rather obvious conclusion that more transparency would *increase* ad effectiveness if the associated data collection practices would be deemed acceptable; unacceptable data collection practices (such as tracking outside of the website or aggressive profiling and inference of user traits) would decrease ad effectiveness.

Figure 1: Distributions of participants' education and background

3 METHODOLOGY

The preliminary set of general causes why advertisements were considered creepy, as discussed in the previous section, served as a basis for the design of an exploratory study, which involved a scenario-based evaluation session and two interviews. The exact setup will be explained in more detail in the remainder of this section.

3.1 Participants

The main criterion for participation was prior exposure to creepy ads, primarily in Facebook and/or Instagram. Respondents were recruited by convenience and snowball sampling.

Twelve respondents participated in the study, of which 50% were female. Two participants were American native English speakers, the remaining ten participants European non-native English speakers. The age ranges from 22 to 50 (*mean* = 29.4, *median* = 25.0, *s.d* = 8.6). They represent a variety of backgrounds, nationalities, occupations and industries – see Figure 1. However, the sample is skewed to experienced Facebook users, with an average of 8.8 years since registration (*median* = 8.5, *s.d.* = 2.4), and to a higher-educated population (83% hold a pre-master's degree or higher).

3.2 Study Setup

Each interview session consisted of three parts, lasting from 40 to 60 minutes. 8 out of 12 participants were interviewed face to face, other sessions were conducted via *Skype for Business*, which allows the interviewee and interviewer to share their screens and see each other during the interview. All interviews were conducted in English.

First, the participants were asked to fill in a pre-interview questionnaire, to collect basic background information. The questionnaire also provided details about the collection, use and processing of the data. Consent to record, transcribe, and analyze the data for research purposes was asked.

The scenario-based evaluation started with a brief warm-up, in which the participants were asked about their general use of Facebook and other social media, their attitudes towards Facebook ads, and creepy ads that they experienced before. Then, participants were instructed to log into their Facebook account[3] and to take screenshots of the first 3 sponsored posts and their corresponding ad-specific explanations. Thoughts and spontaneous comments were encouraged, but not mandatory. Subsequently, participants were asked to share their feelings and opinions about the tasks. As a final part of the observational study, participants were instructed to browse through the general ad explanations and the ad settings.

Finally, participants self-defined their level of privacy concerns and shared their knowledge on privacy-related issues in general and specifically about Facebook.

3.3 Coding

From the literature review, descriptions of possible creepy ad factors were derived and translated into a codebook. Then, all specific ad experiences were transcribed, numbered and coded using the codebook developed earlier. New codes were added if they were not captured by existing codes. With respect to creepiness and user attitude, six types of ads have been derived and analyzed. The ad-specific explanations were assigned to one of the seven possible explanatory goals of recommendation systems, as proposed by Tintarev and Masthoff [18].

Since the pool of participants was not (and was not meant to be) randomly selected, the findings reported are exploratory and interpretive.

4 RESULTS

Most Facebook ads are shown as 'sponsored' posts in between the regular newsfeed. Our participants had mixed feelings about this. Some appreciated the subtlety, which minimizes interruptions, while others expected more distinctions, so that they would not unintentionally click on them. And perhaps because of this low-key presentation, users did not realize they saw so many ads every day until they were interviewed.

> *That's funny, before this interview, I was not aware I get so many ads.* (R09)

> *But also they are a bit invasive. (...) I'm more interested in knowing about or seeing the pages I followed, posts from my friends, but not ads.* (R04)

Generally speaking, users justified the existence of online ads, because the ads 'paid' for the free services. Some also justified the utilization of personalization to make the ad work better for users.

Despite this overall relatively positive attitude, our study participants also reported several incidents that they did not feel comfortable with. These 'creepy' incidents will be discussed in the next subsections.

[3] R08 deactivated her Facebook account shortly before she was invited to this study, so she did not participate in this activity.

4.1 Responses To Creepy Ads

During the pre-interview and the observational study, a total of 45 ad incidents have been discussed with the 12 study participants. Among these, 26 advertisements did not cause any anxiety – but, as will be discussed in Section 4.2, they were not overly positively received either. 19 ads (42%) of the ads discussed were considered creepy – this high percentage is most likely caused by the elicitation method and does *not* represent the amount of creepy ads on Facebook.

Unsurprisingly, the overwhelming majority of creepy ad incidents (17 out of 19) were attributed to **unacceptable/unexpected data collection practices**. It should be noted that these are *perceived* practices, as reported by the user, and not necessarily what actually happened behind the screens.

> *I was telling my parents about it [Radler Beer] on Skype. And all of the sudden it showed up on my news feed on the next day. I thought that was weird because I had not had typed it anywhere.* (R01)

The majority of instances concerned everyday topics, including holidays (2x), beer (2x), flights, shops and tickets. These topics themselves are usually not embarrassing or related to sensitive information, but they are typically *new* and *unfamiliar* to the user. Moreover, in most instances this topic appeared in an ad shortly after the user first encountered it. These ads were perceived as *targeted* and *specific*.

> *I think it's just too much of a coincidence that you cannot exclude that they are tracking your voice at least.* (R12).

Whether data collection practices were considered creepy also depended on their *acceptability* – in line with [6], practices that violated social norms or where no social norms have been developed (yet), were deemed unacceptable.

> *I do not recall typing on any flights to Buenos Aires, but from the conversation somehow maybe this happened. This is still an unsolved mystery that I found pretty weird.* (R04)

Apart from the (perceived) data collection practices, the **context** in which the ad was shown was relevant for user acceptance as well. In one recorded incident, an ad was shown in the participants' Messenger app.

> *I can also see ads in the Facebook Messenger app. I feel that is a bit too much (...) [They] make me feel that Facebook is watching me while I'm talking.* (R07)

As a final factor that influenced the user's perceived creepiness of advertisements, the **sensitivity of data** was mentioned. They were concerned that this information may be undesirably spread:

> *I have [a certain health issue], I think I look into solutions about [the health issue], and then I have the ads on Facebook, that's already something I do not feel comfortable about it.* (R12)

To summarize, the observed reasons for users considering ads creepy are in line with the literature discussed in section 2.4: unexpected or not-accepted data practices and inappropriate use of personal or sensitive information.

4.2 Responses To Non-Creepy Ads

Out of the 45 collected instances, the majority of ads (26) were considered non-creepy. However, 23 of them were either ignored or caused negative emotions.

Most non-creepy ads were deemed *unrelated* to our participants' (current) activities or *irrelevant* and thus uninteresting. A typical user response would be to *ignore* these ads:

> *Unless they are surprising or shocking, I just scan through and scroll through them. And I do not even pay attention (...) I just skip right over them.* (R01)

Even though unrelated or irrelevant ads at first did not cause negative responses, *repetitive appearance* turned out to annoy the participant to a great extent. Six incidents were reported, in which participants complained about 'mechanic repetition'. Frustration about repetition increased when *inferences were deemed incorrect*, especially when they seemed to be based on products that they searched for even *just once*:

> *There are some websites I entered by accident (...) but they keep on reminding me about flights to Hong Kong. I'm not planning to go to Hong Kong. I just wanna to check how much it is.* (R11)

Repetitive incorrect advertisements caused feelings of confusion among our participants, triggering a sense that they were being *stereotyped* [14]:

> *I searched for vegan things it doesn't mean that I'm vegan. They are putting me in buckets, which I do not really feel I'm part of.* (R06)

In several cases, including the incident above, the repetitive advertisement was thought to be the result of logged or inferred *sensitive information*, including those regarded as 'special categories of personal information' in the European GDPR[4]:

> *I've been looking for jobs back in [a US city]. There is one job. This company specialized in travel and cruise and vacations for lesbian women. I've then seen the ads for lesbian cruises on Facebook.* (R04)

Only 3 out of 45 ads were considered *beneficial and harmless*. It was observed that users first assessed whether ads were shown in an appropriate context, did not involve sensitive personal data, followed acceptable ad practices, and only once these conditions were fulfilled, came to assess the *usefulness*. In order to be useful, an ad did not only need to match a user's profile or activities, but also to match a user's actual need or goal. Particularly ads that provided *discount* on a product that a user was considering to buy were likely to be clicked on:

> *It can be convenient [if] I wanna buy a car and then a car thing comes up. Wow, that's convenient. Or the dress is on discount or something.* (R11)

During the study, we noticed that this type of reasoning about acceptability and usefulness did not only apply to online advertisements, but also to *recommendations in general*. For instance, when asked her attitude toward personalization in general, R08 referred to Spotify. She recognized it offers *concrete benefits* ("the music that you like"), believed that only *specific data* ("the music

[4] https://gdpr-info.eu/art-9-gdpr/

you are listening to") was tracked, and *perceived no risk* to share such data. However, when further asked if Spotify also used some of the unacceptable practices that Facebook used, she somehow accepted the possibility that these might happen, in exchange for the clear purpose and benefits that Spotify offers.

4.3 Responses to Ad Explanations

Next to user responses to the online advertisements themselves, we were interested in how they would react to ad-specific explanations. Transcripts of all interviews were examined and coded to one of the seven explanatory goals [18]. Table 1 shows the number of users who mentioned each goal, split into negative and positive remarks.

Table 1: Frequency table of seven explanatory goals and corresponding codes.

Explanatory Goal	Positive	Negative
Transparency	5	5
Scrutability	3	2
Trust	1	6
Persuasiveness	0	1
Effectiveness	0	0
Efficiency	0	2
Satisfaction	5	6

One of the first and most interesting findings about the Facebook ad explanations and settings is that *almost none* of the users were aware of either of them. Participant R03 indicated that he was aware of it, but did not "take the time to read", because "it's pretty obvious [why they show me the ads] most of the time and most of the explanation was pretty obvious". Other participants have not proactively looked for explanations and settings:

> *I like this interview because I've got to explore the pages from Facebook I did not even bother to investigate.* (R04)

The general explanations ('About Facebook Ads') were generally appreciated, although some participants (R04, R06) thought they were "overly simplified": they perceived this "dumped-down" and seemingly "non-invasive, friendly, harmless" and almost "innocent" explanations as *deceiving* and doubted the credibility.

> *I am also curious about what is not explained* (R03)

Participants responded in a similar manner to the ad-specific explanations – see Figure 2 for an example –, commenting upon them as being incomplete, vague and "all the same".

> *I'm not sure if they are completely honest. For example, there is an ad about [a chili sauce]. It says people ages between 18 and 40 who lives in the Netherlands. That's the only explanations and I think they should be more to it that [why] they are targeting me.* (R02)

Several participants demanded *more specific explanations*, especially for ads that they perceived as clearly targeted. For R02, "it really wouldn't be shocking" if the explanation of the chili sauce ad showed that, for instance, five days ago he browsed recipes using that sauce. Such an explanation would make Facebook "*more trustworthy*" to him. Particularly for very specific ads, a *discrepancy* was felt between the ad explanation and the suspected actual targeting strategy.

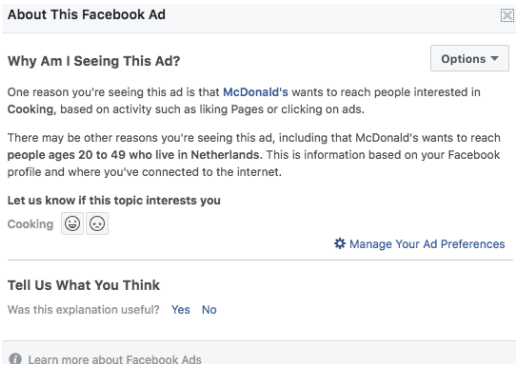

Figure 2: An example ad-specific explanation (for R04)

4.4 General Awareness and Use of Ad Settings

Users were generally aware of many tracking practices, such as collecting data from public posts, public profile, search history, and browsing behavior. Users also believed that these practices were used across the network and the collected data is shared and sold among companies. Cookies were the most-mentioned technology, despite users could seldom articulate how exactly it works.

In the final part of the observational study, during which participants browsed the general ad explanations and ad settings, the discovered practice that they commented most upon was Facebook Pixel[5]. Users found it surprising and concerning because they "do not even need to be logged in" to be tracked.

Other surprising practices included tracking location/GPS data and "profile matching", in which Facebook compares a user with an example profile of "ideal customers" for targeting purposes.

Probably triggered by the detailed examination of ad explanations and settings, as well as by the recapitulation of previous creepy ad incidents in this study, 8 out of 11 participants chose to turn off all personalization options. Two participants did not, indicating that they "do not really care" (R07) and have "nothing to hide" (R02). One participant (R03) felt simultaneously defeated and not willing to spend the time adjusting the ad settings.

Finally, shortly after the study, R08 – who had deactivated her Facebook account just before the study took place – reported that she reactivated her account again, because "all [school name] information is posted on Facebook, so I need an account for staying up-to-date".

5 DISCUSSION AND IMPLICATIONS

Even though users understood and accepted the reasons for advertisements on Facebook, unsurprisingly *not all advertisements were well-received* and about 42% of the ads discussed in this study were considered 'creepy' and the majority of the remaining ads irrelevant or annoying.

It appears that users first assess the acceptability of the ad: is it shown in an appropriate context, isn't the ad based on sensitive personal data, and were the data collection or inference practices acceptable? *Only if these conditions are fulfilled* and the ad is not

[5]https://blog.hootsuite.com/facebook-pixel/

considered 'creepy', users will assess whether the ad is useful or not.

Blending sponsored recommendations or advertisements within the actual content – be it a news feed, search results or a list of 'genuine' recommendations – received mixed responses. They were considered subtle and invasive at the same time. As long as the sponsored recommendations are in line with the user context, they might sufficiently blend in, but as soon as the users sense a *mismatch* [10], and the sponsored content is not labelled as such, users start to feel uneasy and might start to distrust the user context as a whole, particularly if the user context is *considered private*, such as a chat window.

User concerns regarding data collection practices and inferences were not as much related to the sensitivity of data or the outcomes as to the *unexpectedness* of the sponsored recommendation. If a specific item is recommended or advertised that is (from the user's point of view) based on past behavior that the system is *not supposed to know*, and if this item is specific enough, users might start to suspect they are being tracked or observed [11]. Particular when these practices appear not to follow social norms and to be unpredictable, a natural response is to feel anxious and remain alert. This even applies for seemingly harmless data or inferences.

Explanations and transparency seem to reduce users' anxiety and increase trust only to a certain extent – and in some cases lead to more distrust. Even if users are aware of these mechanisms, they usually do not proactively look for them, or expect the explanations to be 'pretty obvious'. In line with [6], we observed that users were dissatisfied with explanations that are *not specific enough*. Particularly in the case of very specific recommendations, users expect that they are targeted to a very specific group and *do not believe* an overly vague or generic explanation (like 'reaching out to a broad audience'). Moreover, they start wondering what aspects are (deliberately) left out of the explanation.

If the conditions discussed above are fulfilled, recommendations or advertisements that are unrelated or irrelevant to the user are usually simply ignored, and relevant, useful recommendations have a fair chance to be taken up, particularly if there is an *incentive*. However, *repetitions* of irrelevant recommendations were regarded as annoying and, moreover, caused feelings of confusion or anxiety, triggering a suspicion that users were incorrectly being stereotyped, being given a label that they do not feel associated with.

6 CONCLUSIONS

In this paper, we reported the results of an in-depth observational study on user acceptance of and response to personalized online advertisements in Facebook. Participants reported about prior experiences, were exposed to three personalized ads and the corresponding explanations, and interacted with Facebook ad explanations and settings.

A limitation of this study is that it is based on a relatively small sample of twelve participants. However, we think the specific insights obtained through in-depth interviews and user responses to solicited and previous experiences with actual online advertisements provide important insightful details on user attitudes and responses that complement the more quantitative and general observations in the existing body of literature.

The results indicate that, in order to be accepted, advertisements or sponsored recommendations should appear in an acceptable context and/or be labelled as such, as perceived unexplained mismatches lead to distrust. Further, user concerns regarding data collection and inferences are often not about the sensitivity of the data, but on the unexpectedness of these practices, which may lead users to suspect they are secretly being monitored. Explanations may reduce these issues, but only if the users perceive them as specific enough – overly generic explanations may have the opposite effect. As a final consideration, if users repetitively encounter the same irrelevant ad, even if at first it was deemed acceptable, the repetition may start users to think that something is going on that does not follow social norms or user expectations.

REFERENCES

[1] Tae Hyun Baek and Mariko Morimoto. 2012. Stay away from me. *Journal of advertising* 41, 1 (2012), 59–76.

[2] Lisa Barnard. 2014. *The cost of creepiness: How online behavioral advertising affects consumer purchase intention.* Ph.D. Dissertation. The University of North Carolina at Chapel Hill.

[3] Lisa Barnard and Daniel Kreiss. 2013. A research agenda for online advertising: Surveying campaign practices, 2000-2012. *International Journal of Communication* 7 (2013), 21.

[4] Robin Burke. 2017. Multisided fairness for recommendation. In *Proc. FAT/ML - 4th Workshop on Fairness, Accountability, and Transparency in Machine Learning, co-located with KDD 2017.*

[5] Arik Friedman, Bart P Knijnenburg, Kris Vanhecke, Luc Martens, and Shlomo Berkovsky. 2015. Privacy aspects of recommender systems. In *Recommender Systems Handbook.* Springer, 649–688.

[6] Tami Kim, Kate Barasz, Leslie K John, and Andrea Morales. 2018. Why Am I Seeing This Ad? The Effect of Ad Transparency on Ad Effectiveness. *Journal of Consumer Research* (2018).

[7] Alfred Kobsa. 2007. Privacy-enhanced web personalization. In *The adaptive web.* Springer-Verlag, 628–670.

[8] Adam DI Kramer, Jamie E Guillory, and Jeffrey T Hancock. 2014. Experimental evidence of massive-scale emotional contagion through social networks. *Proceedings of the National Academy of Sciences* (2014), 201320040.

[9] Mia Liza A Lustria, Juliann Cortese, Seth M Noar, and Robert L Glueckauf. 2009. Computer-tailored health interventions delivered over the Web: review and analysis of key components. *Patient education and counseling* 74, 2 (2009), 156–173.

[10] Francis T McAndrew and Sara S Koehnke. 2016. On the nature of creepiness. *New Ideas in Psychology* 43 (2016), 10–15.

[11] Aleecia McDonald and Lorrie Faith Cranor. 2010. Beliefs and behaviors: Internet users' understanding of behavioral advertising. (2010).

[12] Helen Nissenbaum. 2004. Privacy as contextual integrity. *Wash. L. Rev.* 79 (2004), 119.

[13] Alexandros Paramythis, Stephan Weibelzahl, and Judith Masthoff. 2010. Layered evaluation of interactive adaptive systems: framework and formative methods. *User Modeling and User-Adapted Interaction* 20, 5 (2010), 383–453.

[14] Elaine Rich. 1979. User modeling via stereotypes. *Cognitive science* 3, 4 (1979), 329–354.

[15] Ira S Rubinstein and Nathaniel Good. 2013. Privacy by design: A counterfactual analysis of Google and Facebook privacy incidents. *Berkeley Tech. LJ* 28 (2013), 1333.

[16] Wilbur Schramm. 1997. *The beginnings of communication study in America: A personal memoir.* Sage.

[17] Omer Tene and Jules Polonetsky. 2013. A theory of creepy: technology, privacy and shifting social norms. *Yale JL & Tech.* 16 (2013), 59.

[18] Nava Tintarev and Judith Masthoff. 2011. Designing and evaluating explanations for recommender systems. In *Recommender systems handbook.* Springer, 479–510.

[19] Joseph Turow, Michael X Delli Carpini, Nora A Draper, and Rowan Howard-Williams. 2012. Americans roundly reject tailored political advertising. (2012).

[20] Joseph Turow, Jennifer King, Chris Jay Hoofnagle, Amy Bleakley, and Michael Hennessy. 2009. Americans reject tailored advertising and three activities that enable it. (2009).

[21] 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.* ACM, 4.