

# User Attitudes Towards Controls for Ad Interests Estimated On-device by the Browser

Florian Lachner  
Google  
lachner@chromium.org

Minzhe Yuan Chen Cheng  
Google  
yuanchen@chromium.org

Theodore Olsauskas-Warren  
Google  
sauski@chromium.org

**Abstract**—Online behavioral advertising is a double-edged sword. While relevant display ads are generally considered useful, opaque tracking based on third-party cookies has reached unfettered sprawl and is deemed to be privacy-intrusive. However, existing ways to preserve privacy do not sufficiently balance the needs of both users and the ecosystem. In this work, we evaluate alternative browser controls. We leverage the idea of inferring interests on users’ devices and designed novel browser controls to manage these interests. Through a mixed method approach, we studied how users feel about this approach. First, we conducted pilot interviews with 9 participants to test two design directions. Second, we ran a survey with 2,552 respondents to measure how our final design compares with current cookie settings. Respondents reported a significantly higher level of perceived privacy and feeling of control when introduced to the concept of locally inferred interests with an option for removal.

## I. INTRODUCTION

Almost 30 years ago, a mechanism to sustain states between clients and servers was developed. This mechanism is based on a text string that is placed on a client browser when it accesses a specified server and is commonly known as ‘cookie’ [11].

Cookies are used for a variety of purposes, such as helping users stay signed in to websites, remembering items in shopping carts, or personalizing ads when browsing online. Cross-site ads personalization is enabled by so-called third-party cookies, which, in contrast to first-party cookies, are placed by a domain that is not the visited site. Although the initial intent of cookies was to improve the user experience of online browsing, the term cookie has gained a negative connotation in past years [38]. This has been driven by increased prevalence of online tracking over time, as well as pervasive cross-site tracking based on third-party cookies becoming a norm for Online Behavioral Advertising (OBA) in recent years [44].

Due to the increasing intensity of public discussions about privacy concerns of cookie usage for OBA, established browser vendors have announced in recent years their plans to stop supporting third-party cookies [58]. On one hand, this can be seen as a privacy improvement, on the other hand, blocking third-party cookies by default can have significant implications

on users’ online browsing experiences and disrupt today’s internet infrastructure [32]. Also, prior studies show that people tend to perceive personalized ads as useful [6], [49], [50], [68].

Maintaining relevant ads while improving user privacy is a known challenge. Topics-based personalization is a direction that browsers are moving in [15], [59]. Browsers are also moving away from more traditional forms of cookie-based personalization (see [64]). Given these trends, we wanted to identify and evaluate usable interfaces for topics-based personalization compared with current cookie controls. We build upon the Topics API proposal from Google, i.e., a browser infers a user’s topics of interest and stores them on the user’s device in a privacy-preserving way [16]. Our work is guided by the following question:

**RQ:** *Do users perceive a higher value of personalized ads for users and publishers and a higher level of privacy with browser controls for inferred topics of interest compared to current third-party cookie controls?*

We contribute to the field of Usable Security and Privacy in two ways: First, we draft and compare two interaction patterns for managing browser-facilitated interests for ads personalization. Second, we evaluate user attitudes in a large-scale online survey about the envisioned concept for browser-facilitated ad interests in a world without third-party cookies.

## II. BACKGROUND AND RELATED WORK

People’s perceptions of OBA and tracking limitation mechanisms have been discussed extensively within both academia and industry (see, e.g., [12], [44]). In the following sections we synthesize prior insights about ads personalization and controls.

### A. Online Behavioral Advertising

To increase the relevance of online ads, advertising companies and publishers, i.e., site owners, aim to tailor ads to the interests of website visitors. Throughout the years, tailoring mechanisms have become increasingly nuanced.

A simple approach to tailoring an ad is to choose an ad based on the content of the page, i.e., *contextual advertising*. In contrast, online *behavioral advertising* represents an advanced tailoring approach that chooses an ad based on a user’s online activities, e.g., the sites they visited in the past. This is often referred to as *interest-based advertising* [66].

Another approach is identifying a user’s location to tailor an ad, i.e., *location-based advertising*. An even more direct tailoring mechanism is *retargeting*, i.e., selecting an ad based on a distinct website visit or site interaction [13], [51], [72]. In addition, Wei et al. point out that advertisers can combine different data types for behavioral targeting on, e.g., social media sites [70]. Additionally, recent work has expanded the perspective on behavioral advertising, e.g., by investigating opportunities in rewards-based advertising platforms [53].

With the advancements of online personalization services come inherent and inevitable privacy tradeoffs [2]. In the context of OBA, users are shown relevant ads at the cost of their online behavior being tracked. The complexity of this tradeoff becomes clear when contrasting prior studies, which found diverse user attitudes and perceptions of OBA [8], [73]. Users initially perceive OBA as privacy-intrusive when prompted, particularly if content is perceived as too personalized, however, a majority of users accept personalized ads in exchange for free content and perceive these ads as more likely to be useful than non-personalized ads [46], [49], [50], [66]. In this context, the level of comfort with advertising and tracking depends on the users’ overall privacy attitudes [12]. Unease with OBA can also be reduced when users feel assured that OBA is based on non-personally identifiable information [66].

## B. Tracking Reduction

Research around third-party web tracking dates back to the beginning of the 21st century [44]. In the earlier days, most contributions centered around the reduction of tracking through, e.g., the Do Not Track header and legislation, pop-up blocking, prevention of tracking through social media plugins such as “Like” buttons (see [14], [37], [56]), pseudonymized network layers (see [27]), or same-origin policy enforcements (see [31]).

Prior work has shown that users appreciated it when a browser has built-in tracking reduction functionality while external online opt-out tools were difficult to set up and manage [43]. Further work has studied the application of differential privacy for web analytics (see [3]) and the use of machine learning to counter web tracking (see [7]). Although third-party cookies can directly be controlled within the browser, these controls can be perceived as very technical [39], [62]. Also, site-specific privacy settings on, e.g., social network sites score low on awareness, comprehension, and usability [5].

A common goal of tracking reduction approaches is blocking or limiting how third parties can place and/or read cookies on a user’s device. In contrast, Toubiana et al. have suggested *Adnostic*, a behavioral profiling mechanism based on local interest categorization in the browser [65]. Guha et al. discussed a similar idea as part of their work on *Privad*, a practical private online advertising system [24].

More recently, machine learning approaches have also been investigated to limit cookieless tracking mechanisms, most prominently fingerprinting [30]. Fingerprinting is a covert approach that doesn’t leave cookies or other evidence on the user’s device by inferring a user’s identity based on their browser and operating system properties [44], [56].

Since the European Union passed the General Data Protection Regulation (GDPR) in 2018, discussions about cookies have picked up speed again. Sanchez-Rola et al. conducted a comprehensive study of cookie banners that span the web today [56]. The authors question the effectiveness of the cookie banners when it comes to reducing tracking overall. In addition, many users are irked by these banners [26], [38]. What’s more, the design of cookie banners is often complex which may trick users into a certain behavior through the use of dark patterns [20], [28].

Users have increasingly adopted browser extensions that simply block ads over the past years [10], [18], such as *AdBlock* (see [29]) or auto-respond to cookie banners, such as *I don’t care about cookies* (see [35]). However, many ad-blocker users adopt such tools mainly to improve their online browsing experience, rather than for privacy reasons [48]. Also, while most ad blocking services focus on tracker-based blocking, i.e., who is tracking, Yao et al. point out that users prefer information-based blocking, i.e., what is being tracked [73].

## C. Role of the Browser

Current practices of inferring ads interests are technically complex and often perceived as not intuitive, leading to a decrease in trust in ads personalization [49], [69]. While there seems to be general awareness of the use of cookies in the context of OBA, users have a hard time articulating what cookies are and how they are used [1], [66].

In addition, people’s perceptions of the web ecosystem in the context of OBA have not been studied in-depth. Prior work has yielded four different folk models about the mechanisms of OBA (see [73]). While the goal of Yao et al. was not to evaluate the correctness of people’s folk models, it becomes apparent that mental models vary [73]. In particular, while it is considered a critical aspect for OBA, people have differing perceptions about the relationship between the browser and sites.

## D. User Agency

The continuous opacity of how ads are tailored online has led to increased calls for transparency [17]. Various studies have showcased the importance of and demand for control over user data, which goes hand in hand with transparency [4], [9], [25], [36]. In the context of inferred interests, users who care about the mechanics of personalization have a substantial urge for transparency and control [57]. However, in an online world with numerous ways to store personal data extracted from a variety of sources with a lack of user feedback, agency is not equally distributed between online corporations and users [43], [52], [67]. Privacy transparency tools, e.g., Google’s My Activity, in contrast, have proven to increase trust in online services and increase user agency in the context of web tracking [19], [71]. At the same time, Seberger et al. point out that enhancing users’ feelings of control by giving more privacy choices may lead to empowerment resignation [60]. In addition, the effectiveness of transparency may depend on additional factors, such as data literacy [61].

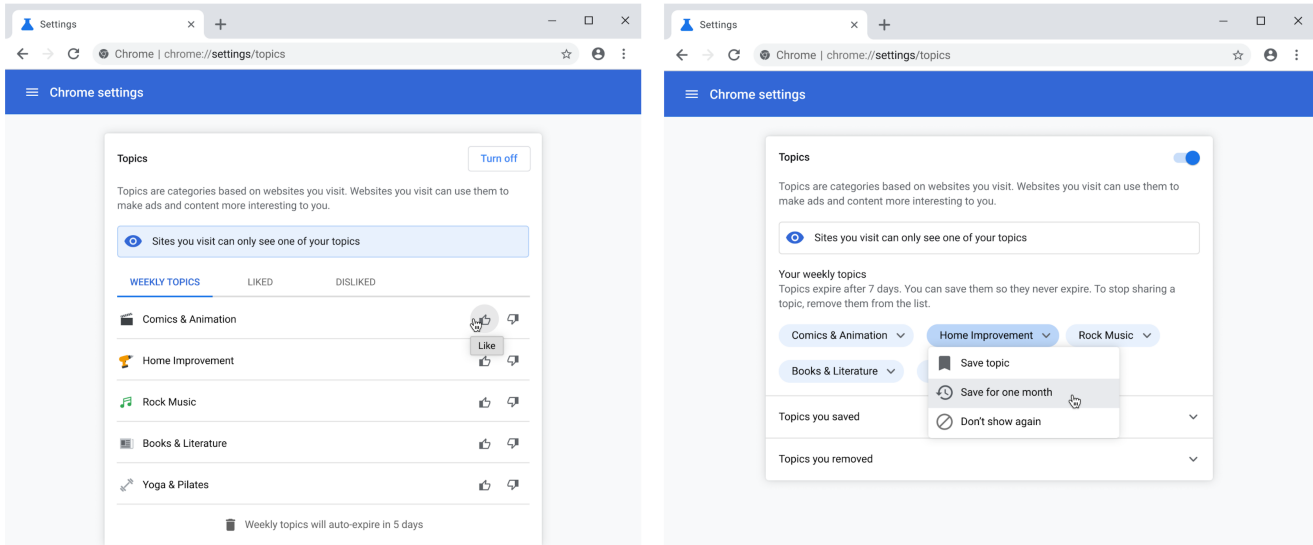


Fig. 1. Prototypes for the qualitative pilot study based on a like/dislike interaction (left) and topic removal (right).

### III. METHODOLOGY

Based on our research question and the analysis of related work in the fields of OBA, tracking reduction, browser mental models, and user agency, we derive the main hypotheses below (H1 - H3). The main hypotheses focus on users' perceptions of controls for inferred topics of interest:

- **H1:** *Users infer a higher level of privacy from topic-based UI compared to cookie-based personalization.*
- **H2:** *Users perceive a higher feeling of control with inferred topics of interest compared to cookie controls.*
- **H3:** *Privacy concerned users perceive a lower value in controls for inferred topics of interest.*

In addition, we derive the following sub-hypotheses (H4 - H6) that refer to aspects related to the interplay between users and publishers, i.e., the sites they visit:

- **H4:** *Users recognize an association between the personal value of personalized ads and the value for publishers.*
- **H5:** *Users expect that sites can financially benefit from site interactions in a topic-based personalization scenario.*
- **H6:** *Users perceive a higher personal value than value for publishers from topic-based personalization.*

We pursued a multi-step approach to validate these hypotheses. The goal of the first three steps A, B, and C was to create a foundation for the design of browser controls for inferred topics of interest. The goal of the final step D was to validate the hypotheses mentioned above:

- Infer technical assumptions from industry approaches.
- Create design directions for topics of interest controls.
- Select one prototype and finalize the design based on user feedback from a qualitative pilot study.
- Validate user attitudes in a large-scale survey study.

#### A. Technical Assumptions

One of the core privacy issues of using third-party cookies for OBA is that companies can track specific sites a user visits and thus re-identify the user and their online activity [56]. As outlined earlier, prior work has suggested privacy-preserving user profiling through inferring topics of interest on a user device without revealing the exact websites, as an alternative approach [24], [65]. Google has recently announced a similar strategy for their Topics API as part of their third-party cookie phaseout plans [15], [59]. We base our technical assumptions for our interface design on the same idea, i.e., the browser locally processes the sites a user visits and infers topics of interest. As described in Google's developer guide, we have applied the following assumptions for our prototypes [16]:

- The browser can share inferred interest with sites.
- The browser cannot share the specific sites a user visited.
- Inferred topics of interests refresh once a week.
- A maximum of five interests are stored per time period.
- Users can control their topics of interests.

#### B. Design Directions

We created and compared two different prototypes, each of which represented a different yet relevant interaction pattern to provide users control over inferred topics of interest (see Fig. 1). For the design of the two conceptual prototypes we leaned into established patterns from ads and privacy settings from browsers and social media sites (see [23]). While we aimed to design for browser-based controls, we believe that users are already used to site-specific controls from social media sites. For both prototypes we added a standard Chrome settings header to emphasize the embedding of our envisioned controls in browser settings.

Our first concept used a list-based visualization and a like/dislike interaction patterns (Fig. 1, left). Our assumption for this prototype was that users feel a sense of agency if they were able to influence the algorithmic, inference model by liking or disliking topics of interest. We also assumed the like/dislike pattern can well explain the limitations of this browser control, meaning users would know if they disliked a category, they would see fewer ads related to it, instead of none. The like/dislike pattern can be found in music streaming apps, e.g. Spotify’s Discover Weekly (see [63]), which allows users to train recommendation algorithms, and on social sites to upvote or downvote comments, e.g., comments on YouTube (see [45]).

Our second concept used a chip-based visualization and a save/remove interaction pattern (Fig. 1, right). Our assumption was users could quickly scan inferred topics due to the compact size of the user interface (UI) component and feel invited to interact with the chips as they stand out visually on the page. Chips UIs are used to enter information, filter content, and dynamically trigger actions, e.g., contact chips on Gmail or category tags in content management systems.

Examples where this pattern is applied are Wordpress or Medium (see [23]). With this design, we wanted to investigate the effect of temporal controls, i.e., saving a topic permanently or for a time, or not showing it again. We discarded additional earlier design ideas such as larger cards or image lists (see [22]) because we wanted to focus on more compact elements that can be embedded in current browser settings UI.

### C. Qualitative Pilot Study

The goal of the pilot study was to identify the final prototype for our survey study. We conducted a qualitative study to better understand conceptual considerations when fine-tuning the design.

1) *Pilot Study Design:* We conducted 45-minute remote interviews with Chrome users from Germany (DE) and the United Kingdom (UK). The study followed a counterbalanced within-subject design with the two design concepts as described above as the independent variable (like/dislike vs. remove/save). All participants were invited to participate voluntarily through an opt-in sign-up form shared via email to the study participant database from our institution and received a customary remuneration in the respective country for their participation. In addition, all participants signed a digital consent before conducting the sessions.

During each session, we talked about the general perception of OBA and then presented the participants with a hypothetical study scenario. In the scenario, we told participants that they should imagine browsing on a furniture site and afterwards visiting a news site where they see an ad about furniture. We then asked about expectations of potential controls in settings if their browser could store topics that they might be interested in based on the sites they visit. Next, we showed and discussed both prototypes. Lastly, we asked participants about the value of potentially adding a topic of interest manually.

TABLE I  
PARTICIPANT PROFILES FOR PILOT INTERVIEWS.

No.	Country	Gender	Age	Ad Blocker	Ads Attitude
1	UK	m	18 - 24	No	Positive
2	UK	m	35 - 44	Yes	Positive
3	UK	f	35 - 44	No	Positive
4	DE	m	35 - 44	Yes	Neutral
5	UK	f	18 - 24	No	Positive
6	UK	f	35 - 44	Yes	Neutral
7	DE	f	35 - 44	No	Negative
8	UK	f	18 - 24	Yes	Negative
9	UK	m	45 - 54	No	Negative

2) *Pilot Study Participants:* Participants were recruited from an existing user study panel that includes both countries. When screening participants, we looked for a balanced mix of browser users in terms of attitude towards personalized ads (7pt. bipolar likert scale) and the usage of ad blockers (see Table I). We invited a total of 10 participants. Due to one no-show we conducted 9 sessions with 4 male and 5 female Chrome users (self-identified gender) ranging between 18 and 54 years old (3x 18-24, 5x 35-44, 1x 45-54). To ensure we had a diverse group in terms of privacy sentiment when selecting participants from the UK and DE, we ended up with seven participants from the UK and two from DE. All participants were native or business fluent English speakers and sessions were conducted in English.

3) *Interview Analysis:* All interviews were recorded, transcribed and qualitatively coded by the first author, who has extensive experience in analyzing qualitative data. Only one author coded the responses until no new themes emerged since our primary goal of the pilot study was to select one design for our survey evaluation. Statistical analysis was not performed for the pilot study since it aimed to identify trends, not statistically significant results.

### D. Survey Study

The goal of the survey study was to evaluate if the general concept of topic-based controls to manage ad interests in the browser represent an improvement over current browser controls. We selected the prototype design that participants in our pilot study preferred and fine-tuned the design based on our learning from the interviews.

1) *Final Prototype:* Before launching the survey, we created our final prototype design so we could evaluate how users rate the current cookie control UI compared with our topic-based personalization UI. Based on our qualitative pilot study we concluded that the remove-prototype aligns better with users’ expectations about managing ad interests in a browser. We selected the respective prototype, simplified the look and feel and modified the explanatory text to better represent the underlying idea of browser-inferred ads interest (see Fig. 2, bottom). Most importantly, we added a note about the option to remove topics to the UI.

2) *Survey Study Design:* We ran a fifteen-minute online survey in the US and UK. Respondents received a customary remuneration in the respective country.

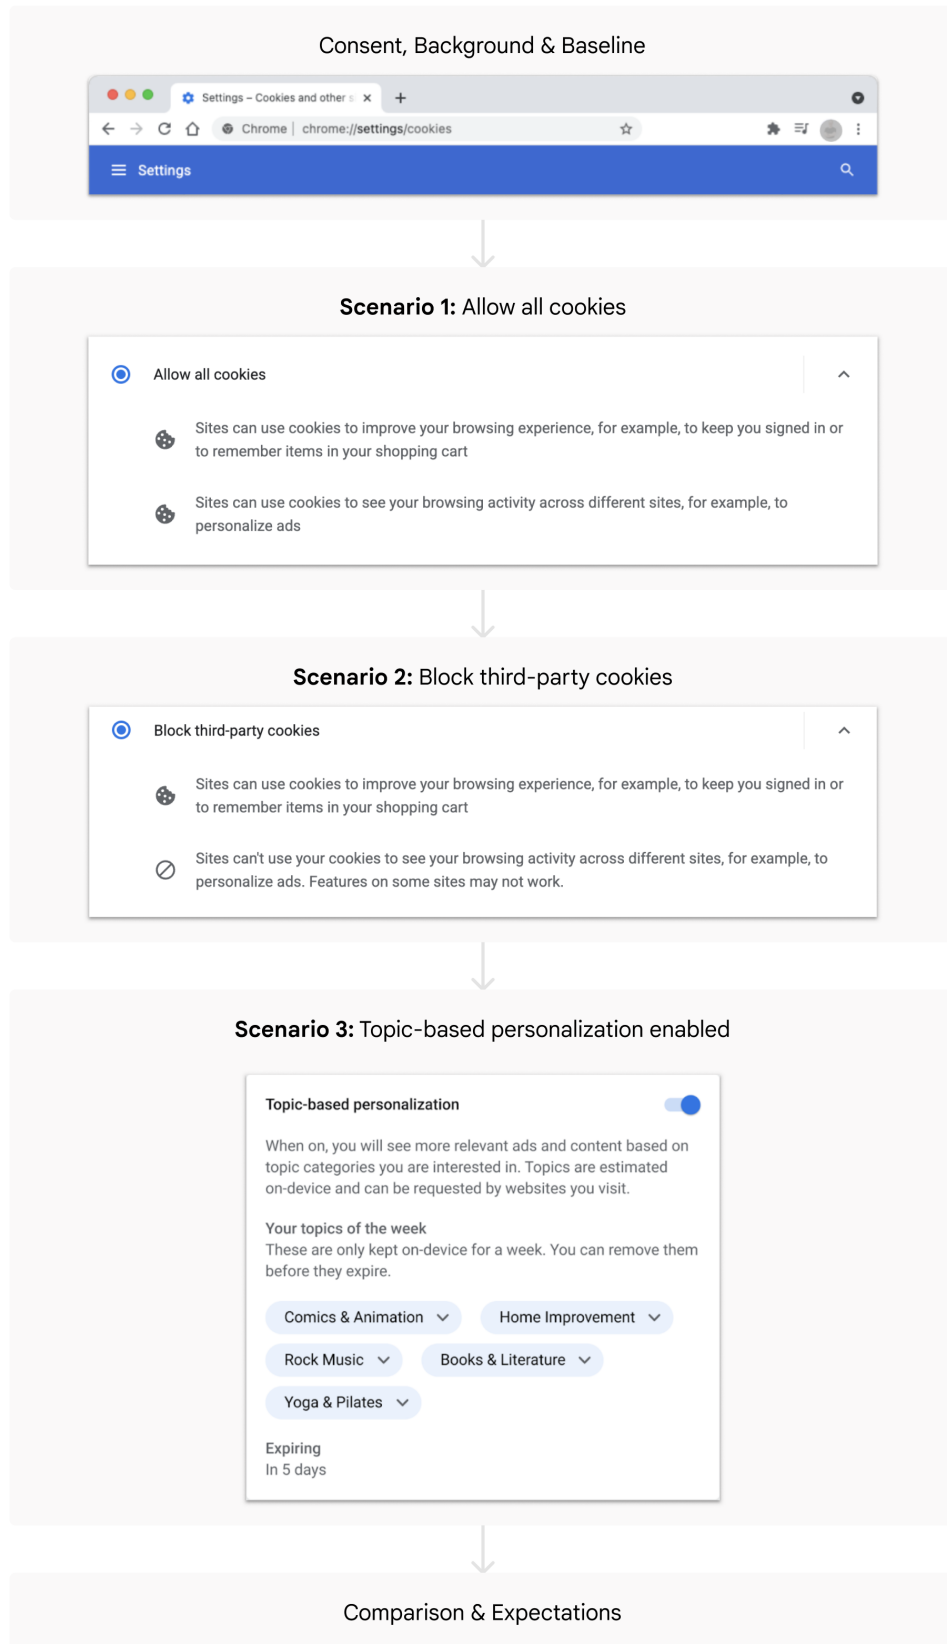


Fig. 2. Survey Flow and Stimuli.

The survey was distributed by the market research agency Qualtrics [54] and data was collected within four weeks between July and August 2021. Before launching the survey, we pre-tested the setup for comprehension, readability, and survey fatigue with five people from the US in moderated 60 minutes remote sessions. Since we did not aim to analyze the pre-test responses, we only screened for Chrome users above 18 years old. Participants for this qualitative pre-test also received an equivalent remuneration for their participation.

3) *Independent Variables*: The questionnaire started with a consent section and demographic background questions. We also included questions to infer respondents’ level of privacy sensitivity and familiarity with OBA. For privacy sensitivity, we used the Westin segmentation [40], which is often used in HCI literature as a proxy for privacy sensitivity [12], [41], [42], [47]. Familiarity with OBA was based on self-reported feedback (5pt. unipolar likert scale). In addition, we asked for the perceived value of personalized ads for both users and publishers as a baseline to compare the provided value based on cookies or topics-based personalization (level of agreement, 7pt. bipolar likert scale).

Next, we introduced our three questionnaire scenarios (see Fig. 2). First, OBA based on third-party cookies (when all cookies are allowed), second, OBA based on third-party cookies (when third-party cookies are blocked), and third, OBA based on local (i.e., on-device) inference of topics of interest. For each scenario, the respondents saw a relevant screenshot/UI mockup, i.e., either an image of current Chrome cookie settings or our prototype UI for topic-based control, as well as textual instructions.

Instructions contained a note that users should imagine that they see a personalized online ad about a pair of shoes on a news website after visiting the site of an online shoe store. Personalized was explained as the concept of an ad being selected for them to see using information about them. We did not counterbalance the cookie-based and topic-based scenarios because the mentioned order represents a realistic sequence if browser vendors phased out third-party cookies and replaced them with an alternative, e.g., locally inferred topics of interest. We also intentionally called out that only one approach per scenario would be active to ensure responses referred to respective mechanisms and the related UI.

4) *Dependent Variables*: After exposing respondents to the prompts and control UI, we collected five sentiments as dependent variables for each of the three scenarios: Personal and publisher value of personalized ads (level of agreement, 7pt. bipolar likert scale), perceived level of privacy (5pt. unipolar likert scale), feeling of control (5pt. unipolar likert scale), and perceptions of how well websites can financially benefit from ad interactions, i.e., support the open web (5pt. unipolar likert scale). At the end of the survey, we asked respondents to directly compare OBA based on third-party cookie controls and based on our topics-based controls, i.e., number of ads shown on sites, perceived level of privacy, and expected relevance of ads (5pt. bipolar likert scale each).

TABLE II  
DISTRIBUTION OF WESTIN SEGMENTS ACROSS SURVEY SAMPLES.

Westin Segment	US (n = 1193)	UK (n = 1359)
Privacy Concerned	40%	31%
Pragmatist	30%	33%
Privacy Unconcerned	30%	36%

5) *Survey Respondents*: We sampled Chrome users and controlled for a representative distribution across age groups (18 years or older). Self-reported gender distribution (US/UK) was 50%/48% woman, 48%/50% man, 2%/2% nonbinary/self-described/not answered.

To allow for a comparison of different privacy sensitivity levels, we sampled equally distributed Westin segments for both US and UK groups. After data cleaning based on a quality control question in which we asked participants to click on a specific item in a single choice question we remained with 40%/31% concerned, 30%/33% pragmatist, 30%/36% unconcerned (US/UK) respondents (see Table II). This differs from earlier and recently identified distributions of Westin segments in the overall population hence we will focus on relative differences and not absolute numbers in this work [40]–[42].

While Westin originally referred to privacy sensitive users as ‘fundamentalists’, we refer to this segment as ‘privacy concerned’ users in this paper. We ended up with 1400 responses in the UK and 1300 responses in the US. Data cleaning as mentioned above resulted in a final sample of 1359 UK respondents and 1193 US respondents.

6) *Survey Analysis*: Statistical analysis was performed only for the survey using IBM SPSS version 25 with significance tested at the 95% level. We applied nonparametric tests for ordinal scales, i.e., Friedman tests for differences between groups, Kruskal Wallis tests for pairwise testing of independent variables of more than 2 groups and Wilcoxon rank tests for dependent variables of 2 groups. We categorize effect sizes of  $r \leq .1 / .3 / .5$  as small /medium/large effects [21]. For ranked variables, we computed Spearman Rank-Order correlations, with effect sizes .0-.19 very weak, .20-.39 weak, .40-.59 moderate, .60-.79 strong, .80-1.0 very strong [55]. Further, we used the Dunn-Bonferroni procedure for post-hoc correction.

#### IV. PILOT STUDY RESULTS

The goal of our qualitative interview study was to get a better sense of what users think about seeing inferred topics of interest in browser settings. In the following sections, we summarize our takeaways in terms of participants’ general sentiment towards the overall approach, specific design feedback for our two UI versions and consequences on the final prototype we eventually incorporated into our survey.

##### A. Expected Control

The participants in our interview study showed mixed sentiment with regards to the value of personalized online ads. One participant said that “they’re extremely helpful” (P3).

In contrast, others mentioned “*to be honest, I’m not very comfortable*” (P6). Across all sentiment levels of ads personalization, participants emphasized that there’s a lack of control:

“*[Companies have a stronger position in personalized ads because...] it’s just one direction of talking.*” (P4)

Once we introduced participants to the idea of browser-inferred topics of interest, we asked them about their expectations of what a related section in browser settings might look like, prior to prompting with our prototypes. Most participants emphasized that they expected to be able to control inferred topics, e.g., “*First, [I would expect to be able to] delete categories of advertisements*” (P4) and “*I would probably imagine that there should be an option to delete which topics you wouldn’t want there*” (P1). In general, there was excitement about the opportunity to gain control over interests:

“*[I would expect...] I can click or unclick specific topics because, for example, I’ve been doing research for a few weeks for someone’s birthday and then I [can] delete that topic, say ‘unselect’ it... It would be fantastic to say I’m not interested anymore.*” (P3)

## B. Privacy

When we showed participants our prototypes, the overall feedback was positive. The majority of participants said that the new level of transparency and associated control would increase user agency and their privacy when browsing online, confirming takeaways from our analysis of related work:

“*[If inferred topics are] shown to me then it increases my privacy because at least I have information of what has driven the personalization ... and it closely ties with the ability to remove topics.*” (P7)

We also indicated that topics would expire and refresh about once a week in our prototypes. While this was seen as an additional privacy benefit since, e.g., “*you’re dropping information [about] me after 7 days so you’re increasing my privacy because you got less information about me*” (P9), others felt like this could simply improve the relevance of ads: “*That would be very helpful to be honest... Because, you know, obviously, there’s a variety of topics that will occur week-to-week. Some things that you may be interested in for many months or years in your life [but] then there are things that are relevant for only a short period of time*” (P8).

## C. Outcome-based Language

In addition to investigating reactions to a topics-based personalization approach, we wanted to identify one design direction that we could use to validate user sentiment in the subsequent online survey. Even though the like/dislike concept looked familiar, e.g., from “*video recommendations based on the topics that I like and dislike*” (P1), most participants preferred the interaction pattern based on topics removal since like/dislike created confusion in the context of inferred ads interest:

“*We are all very used to ‘like’ and ‘dislike’ but I’m ... a bit confused because [I don’t know what it] means [in this context]. If I click ‘like’, will I see more of that, or if I click ‘dislike’, will it disappear? That’s a bit confusing.*” (P7)

In contrast, the concept of ‘removing a topic’ resonated well, since it provided clarity, “*like before... [when] liking or disliking, you would not know what’s happening after you do that. But here [when you click remove] it’s more clear*” (P3). Overall, we had more participants saying that they “*prefer this type of structure*” (P1).

## D. Additional Controls

At the end of our mockup-based pilot testing, we asked participants how they felt about the opportunity to add a topic of interest. We had envisioned that this might increase perceived user agency. However, in our study, the concept of adding a topic of interest did not resonate with participants. Participant 4 initially said “*it depends on what I can add. If I can add just ‘Automotive’ then it would not make sense to me. If [I can be very specific] then it would be good.*” The same participant then suggested a very specific car model to find a good deal which might not align with how we envision on-device inference of general ads interest. In addition, the same participant then mentioned that “*[adding a topic decreases privacy because I give [away more] information [that wasn’t] really collected so far*”. This was also a concern from other participants, i.e., “*[adding topic] would decrease your privacy because you were giving out more information about yourself.*” (P3).

In addition, some participants mentioned that they would probably not visit such an ads topics dashboard regularly. In particular, this was mentioned when seeing the indicated ‘save for a month’ option in our remove/save prototype (Fig. 1, right). While Participant 3 said “*temporary savings... I don’t know. I don’t think I can remember in a month to do that again.*”, Participant 8 called out “*personally [I would not go to such a topic UI] so much [maybe only] if I kept seeing a repetitive ad, then I would interact with it*”.

## E. Summary of Pilot Study Findings

In summary, we found that a key user need for our participants was removing topics that they are not interested in anymore or that are generally not of interest. However, more granular controls in addition to automatic on-device inference, such as saving or adding topics of interest, might be more of a burden than an added value for users.

We synthesized the following takeaways which informed our final concept for step two, i.e., the survey-based concept evaluation (see Fig. 2):

- 1) Seeing inferred interests resonated with participants.
- 2) Outcome-based control such as ‘remove’ provides clarity.
- 3) Saving a topic of interest does not provide a clear value.
- 4) Ability to add a topic of interest may reduce privacy.
- 5) Indicating that topics expire/refresh assures privacy.



TABLE III  
EFFECT OF PRIVACY SENSITIVITY BASED ON WESTIN ON PERCEIVED VALUE OF PERSONALIZED ADS (US).

Value	Westin Segments	US (n = 1193)
Personal	concerned - pragmatist	H = 239.2 (df 2), p <.001
	concerned - unconcerned	
	pragmatist - unconcerned	
Publisher	concerned - pragmatist	H = 46.3 (df 2), p <.001
	concerned - unconcerned	
	pragmatist - unconcerned	n.s.

TABLE IV  
EFFECT OF PRIVACY SENSITIVITY BASED ON WESTIN ON PERCEIVED VALUE OF PERSONALIZED ADS (UK).

Value	Westin Segments	UK (n = 1359)
Personal	concerned - pragmatist	H = 141.0 (df 2), p <.001
	concerned - unconcerned	
	pragmatist - unconcerned	
Publisher	concerned - pragmatist	H = 23.2 (df 2), p <.001
	concerned - unconcerned	
	pragmatist - unconcerned	n.s.

## V. SURVEY STUDY RESULTS

The goal of our survey study was to validate how users feel about OBA when seeing UI for locally inferred topics of interest, compared with third-party cookie controls. In the following sections, we summarize our findings about perceived value of personalized ads, the perception of OBA with cookie vs. topics UI and expected browsing implications.

### A. Baseline of Personalized Ads Value (H4)

Although we sampled above-average-privacy-concerned users based on the Westin segments, more respondents from both countries tended to feel that personalized ads added value to their browsing experience in general (US: AVG 4.0, STD 1.9; UK: AVG 3.9, STD 1.8; 7 pt. Likert scale) (see Fig. 3). At the same time, the value for publishers, i.e., the sites users visit, was rated significantly higher (US:  $z = -14.0$ ,  $p <.001$ ,  $n = 1193$ ,  $r = .41$ , strong effect; UK:  $z = -15.6$ ,  $p <.001$ ,  $n = 1359$ ,  $r = .42$ , strong effect), with an average of 4.8 (STD 1.8) for US and an average of 4.7 (STD 1.7) for UK (see Fig. 4).

In addition, we saw that the perceived personal value and value for publishers are moderately correlated (US:  $r_s(1193) = .514$ ,  $p <.001$ ; UK:  $r_s(1359) = .524$ ,  $p <.001$ ) from which we infer that participants were aware of the interplay between the value of choosing relevant ads and the economics of the open web.

Respondents that fell into the privacy concerned Westin segment rated the value of personalized ads significantly lower than respondents in the privacy unconcerned or pragmatist segment for both personal value and value for publishers (see Tables III and IV). This confirms results from prior studies (see [12]).

Self-reported familiarity with how ads are personalized correlated positively with perceived personal value with a weak/moderate effect (US:  $r_s(1193) = .345$ ,  $p <.001$ ; UK:  $r_s(1359) = .396$ ,  $p <.001$ ). A similar yet weak correlation was found for perceived value for publishers (US:  $r_s(1193) = .291$ ,  $p <.001$ ; UK:  $r_s(1359) = .273$ ,  $p <.001$ ).

The ratings of personal and publisher value of personalized ads changed when we explicitly pointed out that ads were personalized based on cookies (i.e., in scenario one after the baseline questions). Prompting respondents with personalizing ads based on cookies in the respective survey scenario affected the ratings of perceived value of ads. Respondents rated the personal value lower (US:  $z = -3.5$ ,  $p <.001$ ,  $n = 1193$ ,  $r = .1$ ; UK:  $z = -2.2$ ,  $p <.05$ ,  $n = 1359$ ,  $r = .06$ ) and the value for publishers higher (US:  $z = 6.7$ ,  $p <.001$ ,  $n = 1193$ ,  $r = .2$ ; UK:  $z = -8.2$ ,  $p <.001$ ,  $n = 1359$ ,  $r = .22$ ). We saw small to moderate sizes for these effects.

In summary, we accept **H4**, i.e., users recognized an association between the personal value of personalized ads and the value for publishers.

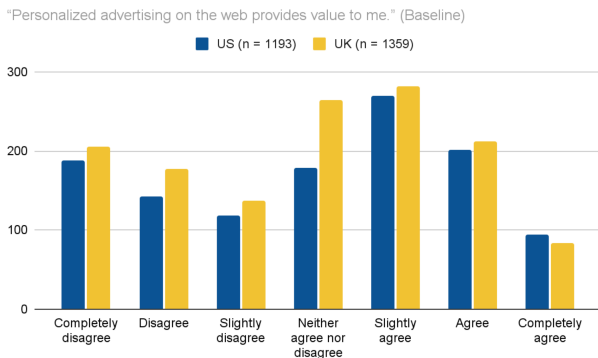


Fig. 3. Personal value of personalized ads (baseline).

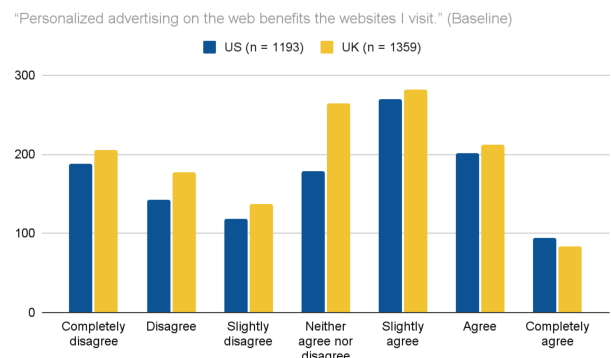


Fig. 4. Publisher value of personalized (baseline).



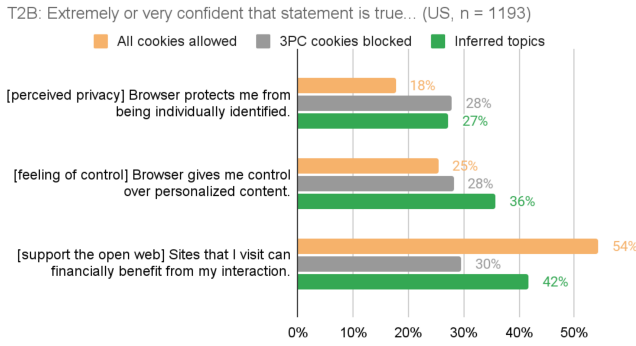


Fig. 5. Top-two-boxes (T2B) for perceived privacy, feeling of control, and support the open web per survey scenario (US).




Fig. 6. Top-two-boxes (T2B) for perceived privacy, feeling of control, and support the open web per survey scenario (UK).

### B. Cookies vs. Topic-based Personalization (H1, H2, H5)

We found statistically significant differences for our dependent variables perceived privacy, feeling of controls and support the open web for between our three survey scenarios:

- **Perceived privacy:** US:  $\chi^2(2) = 309.767, p < .001$ , UK:  $\chi^2(2) = 288.471, p < .001$
- **Feeling of control:** US:  $\chi^2(2) = 136.315, p < .001$ , UK:  $\chi^2(2) = 194.491, p < .001$
- **Support the open web** US:  $\chi^2(2) = 237.399, p < .001$ , UK:  $\chi^2(2) = 186.279, p < .001$

Post hoc analysis with Wilcoxon signed-rank tests with a Bonferroni correction applied was conducted to identify pairwise differences for perceived privacy, feeling of control, and support of the open web.

1) *Perceived Privacy:* For each scenario (all cookies allowed, third-party cookies blocked, inferred topics enabled) we asked participants if they thought the browser protected them from being individually identified (see Fig. 5 and 6). Median ranks were 2 (all cookies allowed), 3 (third-party cookies blocked), and 3 (topics-based personalization) for both US and UK responses. We did not see significant differences between third-party cookies blocked and inferred topics on.

However, we observed significantly higher confidence (strong effect) for *not being individually identified* when the inferred topics mechanism is on compared with all cookies allowed (orange vs. green bars in top part in Fig. 5 and 6; US:  $z = -14.3, p < .001, n = 1193, r = .41$ ; UK:  $z = -14.3, p < .001, n = 1359, r = .39$ ). Almost one third of respondents were extremely or very confident that they are protected from being individually identified for both, third-party cookies being blocked or inferred topics being turned on. Only 18% of respondents in both samples, US and UK, thought the same when all cookies are allowed.

Considering the difference between the scenarios ‘cookies allowed’ and ‘topics-mechanism on’ we partially accept **H1**, i.e., users infer a higher level of privacy from topic-based compared to cookie-based personalization.

2) *Feeling of Control:* When participants were asked about their perceived ability to control personalized content, the inferred topics mechanism induced significantly higher feeling of control compared with both cookie-scenario prompts, i.e., ‘all cookies allowed’ (US:  $z = -11.8, p < .001, n = 1193, r = .34$ , moderate effect; UK:  $z = -13.4, p < .001, n = 1359, r = .36$ , moderate to strong effect) and ‘third-party cookies blocked’ (US:  $z = -6.3, p < .001, n = 1193, r = .18$ , weak to moderate effect; UK:  $z = -4.6, p < .001, n = 1359, r = .12$ , weak effect) (see Fig. 5 and 6).

Median ranks were 3 (scenario one, all cookies allowed), 3 (scenario two, third-party cookies blocked), 3 (scenario three, topics-based personalization) in the US and 2 (scenario one, all cookies allowed), 3 (scenario two, third-party cookies blocked), 3 (scenario three, topics-based personalization) in the UK.

Overall, we accept **H2**, i.e., users perceive a higher feeling of control with inferred topics of interest compared to cookie controls.

3) *Support the Open Web:* Respondents were most confident that sites were able to financially benefit from site interactions when all cookies are allowed (scenario one) and least confident when third-party cookies are blocked (scenario two). In the third scenario, i.e., topics-based personalization, both US and UK responses indicated that inferred topics UI may be perceived as a balance between all cookies are allowed and third-party cookies are blocked (see Fig. 5 and 6).

Respondents rated that sites can financially benefit from interactions significantly higher than when third-party cookies are blocked (US:  $z = -9.4, p < .001, n = 1193, r = .27$ , moderate effect; UK:  $z = -9.1, p < .001, n = 1359, r = .25$ , moderate effect) yet lower compared with ‘all cookies allowed’ (US:  $z = -8.4, p < .001, n = 1193, r = .24$ , moderate effect; UK:  $z = -7.4, p < .001, n = 1359, r = .20$ , moderate effect).

Median ranks were 4 (scenario one, all cookies allowed), 3 (scenario two, third-party cookies blocked), 3 (scenario three, topics-based personalization) in the US.

In the UK, Median ranks were 3 (scenario one, all cookies allowed), 3 (scenario two, third-party cookies blocked), 3 (scenario three, topics-based personalization).

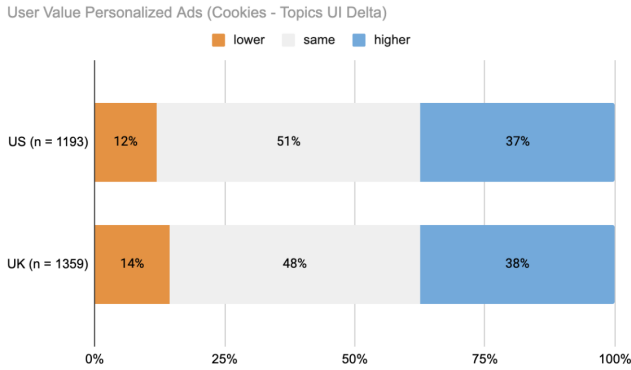


Fig. 7. Delta of self-reported user value of personalized ads when topics-based personalization is enabled instead of third-party cookies. Value (Topics) minus value (Cookies) is either negative (lower), zero (same) or positive (higher).

Since personalization with Topics UI nested between all cookies allowed and third-party cookies blocked, we partially accept **H5**, i.e., users expect that sites can financially benefit from interactions in a topic-based personalization scenario.

### C. Impact of Privacy Sensitivity (H3)

In both US and UK samples, privacy unconcerned respondents inferred a higher level of privacy and a higher feeling of control when being introduced to the inferred topics UI as compared to (third-party) cookies (see Table V). We conclude to accept **H3**, i.e., privacy concerned users perceive a lower value in controls for inferred topics of interest.

For expected implications on how sites can benefit from interactions (i.e., support the open web), we only saw statistically significant differences in the UK sample.

### D. Delta of Personalized Ads Value (H6)

We reported earlier that the perceived user value of personalized ads was considered lower when we called out that personalization is done via third-party cookies.

In contrast, when presenting our topics-based personalization approach and controls, perceived personal value was rated significantly higher (US:  $z = 8.8$ ,  $p < .001$ ,  $n = 1193$ ,  $r = .25$ , moderate effect; UK:  $z = -9.2$ ,  $p < .001$ ,  $n = 1359$ ,  $r = .27$ , moderate effect). At the same time, expected value for publishers was rated lower yet only with a weak effect size (US:  $z = 3.9$ ,  $p < .001$ ,  $n = 1193$ ,  $r = .1$ ; UK:  $z = -4.8$ ,  $p < .001$ ,  $n = 1359$ ,  $r = .14$ ) (see Fig. 7).

Consequently, we accept **H6**, i.e., users perceive a higher personal value than value for publishers from topic-based personalization.

### E. Expected Browsing Implications

After guiding respondents through the cookie scenarios and the topics scenario, we asked participants to compare expected implications on their browsing experience. The majority of respondents expected to see the same amount of ads on a website either based on cookies or inferred topics (US: 56%, UK: 64%). More than one third of respondents in both samples expect a higher level of privacy and higher relevance of ads (i.e., higher level of personalization) (see Fig. 8 and 9).

In the US sample, we did not see statistically significant differences between Westin segments and expected implications on the number of ads, level of privacy, and relevance of ads.

In the UK sample, however, privacy concerned users were more likely to expect more ads (UK:  $H = 7.1$  ( $df 2$ ),  $p < .05$ , pragmatist - concerned  $p < .05$ , pragmatist - unconcerned  $n.s.$ , concerned - unconcerned  $n.s.$ ), yet lower level of privacy (UK:  $H = 17.1$  ( $df 2$ ),  $p < .001$ , unconcerned - concerned  $p < .001$ , unconcerned - pragmatist  $n.s.$ , pragmatist - unconcerned  $n.s.$ ) and lower relevance of ads based on inferred topics (UK:  $H = 9.7$  ( $df 2$ ),  $p < .01$  pragmatist - unconcerned  $p < .05$ , concerned - unconcerned  $p < .05$ , pragmatist - concerned  $n.s.$ ).

## VI. LIMITATIONS

We see our work as a first step towards designing browser controls for topic-based personalization. Our study allowed us to validate that perceived privacy and feeling of control

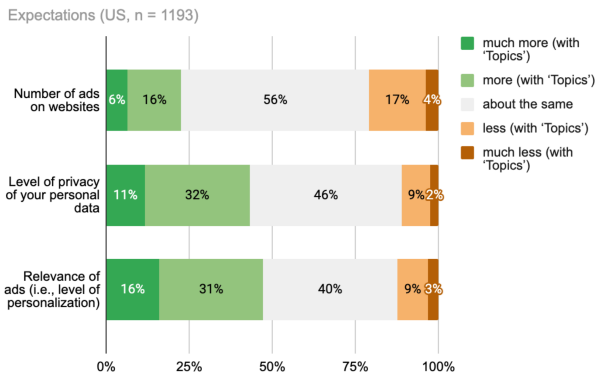


Fig. 8. Expected consequences when topic-based personalization is enabled instead of third-party cookies (US).

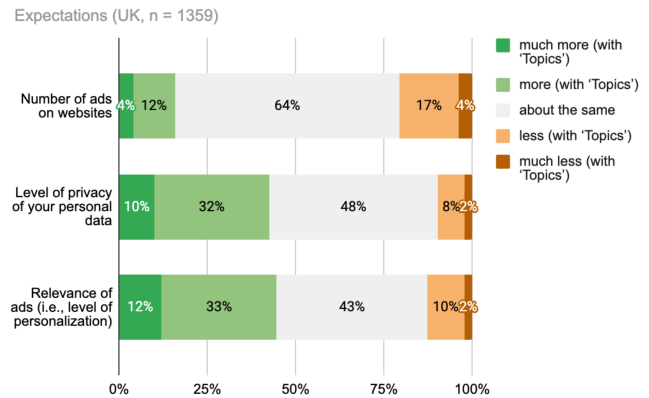


Fig. 9. Expected consequences when topic-based personalization is enabled instead of third-party cookies (UK).

TABLE V  
EFFECT OF PRIVACY SENSITIVITY BASED ON WESTIN ON PERCEIVED PRIVACY, FEELING OF CONTROL AND SUPPORT TO THE OPEN WEB.

Variable	Westin Segments	US (n = 1193)	UK (n = 1359)
Perceived privacy	concerned - pragmatist	H = 92.2 (df 2), p <.001	H = 42.904 (df 2), p <.001
	concerned - unconcerned		
	pragmatist - unconcerned	H = 92.2 (df 2), p <.01	H = 141.0 (df 2), p <.05
Feeling of control	concerned - pragmatist	n.s.	H = 25.914 (df 2), p <.01
	concerned - unconcerned	H = 36.605 (df 2), p <.001	H = 25.914 (df 2), p <.001
	pragmatist - unconcerned	H = 36.605 (df 2), p <.01	n.s.
Support the open web	concerned - pragmatist	n.s.	n.s.
	concerned - unconcerned		H = 20.9 (df 2), p <.001
	pragmatist - unconcerned	n.s.	n.s.

increases when respondents were exposed to our topics-based personalization prototype compared with (third-party) cookies controls. However, the study design also comes along with inherent limitations.

We decided to establish a realistic sequence for a replacement of current cookie-based with topic-based ads personalization controls. Consequently, we did not counterbalance the two scenarios in our survey study. We acknowledge that this may come along priming effects which can be investigated in alternative study setups.

We conducted interviews with participants from Germany and the UK and ran a survey with respondents from the US and UK. Talking to and surveying browser users from different countries might have yielded different feedback and alternative results since sentiment about technological advancements and associated privacy implications differ across countries [34]. Nevertheless, we did not see contradictory results and assessments across the groups we have sampled. Hence, we want to motivate the academic community and the industry to further pursue the path of replacing third-party cookies and associated controls with, e.g., browser-facilitated controls for inferred ads interests.

We wanted to test how people’s privacy attitudes affect the assessment of our designs, as prior work has revealed it affects people’s willingness to share personal data [12]. Similar to previous studies, we saw more critical responses from privacy concerned users. According to Westin, these are users that are more likely to say that “consumers have lost all control over how personal information is collected and used by companies” [40]. In this context, we want to emphasize that actual effects on users’ browsing experiences, e.g., number of ads, level of privacy, and relevance of ads, will depend on how the technology behind our proposed model for locally inferred interest would look like and how marketers will make use of this and/or alternative approaches. We explicitly excluded the investigation and evaluation of a technical implementation for a topics-based concept and perception thereof. Both would require a different approach and setup. Also, in our survey evaluation, we sampled a higher portion of privacy concerned users than previously observed (see [40], [41].) Given the

bias towards different levels of privacy concerns we cannot generalize our results to a larger population. For future work, we recommend investigating the perceptions of representative populations in different countries.

## VII. DISCUSSION

Our work represents an initial exploration about user attitudes towards browser-facilitated controls for ad interests. Our results indicate that such an approach is promising in terms of user attitudes. We see further opportunities for researchers and practitioners to explore alternatives for developing an effective interface and identify opportunities to help users develop effective mental models for interacting with topic-based mechanisms, e.g., through design and UI education.

### A. Transparency and Control Are Key Principles

We have based our work on the assumption that it is possible to improve user privacy while still allowing online companies to show relevant ads based on user interests inferred from visited websites. When discussing third-party cookie alternatives for OBA, one may raise the question if potential utility decreases in behavior targeting are worth the trade-off of inferring high level topics to increase user privacy. Since personalized ads are perceived to be useful and an important pillar of the open web, we consider privacy preserving alternatives to third-party cookies a critical step moving forward [46], [49], [50], [66]. In fact, the web today would have challenges to thrive without alternatives since behavioral targeting represents a substantial aspect of the whole ecosystem [32], [33]. In line with previous assessments we found that transparency and control are key principles for such an approach (see, e.g., [17]). In our pilot study we heard that showing inferred topics of interest is a value added because it creates transparency about what is shared with sites. In our survey study, we saw that respondents had a significantly higher feeling of control (**H2** accepted) when being prompted with a topics-based UI in contrast to current cookie control settings. Given the strongly expressed need for more transparency and control as discussed in the related work section, we want to motivate academic and industry researchers to further investigate topic-based personalization.

### B. Mental Models May Change Over Time

Our survey results indicate that respondents were aware of the trade-off between user value and value for sites they visit in the context of OBA (**H4** accepted).

However, in a scenario where OBA is based on ad interests estimated on-device by the browser, the dynamics of OBA change. Prior work suggests that users' mental models, i.e., how users think about the mechanics of OBA, affects their privacy choices and needs [73]. Adding controls about inferred topics of interest to browser settings may have an impact on established mental models for OBA. Yet, we saw that users already associated a significant role to the browser when it comes to advertising based on their online behavior (see [73]). Since users with different mental models prefer different types of controls we suggest revisiting the analysis of how users think about the mechanics for OBA in the future to derive further needs and requirements for relevant browser controls.

### C. Further Controls Can Increase User Agency

Prior work has showcased that attitudes towards OBA are context-dependent [66]. In our study we used static mocks optimized for the desktop version of a browser. Consequently, we need to better understand what transparency and control means in other contexts, such as on other platforms (e.g., mobile devices) or when needs of a user change over time. Also, we have primarily looked at browser-level controls. An extension of this work would be looking into ad-level controls, similar to [17]. While locally inferred topics of interest primarily serve as a signal for an ad auction, controls about a specific ad shown on a site has the opportunity to provide more details, e.g., what information played into the decision of showing an ad. Both approaches may go hand in hand, yet we still need to understand how users perceive the interplay between new browser-level and ad-level controls. In particular, our results highlight that privacy concerned user saw a lower value in topics-based controls (**H3** accepted). Such users may have different needs than unconcerned users. This is particularly important in the context of empowerment resignation that Seberger et al. discuss in their recent work [60]. In general, when considering further or more granular user controls, it is important to balance the value added of controls and the burden such additional controls put on the end-users. In our pilot interviews, for example, we have heard a preference for automatic inference instead of regularly managing topics of interest.

### D. Inferred Interests Can Balance Requirements

From our survey scenarios, user attitudes can be inferred in terms of privacy towards a world with inferred topics of interest for ads personalization (i.e., without third-party cookies) are comparable to a scenario where third-party cookies are blocked. We also saw that prompting users with the term 'cookies' slightly lowered the self-reported value of personalized online ads. Does that mean that users may primarily react to getting rid of third-party cookies? While this is a fundamental step towards more privacy-preserving solutions, we conclude that

there are further benefits in our proposed direction. In our pilot study interviews, we saw that removing an automatically inferred topic of interest is a key advantage of our proposed design compared with current third-party cookie controls. In addition, in our survey study we saw that respondents associated a statistically significant higher level of feeling of control (**H2** accepted) and provided value of personalized ads for our respondents with the topics-based UI (**H6** accepted). Also, level of privacy was rated higher in the topics-based scenario compared with a scenario in which cookies are allowed (**H1** partially accepted). At the same time, expected financial benefits for publishers in the topics-based scenario was rated between 'all cookies allowed' and 'third-party cookies blocked' (**H5** partially accepted). We can infer from our results that topics-based personalization can suitably balance the needs of users and the web ecosystem. In the future, alternative study setups, e.g., a between-subject design, may complement our learning.

## VIII. CONCLUSION

Our study offers first insights into the perception of browser controls for on-device inference of ads interest. Similar approaches have been discussed in the academic literature and were picked up in industry recently (see [15]). However, we present the first evaluation of user attitudes towards an associated browser control UI.

We conducted a four-step study. First, we designed two different designs inspired by existing browser and social media website controls related to cookies and online advertising. We then talked to browser users from Germany and the UK to see what they think about the proposal in general and about the possible interaction patterns. Second, we selected one design option and ran a survey with 1,193 respondents from the US and 1,359 respondents from the UK to validate how browser users with different levels of online privacy satisfaction perceive such a novel control UI compared to current browser cookie settings.

All in all, we conclude that the shift to locally inferred ads interest and associated browser control is a step in the right direction from a user's point of view. Through the qualitative interviews we identified that a remove-focused interaction pattern aligns well with how participants think about managing ads interest in the context of a browser. Next, via the large-scale online survey, we validated that respondents rate the value of personalized ads, expected level of privacy, and feeling of control significantly higher in comparison to a third-party cookie world. We believe more work in both industry and academia on alternatives to pervasive, cookie-based tracking is necessary to ensure we meet users' expectations and the needs of the open web.

## ACKNOWLEDGMENT

We want to thank all of our study participants for their time and commitment to provide high quality feedback through the study sessions. Further, we want to thank our colleagues and anonymous reviewers for their comprehensive feedback on early manuscripts of this work.

## REFERENCES

- [1] A. Abhyankar, L. Barit, S. Jiwani, R. Sasheendran, M. Sotelo, and L. Cranor. 2022. *The Recipe for the Perfect Batch: Assessing New Cookie Category Terms*. In Eighteenth Symposium on Usable Privacy and Security (SOUPS) 2022. USENIX Association, Boston, MA.
- [2] A. Acquisti, I. Adjerid, R. Balebako, L. Brandimarte, L. Cranor, S. Komanduri, P. Giovanni Leon, N. S., F. Schaub, M. Sleeper. 2017. *Nudges for privacy and security: Understanding and assisting users' choices online*. ACM Computing Surveys (CSUR) 50, 3 (2017), 1–41
- [3] I. E. Akkus, R. Chen, M. Hardt, P. Francis, and J. Gehrke. 2012. *Non-tracking web analytics*. In Proceedings of the 2012 ACM conference on Computer and communications security. 687–698.
- [4] S. R. Alpert, J. Karat, C.-M. Karat, C. Brodie, and J. G. Vergo. 2003. *User attitudes regarding a user-adaptive eCommerce web site*. User Modeling and User-Adapted Interaction 13, 4 (2003), 373–396.
- [5] A. Alqarni and S. Sampalli. 2016. *Privacy-Enhancing of User's Behaviour Toward Privacy Settings in Social Networking Sites*. In Proceedings of the 2016 CHI Conference Extended Abstracts on Human Factors in Computing Systems. 2758–2765.
- [6] S. Banerjee and A. Pal. 2021. *Skipping skippable ads on YouTube: How, when, why and why not?*. In 2021 15th International Conference on Ubiquitous Information Management and Communication (IMCOM). IEEE, 1–5.
- [7] J. Bau, J. Mayer, H. Paskov, and J. C. Mitchell. 2013. *A promising direction for web tracking countermeasures*. Proceedings of W2SP (2013).
- [8] S. C. Boerman, S. Kruikemeier, and F. J. Z. Borgesius. 2017. *Online behavioral advertising: A literature review and research agenda*. Journal of advertising 46, 3 (2017), 363–376.
- [9] L. Brandimarte, A. Acquisti, and G. Loewenstein. 2013. *Misplaced confidences: Privacy and the control paradox*. Social psychological and personality science 4, 3 (2013), 340–347.
- [10] N. H. Brinson, M. S. Eastin, and V. J. Cicchirillo. 2018. *Reactance to personalization: Understanding the drivers behind the growth of ad blocking*. Journal of Interactive Advertising 18, 2 (2018), 136–147.
- [11] A. Cahn, S. Alfeld, P. Barford, and S. Muthukrishnan. 2016. *An empirical study of web cookies*. In Proceedings of the 25th international conference on world wide web. 891–901.
- [12] F. Chanchary and S. Chiasson. 2015. *User perceptions of sharing, advertising, and tracking*. In Eleventh Symposium On Usable Privacy and Security (SOUPS 2015). 53–67.
- [13] Federal Trade Commission et al. 2009. *FTC staff report: Self-regulatory principles for online behavioral advertising*. Federal Trade Commission, Washington, DC.
- [14] M. Dhawan, C. Kreibich, and N. Weaver. 2012. *Priv3: A third party cookie policy*. In W3C Workshop: Do Not Track and Beyond.
- [15] S. Dutton. 2022. *The Topics API Enable interest-based advertising, without resorting to tracking the sites a user visits*. Retrieved September 12, 2022 from <https://developer.chrome.com/en/docs/privacy-sandbox/topics/>
- [16] S. Dutton. 2022. *Topics API: developer guide*. Retrieved November 14, 2022 from <https://developer.chrome.com/en/docs/privacy-sandbox/topics/>
- [17] M. Eslami, S. Rr. Krishna Kumaran, C. Sandvig, and K. Karahalios. 2018. *Communicating algorithmic process in online behavioral advertising*. In Proceedings of the 2018 CHI conference on human factors in computing systems. 1–13.
- [18] J. Estrada-Jiménez, J. Parra-Arnau, A. Rodríguez-Hoyos, and J. Forné. 2017. *Online advertising: Analysis of privacy threats and protection approaches*. Computer Communications 100 (2017), 32–51.
- [19] F. M. Farke, D. G. Balash, M. Golla, M. Dürmuth, and A. J. Aviv. 2021. *Are Privacy Dashboards Good for End Users? Evaluating User Perceptions and Reactions to Google's My Activity*. In 30th USENIX Security Symposium (USENIX Security 21). 483–500.
- [20] M. Fassl, L. T. Gröber, and K. Krombholz. 2021. *Stop the consent theater*. In Extended Abstracts of the 2021 CHI Conference on Human Factors in Computing Systems. 1–7.
- [21] C. O. Fritz, P. E. Morris, and J. Richler. 2012. *Effect size estimates: current use, calculations, and interpretation*. Journal of experimental psychology: General 141, 1 (2012), 2.
- [22] Google. 2022. *Material Design: Image Lists*. Retrieved September 15, 2022 from <https://material.io/components/image-lists>
- [23] Google. 2022. *Material Design: Input Chips*. Retrieved November 14, 2022 from <https://m3.material.io/components/chips/overview>
- [24] S. Guha, A. Reznichenko, K. Tang, H. Haddadi, and P. Francis. 2009. *Serving Ads from localhost for Performance, Privacy, and Profit*. In HotNets. Citeseer, 1–6.
- [25] H. Habib and L. Cranor. 2022. *Evaluating the usability of privacy choice mechanisms*. In Eighteenth Symposium on Usable Privacy and Security (SOUPS 2022). 273–289.
- [26] H. Habib, M. Li, E. Young, and L. Cranor. 2022. *"Okay, whatever": An Evaluation of Cookie Consent Interfaces*. In CHI Conference on Human Factors in Computing Systems. 1–27.
- [27] S. Han, V. Liu, Q. Pu, S. Peter, T. Anderson, A. Krishnamurthy, and D. Wetherall. 2013. *Expressive privacy control with pseudonyms*. ACM SIGCOMM Computer Communication Review 43, 4 (2013), 291–302.
- [28] P. Hausner and M. Gertz. 2021. *Dark patterns in the interaction with cookie banners. Position Paper at the Workshop "What Can CHI Do About Dark Patterns?"* at the CHI Conference on Human Factors in Computing Systems, May 8-13, 2021, Yokohama, Japan.
- [29] Adblock Inc. 2022. *Adblock*. Retrieved September 14, 2022 from <https://getadblock.com/en/>
- [30] U. Iqbal, S. Englehardt, and Z. Shafiq. 2021. *Fingerprinting the fingerprinters: Learning to detect browser fingerprinting behaviors*. In 2021 IEEE Symposium on Security and Privacy (SP). IEEE, 1143–1161.
- [31] C. Jackson, A. Bortz, D. Boneh, and J. C. Mitchell. 2006. *Protecting browser state from web privacy attacks*. In Proceedings of the 15th international conference on World Wide Web. 737–744.
- [32] D. Jaffe and D. O'Brien. 2020. *Statement from the 4A's and ANA responding to Google's announcement regarding third-party cookies*. Retrieved September 12, 2022 from <https://www.aaa.org/statement-from-the-4as-and-ana-responding-to-googles-announcement-regarding-third-party-cookies>
- [33] G. A. Johnson, S. K. Shriver, and S. Du. 2020. *Consumer privacy choice in online advertising: Who opts out and at what cost to industry?* Marketing Science 39, 1 (2020), 33–51.
- [34] P. G. Kelley, Y. Yang, C. Heldreth, C. Moessner, A. Sedley, and A. Woodruff. 2021. *"Mixture of amazement at the potential of this technology and concern about possible pitfalls": Public sentiment towards AI in 15 countries*. IEEE Data Eng. Bull. 44, 4 (2021), 28–46.
- [35] D. Kladnik. 2022. *I don't care about cookies*. Retrieved September 11, 2022 from <https://chrome.google.com/webstore/detail/i-dont-care-about-cookies/fihnjcciajhdofjfbdddfaknhalja?hl=en>
- [36] B. P. Knijnenburg and A. Kobsa. 2013. *Making decisions about privacy: information disclosure in context-aware recommender systems*. ACM Transactions on Interactive Intelligent Systems (TiiS) 3, 3 (2013), 1–23.
- [37] G. Kontaxis, M. Polychronakis, A. D. Keromytis, and E. P. Markatos. 2012. *Privacy-Preserving Social Plugins*. In 21st USENIX Security Symposium (USENIX Security 12). 631–646.
- [38] O. Kulyk, A. Hilt, N. Gerber, and M. Volkamer. 2018. *This website uses cookies: Users' perceptions and reactions to the cookie disclaimer*. In European Workshop on Usable Security (EuroUSEC), Vol. 4.
- [39] O. Kulyk, P. Mayer, M. Volkamer, and O. Käfer. 2018. *A Concept and Evaluation of Usable and Fine-Grained Privacy-Friendly Cookie Settings Interface*. In 2018 17th IEEE International Conference On Trust, Security And Privacy In Computing And Communications/12th IEEE International Conference On Big Data Science And Engineering. IEEE, 1058–1063.
- [40] P. Kumaraguru and L. Cranor. 2005. *Privacy indexes: a survey of Westin's studies*. Carnegie Mellon University.
- [41] M. Kuzmanovic and G. Savic. 2020. *Avoiding the privacy paradox using preference-based segmentation: A conjoint analysis approach*. Electronics 9, 9 (2020), 1382.
- [42] M. Kwasny, K. Caine, W. A. Rogers, and A. D. Fisk. 2008. *Privacy and technology: folk definitions and perspectives*. In CHI'08 Extended Abstracts on Human Factors in Computing Systems. 3291–3296.
- [43] P. Leon, B. Ur, R. Shay, Y. Wang, R. Balebako, and L. 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. 589–598.
- [44] A. Lerner, A. Kornfeld Simpson, T. Kohno, and F. Roesner. 2016. *Internet Jones and the Raiders of the Lost Trackers: An archaeological study of web tracking from 1996 to 2016*. In 25th USENIX Security Symposium (USENIX Security 16).
- [45] Google LLC. 2022. *Youtube*. Retrieved November 14, 2022 from <https://www.youtube.com/>
- [46] M. Malheiros, C. Jennett, S. Patel, S. Brostoff, and M. A. Sasse. 2012. *Too close for comfort: A study of the effectiveness and acceptability of rich-media personalized advertising*. In Proceedings of the SIGCHI conference on human factors in computing systems.

- [47] M. Malheiros, S. Preibusch, and M. A. Sasse. 2013. “Fairly truthful”: *The impact of perceived effort, fairness, relevance, and sensitivity on personal data disclosure*. In International Conference on Trust and Trustworthy Computing. Springer, 250–266.
- [48] A. Mathur, J. Vitak, A. Narayanan, and M. Chetty. 2018. *Characterizing the Use of Browser-Based Blocking Extensions To Prevent Online Tracking*. In Fourteenth symposium on usable privacy and security (SOUPS 2018). 103–116.
- [49] A. Matic, M. Pielot, and N. Oliver. 2017. “OMG! How did it know that?” *Reactions to Highly-Personalized Ads*. In Adjunct Publication of the 25th Conference on User Modeling, Adaptation and Personalization. 41–46.
- [50] K. O’Donnell and H. Cramer. 2015. *People’s perceptions of personalized ads*. In Proceedings of the 24th International Conference on World Wide Web. 1293–1298.
- [51] J. Parra-Arnau, J. P. Achara, and C. Castelluccia. 2017. *Myadchoices: Bringing transparency and control to online advertising*. ACM Transactions on the Web (TWEB) 11, 1 (2017), 1–47.
- [52] S. E. Peacock. 2014. *How web tracking changes user agency in the age of Big Data: The used user*. Big Data & Society 1, 2 (2014), 2053951714564228.
- [53] G. Pestana, I. Querejeta-Azurmendi, P. Papadopoulos, and B. Livshits. 2021. *Themis: A decentralized privacy-preserving ad platform with reporting integrity*. arXiv preprint. arXiv:2106.01940 (2021).
- [54] Qualtrics. 2022. *Qualtrics. Make every interaction an experience that matters*. Retrieved September 10, 2022 from <https://www.qualtrics.com/>
- [55] L. M. Rea and R. A. Parker. 1992. *Designing and conducting survey research: a comprehensive guide*. Jossey-Bass Publishers, San Francisco.
- [56] F. Roesner, T. Kohno, and D. Wetherall. 2012. *Detecting and Defending Against Third-Party Tracking on the Web*. In 9th USENIX Symposium on Networked Systems Design and Implementation (NSDI 12). 155–168.
- [57] S. Schnorf, M. Ortlieb, and N. Sharma. 2014. *Trust, transparency & control in inferred user interest models*. In CHI’14 Extended Abstracts on Human Factors in Computing Systems. 2449–2454.
- [58] J. Schuh. 2020. *News and developments from the open source browser project. Building a more private web: A path towards making third party cookies obsolete*. Retrieved November 11, 2022 from <https://blog.chromium.org/2020/01/building-more-private-web-path-towards.html>
- [59] J. Schuh. 2021. *Privacy Sandbox in 2021: Testing a more private web*. Retrieved November 12, 2022 from <https://blog.chromium.org/2021/01/privacy-sandbox-in-2021.html>
- [60] J. S. Seberger, M. Llavore, N. N. Wyant, I. Shklovski, S. Patil. 2021. *Empowering Resignation: There’s an App for That*. In Proceedings of the 2021 CHI Conference on Human Factors in Computing Systems.
- [61] J. S. Seberger, I. Shklovski, E. Swiatek, S. Patil. 2022. *Still Creepy After All These Years: The Normalization of Affective Discomfort in App Use*. In Proceedings of the 2022 CHI Conference on Human Factors in Computing Systems.
- [62] F. Shirazi and M. Volkamer. 2014. *What deters Jane from preventing identification and tracking on the Web?* In Proceedings of the 13th Workshop on Privacy in the Electronic Society. 107–116.
- [63] Spotify. *Five Ways to Make Your Discover Weekly Playlists Even More Personalized*. Retrieved September 12, 2022 from <https://newsroom.spotify.com/2019-05-02/five-ways-to-make-your-discover-weekly-playlists-even-more-personalized/>
- [64] H. Stewart. 2022. *The Cookieless Future: Why Are Cookies Going Away?* Retrieved September 12, 2022 from <https://learn.g2.com/cookieless-future>
- [65] V. Toubiana, A. Narayanan, D. Boneh, H. Nissenbaum, and S. Barocas. 2010. *Adnostic: Privacy preserving targeted advertising*. In Proceedings Network and Distributed System Symposium.
- [66] B. Ur, P. Giovanni Leon, L. F. Cranor, R. Shay, and Y. Wang. 2012. *Smart, useful, scary, creepy: perceptions of online behavioral advertising*. In Proceedings of the eighth symposium on usable privacy and security. 1–15.
- [67] J. Van Dijck. 2009. *Users like you? Theorizing agency in user-generated content*. Media, culture & society 31, 1 (2009), 41–58.
- [68] K. Varnali. 2021. *Online behavioral advertising: An integrative review*. Journal of Marketing Communications 27, 1 (2021), 93–114.
- [69] J. Warshaw, N. Taft, and A. Woodruff. 2016. *Intuitions, Analytics, and Killing Ants: Inference Literacy of High School-educated Adults in the US*. In Twelfth Symposium on Usable Privacy and Security (SOUPS 2016). 271–285.
- [70] M. Wei, M. Stamos, S. Veys, N. Reitingger, J. Goodman, M. Herman, D. Filipczuk, B. Weinshel, M. L. Mazurek, and B. Ur. 2020. *What Twitter knows: Characterizing ad targeting practices, user perceptions, and ad explanations through users’ own Twitter data*. In 29th USENIX Security Symposium (USENIX Security 20). 145–162.
- [71] B. Weinshel, M. Wei, M. Mondal, E. Choi, S. Shan, C. Dolin, M. L. Mazurek, and B. Ur. 2019. *Oh, the places you’ve been! User reactions to longitudinal transparency about third-party web tracking and inferencing*. In Proceedings of the 2019 ACM SIGSAC Conference on Computer and Communications Security. 149–166.
- [72] W. Xiong, Z. Xiong, and T. Tian. 2022. *Who to show the ad to? Behavioral targeting in Internet advertising*. Journal of Internet and Digital Economics (2022).
- [73] Y. Yao, D. L. Re, and Y. Wang. 2017. *Folk models of online behavioral advertising*. In Proceedings of the 2017 ACM Conference on Computer Supported Cooperative Work and Social Computing. 1957–1969.