

Less is Not More: Improving Findability and Actionability of Privacy Controls for Online Behavioral Advertising

Jane Im
University of Michigan
Ann Arbor, MI, USA
imjane@umich.edu

Ruiyi Wang
University of Michigan
Ann Arbor, MI, USA
ruiyiw@umich.edu

Weikun Lyu
University of Michigan
Ann Arbor, MI, USA
wklyu@umich.edu

Nick Cook
University of Michigan
Ann Arbor, MI, USA
nickcook@umich.edu

Hana Habib
Carnegie Mellon University
Pittsburgh, PA, USA
htq@cs.cmu.edu

Lorrie Faith Cranor
Carnegie Mellon University
Pittsburgh, PA, USA
lorrie@cmu.edu

Nikola Banovic
University of Michigan
Ann Arbor, MI, USA
nbanovic@umich.edu

Florian Schaub
University of Michigan
Ann Arbor, MI, USA
fschaub@umich.edu

ABSTRACT

Tech companies that rely on ads for business argue that users have control over their data via ad privacy settings. However, these ad settings are often hidden. This work aims to inform the design of findable ad controls and study their impact on users' behavior and sentiment. We iteratively designed ad control interfaces that varied in the setting's (1) entry point (within ads, at the feed's top) and (2) level of actionability, with high actionability directly surfacing links to specific advertisement settings, and low actionability pointing to general settings pages (which is reminiscent of companies' current approach to ad controls). We built a Chrome extension that augments Facebook with our experimental ad control interfaces and conducted a between-subjects online experiment with 110 participants. Results showed that entry points within ads or at the feed's top, and high actionability interfaces, both increased Facebook ad settings' findability and discoverability, as well as participants' perceived usability of them. High actionability also reduced users' effort in finding ad settings. Participants perceived high and low actionability as equally usable, which shows it is possible to design more actionable ad controls without overwhelming users. We conclude by emphasizing the importance of regulation to provide specific and research-informed requirements to companies on how to design usable ad controls.

CCS CONCEPTS

• Security and privacy → Usability in security and privacy; • Human-centered computing → Empirical studies in interaction design.

KEYWORDS

Privacy, advertising, ad settings, social media, social platforms, usability, user interface, consent.

ACM Reference Format:

Jane Im, Ruiyi Wang, Weikun Lyu, Nick Cook, Hana Habib, Lorrie Faith Cranor, Nikola Banovic, and Florian Schaub. 2023. Less is Not More: Improving Findability and Actionability of Privacy Controls for Online Behavioral Advertising. In *Proceedings of the 2023 CHI Conference on Human Factors in Computing Systems (CHI '23)*, April 23–28, 2023, Hamburg, Germany. ACM, New York, NY, USA, 33 pages. <https://doi.org/10.1145/3544548.3580773>

1 INTRODUCTION

Major tech companies, such as Meta and Google, rely on online behavioral advertising (OBA) for their revenue; they claim that targeted advertising is also beneficial for users because it enables their products to be free [70]. However, OBA has raised significant privacy concerns among users [6, 46] and researchers have called out its harms to society, in particular increased surveillance [26, 42, 43]. In response, tech companies have argued that users have control over their data via the platforms' ad privacy settings. For example, Meta CEO Mark Zuckerberg wrote to Facebook's users in 2019: "you can find out why you're seeing an ad and change your preferences to get ads you're interested in" [70].

However, contrary to such claims, tech companies' current ad privacy controls are failing to fully support users' privacy needs [21, 29, 65]. Many ad privacy controls lack transparency and do not provide meaningful choices [35, 65]. They are also hard to understand due to the vague explanation of OBA-related concepts [29]. Most importantly, tech companies' ad settings are often difficult to find in the first place, preventing people from using them [29, 35]. Habib et al. [29] observed that many Facebook users are not aware

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than the author(s) must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from permissions@acm.org.
CHI '23, April 23–28, 2023, Hamburg, Germany

© 2023 Copyright held by the owner/author(s). Publication rights licensed to ACM.
ACM ISBN 978-1-4503-9421-5/23/04...\$15.00
<https://doi.org/10.1145/3544548.3580773>

of the platform's existing ad settings and had a hard time finding them because the settings are detached from user experience flows.

We build on design recommendations from prior research [29, 56, 57] to make ad privacy controls more findable by focusing on OBA controls' *location of entry points* and *level of actionability*. Here, an OBA control's entry point means the initial interface a user would click on to find a path that leads to the correct ad setting. In particular, we are interested in entry points in the feed, where users spend most of their time when using social platforms. High actionability means the ad control directly surfaces links to specific ad settings, while low actionability means they point to general settings pages (which is how current ad settings are typically designed). We study the impact of these two on users' behavior, perception of companies, and perceived usability of ad settings. Specifically, our research questions ask how privacy controls' entry point location and level of actionability impact:

RQ1: The findability of controls?

RQ2: Users' effort when finding ad settings?

RQ3: Users' perception of the company that provides the ad settings?

RQ4: Users' perceived usability of ad settings that already exist on the platform?

RQ5: Users' perceived usability of and sentiment towards newly introduced ad control interfaces?

To answer the research questions, we iteratively designed and evaluated ad controls that are easier to find based on Habib et al.'s conclusions and design recommendations [29]. Through a series of formative studies, we explored the design space of alternative advertisement controls. We varied the location of the entry point to the ad settings, placing them within ads (ad menu), at the top of the feed as dashboards, or in the left/top menu bars. Next, we varied the level of actionability at these entry points, providing either only links to general settings pages (low actionability) or more direct links to specific ad controls (high actionability). The findings of our formative study suggested that OBA controls located within ads and at the top of the feed were the most findable. Further, participants liked controls with high actionability because direct links to specific ad settings raised awareness that they exist. Based on the findings, we finalized four designs: 1) *ad menu with low actionability*, 2) *ad menu with high actionability*, 3) *feed dashboard with low actionability*, and 4) *feed dashboard with high actionability*.

To evaluate the final designs, we built a functional prototype as a Chrome extension that augments Facebook's interfaces to allow participants to complete the tasks in the context of their own Facebook account. We conducted a between-subjects, task-based online experiment ($n=110$) to test our designs. Log data generated by the Chrome extension showed that ad controls within ads and at the top of feed both increased the findability of ad settings for the most difficult task of finding advertisers using tailored advertising lists [65], with top of feed being more effective. High actionability increased the findability rate for all tasks and also lessened the number of clicks in attempting each task. For a task related to tailored advertising lists, *feed dashboard with high actionability* increased the findability rate from 22.7% to 63.6%. Participants perceived high actionability and low actionability as equally usable, with high ratings for both. In contrast, participants preferred ad

controls to be situated within ads over being located at the top of the feed. Our designs also positively impacted the perceived usability of Facebook's existing ad settings.

Our findings show that OBA controls' findability can be improved by putting the ad control entry points in ads or by placing settings dashboards at the top of the feed, and especially by providing more actionable options. The results also show that it is possible to design more actionable ad controls without overwhelming users. However, not all participants who found ad settings understood what the settings actually did, suggesting that Facebook's current ad setting interfaces remain confusing. We conclude by emphasizing the importance of platform regulation to provide specific and research-informed requirements to companies on how to design usable OBA controls. Academic researchers can play a critical role here, by auditing companies' ad settings and making concrete design recommendations based on the results. Regulation should also require companies to regularly user test their OBA settings, just as the companies intensively test features relevant to their business models, and publicly release the results to provide accountability.

2 BACKGROUND & RELATED WORK

In many countries, the "notice and choice" model is the primary way through which companies obtain people's consent regarding privacy and data processing. In essence, this model emphasizes providing individuals with information to make informed decisions ("notice") and certain controls for managing how one's information is collected and used ("choice") [51]. Because of this, scholars have used the terms "privacy control" [58] or "privacy self-management" [59] when referring to the current paradigm around privacy protection. In this section, we provide an overview of controls related to OBA, past work related to the design of usable privacy controls, and previous evaluations of OBA controls.

2.1 OBA Controls

In contrast to contextual advertising, where ads are shown based on what the person is seeing on the webpage, online behavioral advertising (OBA) is an advertising technique that leverages personal data to target internet users with highly personalized ads [45]. OBA is primarily enabled by technical mechanisms, such as tracking pixels and cookies, which collect information about users and their online activities [62]. Large technology and ad companies, such as Meta (Facebook), have implemented other complex data collection practices to target ads to users [13]. Prior work has found that while users do not have a full understanding of OBA practices and find some OBA data collection to be privacy-invasive [6, 16, 22, 69], some recognize the utility that OBA can provide [17, 61].

Mechanisms to give users control over the data collected about them for advertising, as well as the ads they see on online platforms, commonly appear on websites [30]. Self-regulatory groups for the advertising industry, such as the Digital Advertising Alliance (DAA) and Network Advertising Initiative (NAI), have implemented websites that allow users to opt out of behavioral advertising from group members [3, 38]. Browser extensions, such as Disconnect¹ and PrivacyBadger,² block ads and tracking mechanisms related

¹Disconnect: <https://disconnect.me/>

²Privacy Badger: <https://privacybadger.org/>

to OBA and support the Global Privacy Control (GPC)³ to communicate a user's preference to limit the sharing or selling of their personal information.

In addition to those third-party control mechanisms, some platforms and websites implement their own settings to allow users some control over the ads they see on the platform. These settings are commonly provided in the service's privacy policy, settings pages, or pages on the website specifically about advertising [30, 56]. Our work seeks to explore how the entry point location and level of actionability of such controls impact their usability.

2.2 Rethinking Usable Privacy Choices

Many scholars have criticized the notice and choice model for failing to actually protect people's privacy [12, 44, 49, 54, 59, 64]. One of the reasons behind this failure is that companies do not provide effective privacy notices and choices [55, 56], mainly due to the lack of incentives and their revenue models [1, 64]. Privacy policies and terms of service are too long and complicated for users to understand [47] and privacy controls are far from being usable and do not provide meaningful choices [23].

Accordingly, privacy experts have increasingly been arguing for rethinking the role of privacy choice [12, 59, 64] and improving privacy choice interface design [23, 55, 57]. While regulations such as the General Data Protection Regulation (GDPR) and the U.S. California Consumer Privacy Act (CCPA) place importance on the quality and the voluntary nature of consent [50, 52], experts have pointed out that there is still a lot of room for improvement in making these regulations concrete about *how* to design privacy choices in a user-centric way. For example, Kretschmer et al.'s work has shown that while GDPR caused online services to provide more ways to opt out of data processing, many did not give users convenient ways to achieve it [40].

Human-Computer Interaction (HCI) researchers can play an important role here by providing specific design recommendations based on research to regulators [25]. For instance, Habib & Zou et al.'s work directly impacted the CCPA regulations by designing and testing icon-link text pairings for the CCPA's do-not-sell opt-out that minimized users' misconceptions [31].

We contribute to this line of work by developing findable ad privacy controls, and studying and measuring their impact on users to provide implications for design and policy. We used Habib and Cranor's [28] framework for privacy choice mechanisms' usability to explore designs. The framework provides seven aspects of usability for evaluating privacy choices: 1) user needs, 2) user ability and effort, 3) user awareness, 4) user comprehension, 5) user sentiment, 6) decision reversal, and 7) nudging patterns. We focus on *user awareness*, as being aware of the existence of ad controls is a precursor to finding and customizing them [29]; without such awareness the users are unlikely to customize any software [7]. Although Habib and Cranor [28] noted that much of users' effort in using privacy choices is in *finding* them, they considered *user ability and efficiency* as a separate usability aspect for measuring users' effort in *making* privacy choices (e.g., required users' actions, such as clicks and scrolls). This is because users may make other errors

once they have found privacy choices. Thus, we also considered *user ability and efficiency* to ensure our designs are usable.

2.3 Evaluating and Redesigning OBA Controls

Extensive research has shown that current OBA controls are problematic in many ways. Research based on a combination of surveys and interviews has shown that they are hard to find and access in the first place because they are often detached from user experience flows on platforms [29, 35]. Habib et al. observed that despite Facebook's existing controls meeting some user needs—especially those related to controlling ad content, advertisers, or information used in targeting—their findability was low [29]. In particular, study participants had a hard time finding the Ad Preference page,⁴ where most of Facebook's OBA controls are located [29]. Hsu et al.'s work also showed study participants were largely not aware of Facebook's ad personalization settings and struggled to find them [35]. Studies have also uncovered that OBA controls' explanations are hard to understand, especially for those related to collected data [29], provide misleading facts [60], are vague [20], or potentially inaccurate [4]. Furthermore, there is no evidence that platforms actually respect users' expressed preferences via ad controls. Prior audits have shown that Facebook's algorithms eventually showed ads with the same sensitive content over time, despite the user having chosen to see fewer ads about the topic [27].

Researchers have also focused on improving OBA privacy controls by using system-building approaches. Weinschel et al.'s browser extension visualized information that data trackers may have collected about users for a long period of time, which helped participants have a more accurate understanding of tracking and increased their interest in privacy-protecting measures [66]. Jin et al. [39] built a system that used flowcharts to show users why an ad has been shown based on one's Facebook profile, and also let users edit their profile. Based on the findings, the authors argued that giving users more transparency and control will help users become more receptive to OBA [39]. Similarly, Barbosa et al. built and deployed a system that considers both online and offline behaviors and lets users review and interact with their ad profiles [8]. The authors argued that letting users interact with one's profile helps them have more trustworthy experiences with targeted ads [8].

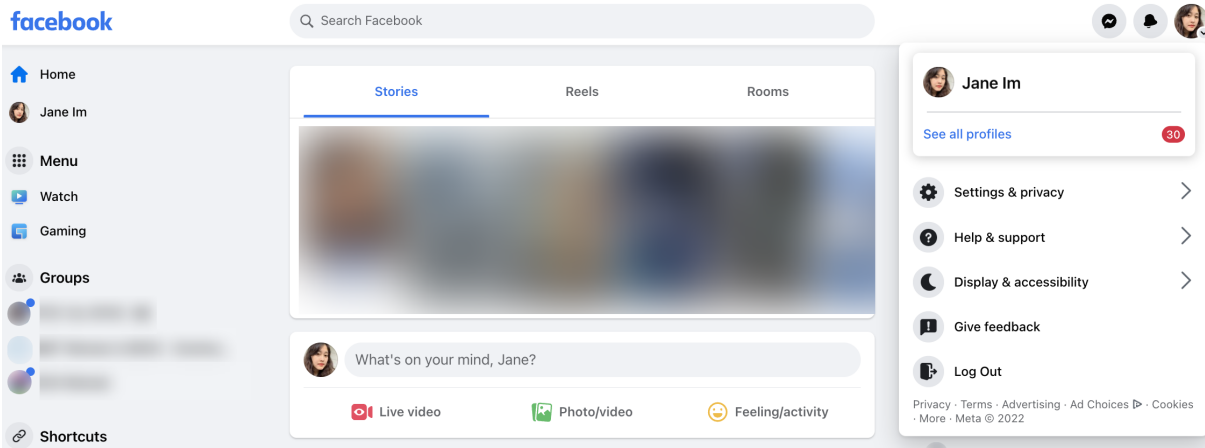
Complementing prior research on improving ad controls, our work focuses on the findability of ad settings rather than adding new kinds of settings functionalities. In doing so, we contribute an understanding of the *impact of interface design on existing OBA controls' findability and user experience*. While our study has been conducted in the context of Facebook because of its extensive tracking and major role in the OBA ecosystem [33, 62], our intention is to broadly inform the design of OBA controls to improve their findability, and study findable controls' impact on users' behavior, perception of companies, and perceived usability of ad settings.

3 FORMATIVE STUDY AND DESIGNS

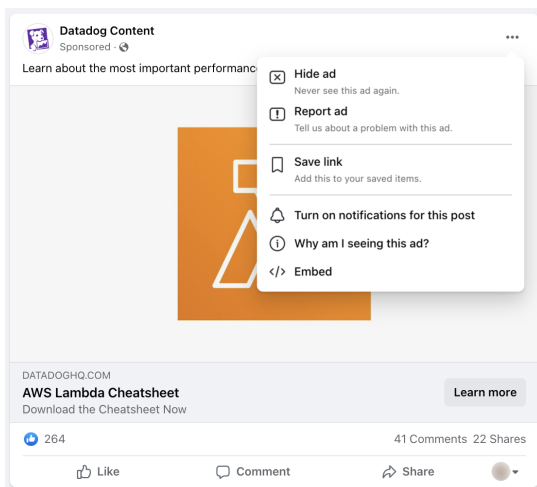
We first conducted formative user studies to iteratively design ad control interfaces with a focus on improving entry points and actionability of ad privacy settings. We used Habib and Cranor's framework [28] for privacy choice mechanisms' usability to inform

³Global Privacy Control: <https://globalprivacycontrol.org/>

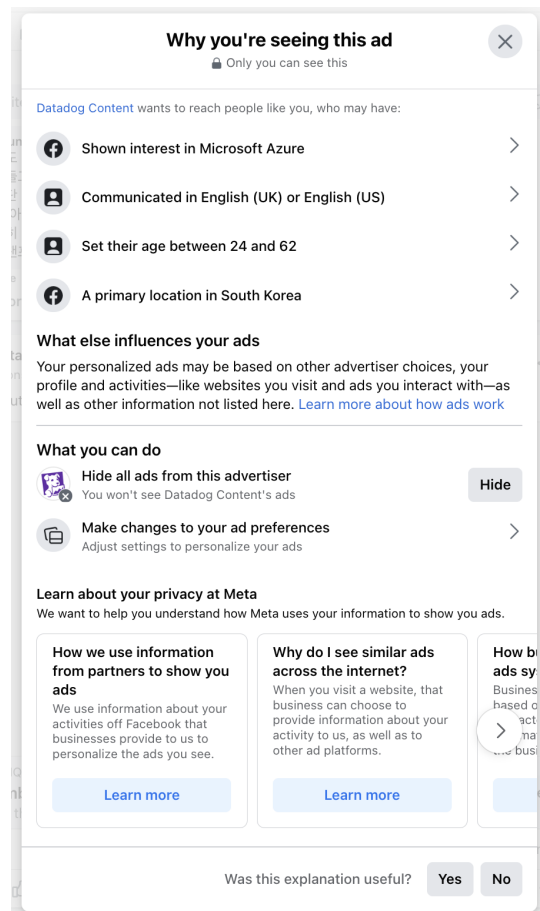
⁴<https://www.facebook.com/adpreferences>



(a) The user's profile image in the top right corner of the menu bar is one of the entry points for reaching the general ad settings page.



(b) Facebook's ad contextual menu that can be reached by clicking on the three dots at the top right corner of an ad. Users can click "Why am I seeing this ad?" for more information and controls.



(c) "Why you're seeing this ad" popup that can be reached from the ad contextual menu by clicking "Why am I seeing this ad?"

Figure 1: Control condition (Facebook's original interface)'s path to ad setting functionalities.

our design directions. In particular, we focused on *user awareness*, as users should be able to find advertisement controls to use them [29]. We also considered *user ability and effort* because if ad controls require too much effort (e.g., time, number of clicks) for users to make privacy choices, they will not be usable. In the following subsections, we describe Facebook’s existing OBA controls, and how we identified initial designs and selected the designs for the main study.

3.1 Existing Facebook OBA Controls

Habib et al. [29] identified 36 interactions/paths that lead users to controls related to advertising on Facebook.⁵ As shown in Figure 1, the primary locations are (1) the contextual menu that can be reached by clicking three dots at the top right corner of each ad (Figure 1-b), (2) the Ad Preferences page, which can be reached via the global account setting⁶ that is linked from the top menu bar’s dropdown (Figure 1-a) and the contextual menu’s “Why am I seeing this ad?” (Figure 1-b), and (3) the Off-Facebook Activity page,⁷ which is in Your Facebook Information,⁸ a tab located in the global account settings. In particular, the Ad Preferences page includes Ad Settings,⁹ where most of Facebook’s ad controls related to privacy and data collection are located at.

As Habib et al.’s work has shown, the entry points of Facebook’s OBA controls are hidden and not directly surfaced to users [29]. For example, in order for a user to go to Ad Settings, they have to click on their profile image at the top menu bar to open the dropdown menu (Figure 1-a), then click Settings & privacy, click Settings, scroll down to find and click on the Ads tab, and then click Ad Settings. While users can also access Ad Settings via an ad’s contextual menu where Ad Preferences is linked, the link to Ad Preferences is hidden inside “Why you’re seeing this ad,” (Figure 1-c), which is not apparent to users [28]. In our work, we consider exposing the link to Ad Settings in the main feed by creating *entry points* that are integrated into user experience flows [23]. At the same time, even if a user finds the link to Ad Settings, there is still a risk of users not knowing what to expect and find within the page [28]. Therefore, we also considered improving *actionability* by directly surfacing specific ad setting functionalities.

3.2 Initial Designs

Here, we briefly describe our initial designs that focus on location of entry point and level of actionability. We explored the design space by creating mockups using Google Slides and AdobeXD.

3.2.1 Location of entry point. First, we focused on the location of entry points as we considered them important for user awareness and also for the control’s efficiency. Table 1 shows an overview of major entry points that were explored during the formative study: contextual menu within ads (*ad menu*), top menu bar’s dropdown link (*top menu dropdown link*), icon shortcut next to the Notifications icon (*left menu icon*), and link in the feed’s left side (*left menu link*) are based on Habib et al.’s study [29], while the dashboard at

the feed’s top (*feed dashboard*) was inspired by Facebook’s prior design (described below).

Ad button and dropdown menu. Facebook provides a contextual menu for each ad that can be reached by clicking on the three dots at the top right of the ad. In our design, we made the entry point more prominent by changing the three dots to a button that says “Ad settings.” We describe the options surfaced in the dropdown menu in Section 3.2.2.

Feed dashboard. For this design, we added a “privacy dashboard” [21] at the top of the main feed (Figure 2-b). The rationale behind this design is that Facebook has historically surfaced privacy-related reminders on the main feed (e.g., Figures 20 and 21 in Appendix).

Left menu icon button and tab. Facebook provides icon-based buttons to access Messenger and Notifications at the top of the page next to the user’s profile image (see top right of Figure 5). We tested including an icon that opened a tab with ad setting functions right near the Messenger icon (Figure 2-c). This design was proposed by Habib et al.’s study participants [29].

Top/Left menu link to settings page. In this design, we brought the link to Ad Settings earlier in the top menu dropdown, right after the user clicks on Settings & privacy menu (*top menu dropdown link*; Figure 2-d), also a design recommendation made by Habib et al.’s study participants [29]. We also considered other locations for shortcuts, such as the *left menu link* (Figure 2-e) because Facebook was A/B testing a layout where the icon shortcuts (e.g., Notification icon) were located in the left menubar, instead of at the top.

3.2.2 Actionability. We also explored the level of actionability provided by the interface as it impacts user awareness and ad controls’ efficiency. Typical ad settings provided by tech companies have low actionability (e.g., entry points link to general settings pages without surfacing specific privacy choices). We wanted to assess the utility and usability of highly actionable ad control interfaces.

For designs with **low actionability** (*top menu dropdown link*, *left menu link*), we provided a link to Ad Settings,¹⁰ as this page has most of the OBA controls related to data collection and privacy.

For designs with **high actionability** (e.g., *ad menu*, *feed dashboard*, *left menu icon*), we directly surfaced options related to ad interests and relevance (e.g., Ad Topics¹¹) and options related to data collection and privacy (e.g., Off-Facebook activity, Data about your activity from partners). This is because users have varied goals when it comes to ad controls [29]. While our focus was on surfacing controls that are related to privacy, we also wanted our designs to capture the attention of users that care about ad personalization (described as “advertising curators” by Habib et al. [29]), as those users might benefit from, and even be interested in, thinking more about their privacy. One main goal of the formative study was to see if there was indeed such a need for highly actionable ad controls.

3.2.3 Variation and exploration of designs. While the location of entry points remained the same (i.e., Table 1 lists all entry points we explored), we note that we explored variations of the designs described above, which is a common human-centered design approach for identifying promising designs [67]. We tested different

⁵Habib et al.’s appendix materials document the entire set of paths and settings [29].

⁶<https://www.facebook.com/settings>

⁷https://www.facebook.com/off_facebook_activity/

⁸https://www.facebook.com/settings?tab=your_facebook_information

⁹https://www.facebook.com/adpreferences/ad_settings

¹⁰https://www.facebook.com/adpreferences/ad_settings

¹¹https://www.facebook.com/adpreferences/ad_topics

Name	Entry point	Level of actionability	Type of interfaces
<i>ad button & dropdown menu</i>	within ads	high	text-based button(s) with dropdown menu
<i>feed dashboard</i>	top of the feed	high	dashboard with text-based links
<i>left menu icon button & tab</i>	left-side menubar	high	icon-based button opening a tab of menu bar
<i>top menu link to settings page</i>	top menu bar's profile image > dropdown	low	text-based link to Ad Settings page
<i>left menu link to settings page</i>	left-side menubar	low	text-based link to Ad Settings page

Table 1: Overview of major design ideas explored in the formative study.

numbers of buttons/links and types of buttons (e.g., text-based, toggle) (designs shown in Figure 24). Here, we described a subset of the most fruitful designs that illustrate regions of the design space that we decided to focus on.

3.3 Formative Study

After deciding on initial designs intended to improve on the usability factors user awareness and ability & effort, we created low-fidelity, interactive prototypes using Adobe XD. Using the prototypes, we conducted a formative evaluation study to refine and select the final designs for the main study. The goals were to inform: (1) which entry point locations to focus on, (2) how to provide higher actionability, (3) decide which ad settings to surface in our interface, (4) test our interface design and wordings, and (5) further refine the designs. The study was reviewed and approved as exempt from ongoing oversight by the Institutional Review Board (IRB) at the University of Michigan.

3.3.1 Study Protocol. Before the study, all participants reviewed and agreed to a consent form via an online survey. All interviews were conducted remotely via Zoom and recorded and transcribed automatically with the participant's consent. Following Nielsen's recommendation [48], we conducted our study over multiple iterations with five participants per cycle. Participants were compensated with a \$15 Amazon gift card; sessions lasted around an hour (studies that included more designs tended to take longer).

During each session, we first asked about participants' experiences with and thoughts about ads on Facebook in order to understand participants' needs regarding ad settings and inform follow-up questions. Then, we showed participants interactive prototypes of ad control interfaces. We used a Wizard of Oz prototyping approach (i.e., the interviewer controlled the flow of the mockup based on where participants wanted to click) and participants were shown three to six designs in randomized order.

In the first round, we introduced the designs and focused on asking about participants' impressions to gain initial feedback. In subsequent rounds, we asked participants to find a Facebook ad setting that meets their needs in order to observe whether participants noticed the ad control. When participants could not find an interface after multiple attempts, we provided hints and eventually guided them to the control. After the participant found the control, we asked where they would click and what they expected would happen to understand their comprehension of the interface. Then, we showed them the next screen depicting how the interface would change after a user's click. At the end of the session (when participants were finished with interacting with all ad control interfaces),

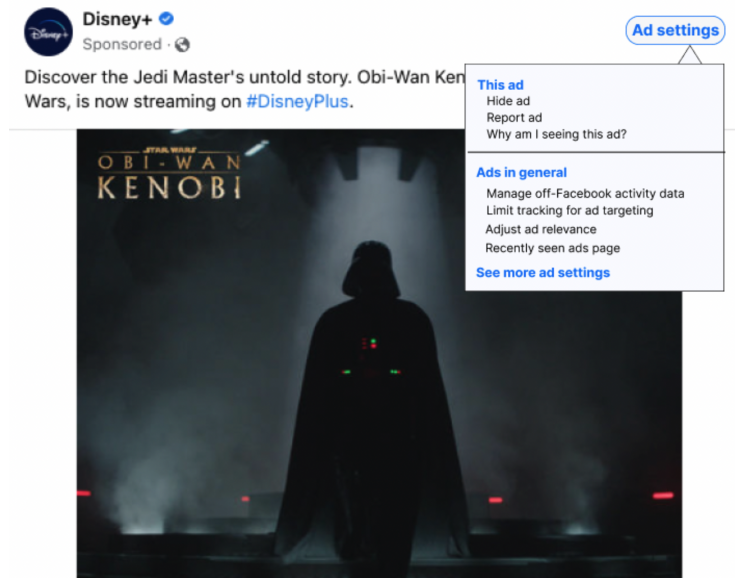
we asked them to rank the designs by their preference and explain their reasons for the ranking. We also asked if they would want to add or remove any functionalities from the interfaces and whether they had any other feedback. Finally, we asked them to complete a post-study demographic survey form online.

3.3.2 Recruitment and Participants. We recruited participants for the formative study by posting a screening survey on Craigslist, following Habib et al.'s approach [29]. The screening survey was advertised in Ann Arbor, Chicago, and South Bend–Mishawaka Metropolitan Statistical Area, and included questions about participants' frequency of using Facebook, how often they click or comment on ads, how often they buy products after seeing ads on Facebook, and their age. Other demographic questions were in the post-study survey. Both surveys are available in this paper's Online Supplementary Materials.

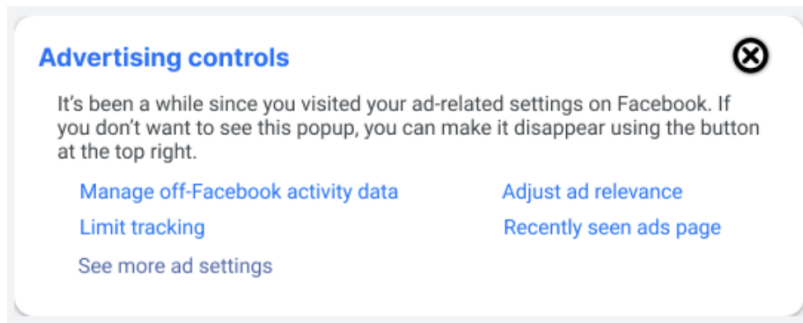
We recruited 20 participants (five participants per design iteration) who use Facebook daily or weekly. We purposefully included a wide range of participants from those who rarely interact with Facebook ads to those who frequently do. This was in order to gain diverse perspectives from Facebook user groups who are privacy-concerned and/or do not care about ads, as well as those who find Facebook ads useful. Nine participants reported buying products at least once per month, six participants reported buying less than once per month but at least once in the last year, and five participants reported not buying products at all after seeing ads on Facebook every day. We stopped recruiting participants after we reached saturation of findings regarding users' awareness of and preferences towards controls.

Among the 20 participants, 15 submitted the post-study demographic survey form. Based on the 15 participants' responses, nine participants had a Bachelor's degree, four participants had a Master's degree, two participants had some college experience but no degree, and one participant had a high school or equivalent degree. Nine participants self-identified as Black, five participants as White, one as Asian, and one participant as having mixed race. Participants' age ranged from 25 to 49 (median=30.5; $n = 12$ as some participants did not submit the post-study survey or preferred to not disclose their age). Five participants self-identified as women, while ten self-identified as men.

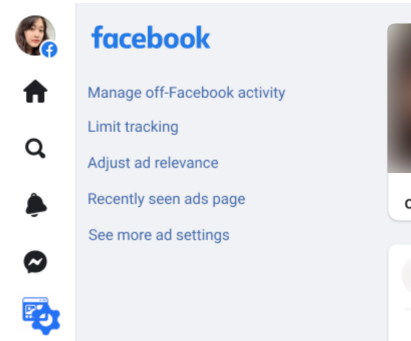
3.3.3 Adapting to Facebook's interface changes and A/B testing. During the period we conducted the formative study, Facebook made several changes to the main feed's interface. For example, for one layout they were A/B testing, Facebook removed links to certain settings in the left-side menu bar. This was one of the reasons for not selecting *left menu link* (Figure 2-e) for the final



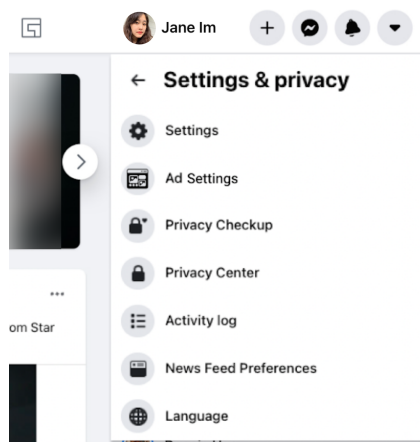
(a) ad button and dropdown menu



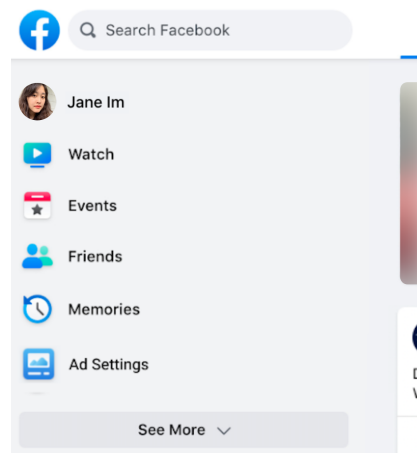
(b) feed dashboard



(c) left menu icon button & tab



(d) top menu link to Ad Settings page



(e) left menu link to Ad Settings page

Figure 2: Mockups that show major designs that we explored during the formative study. We used AdobeXD to add new OBA control designs to screenshots of Facebook's interface in May-June 2022. Variations are included in the Appendix (Figure 24).

study. Facebook also changed the entry point to the dropdown linking Settings & privacy from a dropdown icon (Figure 19 in Appendix) to the users' profile image for one layout (Figure 1-a). This was later introduced in the other layout as well. We updated the mockup accordingly after noticing the change. We checked and made sure that Facebook's interface changes did not occur or have a major impact on the main study, which we describe in Section 5.

3.3.4 Analysis. For each round of formative testing, the first author read the transcripts and notes taken during sessions and conducted deductive coding. Codes mainly focused on: (1) participant's ease/difficulty of finding the ad control interface (awareness), (2) participant's sentiment towards the ad control (e.g., whether the participant liked the entry point), (3) participant's expectation of what would happen after clicking/interacting with the interface, (4) participants' perception of Facebook's existing ad settings, and (5) other feedback. Findings were discussed among the team in weekly meetings to decide on refinements and what designs to further test in the next round of interviews.

3.4 Findings from Formative Study

The formative study's multiple rounds provided us with insights for refining entry point and actionability designs for ad settings.

3.4.1 Location of entry points.

Ad button and dropdown menu. The *ad menu* design's findability rate was consistently high (Figure 2-a). Most participants also preferred it over the rest of the other entry points when asked to rank designs they saw during the session. Participants also liked having direct links to specific functionalities within the dropdown menu (Figure 2-a). They acknowledged that even if some options are not related to just one particular ad (e.g., Manage off-Facebook activity data), the links served as a good reminder that general ad settings exist.

Feed dashboard. In the formative study, almost all participants noticed the *feed dashboard* design immediately (Figure 2-b). However, participants' preferences regarding the dashboard were mixed. Some liked that it was easily accessible, as they could just scroll to the top, and also that it provided many options. For instance, a participant said *"I mean, this is kind of what I probably like best because you really can't avoid it. And ... my hypothesis is that there are people who wouldn't know that these controls exist much less."* But others thought it took up too much space or they did not like that it was the first thing they had to see. One participant also noted that it is easy for users to scroll past it if it is not pinned to the top of the page. However, as many participants thought that the dashboard took up too much space; we ended up not proceeding with and testing a pinned version. Instead, we made the final version collapsible (Figure 22).

Left menu icon button & tab. Many participants had a hard time finding the *left menu icon*. Some participants explained that it was not clear that the icon was related to ads. This finding echoes prior research that text-based links or combinations of icons and text can be more effective than just an icon [31]. Other participants noted that they generally did not pay attention to icon buttons on the feed. For example, a participant said *"I really did not notice.... People*

like me that really don't pay attention to details probably make the same mistake." Because of the low findability, we abandoned the icon button as a potential entry point for ad settings and did not include it in our main study.

Top menu link to settings page. Many participants had a hard time finding Ad Settings from the dropdown menu that can be accessed at the top menu bar (Figure 2-d), which aligns with Habib et al.'s findings [29]. We found that participants struggled even more after Facebook changed the dropdown icon (Figure 19 in Appendix) to the user's profile image (Figure 1-a). But some participants also thought the location was intuitive, as it was under Settings & privacy. However, we did not consider it in our final designs due to the low findability.

3.4.2 Level of actionability. One major finding from the formative study was that many participants preferred or were open to the idea of surfacing direct links to ad setting functionalities, instead of only providing the link to the Ad Settings page. This was also the case for participants who found ads useful or were not privacy-concerned. These participants said while they find Facebook ads useful and would prefer to keep having personalized ads based on their online activity, they would like to have the freedom in being able to choose from available ad settings options, including privacy features. Thus, we decided to vary the level of actionability for the final study to understand the impact of high actionability on users. Of course, when there were too many buttons or options in the dropdown (e.g., Figure 24), participants preferred having fewer options, indicating the importance of finding the right balance to not overwhelm users.

3.4.3 Facebook's explanation of ad settings. Many participants commented that they found Facebook's explanation of the ad setting functionalities related to data and privacy (e.g., Data about your activity from partners) vague or confusing, which echoes prior work [20, 29]. Participants wanted more information about how their choices impacted their data and what "business partners" actually meant. Because our work is focused on findability, we did not alter Facebook's explanation provided in the settings for the main study. However, we added brief explanations about each functionality that was surfaced (Figures 4 and 6).

Based on the results of the formative study, we decided to place the controls' entry points within ads and at the top of the main feed. The major reason was their high findability rate. We also considered participants' preferences. Almost all participants liked having controls within ads, but participants' reaction to the dashboard in the feed was mixed, as described above. We decided to include the dashboard because it is highly actionable and resembles ways Facebook has been known to actually surface its privacy-related features (Figures 20 and 21 in Appendix). We tested variants in which we included multiple buttons at the top of the ad to take people directly to actions available through the dropdown menu (Figure 24). However, many participants preferred the single button so we decided to focus on one ad settings button to open the dropdown *ad menu*.

condition	Location of entry points	Actionability
control condition (Facebook)	Facebook’s entry points (profile image in feed’s menu bar or three dots in ad/dropdown menu)	low
ad menu & low actionability	Facebook’s entry points + button in ad/dropdown menu	low
ad menu & high actionability	Facebook’s entry points + button in ad/dropdown menu	high
feed dashboard & low actionability	Facebook’s entry points + dashboard at top of feed	low
feed dashboard & high actionability	Facebook’s entry points + dashboard at top of feed	high

Table 2: Overview of the between-subjects online experiment’s five conditions.

3.5 Iterative Pilot Testing and Development of Chrome Extension

After the formative study, we built a Chrome extension to add our design elements to a user’s Facebook experience (we describe the system details in Section 5.2). During development, we continued to pilot-test the extension with colleagues and alumni of our institution. We pilot-tested the final prototype with remote user/usability studies via Zoom. Participants were asked to complete tasks that required finding different Facebook ad controls while thinking aloud and sharing their screen. Sessions were not recorded to minimize collecting personal data (as participants shared their screen showing their Facebook account). The sessions typically lasted around 50 minutes and all participants were compensated with Amazon gift cards (\$15).

The evaluations were iterative: we recruited 3-5 participants per batch following Nielsen’s recommendation [48], discovered a set of usability issues or technical problems from each batch, fixed the system to address the problems, and then evaluated with another batch of participants. We repeated this until we reached saturation of issues after four iterations, which resulted in recruiting a total of 15 participants. This iterative approach helped ensure that the augmented interfaces would not hinder participants’ Facebook usage in the main study.

The pilot tests uncovered some important usability issues that were not previously found in the formative interviews. For example, initially, the *ad menu* and *feed dashboard* designs had white backgrounds that resembled more closely how Facebook designed its buttons and popups (e.g., Figures 25 and 26). However, during the pilot tests with the prototype, we discovered that many participants did not notice the ad buttons right away, even when the participants were instructed to find two ads on their feed (tasks described in Section 5.5). This was because participants were used to scrolling the feed very quickly and did not glance at the ad’s right corner because the advertiser’s name is on the left. Participants commented that the buttons “blended in” with Facebook’s ads and overall interface, mainly due to their white background color. Thus, we changed the designs so that they are more noticeable mainly by changing the color of the ad button and the dashboard’s top to blue. We also added a gear icon that is used on Facebook’s setting page to the ad button. Mimicking Facebook’s style, we used the same blue color as Facebook’s logo. In the next round of pilot tests, participants indeed noticed our user interface components more. Many commented that they thought while the button and dashboard stood out, they still followed Facebook’s color scheme

and “looked sleek enough.” Once we finalized the colors, we did not change them during the main study.

Another important finding was related to the dashboard’s location. While Facebook typically puts reminders between Stories and posts (e.g., reminders about passwords as shown in Figure 21), our pilot studies showed that a few participants scrolled down the feed immediately after installing the extension. A few participants also noted that because they jump straight to the posts, putting the dashboard at the feed’s top is better for grabbing their attention. Therefore, we decided to add it at the feed’s very top.

We addressed all usability issues and technical problems discovered during the pilot studies before the main study. The last batch of pilot studies showed that the augmented interfaces did not hinder users’ experience on Facebook.

4 FINAL DESIGNS

Here, we describe the finalized designs that we used in the main study (Section 5), based on the findings from the formative study and iterative pilot testing of the Chrome extension described above.

4.1 Entry Point Location

4.1.1 Ad button and dropdown menu. The final ad button we used had a blue background with a gear icon to make it noticeable (Figures 3 and 4). When a user clicks on the button, the contextual menu appears, with Facebook’s original options displayed under the header “For this ad” and our augmented options underneath “For all ads.”

4.1.2 Feed dashboard. Similar to the ad button, we made the dashboard noticeable by making the background partially blue (Figures 5 and 6). Our formative study’s participants were concerned about the dashboard taking up too much space, so we added a dropdown button at the top right of the dashboard for collapsing it (Figure 22).

4.2 Actionability

4.2.1 High actionability. For interfaces with a high level of actionability (Figures 4 and 6), we included entry points to Facebook’s existing settings that are a mix of data privacy settings and ad content curation settings. They were chosen based on our formative study’s finding that many participants wanted both types, as well as prior literature [29, 53, 65]: Data about your partners, Off-Facebook activity, and Audience-based Advertising are closely related to data privacy, and Ad Topics is more related to interpersonal privacy and ad content (Table 3). The settings’ locations, reasons for including each setting grounded in prior literature, and relevant tasks

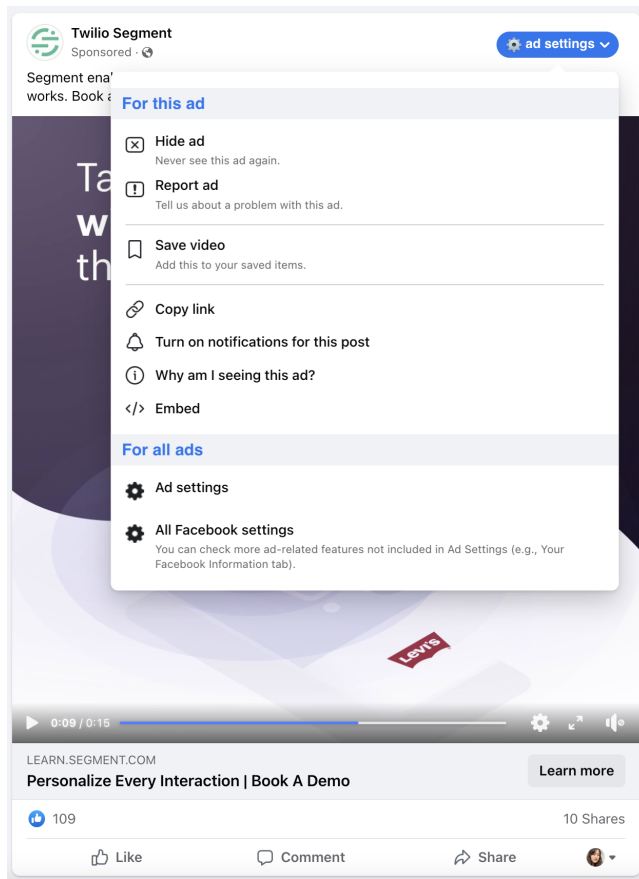


Figure 3: Ad button/menu with low actionability. The ad button is located within ads and the menu surfaces links to the Ad Settings and Facebook's general settings page.

are further described in Section 5.5. We also included Ad Settings because it is a general page for OBA privacy settings, including further settings that we did not directly surface. We included a short explanation for each functionality as our formative studies revealed that many participants had a hard time understanding Facebook's explanations included in settings related to data collection (e.g., Data about your activity from partners, Off-Facebook activity). We iterated on the short explanations during the formative study and pilot tests of the extension.

4.2.2 Low actionability. For interfaces with low actionability (Figures 3 and 5), we included two links: one to the Ad Settings tab of Ad Preferences where many ad privacy settings are located, and another link to the general settings page. The reason for including the general settings is that the Off-Facebook activity is not located in Ad Settings. It is instead located in Your Facebook Information, a tab that is located on the general settings page. The others (Data about your partners, Ad Topics, Audience-based Advertising) are all located within Ad Settings or Ad Preferences. We also consider Facebook's original interface (i.e., without any design changes), which served as the control condition, as having low actionability.

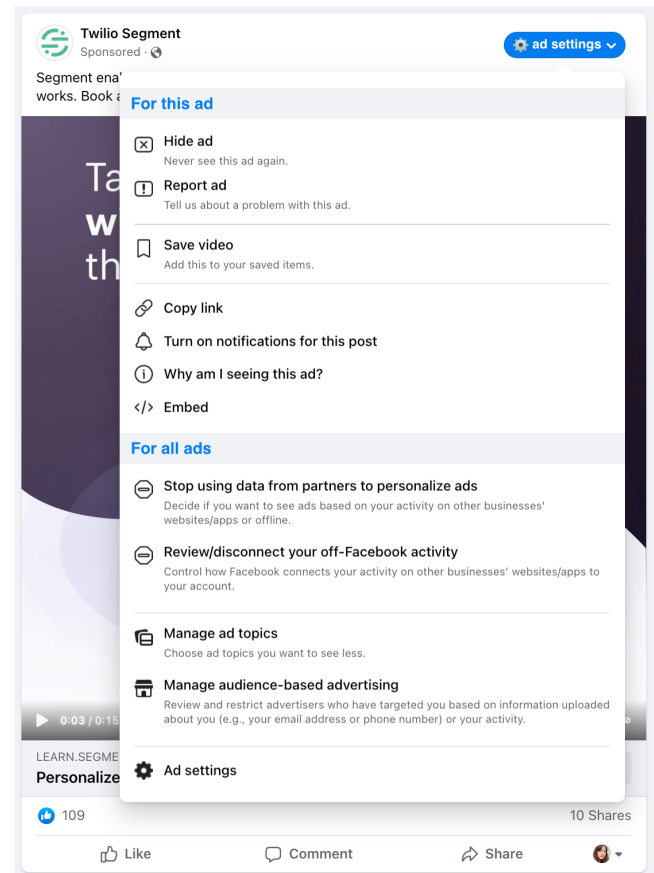


Figure 4: Ad button/menu with high actionability. Compared to the *ad menu with low actionability*, this condition provides direct links to privacy controls that are currently located on Ad Settings and Facebook's general settings page.

5 MAIN STUDY: BETWEEN-SUBJECTS ONLINE EXPERIMENT

To understand our designs' impact on users' behavior and sentiment, we conducted an online between-subject experiment with 110 participants recruited via Prolific.¹² In this section, we present our research questions and hypotheses, our Chrome extension for the experiment, study protocol, analysis approach, and limitations of our study.

5.1 Method and Hypotheses

We conducted a remote, online experiment (as opposed to an in-person study) to gather behavioral data while participants completed tasks in a realistic setting.

5.1.1 Measurement. Using the final designs described in Section 4, we varied the location of ad controls' entry points (*ad menu, feed dashboard, control: Facebook's original entry points*) and level of actionability (*low actionability, high actionability*) (Table 2), and measured whether participants found each ad setting related to the

¹²<https://www.prolific.co/>

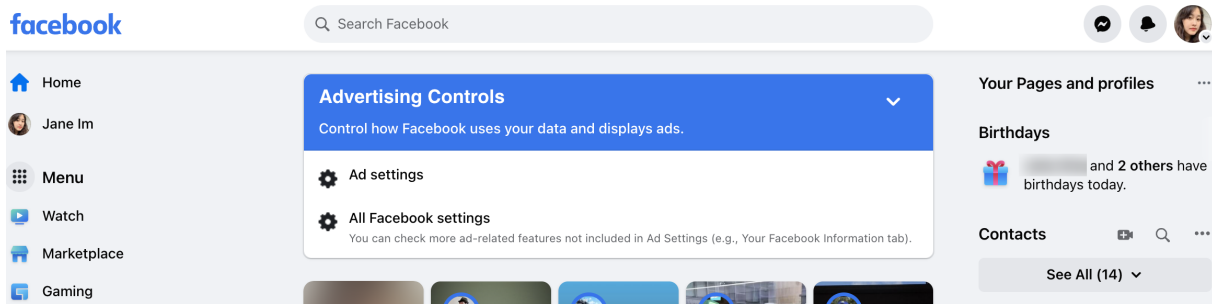


Figure 5: Dashboard at the top of the main feed with low actionability. The dashboard is located at the top of Facebook’s main feed and surfaces links to the Ad Settings and Facebook’s general settings page.

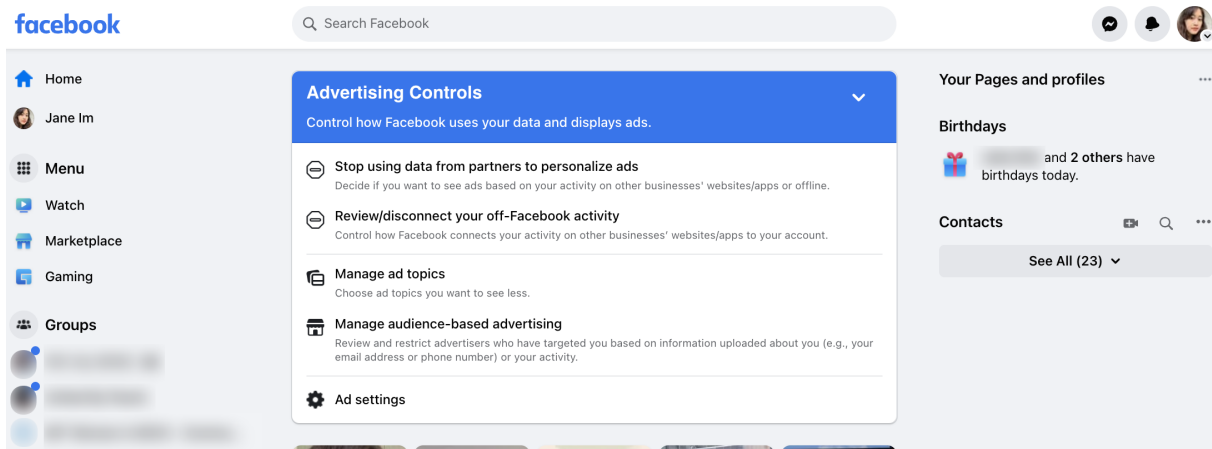


Figure 6: Dashboard at the top of the main feed with high actionability. Compared to the *feed dashboard with low actionability*, this design provides direct links to privacy controls that are currently located on Ad Settings and Facebook’s general settings page.

tasks, number of clicks during tasks, perception of Facebook after completing tasks, and perceived usability of and sentiment towards the interface they tested.

5.1.2 Hypotheses. Below, we describe our hypotheses and how they help us answer each research question, along with the measurement for testing the hypotheses. In each hypothesis, “control condition” and “baseline” refers to Facebook’s original interface.

H1: Our intervention will increase the findability of controls compared to the baseline.

- H1.1: More participants given *ad menu* (Figures 3 and 4) or *feed dashboard* (Figures 5 and 6) interfaces will find the privacy controls than participants in the control condition.
- H1.2: More participants given interfaces with *high actionability* (Figures 4 and 6) will find the ad setting controls compared to groups with *low actionability* controls (Figures 1, 3, and 5).

These hypotheses help us answer *RQ1: How do privacy controls’ location of initial entry point and level of actionability impact the findability of control?* We used log data of participants’ path history

and click behavior on Facebook to measure ad settings’ findability rates.

H2: Our intervention will decrease the number of clicks (a proxy for effort [28]) required to find ad controls compared to the baseline.

- H2.1: Participants given *ad menu* or *feed dashboard* interfaces will find the settings with fewer clicks compared to the control condition.
- H2.2: Participants given interfaces with *high actionability* will find the settings with fewer clicks compared to groups with *low actionability* conditions.

These hypotheses help us answer *RQ2: How do privacy controls’ location of initial entry point and level of actionability affect the users’ effort when finding ad settings?* We hypothesized H2.1 because we expected that treatment groups will more easily find the settings and thus require fewer clicks. The reasoning behind H2.2 is that the length of the path from the entry point to the controls is shorter for *high actionability*. We measured the number of clicks using the Chrome extension.

H3: Our intervention will positively impact users' perception of Facebook compared to the baseline.

- H3.1: Compared to the control condition, groups given *ad menu* or *feed dashboard* will be more positive towards Facebook.
- H3.2: Groups with *high actionability* interfaces will have a more positive perception of Facebook compared to those in *low actionability* conditions.

These hypotheses help us answer *RQ3: How do privacy controls' location of initial entry point and level of actionability impact users' perception of Facebook?* Our reasoning for the above hypotheses is that more participants in the treatment groups would find and learn (or be reminded) that Facebook provides them with various ad settings compared to the control condition. We measured participants' perception of Facebook via their answers to survey questions.

H4: Our intervention will positively impact users' perceived usability of Facebook's existing ad settings compared to the baseline.

- H4.1: Participants in the *ad menu* and *feed dashboard* conditions will rate Facebook's existing ad settings as more usable than the control condition.
- H4.2: Participants in *high actionability* conditions will rate Facebook's existing ad settings as more usable than *low actionability* conditions.

These hypotheses help us answer *RQ4: How do privacy controls' location of initial entry point and level of actionability impact users' perceived usability of ad settings existing on Facebook?* The reasoning behind the hypotheses is that treatment groups' participants may be able to find and use the settings more easily without frustration compared to the control group, which in turn could impact their perceived usability of the settings. We measured participants' perceived usability of ad settings via their survey responses.

H5: Location of entry points and level of actionability will impact users' perceived usability of newly introduced ad controls.

- H5.1: Participants will perceive the *ad menu* designs as more usable and desirable than the *feed dashboard*.
- H5.2: Participants will rate *high actionability* designs as more usable and desirable compared to *low actionability* interfaces.

These hypotheses help us answer *RQ5: How do privacy controls' entry point location and level of actionability impact users' perceived usability of and sentiment towards newly introduced ad control interfaces?* The reasoning behind H5.1 is that the *feed dashboard* is a more intrusive design than the *ad menu* due to its location and size. The reasoning behind H5.2 is that the *high actionability* interfaces should provide more direct access to the ad setting functionalities than *low actionability*. We test these hypotheses by asking how participants perceived the new controls via survey questions.

5.2 Chrome Extension Development

To enable participants to experience the treatment conditions in the context of their own Facebook account, we built a Chrome browser extension and server backend that augmented the Facebook feed according to a participant's assigned condition. We tried to mimic Facebook's style as closely as possible. For example, all of the icons and colors used corresponded to Facebook's design. We

also re-used Facebook's class names to render most of the HTML elements, especially to ensure they worked both in "light" and "dark" mode.¹³ The extension was designed to be both MacOS and Windows compatible, which required supporting different HTML element class names for some ad settings. The server was built with Django, Unicorn, and Nginx.

Once installed, the Chrome extension silently added the new interface elements to Facebook's main feed for treatment groups, i.e., the extension did not explicitly alert users about interface changes in any way. This allowed us to measure how many participants actually discovered and used the controls without explicit nudging.

At the beginning of their study session, participants needed to click a "start session" button in the extension (Figure 27), to allow the extension to start tracking where the participant clicked within Facebook pages, the URLs of Facebook pages they visited, the titles of ad-related popups they saw, along with timestamps. The extension logged this information using attributes (e.g., class, role, name) and inner text of the HTML elements.

5.2.1 Adapting to Facebook's interface changes and ensuring they do not impact the experiment results. During the entire project duration (i.e., from the start of the formative interview study until the completion of final experiment), the first author accessed Facebook daily using two accounts (each had different layouts due to Facebook's A/B testing) to keep track of changes. We briefly note that right before launching the main study (when the authors were wrapping up the pilot studies), Facebook made changes to a few HTML elements' class names that broke the Chrome extension, which led us to update the system. However, during the final study, we did not notice any changes that could have impacted the results. There were also no reports from participants about respective issues during the experiment.

5.3 Ethical Considerations

We built the Chrome extension to minimize personal data collection. That is, compared to methods that require participants to video record their screens (which would capture a lot of revealing information, such as photos of participants' friends or other non-consenting individuals), gathering log data of clicks and page history on Facebook reduces the amount of personal information being collected. The IRB at the University of Michigan reviewed and determined that the main study is exempt from ongoing IRB oversight, per the federal exemption category of being a benign behavioral intervention.¹⁴

However, a possibility of gathering personally identifiable information remained, depending on where participants click. For example, if a participant clicks on their name that serves as the link to their profile page, the extension would save the name. We were transparent in the consent form of what log data would be collected and emphasized that personal information would be de-identified. All data was stored on our institution's server which only the first author (and the school's ITS team) had access to. After the experiment, the first author manually went through the log

¹³Explanation of dark mode on Facebook: <https://www.facebook.com/help/282686829455697>

¹⁴<https://www.hhs.gov/ohrp/regulations-and-policy/regulations/45-cfr-46/common-rule-subpart-a-46104/index.html>

data and de-identified personal information. For instance, if the log data included a name, the first author changed it to '[name]'.

5.4 Recruitment and Participants

To recruit participants, we first posted a screening survey on Prolific. In the survey, potential participants first read the consent form and were asked whether they consented. Then, they were asked about their Facebook, browser, and ad blocker usage. All potential participants were compensated \$1.25 for completing the screening survey. We invited participants who (1) used Facebook weekly/daily, (2) used a desktop or laptop computer to access Facebook weekly/daily, (3) are at least 18 years old, (4) mostly/always use Google Chrome, (5) do not use ad blockers on Chrome, and (6) have English set as their Facebook's language. We decided to recruit people who use desktop or laptop computers for accessing Facebook daily/weekly in order to ensure that familiarity with the desktop version of Facebook was not a confounding variable. We deliberately recruited participants who did not use ad blockers because it was important that all participants would see ads during the study. Our early pilot studies showed that sometimes participants were unaware of their ad blockers and later got confused about the instructions (described in Section 5.5) because they could not see any ads on Facebook. We distinguished who had ad blockers installed on Chrome by including a question in the screening survey that showed three red boxes with an ad about real estate injected in the right one (those with ad blockers could not see the ad in the first place). Potential participants were asked in which box they see an ad about real estate. There were four choices: left box, middle box, right box, I am not sure. We only let those who answered "right box" participate in the main study.

Participants were randomly placed in an experimental condition and a total of 110 participants completed the study, with an average completion time of 41.8 minutes (*median* = 34.8 minutes). All participants were compensated \$15 via Prolific. Participants' age ranged from 19 to 79, with a median age of 39 ($n = 107$; three participants preferred to not disclose their age). 56.4% identified as women, while 41.8% identified as men, and 1.8% identified as non-binary. 80.9% identified as White, 8.2% identified as Black, 1.8% identified as Asian, and 6.4% identified as having mixed race. Regarding education, 34.5% of participants had a 4-year college degree, 24.5% had some college experience but no degree, 13.6% had a Master's degree, 12.7% had an Associates degree, 10.9% graduated from high school, 1.8% had a doctoral degree, 0.9% had less than high school degree, and 0.9% had a professional degree (JD, MD). 72.7% reported being employed, 8.2% identified as homemakers, 4.5% identified as being out of work and looking for work, 3.6% identified as students, 2.7% identified as being retired, and 0.9% reported being out of work but not looking for jobs.

5.5 Tasks, Procedures, and Relevant Ad Settings

Data was collected for almost two weeks in August 2022. Participants were guided through the study via a Qualtrics survey (included in this paper's Online Supplementary Materials). First, participants were instructed to install the Chrome extension. For all conditions, participants were told that the Chrome extension may

add new interfaces to Facebook's main feed for some participants, but it was not disclosed what those might be.

We asked participants to complete tasks via instructions because our early batch of pilot studies showed participants immediately and quickly scrolled down the feed after logging in. This caused many to not notice ads (likely due to "banner blindness"¹⁵) and our designs without further instructions. Thus, after participants installed the Chrome extension, we first asked them to briefly report if they noticed any change on Facebook's interface at a first, quick glance of the main feed. Then, we asked participants to find two ads on their feed and (1) briefly summarize their topics, and (2) rate each ad's relevance to their interests, which mimics real-world situations when users notice ads in their feed. The tasks were intended to carefully draw participants' attention to the location where the new interfaces are, if they were present (i.e., not in the control condition). Furthermore, these tasks did not put the control condition participants to a disadvantage, as they were also instructed in the exact same way and could glance at Facebook's existing ad contextual menu (Figure 1-b) or the feed's menu bar (Figure 1-a).

Next, participants were asked to complete three tasks that involved finding a specific ad setting (Table 3). Using Qualtrics, we showed the tasks to each participant in a randomized order. Below, we describe each task, reasons for including each relevant ad setting grounded in prior work [29, 53, 65], and the path to each setting.

5.5.1 Task: Manage ad topics. In this task, we asked participants to find a way to see fewer ads about a certain topic. Using a feature provided by Qualtrics, the description of the topic was populated based on the participant's response to an earlier question. This is because we wanted to ensure the scenario mentions an ad that the participant has recently seen (instead of including an arbitrary one that could be irrelevant). We chose this task because prior work has shown users express a strong need for having more control over what kind of ads they see [29]. Furthermore, ad topics are deeply related to interpersonal privacy, as ads can reveal a lot about oneself to others [53]. The respective setting for this task is the Ad Topics page,¹⁶ which was redesigned by Facebook in early 2022 (paths to it are described in Table 4). Once a participant reached Ad Topics, they had to search for the topic, choose one from the search results, and then click "See Less."

5.5.2 Task: Stop personalized ads based on online activity on other websites and apps. This task asked participants to stop Facebook from targeting them with personalized ads based on online activity on other websites and apps other than Facebook's products (i.e., websites/apps other than Facebook and Instagram). We chose this task because prior research has shown privacy-concerned users mostly want ways to prevent data sharing and tracking [29]. While Facebook does not provide a way to completely opt out of tracking, this task is the closest to meeting the needs, as it prevents Facebook from using online behavior from other apps/websites to show personalized ads [29]. The privacy actions for this task were either going to Off-Facebook activity page¹⁷ and clicking "Disconnect

¹⁵Banner blindness refers to internet users avoiding paying attention to ad banner-like information [14].

¹⁶https://www.facebook.com/adpreferences/ad_topics

¹⁷https://www.facebook.com/off_facebook_activity

Relevant ad setting	Scenario
Ad Topics page	Imagine that as you're scrolling on Facebook's main page, you wish to see fewer ads about [topic of an ad that the participant saw before the task]
Data about you from Partners Off-Facebook activity	Imagine you went to a few travel websites recently, and you don't like that the ads you're now seeing on Facebook are all related to hotel deals. Imagine you want to stop Facebook from showing you advertisements based on the websites/apps you have visited.
Audience-based advertising	Imagine you recently read a news article about advertisers being able to upload or use a list of information such as email addresses or phone numbers to show you (or exclude you from seeing) certain ads on Facebook. Imagine you want to review which companies reached you on Facebook using such lists.

Table 3: Scenario tasks. The three tasks were shown to each participant in a randomized order using Qualtrics' functionality.

Relevant ad setting	Paths
Ad Topics page	Click the profile image located in the main feed's top menu bar > Settings & Privacy > Settings > Ads > Ad Preferences > Ad Topics Click the three dots in an ad's top right corner > Why you're seeing this ad > Shown interest in [topic] > Manage your ad topics
Data about you from Partners	Click profile image located in the main feed's top menu bar > Settings & privacy > Settings > Ads > Ad Preferences > Ad Settings > Data about your activity from partners Click the three dots in ads > Why you're seeing this ad > How we use information from partners to show you ads (or Why do I see similar ads on the internet?) > Data about your activity from partners Click the three dots in ads > Why you're seeing this ad > Make changes to your ad preferences > Ad Settings > Data about your activity from partners
Off-Facebook activity	Click profile image located in the main feed's top menu bar > Settings & privacy > Settings > Your Facebook Information
Audience-based advertising	Click profile image located in the main feed's top menu bar > Settings & privacy > Settings > Ads > Ad Preferences > Ad Settings > Audience-based advertising

Table 4: Paths to ad controls needed to complete the scenario tasks.

off-Facebook activity," or finding the "Data about you from partners popup" and clicking the toggle button to opt out. Off-Facebook activity is a separate page, while the "Data about your activity from partners" popup is located in Ad Settings. Table 4 describes the paths to each.

5.5.3 Task: Find advertisers that used tailored audience list. We also asked participants to find one advertiser that targeted to them using audience lists. We included this task because audience lists are an advertising technique that not many people are aware of, and are considered very invasive once users learn about them [65]. The right setting to find was the "Audience-based advertising" popup, which shows a list of advertisers that chose to display their ads to certain audiences that included the user. Then, the participant had to peruse through advertisers to find one that has reached them by using a tailored list. The link to the popup is located on the Ad Settings page. We considered this task as the most difficult because it had the fewest entry points compared to the prior two (the path to it is described in Table 4).

5.5.4 Post-task questions. After the three tasks, participants answered five questions about their perception of Facebook. These

questions asked how participants felt about the control that Facebook gives them over their data, how likely they thought it was that Facebook would keep its promises about not selling users' data to other companies, how likely they thought it was that Facebook would respect their ad-related privacy choices on the platform, and how likely they thought it was that Facebook conforms to its privacy notice. Then, for treatment groups, participants were shown a screenshot of the augmented ad control interface (according to their condition) and asked whether they noticed or used the interface. They then answered questions asking about their sentiment and perceived usability of the interface.

We showed all participants screenshots of Off-Facebook activity, Data about your activity from partners, Audience-based advertising, Ad Topics, and Ad Settings page. For each screenshot, participants were asked whether they saw this interface during the tasks, whether they saw it before the study, and how difficult they felt it was to find the setting as well as other usability-related questions (if they saw it during the tasks). Specifically, we asked a subset of the System Usability Scale (SUS) questions [10], which measured participants' perceived complexity of the interface, ease of understanding the interface, and ease of use. We also asked about

participants' sentiment towards the control's location. Lastly, participants answered questions about their demographics and level of privacy concern [15].

5.6 Analysis

We used our extension's log data for measuring the findability of ad settings. We also analyzed participants' qualitative descriptions about how they attempted to complete each task using open coding. These were useful for gauging whether participants recognized the right setting after finding them. For example, a user can find a setting, but close it and move on to another one due to not understanding its functionality. The first author manually went through and read the log traces of each participant to get a sense of the data, and then created a codebook. The codebook contained criteria for deciding whether the participant recognized what was the right ad setting for each task. Then, using the codebook, the first author and second author independently coded a random subset of 33 participants' responses for the three tasks. In order to understand the context of the participant's description, both authors also referenced whether the participant came across the right setting during the task and made privacy choices on the page/popup based on the log data. The inter-rater reliability (IRR, Cohen's κ) for each task was: 0.72 (manage ad topics), 0.94 (stop personalized ads based on online activity on other websites/apps), and 0.81 (find advertisers that used a tailored audience list). Then, the two authors discussed the disagreements, updated the codebook, and annotated the rest ($n = 77$). The IRR (κ) for each task in the second round was: 0.80 (manage ad topics), 0.86 (stop personalized ads based on online activity on other websites/apps), 0.73 (find advertisers that used tailored audience list). The two authors also resolved all disagreements after the final round.

Because our data is not normally distributed, we performed Align Rank Transform (ART) [68] before running ANOVAs, and then performed post-hoc pairwise analysis using ART-c [18] with Holm-Bonferroni correction. We ran separate tests for entry points and actionability. For example, ad menu with low actionability and ad menu with high actionability were grouped into *ad menu*, and dashboard with high actionability and ad menu with high actionability were grouped into *high actionability*. Our *a priori* power analysis ($\alpha=0.05$; $1-\beta=0.8$) showed that our study required 110 participants (22 per condition) to detect a medium-sized effect.

In our analysis, we measured both findability and discoverability rates of ad settings. The discoverability rate is different from findability in the sense that it measures how many participants discovered a setting *for the first time*, whereas the findability rate reflects how many participants found the setting at all, including both participants who had seen the setting before and those who discovered and found it for the first time. To make this distinction, participants were asked after completing the tasks whether they had seen each setting *before the study*.

5.7 Limitations

Our study design has the following limitations.

5.7.1 Participant pool. We recruited participants from the United States who are fluent in English, limiting our study to a specific stakeholder group in the North American context. Thus, our results

may not generalize to non-English speakers from other countries, especially if Facebook uses different interfaces for different regions.

5.7.2 Desktop version of Facebook, Chrome, and No Ad Blockers. Another limitation of our study is that participants used only desktop versions of Facebook. As described in Section 5.4, we recruited participants who frequently use desktops/laptops to access Facebook so that unfamiliarity with Facebook's desktop version does not impact the experiment results. While this was because of the technical difficulty of augmenting interfaces on mobile apps, we concluded using Chrome extensions is reasonable considering prior work's approaches and findings [9, 37]. Furthermore, our designs (*ad menu*, *feed dashboard*) could be applied to mobile versions (e.g., the mobile Facebook app also has ad menus, and feed's top is used for surfacing information). However, future work should examine how the results hold for mobile apps, considering many users access Facebook using phones. We also required that participants used Windows and MacOS desktop or laptop computers to complete their tasks. We did this because the class names of Facebook's HTML elements are different across operating systems. Thus, our results may not immediately generalize to the users of Facebook's mobile version or users of other platforms (e.g., ChromeOS, Ubuntu).

Next, since our system only worked on Chrome, we recruited participants who mostly or always use the Chrome browser. While we believe it is a reasonable first step considering many internet users use Chrome, future studies should consider how the results hold for those who primarily use other browsers. Also, as described in Section 5.4, we recruited participants who do not have ad blockers as all participants needed to see ads to complete the tasks (our pilot studies showed the need for this requirement).

5.7.3 How users set up their accounts and user interface elements. To contextualize our results, we also note that 33.6% of the participants completed the task while being in "dark" mode, and due to Facebook's A/B testing, participants saw a slightly different user interface (UI) for a few elements. However, these were minor discrepancies unrelated to ads (e.g., icons and links at the top/left of the feed, as shown in Figure 23). The overall layout and entry points to ad controls, which is the focus of our study, were the same across layouts (i.e., dropdown link in the top menu bar and contextual menu within ads as shown in Figures 1 and 23). Our pilot studies showed there was not much of a difference for light versus dark modes, as well as for the UI differences. Furthermore, interface designers typically do not have control over such factors (e.g., which mode users choose to have). To account for this, we randomly sampled participants and randomly assigned them to conditions to account for any kinds of random effects.

5.7.4 Intervention designs. We acknowledge there is a possibility that the interventions' color, location, and options could have primed participants. As described in Section 3.5, we chose the final color to be Facebook's trademark blue, which is atypical considering platforms' existing ad control designs. We made this decision after participants in early pilot tests did not notice the interventions when they were white or gray (i.e., it would be infeasible to conduct the experiment if all participants cannot find the interventions). The interventions, especially the *feed dashboard*, are also located at less subtle places compared to the controls companies typically provide

condition	Reported noticing and using	Reported noticing but not using	Reported not noticing	Clicked interface at least once	Did not click interface at all
ad menu with low actionability	81.82% (18)	4.55% (1)	13.64% (3)	81.82% (18)	18.18% (4)
ad menu with high actionability	86.36% (19)	0.0% (0)	13.64% (3)	86.36% (19)	13.64% (3)
feed dashboard with low actionability	63.64% (14)	31.82% (7)	4.55% (1)	68.18% (15)	31.82% (7)
feed dashboard with high actionability	95.45% (21)	4.55% (1)	0.0% (0)	90.91% (20)	9.09% (2)

Table 5: Ratio of participants (total number in parentheses) per condition who self-reported noticing or/and using the augmented interfaces, as well as whether they clicked on the interface at least once. Each condition had 22 participants.

to users. Results may also differ when participants are instructed to find features that are not directly surfaced in the interface.

5.7.5 Study context. Our study focused on the particular context of Facebook at the time of our study. However, we believe our findings will still have broader relevance because ad contextual menus are widely used by social media (e.g., Twitter, Youtube, Reddit) and the top of content feeds is a common place for platforms to surface something to their users—an approach already used by Facebook in the past for their privacy checkup feature (Figure 20).

6 RESULTS

In this section, we report the findings from our main experiment.

6.1 Whether participants noticed and/or clicked on the newly augmented interfaces

We first briefly report how many participants in each condition noticed and/or clicked the newly added interfaces. The *feed dashboard* conditions had over 95% of participants self-reporting noticing the interface. Among them, 95.45% (21/22) of the participants in the *feed dashboard with high actionability* condition self-reported having used the interface (although the log data showed that actually 90.91% (20/22) of participants clicked on the interface). 63.64% (14/22) of the participants in the *feed dashboard with low actionability* condition self-reported using the interface. Both *ad menu with high actionability* and *ad menu with low actionability* conditions had over 80% participants self-reporting noticing and using it, with the log data confirming these numbers. For all newly augmented interfaces, not all participants clicked or used them.

At the same time, we note that the controls' locations, colors, and options could have impacted the results, as acknowledged in Section 5.7. For instance, all participants in the *feed dashboard with high actionability* condition self-reported noticing the design; this could be because it is at the feed's top and is larger than the *feed dashboard with low actionability*—which is due to including more direct options.

6.2 RQ1: Findability and Discoverability of Facebook's Ad Setting Functionalities

We report the *findability* and *discoverability* of Facebook's ad settings per task. Figures 7, 8, and 9 show the discoverability and findability rates for each of the three tasks.

6.2.1 Settings findability per task. We included all participants for measuring findability regardless of whether they self-reported having seen a setting before (as described in Section 5.6). Overall,

controls were significantly more findable in *high actionability* conditions for all three tasks, while the entry point was only found to have a significant impact on findability for the task on finding Audience-based advertising.

Manage ad topics. For the task on managing ad topics (Figure 7), our tests did not find any statistically significant difference in terms of findability between the control condition and the *ad menu* or *feed dashboard* conditions ($F_{(2,105)} = 1.22, p = .30, \eta_p^2 = 0.02$). However, significantly more participants in the *high actionability* conditions found the Ad Topics than in the *low actionability* conditions ($F_{(1,108)} = 4.26, p < .05, \eta_p^2 = 0.04$).

Stop personalized ads based on online activity on other websites and apps. For the task of stopping personalized ads based on off-Facebook activity (Figure 8), we did not find any statistically significant difference in findability between the control condition and the *ad menu* or *feed dashboard* conditions ($F_{(2,105)} = 2.15, p = .12, \eta_p^2 = 0.04$). However, there was again a statistically significant difference between the actionability conditions, with a significantly higher findability rate for *high actionability* over *low actionability* ($F_{(1,105)} = 6.13, p < .05, \eta_p^2 = 0.06$).

Find advertisers that used tailored audience lists. For this task, Figure 9 shows a significant difference between the findability rates of conditions with different entry point locations ($F_{(2,105)} = 8.24, p < .0005, \eta_p^2 = 0.14$). The findability rate of the *feed dashboard* conditions was significantly higher than the control's ($p < .05, d = 0.65$). The *ad menu* conditions' findability rate was also significantly higher than the control's ($p < .05, d = 0.60$). Furthermore, findability was significantly higher for the *high actionability* conditions than for *low actionability* ($F_{(1,105)} = 26.49, p < .0001, \eta_p^2 = 0.20$).

One possible explanation why both factors (entry point, actionability) impacted findability for this task is that the *Audience-based advertising* control is hidden more deeply within Facebook's settings compared to the others, as it is not linked from the contextual menu (see sections 5.5.1 and 5.5.2). Furthermore, it could be that participants found the setting the most difficult to understand. Prior research has shown that internet users are largely unaware of tailored audience lists [65], which may have impacted their ability to find the setting.

6.2.2 Settings discoverability per task. The majority of participants across conditions self-reported seeing the ad settings for the first time during our study. Looking only at these participants' discoverability rate, we see that results largely mirror the findability results reported above, with the difference that entry point location had a

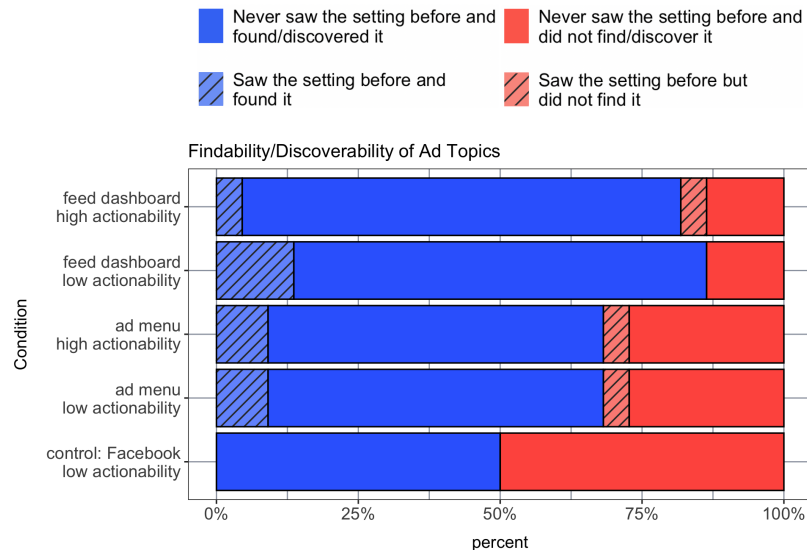


Figure 7: Findability and discoverability rate of the task “manage ad topics.”

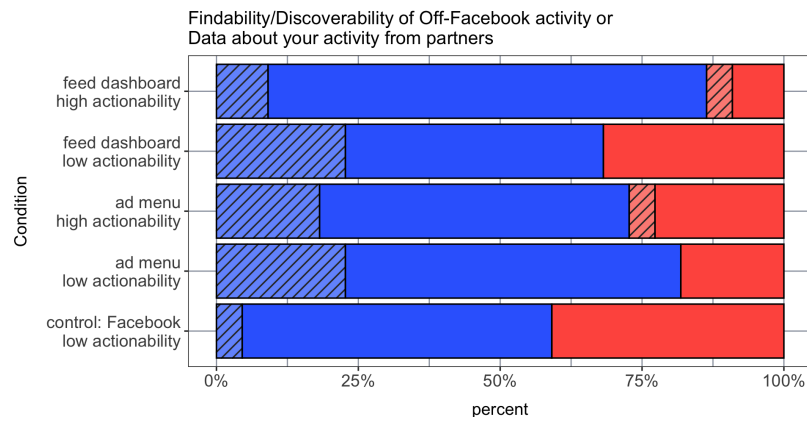


Figure 8: Findability and discoverability rate of the task “stop personalized ads based on online activity on other apps/websites.”

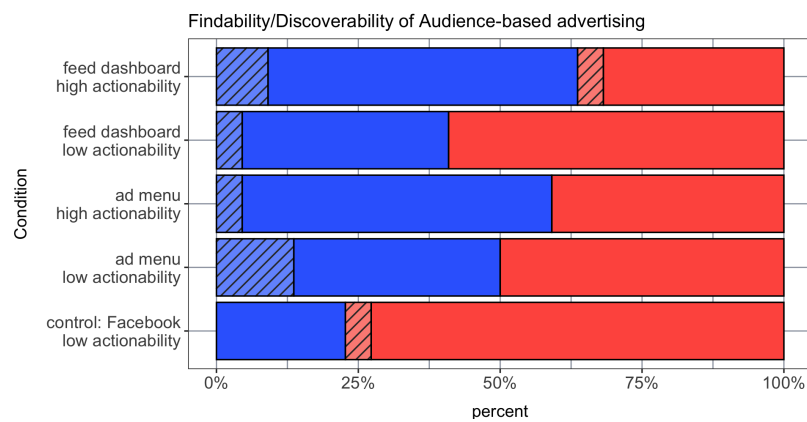


Figure 9: Findability and discoverability rate of the task “find advertisers that used tailored audience lists.”

significant effect for the ad topics task in addition to the audience list task. In short, surfaced entry points to ad settings could have been of more help to first-timers than those who have already seen them before.

Manage ad topics. 90% (99/110) of participants self-reported seeing the Ad Topics page for the first time during the study. For them, there is a significant difference in discoverability between groups with different entry points ($F_{(2,94)} = 9.87, p < .0005, \eta_p^2 = 0.17$). Post-hoc tests showed that the *feed dashboard* conditions' discoverability rate was significantly higher than the control group's ($p < .05, d = 0.78$). *Higher actionability* groups also had a significantly higher discoverability rate than *lower actionability* ($F_{(1,94)} = 4.30, p < .05, \eta_p^2 = 0.04$).

Stop personalized ads based on online activity on other websites and apps. 84.54% (93/110) of the participants indicated they saw "Data about your activity from partners" for the first time in the study, while 95.45% (105/110) answered they saw "Off-Facebook activity" for the first time. For these participants, *high actionability* groups' discoverability rate was significantly higher than for *lower actionability* ($F_{(1,86)} = 12.36, p < .0001, \eta_p^2 = 0.13$). Differences between entry points were not significant ($F_{(2,86)} = 0.99, p = .38, \eta_p^2 = 0.02$).

Find advertisers that used tailored audience lists. Lastly, 91.82% (101/110) of the participants self-reported they saw "Audience-based advertising" for the first time during the study. Significantly more of these participants in the *high actionability* groups discovered Audience-based advertising than in the *lower actionability* condition ($F_{(1,96)} = 28.30, p < .0001, \eta_p^2 = 0.23$). While the omnibus test indicated a difference across entry point conditions ($F_{(2,96)} = 6.08, p < .005, \eta_p^2 = 0.11$), post-hoc tests were not significant.

Summary. For two out of three tasks, the interventions' entry points significantly impacted findability or discoverability, showing support for H1.1. Specifically, both *ad menu* and *feed dashboard* significantly improved *Audience-based advertising's* findability compared to the baseline. In terms of discoverability, *feed dashboard* improved *Ad Topics's* discoverability, and while the omnibus test indicated a difference across entry point conditions for *Audience-based advertising* task as well, post-hoc tests were not significant. H1.2 was supported across all tasks.

6.2.3 Recognition of and engagement with OBA controls after finding them. Although not a major focus, using self-report data, we gauged how many participants recognized the control was the right one per task after finding them. Perhaps unsurprisingly, not all participants recognized the setting they found was the correct one.

Manage ad topics. Out of 78 participants who found the Ad Topics page, 68 (87.18%) of them recognized it as the right setting for the task. A few participants noted that they were surprised the search query returned so many granular topics ("There were a LOT that were art related, I didn't turn 'off' ('see less') every single one").

Stop personalized ads based on online activity on other websites and apps. Among the 81 participants who found either the Off-Facebook activity or Data about your activity from partners, 77 (95.06%) of them recognized either of them as the right setting.

Interestingly, among 50 participants who found the Data about our partners and clicked on the toggle button to make a privacy choice, 5 participants seemingly unintentionally opted in to see personalized ads based on online activity on websites/apps other than Facebook/Instagram, instead of opting out. For instance, a participant wrote "I went to the new add-on interface, clicked on the 'Stop using data from partners to personalize ads' sub-tab, and clicked on 'Not Allowed.'" However, 'Not Allowed' meant the participant had already opted out and thus clicking on it again means they would opt in. But the participant thought by clicking 'Not Allowed', they would opt out (Figure 16).

Find advertisers that used tailored audience list. Out of 52 participants who found Audience-based advertising, 41 (78.84%) recognized it as the right control for the task. Among the 41, 18 participants were able to find one advertiser that targeted them via tailored advertising lists. We note that for some participants, they could have had more difficulty in finding one and eventually gave up, as when a user clicks "See all businesses," the setting shows a long list which also includes advertisers who targeted them via other ways (Figure 18). For example, a participant wrote "I was very surprised by the length of the list, so I worked on several companies and I will go back and finish the list later."

These findings show that the usability issues of Facebook's interface hindered some participants from making the right privacy choice, even if they found the ad setting relevant to the task.

6.3 RQ2: Participants' effort in finding controls: Number of clicks

To understand participants' effort in finding controls, we analyzed participants' number of clicks during each task based on the log data. The extension only logged clicks made on Facebook and not other websites. Number of clicks is a more reliable metric than time as participants could have multitasked during the tasks [28]. For all tasks, participants in *high actionability* conditions required significantly fewer clicks (see Figure 10).

Manage ad topics. Compared to participants in *low actionability* groups (*median* = 16), those in *high actionability* groups (*median* = 11) needed significantly fewer clicks during the task ($F_{(1,105)} = 4.51, p < .05, \eta_p^2 = 0.04$). Our tests also showed a significant difference in number of clicks across entry point locations ($F_{(2,105)} = 3.29, p < .05, \eta_p^2 = 0.06$), but post-hoc tests showed only marginal differences between *feed dashboard* and the control ($p = .07, d = 0.60$), as well as between *ad menu* and *feed dashboard* ($p = .098, d = 0.42$). The difference in clicks between *ad menu* and control was not significant ($p = .50$).

Stop personalized ads based on online activity on other websites and apps. For this task, the *high actionability* groups (*median* = 9) needed significantly fewer clicks than the *low actionability* groups (*median* = 14.5) ($F_{(1,105)} = 5.20, p < .05, \eta_p^2 = 0.05$). Our tests showed there was only a marginal difference in the number of clicks across conditions with different entry points ($F_{(2,105)} = 2.53, p = .08, \eta_p^2 = 0.05$).

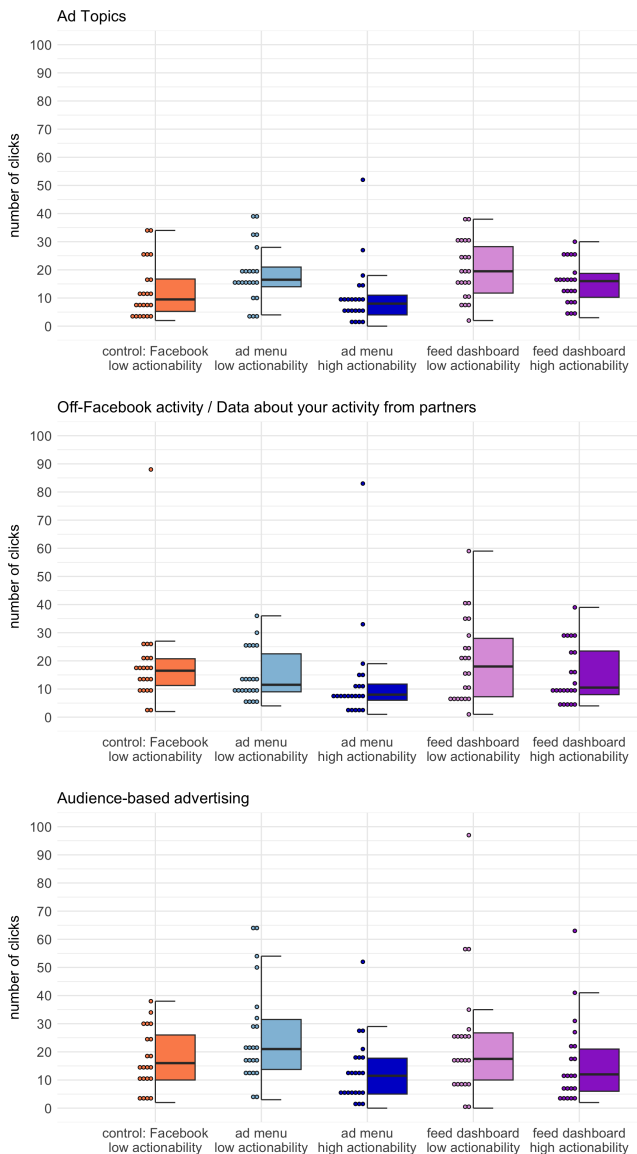


Figure 10: Participants' number of clicks during each task.

Find advertisers that used tailored audience lists. Similarly, for this task, the *high actionability* groups had significantly fewer clicks (*median* = 12) than the *low actionability* groups (*median* = 18) ($F_{(1,105)} = 5.26, p < .05, \eta_p^2 = 0.05$). Differences between entry points were not significant ($F_{(2,105)} = 0.01, p = .98, \eta_p^2 = 0.0003$).

Summary. Our test could only find marginally significant differences or could not find any significant difference in number of clicks between the control condition and *ad menu/feed dashboard* for all three tasks. Thus, our test for H2.1 was inconclusive. H2.2 was supported across all tasks. Overall, the findings show that high level of actionability was effective in reducing users' effort in finding ad settings.

6.4 RQ3: Participants' perception of Facebook

After participants completed the three tasks, they were asked to answer five questions related to their perception of Facebook (see Figure 11). For all questions, our tests could not find any significant difference between the control group and treatment groups' ratings (for both location of entry point and level of actionability). While our tests showed a significant difference for the question "How likely do you think Facebook will respect your ad-related privacy choices made on their platform?" ($F_{(2,105)} = 4.33, p < .05, \eta_p^2 = 0.08$), post-hoc tests showed that those in *dashboard* condition (*median* = 4) thought it is more likely Facebook would respect their choices compared to the *ad menu* condition (*median* = 3; $p < .05, d = 0.54$). We also found an interaction between *location* and *actionability* ($F_{(1,105)} = 9.44, p < .005, \eta_p^2 = 0.08$). The post-hoc tests showed that those in *dashboard with high actionability* condition (*median* = 4; $p < .05, d = 0.93$) and *dashboard with low actionability* condition (*median* = 4; $p < .05, d = 1.01$) both significantly replied it is more likely that Facebook would respect their ad-related privacy choices compared to the *ad menu with high actionability* condition (*median* = 2).

We found an interaction between *location* and *actionability* regarding participants' responses to the question "How trustworthy is Facebook as a company?" ($F_{(2,105)} = 6.17, p < .05, \eta_p^2 = 0.06$). However, our post-hoc tests could not find pairwise differences.

Summary. Although our tests found a significant difference between *ad menu* and *feed dashboard* groups' perception of Facebook for one question, our tests could not find any significant differences when comparing treatment conditions to the control condition. Thus, our tests for H3.1 and H3.2 that the interventions would be different from the control were inconclusive. Regardless, it is worth noting that overall participants' trust in Facebook was low across all groups. One speculation behind *ad menu* and *feed dashboard* groups' difference could be that participants tended to more positively perceive companies that transparently disclose their advertisement controls and practices via more radical designs like *feed dashboard*, than the *ad menu* (which is an adaptation of Facebook's existing ad contextual menu).

6.5 RQ4: Participants' perceived usability of Facebook's existing ad settings

For participants who encountered the relevant ad settings during the study, we also asked about their perceived usability of each setting using Likert-scale questions (described in Section 5.5.4). Interestingly, our new entry points (*ad menu, feed dashboard*) and high actionability both positively impacted participants' perceived usability of existing Facebook ad settings, although it varied across settings (Figure 12). Here, we report the results that were found as statistically significant by our tests.

Ad Topics. Our tests found that across groups with different entry point locations, there was a significant difference in participants' perception of Ad Topics regarding the interface's complexity ($F_{(2,78)} = 5.25, p < .01, \eta_p^2 = 0.12$) and how understandable it was ($F_{(2,78)} = 6.03, p < .005, \eta_p^2 = 0.13$), as well as participants' sentiment towards the functionality's location ($F_{(2,78)} = 6.55, p < .005, \eta_p^2 = 0.14$). Post-hoc tests showed that participants

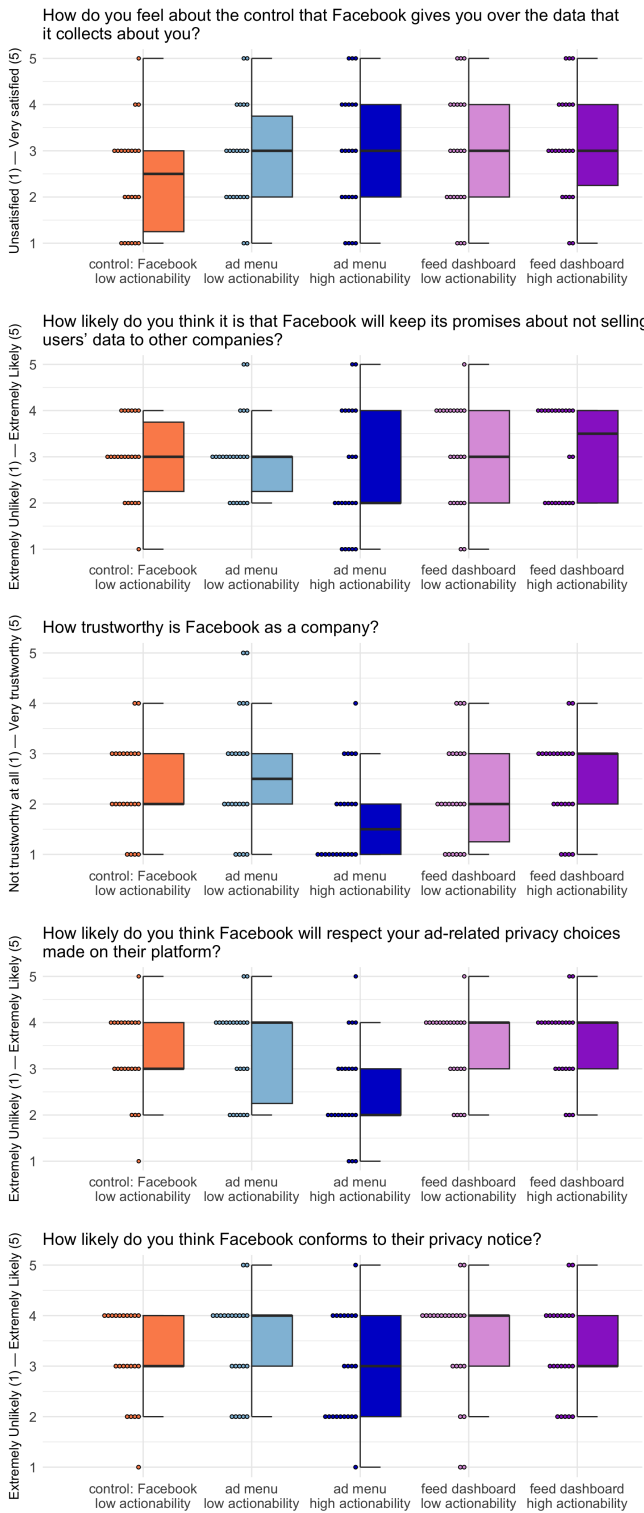


Figure 11: Participants' answers to questions asking about their perception of Facebook.

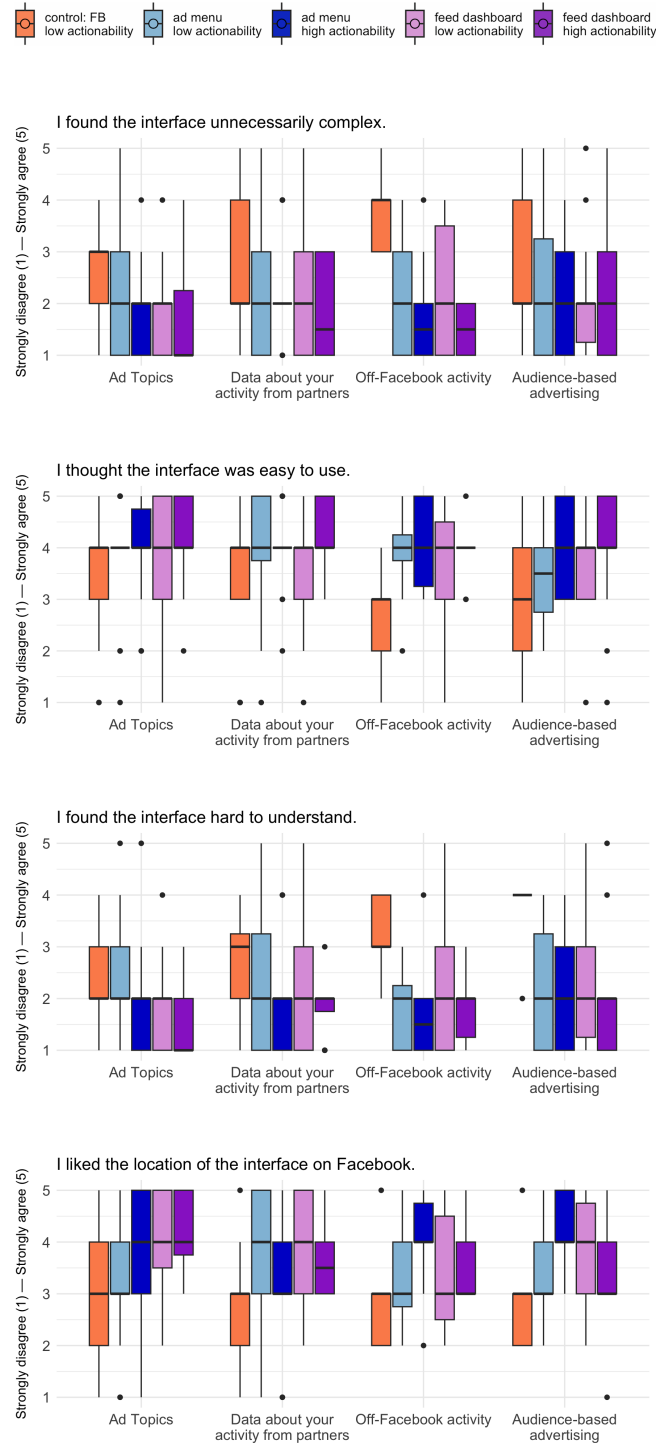


Figure 12: Participants' answers to usability-related Likert-scale questions for Ad Topics page, Data about your activity from partners, Off-Facebook activity page, and Audience-based advertising.

in the *feed dashboard* condition who found Ad Topics (*median* = 4) significantly liked its location more compared to the control group (*median* = 3; $p < .005, d = 1.03$). Interestingly, participants in the *feed dashboard* condition perceived Ad Topics as less complex (*median* = 1) compared to the control group (*median* = 3; $p < .05, d = 0.82$), and also answered it was easier to understand (*median* = 2) compared to the control group (*median* = 1; $p < .05, d = 0.86$). One possible explanation for this finding could be that the treatment groups' reduced effort in finding the setting increased the setting's perceived usability, as participants had less frustration in finding it. Our tests also found that participants in the *high actionability* condition (*median* = 4) liked the Ad Topic's location compared to those in *low actionability* (*median* = 3; $F_{(2,78)} = 4.29, p < .05, \eta_p^2 = 0.05$), and also thought the page was easier to understand (*median* = 2) than the *low actionability* (*median* = 2; $F_{(2,78)} = 5.38, p < .05, \eta_p^2 = 0.06$). It could be that the short explanation provided in the *high actionability* interfaces helped participants understand the setting more easily.

Data about your activity from partners. For *Data about your activity from partners*, our tests found a significant difference in participants' responses to "I found the interface hard to understand" across groups with different entry points ($F_{(2,77)} = 4.37, p < .05, \eta_p^2 = 0.10$). Post-hoc test showed that participants in the *feed dashboard* (*median* = 2) condition perceived it as less difficult to understand compared to the control group (*median* = 3; $p < .05, d = 0.75$).

Off-Facebook activity. Across groups with different locations of controls' entry points, our tests found a significant difference in participants' thoughts on how complex the Off-Facebook activity page was ($F_{(2,39)} = 5.48, p < .01, \eta_p^2 = 0.22$), how easy it was to understand ($F_{(2,39)} = 5.16, p < .05, \eta_p^2 = 0.21$), and how easy it was to use ($F_{(2,39)} = 4.10, p < .05, \eta_p^2 = 0.17$). Post-hoc tests showed that the *feed dashboard* (*median* = 2; $p < .01, d = 1.66$) and *ad menu* (*median* = 2; $p < .01, d = 1.71$) groups perceived the page as less complex than the control group (*median* = 4). Both groups (*ad menu: median* = 2, $p < .05, d = 1.51$; *feed dashboard: median* = 2, $p < .05, d = 1.21$) also perceived it as less hard to understand compared to the control condition (*median* = 3). Those in the *ad menu* groups also thought it was easy to use (*median* = 4; $p < .05, d = 1.36$) compared to the control condition (*median* = 3). Our tests also found that those in *high actionability* conditions (*median* = 1.5) perceived the page as less complex than those in *low actionability* conditions (*median* = 3; $F_{(2,39)} = 9.02, p < .005, \eta_p^2 = 0.19$). *High actionability* groups (*median* = 4) also perceived Off-Facebook activity as easy to use compared to *low actionability* (*median* = 4; $F_{(1,39)} = 5.39, p < .05, \eta_p^2 = 0.12$). One reason behind such findings could be that the participants had less frustration in finding the setting compared to the control group.

Audience-based advertising. Our tests found that across groups with different control entry points, there was a significant difference in how participants found the Audience-based advertising setting easy to understand ($F_{(2,58)} = 4.06, p < .05, \eta_p^2 = 0.12$). Post-hoc tests showed that the control group (*median* = 4) perceived it as more difficult to understand than both the *ad menu* (*median* = 2; $p < .05, d = 1.18$) and *feed dashboard* (*median* = 2; $p < .05, d = 1.24$) groups. Our test also found that participants

in *high actionability* groups (*median* = 2) perceived the setting as easier to understand compared to those in *low actionability* groups (*median* = 2; $F_{(2,58)} = 5.78, p < .05, \eta_p^2 = 0.09$). As mentioned above, one possible reason could be that the short description provided in the interface for *high actionability* groups helped participants' understanding. Lastly, while there was an interaction effect between *location of entry point* and *level of actionability* regarding participants' sentiment towards the location of Audience-based advertising ($F_{(1,58)} = 5.55, p < .05, \eta_p^2 = 0.09$), post-hoc tests did not detect significant differences between combinations of different entry points and level of actionability.

Summary. H4.1 was supported, as entry points impacted participants' perceived usability and sentiment towards the settings interfaces across all tasks (although the questions to which ratings showed difference slightly varied). H4.2 was partially supported; *high actionability* interfaces positively impacted participants' perception of ad setting functionalities compared to *low actionability*, except for *Data about your activity from partners* for which the test could not find any significant difference.

6.6 RQ5: Participants' sentiment towards and perceived usability of the new ad controls

We were further interested in understanding if there were differences in users' perceived usability of and sentiment towards our new ad control interfaces (Figures 13 and 14). When asked about how they felt about the interface overall (Figure 13), participants significantly preferred the *ad menu* designs (*median* = 4) over the *feed dashboard* designs (*median* = 4; $F_{(1,84)} = 5.10, p < .05, \eta_p^2 = 0.03$), although the median values for both were four. Our tests could not find significant difference between *low actionability* (*median* = 4) and *high actionability* (*median* = 5) conditions ($F_{(1,84)} = 0.60, p = .44, \eta_p^2 = 0.01$). Similarly, regarding how participants felt about the options provided in the interface (Figure 13), our tests could not find any significant difference between *low actionability* (*median* = 4) and *high actionability* (*median* = 5) conditions ($F_{(1,84)} = 2.90, p = .09, \eta_p^2 = 0.03$).

Participants were also asked four Likert-scale questions about the new interfaces' usability and location after completing the tasks (Figure 14). Participants significantly preferred *ad menu's* location (*median* = 5) compared to the *feed dashboard's* (*median* = 4; $F_{(1,84)} = 10.48, p < .005, \eta_p^2 = 0.10$). Furthermore, participants rated the *ad menu* as easier to use (*median* = 5) than the *feed dashboard* (*median* = 4; $F_{(1,84)} = 9.49, p < .005, \eta_p^2 = 0.11$).

When asked about how likely they would use the new interface if Facebook implemented it on its platform, overall participants across all treatment groups answered that they would likely use it (Figure 13). Our tests could not find any significant difference between *ad menu* (*median* = 5) and *feed dashboard* conditions (*median* = 4; $F_{(1,84)} = 2.65, p = .11, \eta_p^2 = 0.03$), nor between *low actionability* (*median* = 4) and *high actionability* (*median* = 5) conditions ($F_{(1,84)} = 2.21, p = .14, \eta_p^2 = 0.03$).

Summary. H5.1 was supported, as the results clearly show that participants preferred *ad menu* over *feed dashboard*. However, the overall rating of *feed dashboard* was positive as well (Figures 13 and 14). Our tests could not find any significant difference between

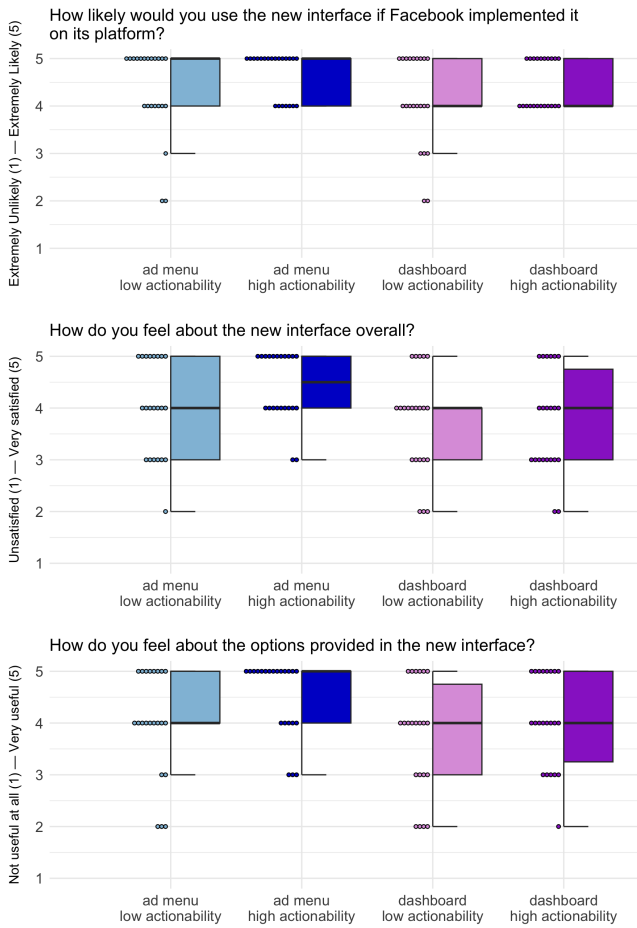


Figure 13: Participants' answers to questions asking about their sentiment towards the new ad control interfaces.

high actionability and *low actionability*. Thus, our test for H5.2 was inconclusive. But similarly, participants' ratings were overall high for both types of interfaces (Figures 13 and 14).

7 DISCUSSION

In this section, we discuss specific design recommendations for improving OBA controls' findability and our findings' implications for regulation. In particular, we emphasize the importance of integrating specific user research and design requirements into platform regulation, and the role academics can play.

7.1 Adjusting Entry Points and Actionability Results in Findable and Usable OBA Controls

Our findings show that the entry point location significantly impacted ad controls' findability or discoverability for two out of three tasks. The *feed dashboard* design significantly increased the discoverability rate for the Ad Topics page. Both *ad menu* and *feed dashboard* designs helped participants find the audience-based

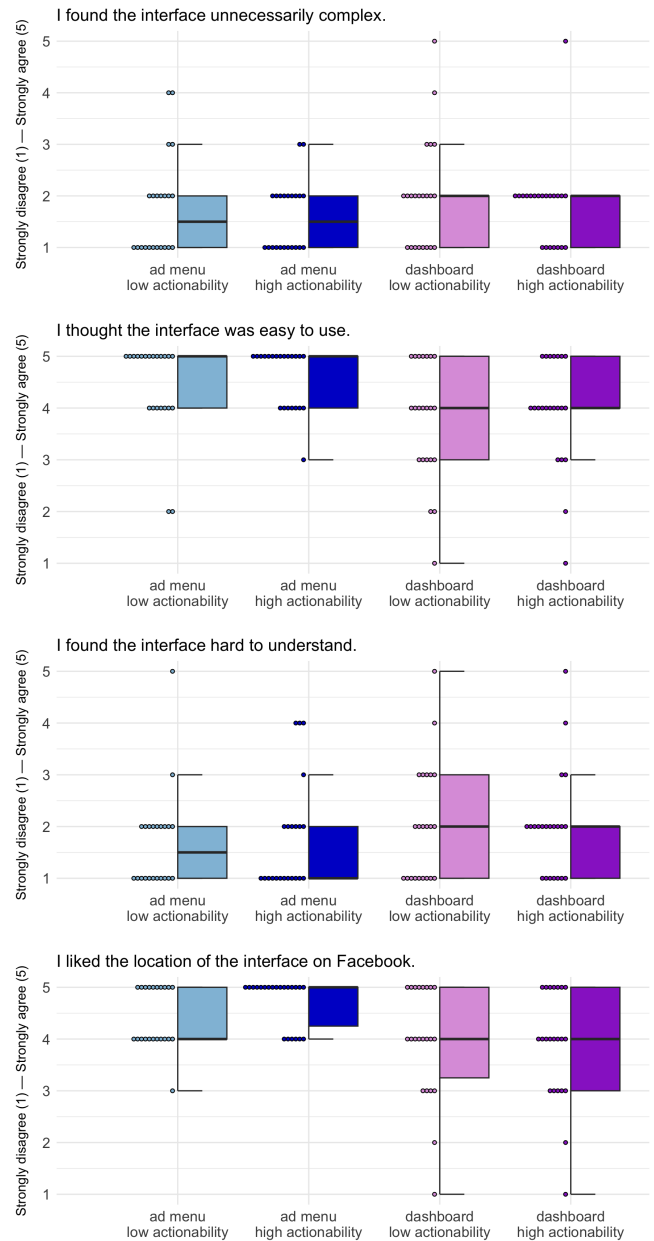


Figure 14: Participants' answers to usability-related Likert-scale questions for the new ad control interfaces.

advertising control. In particular, many participants gave high usability ratings for the *ad menu* designs (Figure 13 and 14). It is worth noting that the ad button/menu we designed is far from what platform companies typically provide users. The button is larger, has a blue background, and even has an icon, which is quite different from the three small dots that Facebook provides (Figure 1). Further, the newly designed contextual menu includes general ad settings, and not just controls for one particular ad. Although our designs may be atypical, participants in our study gave them high usability

ratings. In short, our findings show there are benefits to making OBA controls more prominent in ads and including more options within ad contextual menus.

While participants preferred the *ad menu* design over the *feed dashboard* design, the *feed dashboard* was actually more effective in increasing the findability rate for some tasks. For example, the findability rate of Audience-based advertising increased from 22.7% to 63.6% when participants were given the *feed dashboard with high actionability*. However, participants seemed to prefer it less than the *ad menu*, mainly due to it taking up space in the main feed, where they are used to seeing posts first. This suggests that surfacing ad settings and possibly other settings in a dashboard embedded in the feed is a promising approach, but that the dashboard has to be designed carefully to not frustrate the user experience, for example by giving users the option to hide the dashboard or by only having it appear periodically, for what Im et al. call “periodic checks” [36].

Our study’s results also show that actionability is crucial for ensuring that ad controls are findable, as *high actionability* increased ad settings’ findability across all tasks. This is encouraging when also considering participants’ high usability ratings for the *high actionability* interfaces (Figure 13). Conventionally, users, and even researchers, are used to major tech companies providing many privacy functions on the general settings page, with vague descriptions about what a user can find there. From a user experience perspective, this could seem intuitive, as it is important to not overwhelm users with too many functions. However, participants’ response to *high actionability* interfaces were overall positive and our tests did not find any statistical significance between participants’ preference between *low actionability* and *high actionability*. In short, our study’s results show that surfacing more actionable options *can benefit users without overwhelming them*.

Simultaneously, it should be noted that our interfaces with *high actionability* included options that were related to study tasks, which might have contributed to high perceptions of usefulness. For future work, researchers and platforms should consider how different ad setting functionalities should be surfaced in a findable and actionable way for different groups of users. Habib et al. [29] have identified four user groups (“The Privacy Concerned,” “The Advertising Curators,” “The Advertising Irritated,” “The Advertising Disengaged”) when it comes to ad control needs, and have argued that a “one-size-fits-all” approach will likely result in not meeting some users’ needs. Platforms could potentially design a feature that asks users what kind of ad-related needs they have, and then populate *ad menu* and *feed dashboard* interfaces with the desired options in an actionable way.

Lastly, it is important to acknowledge that even if ad controls are findable and actionable, it is still likely that default values are important. For example, Liu et al. [41] showed that 36% of content on Facebook (from 200 users) was shared with the default visibility settings. Furthermore, because default settings are closely tied to the platforms’ business, they may not be what some or many users expect [41]. Thus, while it is important for ad controls to be findable and actionable so that users can mark their preferences whenever they want to, regulators, researchers, and industry practitioners should understand the importance of the controls’ default values.

7.2 Ad Controls’ Entry Points and Actionability Impact Users’ Sentiments towards Ad Settings

Varying ad controls’ entry points and level of actionability also impacted users’ perceived usability of and sentiment towards existing ad settings on Facebook. One possible interpretation is that how the ad setting’s entry point is designed impacts how users perceive the final destination of the function because users are less frustrated by the search for specific settings. In particular, findings show it is likely the explanations provided in interfaces with high actionability helped users better understand what actions they can actually take in the setting. Currently, the wordings that companies provide in OBA settings are not action-oriented. For example, Facebook’s links to OBA settings rarely start with verbs (e.g., “Data about your activity from partners,” “Audience-based advertising”). Formulating the option explanations so that they include what actions users can take (e.g., “Decide if you want to see ads based on your activity on other businesses’ websites/apps or offline”) can help users find the right settings and also increase user satisfaction.

On the other hand, our tests did not find that our interventions significantly impacted participants’ perception of Facebook compared to the baseline. This could be because of users’ fatigue in Facebook’s history of privacy controversies [24, 29] and the short duration of our experiment. That is, a short interaction with the interventions could not have been enough to impact users’ perceived image of Facebook, but enough to more positively evaluate the settings, especially considering many participants newly discovered them via our interventions (Figures 7-9).

Thus, one important future research direction is to understand the longitudinal impact of such changes in users’ sentiment [39]. While participants’ sentiment towards platforms’ ad settings is relevant to user experience, it is an open question what that means for users’ perception of companies in the long run. As Hohnhold et al.’s study [34] has shown, what may immediately seem contradictory to a company’s business may not be so in the end, when considering user satisfaction is pivotal for long-term revenue.

One intervention that should be explored in future longitudinal studies is strongly signaling to users why the company is providing findable and actionable ad controls (e.g., the company aims to rebuild users’ trust by providing findable and actionable privacy mechanisms, an approach researchers have suggested [29]). This is different from our study where we silently augmented our new designs without any explanation on what they mean for users and the company (e.g., company’s intention). If such studies show that findable and actionable ad controls actually lead to a positive change in users’ perception of companies or online services, it means companies do have an incentive to design such ad controls.

7.3 Integrating Interface Design and Platform Regulation

Our findings also highlight that discovering ad privacy controls does not necessarily mean users recognize and understand their functionality. Perhaps unsurprisingly, a few participants who came across the appropriate setting for a given task did not recognize it as such. For the task that was most difficult for participants, finding advertisers that used tailored audience lists, it could be that the

concept of advertising list was hard to grasp because many users are not aware of this practice [65]. Inspecting the log data showed that many participants who failed to recognize *Audience-based advertising* ended up going to other settings, such as *Data about your activity from partners*. Even for those who recognized the right ad setting, comparing the log data and participants' survey responses revealed that some participants ended up making privacy choices inconsistent with the assigned task. For example, a few participants who found *Data about your activity from partners* actually opted in to see personalized ads based on their online activity on other apps and websites, instead of opting out. One major reason was due to Facebook's confusing interface, as participants could not understand whether they were currently opted in or out.

Therefore, it is critical to take a holistic perspective in considering both interface design and regulation [25]. Many regulations across countries say that companies should provide OBA controls to users, but they do not concretely specify *how* they should be designed. Our findings indicate the need for regulators to provide specific and research-informed guidance and requirements to companies on how to design OBA controls. Regulation should also require companies to user-test their OBA settings, just as they intensively test features related to their traditional business models. Companies should also be required to publicly release the results, similar to reports in other sectors (e.g., environmental impact, product safety), given the impact of the OBA ecosystem on society [32, 63].

Lastly, just as researchers have argued for the importance of auditing algorithms on platforms [5, 11, 19], there is value in auditing how tech companies create and change their ad control interfaces. Based on the audits, researchers can provide concrete design insights for regulation. We argue that academic researchers, as well as organizations that are independent from major platforms, can play a central role in this. At the same time, it is important to acknowledge that despite the importance of studying and designing ad control interfaces, there is also anecdotal evidence that platforms may not honor users' choices marked via such controls [27]. We thus acknowledge that it is also critical for researchers and organizations to audit the efficacy of ad settings. In particular, longitudinal studies are needed [2]. This is because how the platforms' algorithms surface ads in the long run may differ from what the user sees immediately after making changes in their settings [27].

8 CONCLUSION AND FUTURE WORK

While tech companies that rely on ads for revenue argue users have control over their data via ad privacy settings, prior research has shown that they are often difficult to find. In this work, we explored how to design findable and actionable ad controls and studied their impact on users' behavior and sentiment. We first conducted a formative study to iteratively design ad control interfaces. Our designs varied in the setting entry point (within ads, at the feed's top) and setting's level of actionability, where high actionability directly surfaced links to specific advertisement settings, and low actionability pointed to general settings pages. Then, we built a Chrome extension to augment the designs on Facebook and conducted a between-subjects online experiment with 110 participants. Our findings show that entry points within ads or at the feed's top,

as well as interfaces with high actionability, increased ad settings' findability and discoverability. Controls with high actionability also reduced users' effort to find ad settings. Furthermore, participants gave high usability ratings to controls with high actionability, with no significant difference from low actionability, which shows there is potential to design more actionable ad controls. We conclude with specific recommendations for designing findable ad privacy settings and respective regulation. In particular, we emphasize the importance of regulation on platforms to provide research-informed requirements to companies on designing usable OBA controls. Future work should explore the possibility of measuring the long-term impact of findable ad privacy control interfaces as well as audit how platforms' ad setting interfaces evolve.

ACKNOWLEDGMENTS

We thank the Security Privacy Interaction Lab at the University of Michigan, especially Sean Scarnecchia, Annie Chen, Abraham Mhaidli, Yixin Zou, and Brian Tang, for their feedback on the work. Sean Scarnecchia and Annie Chen provided valuable help on improving the Chrome extension. Abraham Mhaidli and Yixin Zou gave detailed advice on the survey and experiment design. Brian Tang gave helpful tips on refining the survey and using Qualtrics. We also thank members of the Computational HCI Lab, especially Sumit Asthana, Anindya Das Antar, Divya Ramesh, Snehal Prabhudesai, Tsedeniya Amare, and Nel Escher, for pilot testing the Chrome extension or/and giving feedback on an early draft of the paper. Im also thanks them for the mental support, which made the paper-writing process much more enjoyable. We also thank Jiyeon Kim for giving immensely helpful feedback on the designs and Jaewon Kim for providing comments that improved the survey wordings. We would also like to thank the reviewers for their constructive and insightful comments. Im would especially like to thank Chaerin Im, Jiyeon Kim, and Jeeyoon Hyun for their encouragement and feedback, and Nayoung Kim, Seyoung Im, and Jungsook Cheon, for their wise advice. Im was supported by a University of Michigan Barbour Scholarship.

REFERENCES

- [1] Alessandro Acquisti, Idris Adjerid, Rebecca Balebako, Laura Brandimarte, Lorie Faith Cranor, Saranga Komanduri, Pedro Giovanni Leon, Norman Sadeh, Florian Schaub, Manya Sleeper, et al. 2017. Nudges for privacy and security: Understanding and assisting users' choices online. *ACM Computing Surveys (CSUR)* 50, 3 (2017), 1–41.
- [2] Muhammad Ali, Angelica Goetzen, Piotr Sapiezynski, Elissa Redmiles, and Alan Mislove. 2022. All Things Unequal: Measuring Disparity of Potentially Harmful Ads on Facebook. In *Proceedings of the Workshop on Technology and Consumer Protection (ConPro'22)*. San Francisco, CA, USA.
- [3] Digital Advertising Alliance. 2022. YourAdChoices. <https://optout.aboutads.info>.
- [4] Athanasios Andreou, Giridhari Venkatadri, Oana Goga, Krishna Gummadi, Patrick Loiseau, and Alan Mislove. 2018. Investigating ad transparency mechanisms in social media: A case study of Facebook's explanations. In *NDSS 2018-Network and Distributed System Security Symposium*. 1–15.
- [5] Joshua Asplund, Motahare Eslami, Hari Sundaram, Christian Sandvig, and Karrie Karahalios. 2020. Auditing race and gender discrimination in online housing markets. In *Proceedings of the International AAAI Conference on Web and Social Media*, Vol. 14. 24–35.
- [6] Brooke Auxier, Lee Rainie, Monica Anderson, Andrew Perrin, Madhu Kumar, and Erica Turner. 2019. Americans and privacy: Concerned, confused and feeling lack of control over their personal information. (2019).
- [7] Nikola Banovic, Fanny Chevalier, Tovi Grossman, and George Fitzmaurice. 2012. Triggering Triggers and Burying Barriers to Customizing Software. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems* (Austin, Texas,

- USA) (CHI '12). Association for Computing Machinery, New York, NY, USA, 2717–2726. <https://doi.org/10.1145/2207676.2208666>
- [8] Natá M Barbosa, Gang Wang, Blase Ur, and Yang Wang. 2021. Who Am I? A Design Probe Exploring Real-Time Transparency about Online and Offline User Profiling Underlying Targeted Ads. *Proceedings of the ACM on Interactive, Mobile, Wearable and Ubiquitous Technologies* 5, 3 (2021), 1–32.
 - [9] Md Momen Bhuiyan, Michael Horning, Sang Won Lee, and Tanushree Mitra. 2021. NudgeCred: Supporting News Credibility Assessment on Social Media Through Nudges. *Proc. ACM Hum.-Comput. Interact.* 5, CSCW2, Article 427 (oct 2021), 30 pages. <https://doi.org/10.1145/3479571>
 - [10] John Brooke. 1995. SUS: A quick and dirty usability scale. *Usability Eval. Ind.* 189 (11 1995).
 - [11] Shea Brown, Jovana Davidovic, and Ali Hasan. 2021. The algorithm audit: Scoring the algorithms that score us. *Big Data & Society* 8, 1 (2021), 2053951720983865.
 - [12] Fred H Cate. 2006. The failure of fair information practice principles. *Consumer protection in the age of the information economy* (2006).
 - [13] Facebook Help Center. 2022. How Ads Work on Facebook. https://www.facebook.com/help/516147308587266/?helpref=hc_fnav.
 - [14] Patrali Chatterjee. 2008. Are Unclicked Ads Wasted? Enduring Effects of Banner and Pop-Up Ad Exposures on Brand Memory and Attitudes. *Journal of electronic commerce Research* 9, 1 (2008).
 - [15] Jessica Colnago, Lorrie Faith Cranor, Alessandro Acquisti, and Kate Hazel Jain. 2022. Is it a concern or a preference? An investigation into the ability of privacy scales to capture and distinguish granular privacy constructs. In *Eighteenth Symposium on Usable Privacy and Security (SOUPS 2022)*. 331–346.
 - [16] Kovila PL Coopamootoo, Maryam Mehrzad, and Ehsan Toreini. 2022. "I feel invaded, annoyed, anxious and I may protect myself": Individuals' Feelings about Online Tracking and their Protective Behaviour across Gender and Country. *arXiv preprint arXiv:2202.04682* (2022).
 - [17] Claire Dolin, Ben Weinschel, Shawn Shan, Chang Min Hahn, Euirim Choi, Michelle L Mazurek, and Blase Ur. 2018. Unpacking perceptions of data-driven inferences underlying online targeting and personalization. In *Proceedings of the 2018 CHI conference on human factors in computing systems*. 1–12.
 - [18] Lisa A Elkin, Matthew Kay, James J Higgins, and Jacob O Wobbrock. 2021. An aligned rank transform procedure for multifactor contrast tests. In *The 34th Annual ACM Symposium on User Interface Software and Technology*. 754–768.
 - [19] Nel Escher and Nikola Banovic. 2020. Exposing Error in Poverty Management Technology: A Method for Auditing Government Benefits Screening Tools. *Proceedings of the ACM on Human-Computer Interaction* 4, CSCW1 (2020), 1–20.
 - [20] Motahhare Eslami, Sneha R Krishna Kumaran, Christian Sandvig, and Karrie Karahalios. 2018. Communicating algorithmic process in online behavioral advertising. In *Proceedings of the 2018 CHI conference on human factors in computing systems*. 1–13.
 - [21] Florian M Farke, David G Balash, Maximilian Golla, Markus Dürmuth, and Adam 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.
 - [22] Lisa Farman, Maria Leonora Comello, and Jeffrey R Edwards. 2020. Are consumers put off by retargeted ads on social media? Evidence for perceptions of marketing surveillance and decreased ad effectiveness. *Journal of Broadcasting & Electronic Media* 64, 2 (2020), 298–319.
 - [23] Yuanyuan Feng, Yaxing Yao, and Norman Sadeh. 2021. A design space for privacy choices: Towards meaningful privacy control in the internet of things. In *Proceedings of the 2021 CHI Conference on Human Factors in Computing Systems*. 1–16.
 - [24] Sheera Frenkel and Cecilia Kang. 2021. *An ugly truth: Inside Facebook's battle for domination*. Hachette UK.
 - [25] Colin M Gray, Cristiana Santos, Nataliia Bielova, Michael Toth, and Damian Clifford. 2021. Dark patterns and the legal requirements of consent banners: An interaction criticism perspective. In *Proceedings of the 2021 CHI Conference on Human Factors in Computing Systems*. 1–18.
 - [26] Gabriel Grill. 2021. Future protest made risky: Examining social media based civil unrest prediction research and products. *Computer Supported Cooperative Work (CSCW)* 30, 5 (2021), 811–839.
 - [27] Dorota Glowacka and Karolina Iwańska. 2021. Algorithms of trauma: new case study shows that Facebook doesn't give users real control over disturbing surveillance ads. Retrieved December 9, 2022 from <https://en.panoptykon.org/algorithms-of-trauma>
 - [28] Hana Habib and Lorrie Faith Cranor. 2022. Evaluating the usability of privacy choice mechanisms. In *Eighteenth Symposium on Usable Privacy and Security (SOUPS 2022)*. 273–289.
 - [29] Hana Habib, Sarah Pearman, Ellie Young, Ishika Saxena, Robert Zhang, and Lorrie Faith Cranor. 2022. Identifying User Needs for Advertising Controls on Facebook. *Proc. ACM Hum.-Comput. Interact.* 6, CSCW1, Article 59 (apr 2022), 42 pages. <https://doi.org/10.1145/3512906>
 - [30] Hana Habib, Yixin Zou, Aditi Jannu, Neha Sridhar, Chelse Swoopes, Alessandro Acquisti, Lorrie Faith Cranor, Norman Sadeh, and Florian Schaub. 2019. An Empirical Analysis of Data Deletion and {Opt-Out} Choices on 150 Websites. In *Fifteenth Symposium on Usable Privacy and Security (SOUPS 2019)*. 387–406.
 - [31] Hana Habib, Yixin Zou, Yaxing Yao, Alessandro Acquisti, Lorrie Cranor, Joel Reidenberg, Norman Sadeh, and Florian Schaub. 2021. Toggles, dollar signs, and triangles: How to (in) effectively convey privacy choices with icons and link texts. In *Proceedings of the 2021 CHI Conference on Human Factors in Computing Systems*. 1–25.
 - [32] Brittan Heller and Avi Bar-Zeev. 2021. The problems with immersive advertising: in AR/VR, nobody knows you are an ad. *Journal of Online Trust and Safety* 1, 1 (2021).
 - [33] Anne Helmond, David B Nieborg, and Fernando N van der Vlist. 2017. The political economy of social data: a historical analysis of platform-industry partnerships. In *Proceedings of the 8th international conference on social media & society*. 1–5.
 - [34] Henning Hohnhold, Deirdre O'Brien, and Diane Tang. 2015. Focusing on the long-term: It's good for users and business. In *Proceedings of the 21th ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*. 1849–1858.
 - [35] Silas Hsu, Kristen Vaccaro, Yin Yue, Aimee Rickman, and Karrie Karahalios. 2020. Awareness, navigation, and use of feed control settings online. In *Proceedings of the 2020 CHI Conference on Human Factors in Computing Systems*. 1–13.
 - [36] Jane Im, Jill Dimond, Melody Berton, Una Lee, Katherine Mustelie, Mark S. Ackerman, and Eric Gilbert. 2021. Yes: Affirmative Consent as a Theoretical Framework for Understanding and Imagining Social Platforms. In *Proceedings of the 2021 CHI Conference on Human Factors in Computing Systems (Yokohama, Japan) (CHI '21)*. Association for Computing Machinery, New York, NY, USA, Article 403, 18 pages. <https://doi.org/10.1145/3411764.3445778>
 - [37] Jane Im, Sonali Tandon, Eshwar Chandrasekharan, Taylor Denby, and Eric Gilbert. 2020. Synthesized Social Signals: Computationally-Derived Social Signals from Account Histories. In *Proceedings of the 2020 CHI Conference on Human Factors in Computing Systems (Honolulu, HI, USA) (CHI '20)*. Association for Computing Machinery, New York, NY, USA, 1–12. <https://doi.org/10.1145/3313831.3376383>
 - [38] Network Advertising Initiative. 2022. NAI Consumer Opt Out. <https://optout.networkadvertising.org/>.
 - [39] Yucheng Jin, Karsten Seipp, Erik Duval, and Katrien Verbert. 2016. Go with the flow: effects of transparency and user control on targeted advertising using flow charts. In *Proceedings of the international working conference on advanced visual interfaces*. 68–75.
 - [40] Michael Kretschmer, Jan Pennekamp, and Klaus Wehrle. 2021. Cookie banners and privacy policies: Measuring the impact of the GDPR on the web. *ACM Transactions on the Web (TWEB)* 15, 4 (2021), 1–42.
 - [41] Yabing Liu, Krishna P. Gummadi, Balachander Krishnamurthy, and Alan Mislove. 2011. Analyzing Facebook Privacy Settings: User Expectations vs. Reality. In *Proceedings of the 2011 ACM SIGCOMM Conference on Internet Measurement Conference (Berlin, Germany) (IMC '11)*. Association for Computing Machinery, New York, NY, USA, 61–70. <https://doi.org/10.1145/2068816.2068823>
 - [42] Natasha Lomas. 2021. UK privacy watchdog warns adtech the end of tracking is nigh. <https://techcrunch.com/2021/11/25/ico-adtech-reform-opinion/>
 - [43] Natasha Lomas. 2022. Give users genuine control over ad targeting, MEPs urged. <https://techcrunch.com/2022/01/13/joint-statement-surveillance-ads-inferred-data/>
 - [44] Mary Madden, Michele Gilman, Karen Levy, and Alice Marwick. 2017. Privacy, poverty, and big data: A matrix of vulnerabilities for poor Americans. *Wash. UL Rev.* 95 (2017), 53.
 - [45] Miguel Malheiros, Charlene Jennett, Sneha Patel, Sacha Brostoff, and Martina Angela 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 (Austin, Texas, USA) (CHI '12)*. Association for Computing Machinery, New York, NY, USA, 579–588. <https://doi.org/10.1145/2207676.2207758>
 - [46] Aleecia McDonald and Lorrie Faith Cranor. 2010. Beliefs and behaviors: Internet users' understanding of behavioral advertising. *Tprc*.
 - [47] Aleecia M McDonald and Lorrie Faith Cranor. 2008. The cost of reading privacy policies. *Isjlp* 4 (2008), 543.
 - [48] Jakob Nielsen. 1995. How to conduct a heuristic evaluation. *Nielsen Norman Group* 1, 1 (1995), 8.
 - [49] Helen Nissenbaum. 2011. A contextual approach to privacy online. *Daedalus* 140, 4 (2011), 32–48.
 - [50] Midas Nouwens, Ilaria Liccardi, Michael Veale, David Karger, and Lalana Kagal. 2020. Dark patterns after the GDPR: Scraping consent pop-ups and demonstrating their influence. In *Proceedings of the 2020 CHI conference on human factors in computing systems*. 1–13.
 - [51] Jonathan A Obar and Anne Oeldorf-Hirsch. 2020. The biggest lie on the internet: Ignoring the privacy policies and terms of service policies of social networking services. *Information, Communication & Society* 23, 1 (2020), 128–147.
 - [52] Eugenia Politou, Efthimios Alepis, and Constantinos Patsakis. 2018. Forgetting personal data and revoking consent under the GDPR: Challenges and proposed solutions. *Journal of Cybersecurity* 4, 1 (2018), tty001.
 - [53] Emilee Rader. 2014. Awareness of behavioral tracking and information privacy concern in facebook and google. In *10th Symposium On Usable Privacy and Security (SOUPS 2014)*. 51–67.

- [54] Joel R Reidenberg, Travis Breaux, Lorrie Faith Cranor, Brian French, Amanda Grannis, James T Graves, Fei Liu, Aleecia McDonald, Thomas B Norton, and Rohan Ramanath. 2015. Disagreeable privacy policies: Mismatches between meaning and users' understanding. *Berkeley Tech. LJ* 30 (2015), 39.
- [55] Florian Schaub, Rebecca Balebako, and Lorrie Faith Cranor. 2017. Designing effective privacy notices and controls. *IEEE Internet Computing* 21, 3 (2017), 70–77.
- [56] Florian Schaub, Rebecca Balebako, Adam L Durity, and Lorrie Faith Cranor. 2015. A design space for effective privacy notices. In *Eleventh symposium on usable privacy and security (SOUPS 2015)*. 1–17.
- [57] Florian Schaub and Lorrie Cranor. 2020. *Usable and Useful Privacy Interfaces*. IAPP.
- [58] Paul M Schwartz. 1999. Internet privacy and the state. *Conn. L. Rev.* 32 (1999), 815.
- [59] Daniel J Solove. 2012. Introduction: Privacy self-management and the consent dilemma. *Harv. L. Rev.* 126 (2012), 1880.
- [60] Michael Carl Tschantz, Serge Egelman, Jaeyoung Choi, Nicholas Weaver, and Gerald Friedland. 2018. The accuracy of the demographic inferences shown on Google's Ad Settings. In *Proceedings of the 2018 Workshop on Privacy in the Electronic Society*. 33–41.
- [61] 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*. 1–15.
- [62] Fernando N Van der Vlist and Anne Helmond. 2021. How partners mediate platform power: Mapping business and data partnerships in the social media ecosystem. *Big Data & Society* 8, 1 (2021), 20539517211025061.
- [63] José Van Dijck, David Nieborg, and Thomas Poell. 2019. Reframing platform power. *Internet Policy Review* 8, 2 (2019), 1–18.
- [64] Ari Ezra Waldman. 2021. *Industry unbound: The inside story of privacy, data, and corporate Power*. Cambridge University Press.
- [65] Miranda Wei, Madison Stamos, Sophie Veys, Nathan Reitingger, Justin Goodman, Margot Herman, Dorota Filipczuk, Ben Weinschel, Michelle L Mazurek, and Blase 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.
- [66] Ben Weinschel, Miranda Wei, Mainack Mondal, Euirim Choi, Shawn Shan, Claire Dolin, Michelle L Mazurek, and Blase 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.
- [67] James Wilson and Daniel Rosenberg. 1988. Rapid prototyping for user interface design. Elsevier, 859–875.
- [68] Jacob O Wobbrock, Leah Findlater, Darren Gergle, and James J Higgins. 2011. The aligned rank transform for nonparametric factorial analyses using only anova procedures. In *Proceedings of the SIGCHI conference on human factors in computing systems*. 143–146.
- [69] Yaxing Yao, Davide Lo Re, and Yang 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.
- [70] Mark Zuckerberg. 2019. The Facts About Facebook. <https://www.wsj.com/articles/the-facts-about-facebook-11548374613>

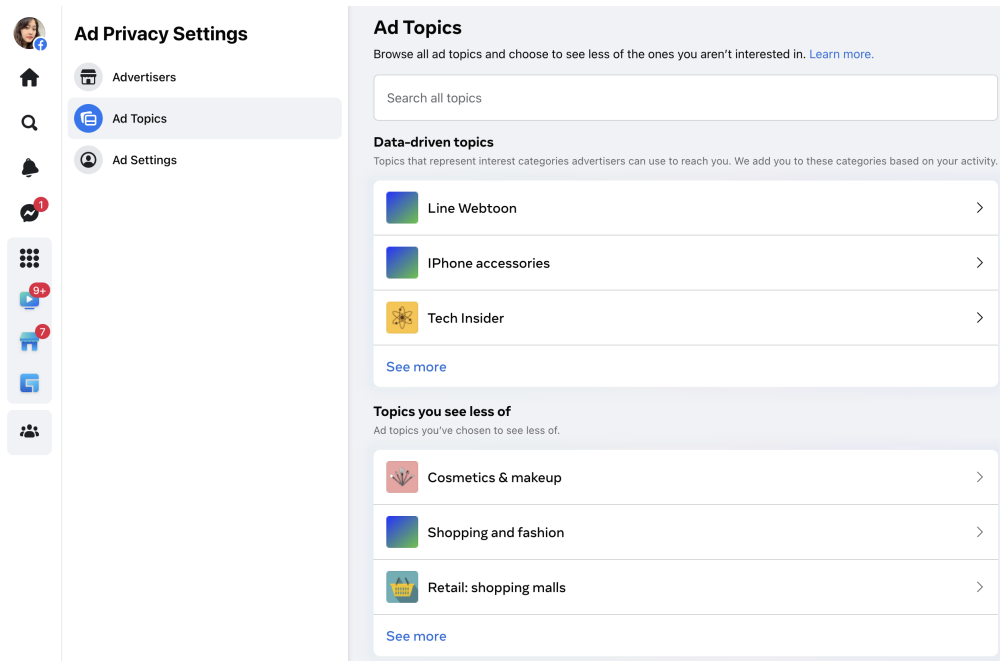


Figure 15: Screenshot of Facebook’s Ad Topics page during the experiment.

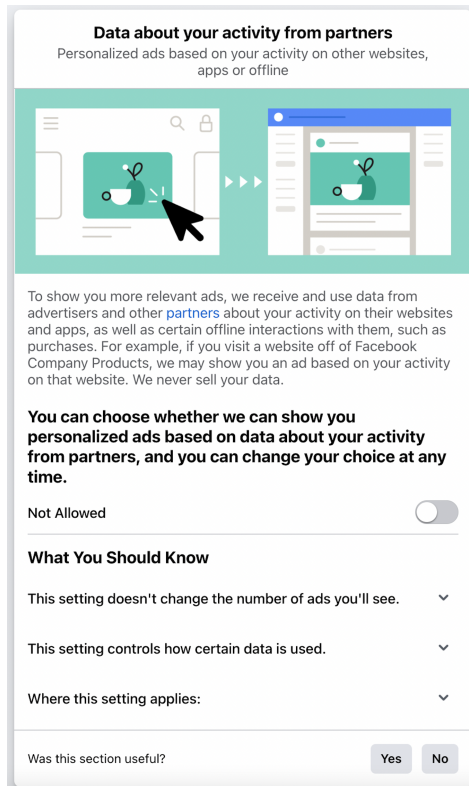


Figure 16: Screenshot of Facebook’s Data about your activity from partners during the experiment. There were slight variations of the opt-in/out wording due to A/B testing (e.g., “Allowed/Not Allowed”, “Use Data from Partners”), but otherwise the popup was largely the same.

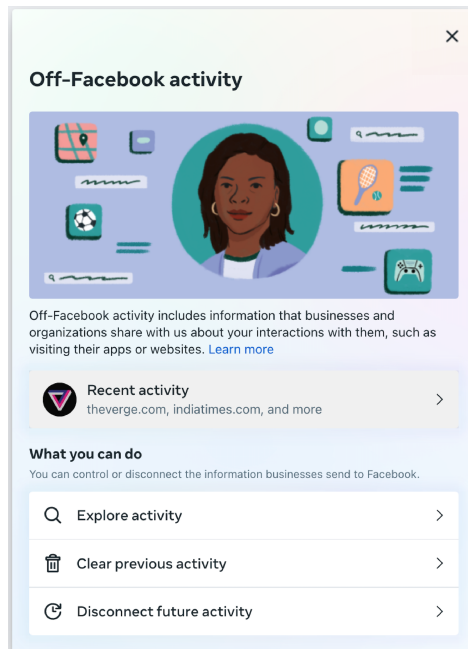


Figure 17: Screenshot of Facebook's Off-Facebook activity page during the experiment.

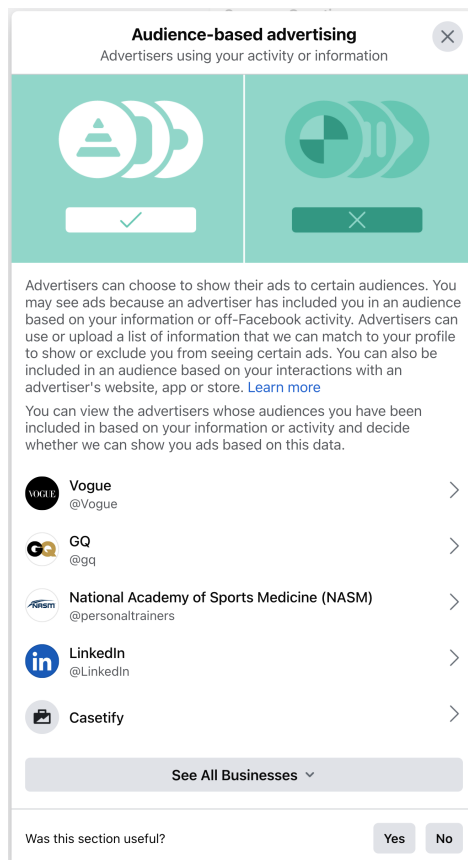


Figure 18: Screenshot of Facebook's Audience-based advertising during the experiment.

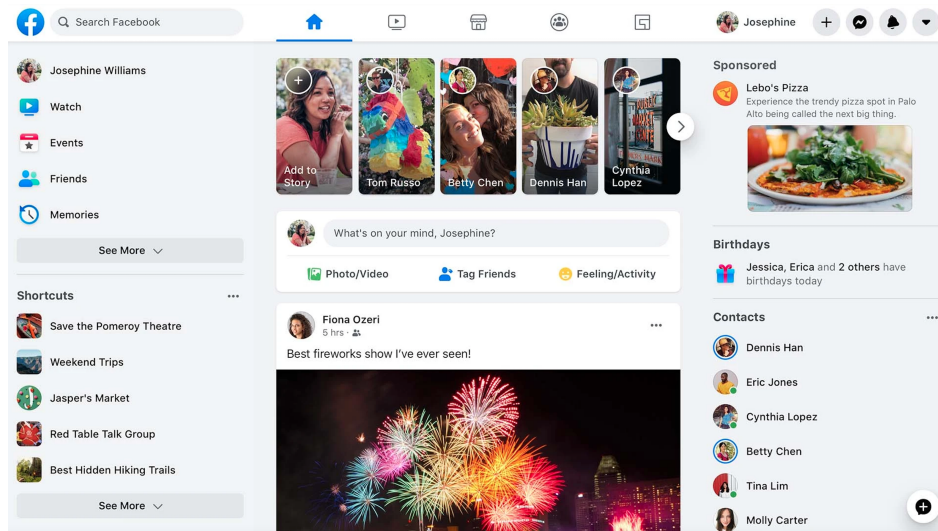


Figure 19: Image of Facebook’s main feed, provided by Meta in July 2020 (source: <https://engineering.fb.com/2020/07/30/web/facebook-com-accessibility/>). The dropdown icon was later changed around June 2022 to a profile image (example: Figure 23). At the top right, there is a dropdown icon which opens a menu that shows options including Settings & privacy.

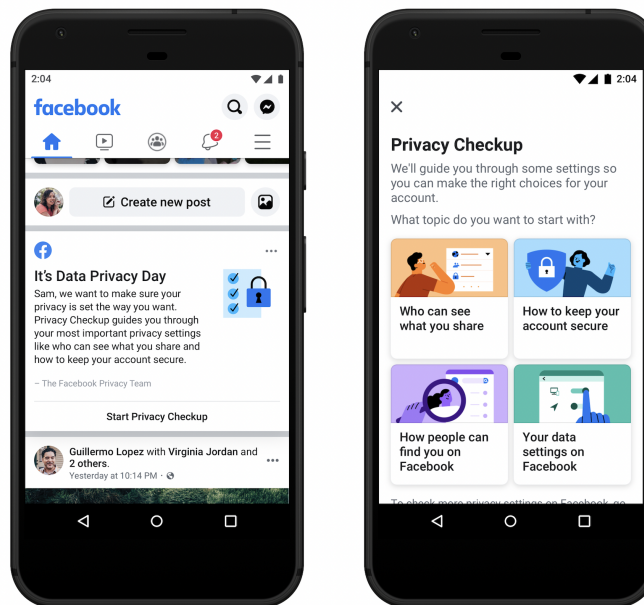


Figure 20: Image of Facebook’s Privacy Checkup box appearing in the main feed, provided by Meta in January 2020 (source: <https://about.fb.com/news/2020/01/data-privacy-day-2020/>).

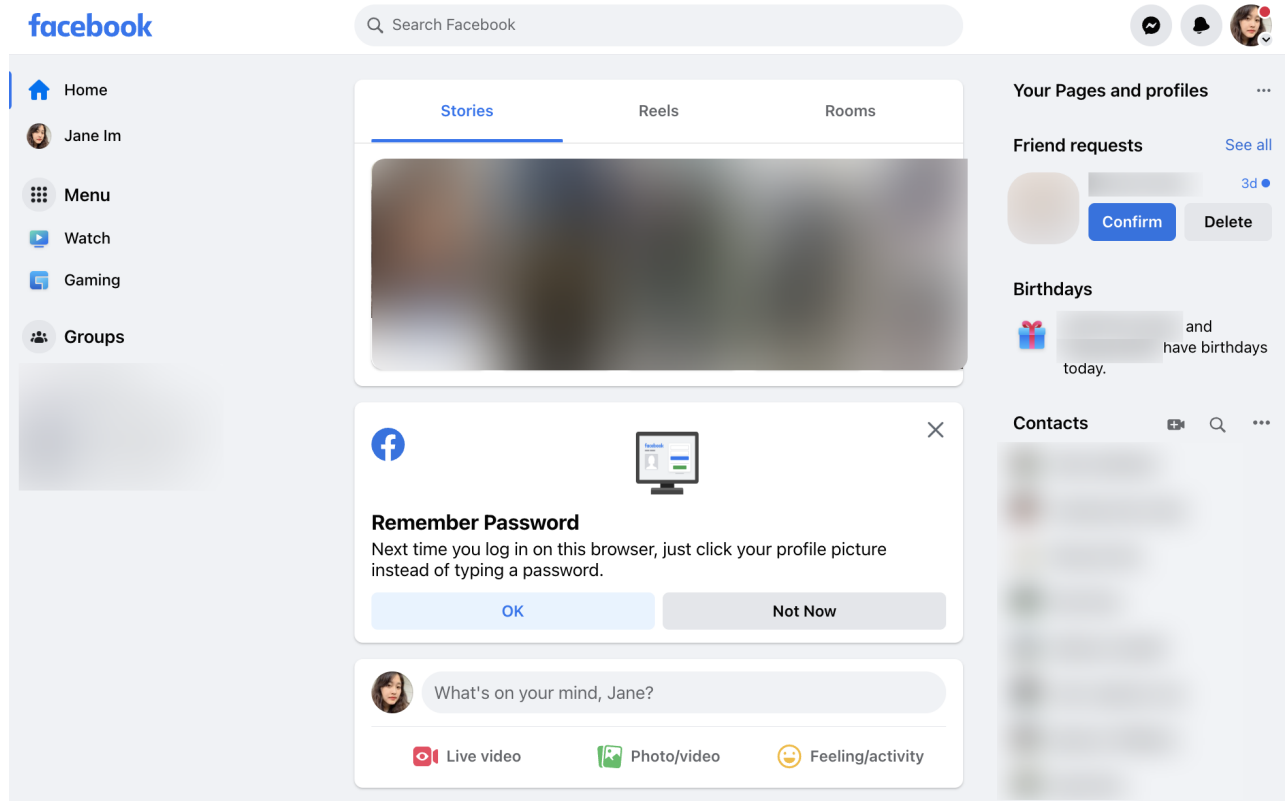


Figure 21: Image showing Facebook’s reminder about passwords. It is located between Stories and posts.

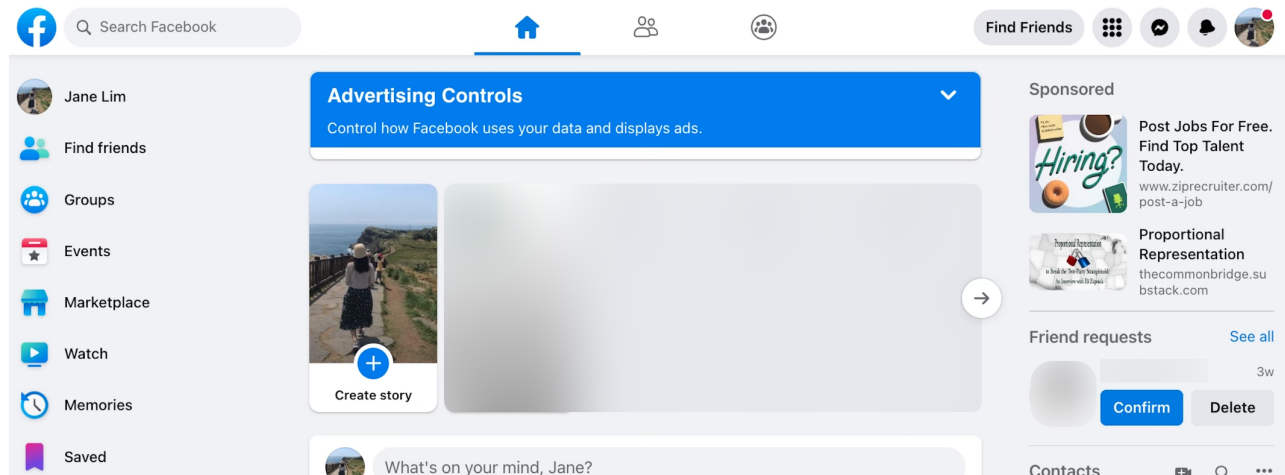
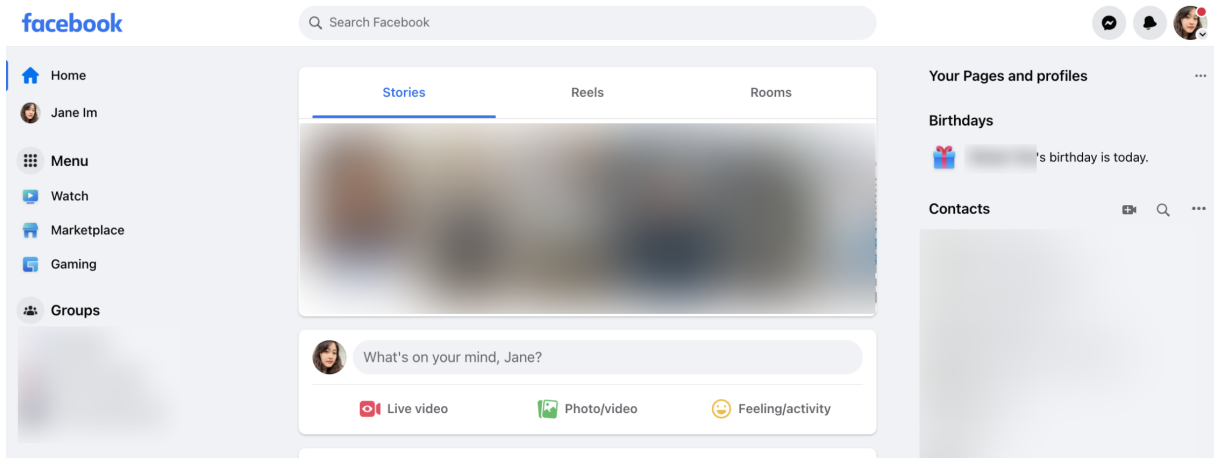
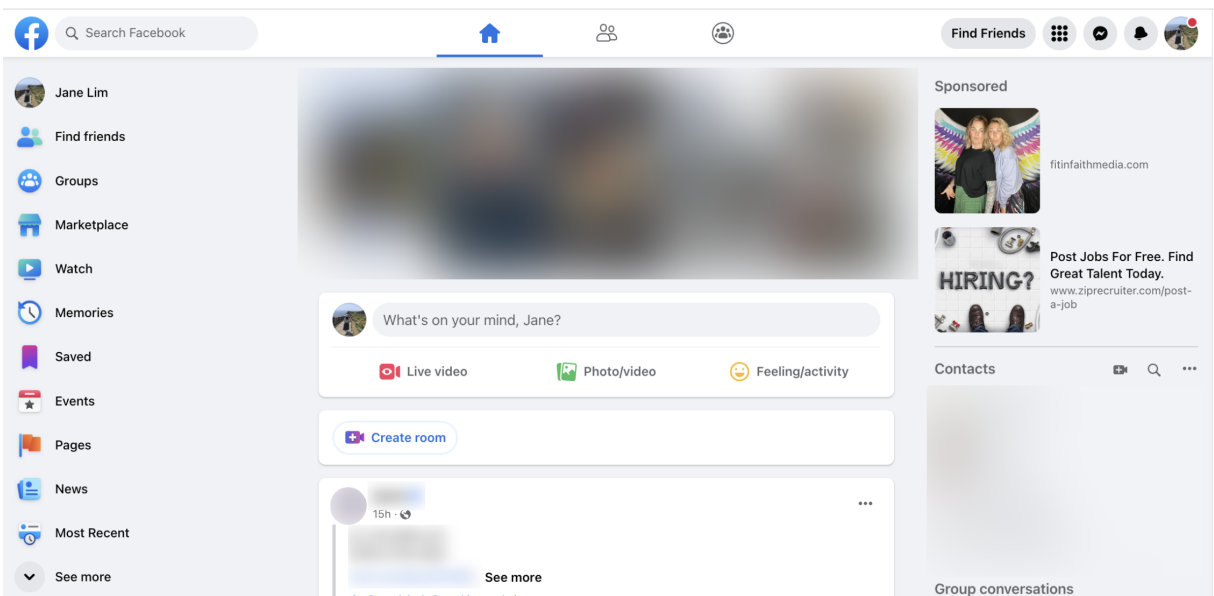


Figure 22: Participants were able to collapse the *feed dashboard* (via the button at the top right). Uncollapsed versions are shown in Figures 5 and 6.



(a)



(b)

Figure 23: Two layouts of Facebook’s desktop version that the authors found Meta was A/B testing during the experiment.

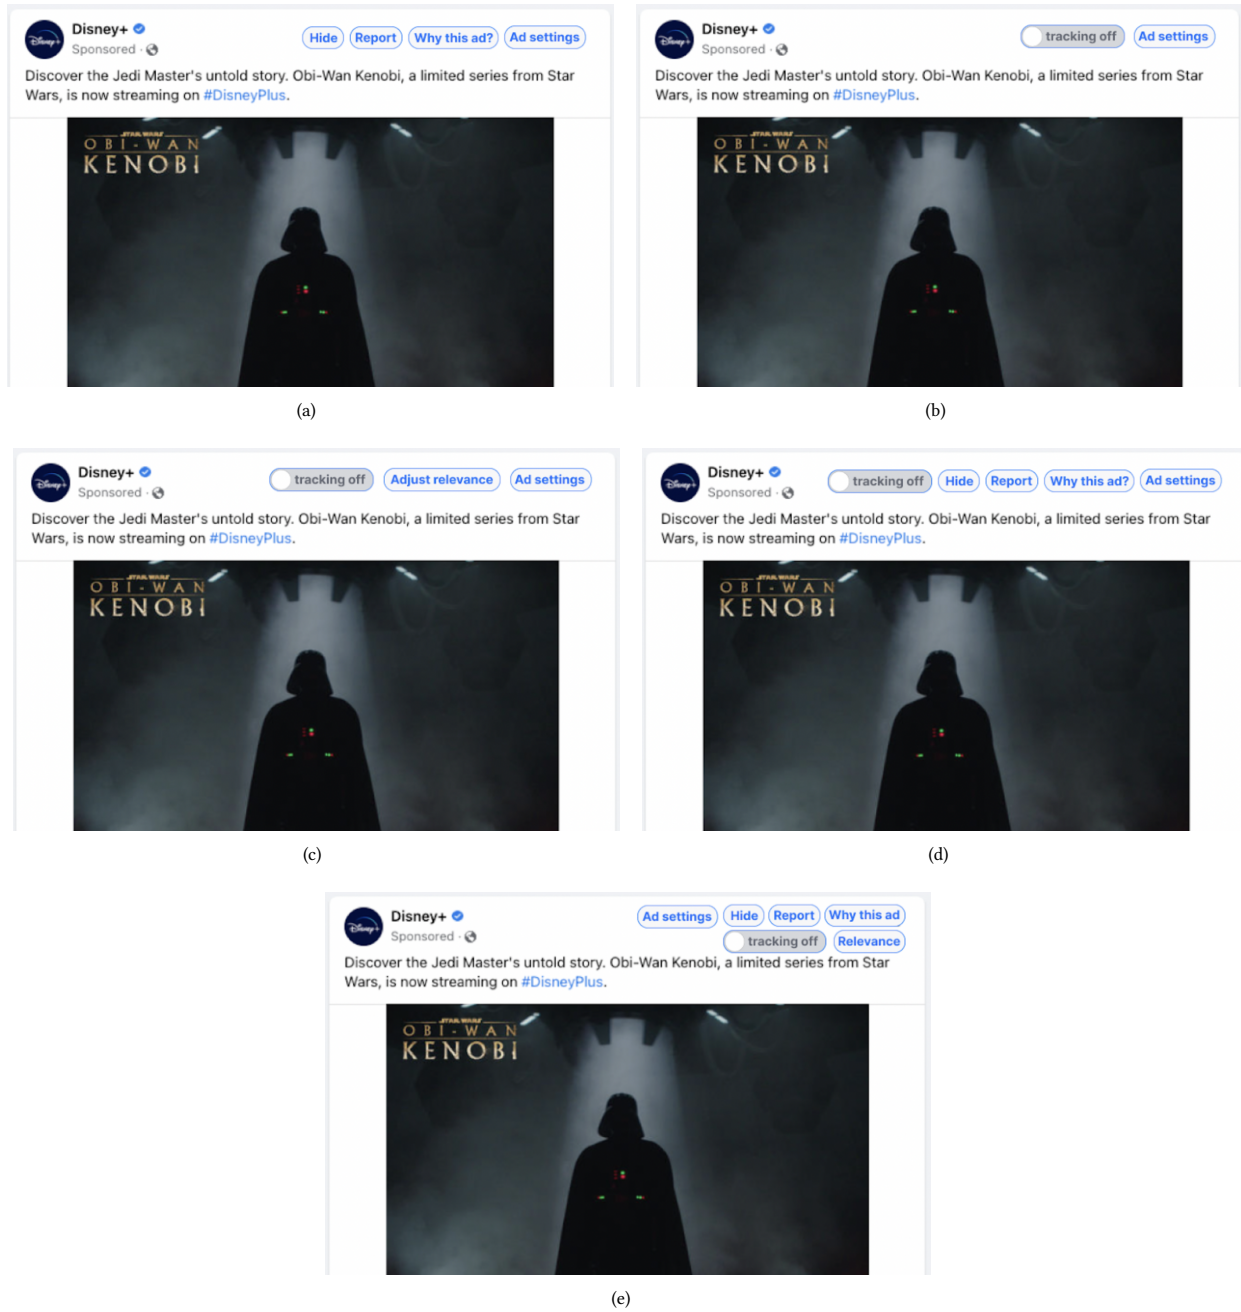


Figure 24: Variations of *ad menu*. We experimented with numbers and types of buttons.

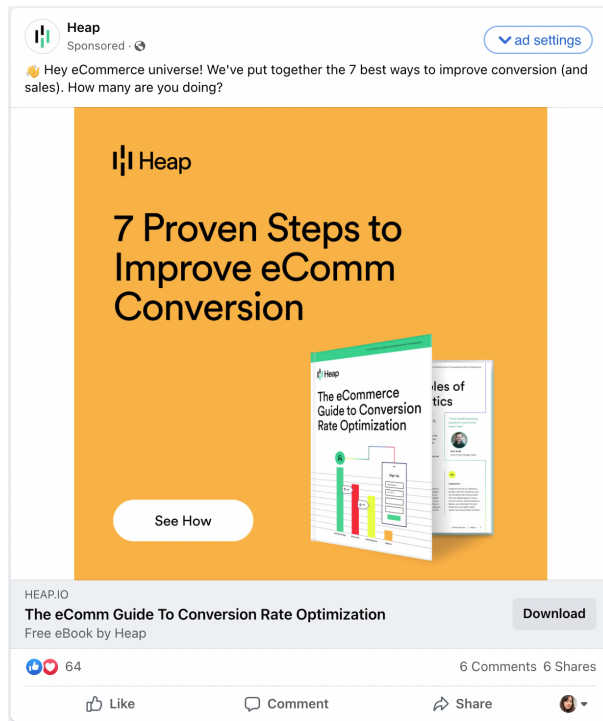


Figure 25: Screenshot of an early version of the ad button.

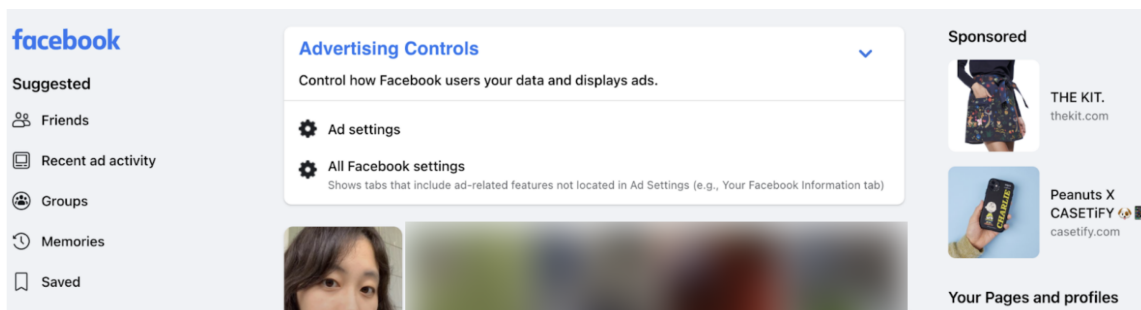


Figure 26: Screenshot of an early version of the feed dashboard (low actionability).

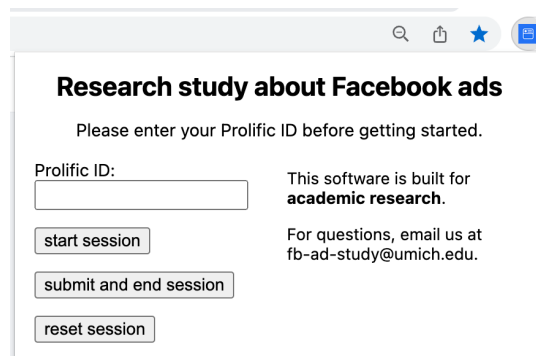


Figure 27: Screenshot of Chrome extension's popup.