

Characterizing the Usability and Usefulness of U.S. Ad Transparency Systems

Kevin Bryson, Arthur Borem, Phoebe Moh[†], Omer Akgul[‡], Laura Edelson[§],
Tobias Lauinger[¶], Michelle L. Mazurek[†], Damon McCoy[¶], Blase Ur
University of Chicago, [†]*University of Maryland*, [‡]*Carnegie Mellon University*,
[§]*Northeastern University*, [¶]*New York University*

Abstract—Online targeted ads are those shown only to certain users based on interests, demographics, or behaviors. Because targeted ads raise many privacy concerns, many platforms provide ad transparency systems (ATSS) to inform users about this practice. To better understand what current ATSS are communicating to users—and how—we first taxonomized the design and content of 22 of the most popular English-language websites’ ATSS as presented to users in the United States. We found substantial differences across ATSS in both the prevalence of transparency-enhancing features (e.g., whether they show users what has been inferred about them) and the presentation of information (e.g., the terminology used, where settings are located). Across all platforms, however, we observed consistent ambiguity about what data is used to target ads and the actual impact of altering settings. To gauge how these different design choices impact users, we conducted an online user study in which 198 participants used their own account to explore the ATS of one of eight representative platforms. We found that many of the questions participants hoped the ATS would answer remained unanswered after exploring the ATS. More broadly, participants found current ATSS simultaneously complex and lacking key details. We pinpoint ATS design decisions that best support users.

1. Introduction

Many online advertisements are *targeted ads*, which are shown only to specific users based on advertiser-provided criteria and algorithms [1], [2]. Targeting can leverage a wealth of data, such as users’ inferred interests, demographics, and activities [3], [4]. Academics and governments have highlighted dangers of hyper-targeted advertising, including potential impacts on credit access, housing, employment, and politics [5]–[9]. For instance, the U.S. DOJ found that targeted housing ads on Facebook violated the Fair Housing Act by using discriminatory characteristics [5].

Spurred by potential privacy and fairness harms [10], [11], as well as consumer demand, legislation like the GDPR, DSA, and CCPA have sought to establish user rights to transparency about ad targeting and the usage of user data [12]–[14]. Such legislation, however, has not specified exactly how the user interfaces should be designed or exactly how the mechanisms ought to behave. Without such oversight, platforms have designed and defined ad transparency

themselves. Many popular platforms now provide *ad transparency systems* (ATSS), which we define as a collection of controls and affordances on a platform that: (i) give users information about how targeted advertising works on that platform; (ii) let users manage settings about what data is used to make inferences about them; and (iii) let users modify their experience with targeted ads.

To date, researchers have studied a few major platforms’ ATSS in isolation, including Google [15], [16], Facebook [17], and Twitter [3]. However, researchers have not taken a holistic view of the large range of current ATSS, so the community lacks a systematic understanding of the affordances popular websites’ ATSS currently implement.

This paper aims to fill this gap by analyzing the user-facing design of current ATSS and how these design decisions impact users. We first investigated **what features and affordances different ATSS currently offer**. Specifically, we developed a taxonomy of the features available across U.S. ATSS on 22 popular platforms. We found significant differences across platforms in the prevalence of transparency-enhancing features like ad explanations and the disclosure of inferred ad topics. We also found significant differences in user experience and terminology used to describe platforms’ practices. Across all platforms, however, we observed consistent ambiguity around what data is used for ad targeting, as well as the impact of changing settings.

Having shown that platforms’ ATSS differ in design, we then wanted to know how these design choices impact usability and usefulness for users. Specifically, we wanted to learn **what participants want to do with ATSS** and **what users are able to do with current ATSS**. To this end, we conducted an online user study with 198 participants across eight representative platforms. We found that many of the questions participants hoped the ATS would help them answer—especially those about data collection and storage, third party data selling, and disabling specific categories of ads—remained unanswered after exploring the ATS. Furthermore, participants had difficulty navigating ATS interfaces that span multiple pages and understanding what some key ATS controls actually do.

Taken together, these findings indicate that users find current ATSS to be insufficient, as they are simultaneously complex while lacking detail in key areas. We recommend that platforms: (i) fully describe their data provenance and how inferences are made, (ii) provide more options for users

to manually specify their interests, and (iii) standardize the language used in ATSS’ controls and descriptions.

2. Background and Related Work

Ad-delivery. Targeted ads are served to users from a combination of advertiser-specified characteristics (e.g., demographics, audience lists, website visits) and by platform-specific ad-delivery algorithms. These delivery systems are opaque; where user data comes from, how data is used, and how ads are targeted are obfuscated [1], [18]. Studies of ad-delivery algorithms have brought to light problematic behaviors these algorithms can propagate, such as discrimination in housing and employment opportunity through Facebook’s ad-delivery system [19]–[21]. Indeed, ad-delivery systems have been used to target groups of users with political messaging ahead of elections, the Cambridge Analytica scandal being the most notorious example [8], [9].

Ad transparency. Many prominent online platforms, such as Facebook, Google, and Twitter, began to implement ATSS in response to political and legal pressure [10], [11]. Crucially, these platforms developed their ad transparency measures independently, without shared design guidelines. As such, different measures and affordances within ATSS have likely emerged and diverged from one another. The ATSS of large platforms, like Google and Facebook, have been the subject of extensive study [15], [17]. However, the ATSS on myriad smaller platforms are largely understudied.

Transparency and privacy affordances. One affordance in the ad transparency ecosystem is *ad explanations*, encompassing why an ad was served, how user data was collected, and what information is ultimately presented to users [22]. Empirical evaluations of Facebook’s ad explanations has shown that inferences shown to users are often incomplete, vague, or misleading [17]. While developing more informative and usable ad explanations is possible, platforms are likely disincentivized from implementing them because they can decrease users’ trust in the advertiser [3].

Ad controls, the ability for users to modify how often they see an ad, are another facet of affordances. Castleman et al. demonstrate that Facebook’s “See less” ad control and ad explanations are ineffective for certain ad-targeting strategies due to an algorithmically-mediated delivery system [2].

In another type of transparency and privacy affordance, ATSS reveal data to a user, allowing them to see what data a platform has collected and how a platform uses it. Prior work has examined how users interact with and perceive these mechanisms. Users’ concerns with data tracking decreased and their perceived benefit increased after exposure to Google’s privacy dashboard [15]. However, Google users also reported more concern when they connected their activities to specific inferences compared to when they were only shown generic information [16]. When examining Twitter’s targeted advertising mechanisms, users found more generic interests less concerning, but found more accurate and specific information more concerning [3]. Users also expressed

privacy concerns when transparency tools revealed practices that violate norms about information flows, such as by using inferred information [23]. Finally, large gaps exist between data used by companies and what users see in transparency dashboards, which often lack specificity [24].

Usability of privacy and ad transparency systems. Several studies have highlighted usability challenges in ATSS. In particular, discoverability represents a significant barrier for users. Users often struggle with controls for opting out of targeted advertising because they have difficulty finding and understanding these controls, or may be generally unaware of them [25], [26]. However, Im et al. found that by making controls and information more actionable and available in different contextual modalities, users were more likely to successfully find the controls they were searching for [27]. While some ATSS have hard to find or understand controls, others lack them entirely. Users find low actionability interfaces, which do not allow participants to modify their ad experience in situ, less useful [2], [26]–[29]. Wu et al. emphasize the harm done by the ubiquitousness of targeted advertising, finding that users feel a “constriction of behavior” that makes attempts at managing ad transparency and privacy settings futile [30]. Marginalized users have different needs and desires from ATSS, namely around refusal and greater control over the inferences applied to them [31].

This body of work largely has examined affordances offered by large platforms such as Facebook and Google. It remains to be seen how common these affordances are outside of these major players and how smaller platforms implement them, if at all. In addition, these affordances are often studied in isolation from the rest of the ATSS, and we analyze the effect of these affordances on usability and users’ ability to carry out transparency goals.

3. Taxonomy Methodology

We first describe how we taxonomized current ATSS.

3.1. Platform Selection

Because we aimed to characterize large platforms’ ATSS, we began with the 150 most popular domains on the Tranco list [32] as of November 9th, 2022. We attempted to locate each platform’s ATSS in three ways: accessing the page manually while not logged in, searching for the name of the platform alongside the term “transparency” in a search engine, and (when permitted) creating an account and logging in before searching settings pages. Our inclusion criteria required that ATSS: (1) be accessible in English from a desktop browser; (2) let users see information or data used to personalize ads; and (3) let users manage related settings.

Using these criteria, we found 22 platforms with ATSS among the Tranco Top 150. We excluded the remaining 128 domains because their sites were inaccessible (33 domains), provided no controls or information about ads (24), duplicate another site in the list (24), only provide general information about ads (21), were not in English (12), offered no way to

make an account (12), or were an adult website (2). Note that many of the platforms that were inaccessible seemed to be CDNs, while the platforms that provided only general information about ads typically displayed information related to platform governance rather than ad transparency specifically (e.g., zoom.us) or had no transparency-relevant content at all. For the 22 ATSS we studied, we initially examined the platform through a desktop browser (Google Chrome), a mobile browser (Google Chrome on Android), and (when offered) through that platform’s Android app. Because we found that the desktop versions of these ATSS often had either more features or the same number of features as their mobile counterparts, we chose to analyze all ATSS through their desktop interface. For each ATS, between November 2022 and October 2023 we repeatedly collected screenshots from the desktop interfaces of each part of the ATS.

3.2. Qualitative Analysis Approach

To characterize the features of the 22 ATSS in a rigorous way, we performed qualitative coding of the screenshots we collected. Members of the research team initially performed open coding on a selection of popular platforms’ ATSS and then used the resultant codebook to train additional coders. We centered our analysis on what we termed ATS “units.” As exemplified by Figure 10 in the appendix, we defined an ATS unit as a self-contained aspect of the ATS with a single purpose: for either providing information to the user or enabling the user to change some setting. Two members of the team first independently defined each unit, collaboratively resolving all disagreements on the scope of each ATS unit. Two members of the research team then independently coded each unit, again resolving all disagreements. Regarding our team’s positionality, we are a group of privacy and HCI researchers based in the US. All members of the research team who performed coding regularly use multiple of the platforms studied. Many of us set our personal accounts to limit ad personalization.

3.3. Limitations

Our observations are limited to elements in the immediate context of the ATS. This includes all controls and surrounding text, but not links outside the system to FAQs or other pages. Because screenshots were our primary unit of observation for coding consistency, we could not systematically describe the breadth of data shown to different users, nor dynamic aspects of the ATS interface. While we considered using short videos showcasing interaction with ATSS, screenshots simplified the coding process and let us focus on the text and controls provided. Further, we could not reliably describe default settings on ATSS—the settings seemed to vary based on when the account was created.

When we started collecting screenshots in 2022, many platforms had substantial variance between their desktop and mobile interfaces, and sometimes even between iOS and Android devices. Since then, we have observed changes primarily to mobile interfaces that most frequently unify

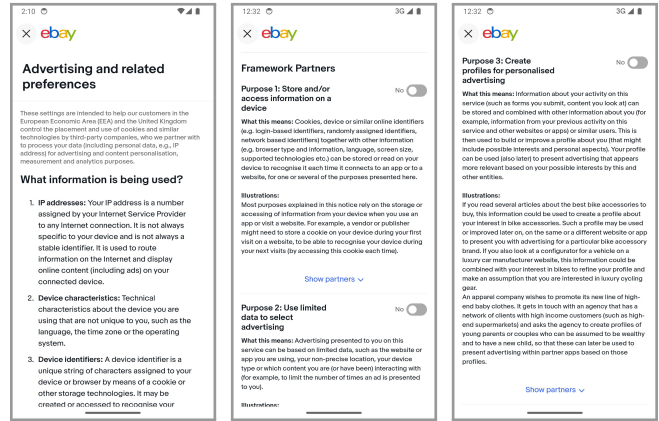


Figure 1: A partial snapshot of eBay’s ATS.

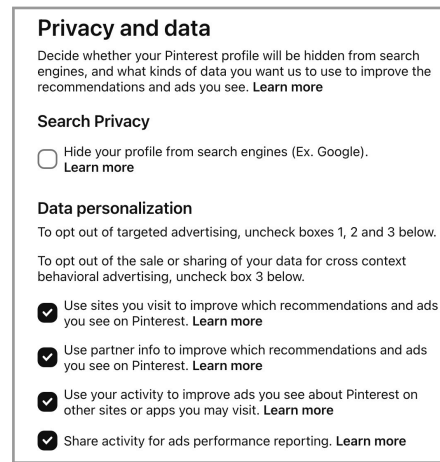


Figure 2: Pinterest’s full ATS.

experiences for desktop and mobile users, both visually and functionally. However, Twitter and TikTok prominently buck this trend. Twitter’s mobile app ATS was similar to its desktop counterpart until late 2023, when several features were obfuscated in the mobile version, including the ability to control inferred interests. Currently, these features require a user to click on a different button than the desktop interface. Doing so opens a web browser prompting users to sign in to Twitter before displaying the full set of settings available on the web interface. On TikTok, the web interface has no ATS, possibly because the web interface does not show ads.

4. Taxonomy Results

In this section, we describe the types of affordances we identified and their prevalence across the 22 platforms (Table 1). This analysis leads us to identify positive and negative trends for users across ATSS to address our initial research question, **what features are available in ATSS?**

TABLE 1: Taxonomy of the ATS features and affordances for the 22 platforms we studied.

Total	Tranco	19	6	3	5	7	7	9	10	9	3	3	3	8	15	6	5	6	13	4	4	10	10	13
		Platform personalization	Ad explanation	Ad related activity logs	Advertisers associated	Inferred topics	Advertiser information	Data shared by user	Data inferred by platform	Data sharing ambiguous	Sorting	Searching	Filtering	Report ad	Personalization toggle	Topic frequency	Advertiser frequency	Ad frequency	Web activity	Location	Specific data types	Audience lists	Data outside platform	Data within platform
Google	1	●	●	●	●	●	●	●	●	●	●	●	●	●	↑↓	↑↓	↑↓	↑↓	↑↓	↑↓	↑↓	↑↓	↑↓	↑↓
Facebook	4	●	●	●	●	●	●	●	●	●	●	●	●	●	↑↓	↑↓	↑↓	↑↓	↑↓	↑↓	↑↓	↑↓	↑↓	↑↓
Microsoft	5	●	●	●	●	●	●	●	●	●	●	●	●	●	↑↓	↑↓	↑↓	↑↓	↑↓	↑↓	↑↓	↑↓	↑↓	↑↓
X	8	●	●	●	●	●	●	●	●	●	●	●	●	●	↑↓	↑↓	↑↓	↑↓	↑↓	↑↓	↑↓	↑↓	↑↓	↑↓
Instagram	9	●	●	●	●	●	●	●	●	●	●	●	●	●	↑↓	↑↓	↑↓	↑↓	↑↓	↑↓	↑↓	↑↓	↑↓	↑↓
Apple	12	●	●	●	●	●	●	●	●	●	●	●	●	●	↑↓	↑↓	↑↓	↑↓	↑↓	↑↓	↑↓	↑↓	↑↓	↑↓
LinkedIn	14	●	●	●	●	●	●	●	●	●	●	●	●	●	↑↓	↑↓	↑↓	↑↓	↑↓	↑↓	↑↓	↑↓	↑↓	↑↓
Yahoo	19	●	●	●	●	●	●	●	●	●	●	●	●	●	↑↓	↑↓	↑↓	↑↓	↑↓	↑↓	↑↓	↑↓	↑↓	↑↓
Amazon	22	●	●	●	●	●	●	●	●	●	●	●	●	●	↑↓	↑↓	↑↓	↑↓	↑↓	↑↓	↑↓	↑↓	↑↓	↑↓
Warner Bros	24	●	●	●	●	●	●	●	●	●	●	●	●	●	↑↓	↑↓	↑↓	↑↓	↑↓	↑↓	↑↓	↑↓	↑↓	↑↓
Pinterest	34	●	●	●	●	●	●	●	●	●	●	●	●	●	↑↓	↑↓	↑↓	↑↓	↑↓	↑↓	↑↓	↑↓	↑↓	↑↓
Reddit	44	●	●	●	●	●	●	●	●	●	●	●	●	●	↑↓	↑↓	↑↓	↑↓	↑↓	↑↓	↑↓	↑↓	↑↓	↑↓
TikTok	74	●	●	●	●	●	●	●	●	●	●	●	●	●	↑↓	↑↓	↑↓	↑↓	↑↓	↑↓	↑↓	↑↓	↑↓	↑↓
Tumblr	85	●	●	●	●	●	●	●	●	●	●	●	●	●	↑↓	↑↓	↑↓	↑↓	↑↓	↑↓	↑↓	↑↓	↑↓	↑↓
Spotify	89	●	●	●	●	●	●	●	●	●	●	●	●	●	↑↓	↑↓	↑↓	↑↓	↑↓	↑↓	↑↓	↑↓	↑↓	↑↓
eBay	113	●	●	●	●	●	●	●	●	●	●	●	●	●	↑↓	↑↓	↑↓	↑↓	↑↓	↑↓	↑↓	↑↓	↑↓	↑↓
Canva	118	●	●	●	●	●	●	●	●	●	●	●	●	●	↑↓	↑↓	↑↓	↑↓	↑↓	↑↓	↑↓	↑↓	↑↓	↑↓
SoundCloud	123	●	●	●	●	●	●	●	●	●	●	●	●	●	↑↓	↑↓	↑↓	↑↓	↑↓	↑↓	↑↓	↑↓	↑↓	↑↓
Fandom	128	●	●	●	●	●	●	●	●	●	●	●	●	●	↑↓	↑↓	↑↓	↑↓	↑↓	↑↓	↑↓	↑↓	↑↓	↑↓
The Guardian	136	●	●	●	●	●	●	●	●	●	●	●	●	●	↑↓	↑↓	↑↓	↑↓	↑↓	↑↓	↑↓	↑↓	↑↓	↑↓
AliExpress	140	●	●	●	●	●	●	●	●	●	●	●	●	●	↑↓	↑↓	↑↓	↑↓	↑↓	↑↓	↑↓	↑↓	↑↓	↑↓
Stack Overflow	142	●	●	●	●	●	●	●	●	●	●	●	●	●	↑↓	↑↓	↑↓	↑↓	↑↓	↑↓	↑↓	↑↓	↑↓	↑↓

●, ↑↓, and ↑↓ indicate presence of affordance while ↑↓ indicates the affordance is a control, ↑↓ indicates a divergence from expected behavior or design.

4.1. User and platform information

We define **user information** as the *descriptions* of data provided by or associated with users (e.g., gender). eBay provides a lengthy description of user data in the leftmost images in Figure 1. Sometimes a description of data provenance is included, indicating where the data comes from (e.g., *shared by the user* or *inferred by the platform*).

However, across nine platforms (*Data sharing ambiguous*, Table 1), we found that those descriptions of data provenance are **ambiguous**. This occurs most frequently in descriptions of platform personalization, ad explanations and data controls (Section 4.3). Rather than specifying types of data that *are* used, platforms commonly mention several types of data that *may* be used:

“Tumblr may use information you provide, such as your age, to tailor ads. . . Tumblr may also use this information, including personally identifiable information, to tailor ads when you are using our services. . . We also select ads to show you based on information that we collect from other sources, such as your age or location, searches you conduct, apps on the device, or based on your other interests.”

In this excerpt, age is both shared by the user and collected from other sources. Personally identifiable information has unmentioned origins and location, searches you conduct, apps on the device, and other interests are collected from an unspecified other source. Ambiguous descriptions are a common deficiency in transparency. Excluding The Guardian and AliExpress, every ATS that describes data provenance includes ambiguous descriptions of data origins (Table 1).

We also analyzed platform information, which includes explanations of a **platform’s personalization** practices like,

“Information about your activity on this service. . . can be stored and combined with other information about you” (eBay) (Figure 1). While most ATSs (19 of 22) convey platform personalization (Table 1), they do so using different strategies to describe their practices, often with different verbiage. eBay’s interface (Figure 1), for example, provides paragraph-long explanations for each of the seven controls offered, while Pinterest (Figure 2) opts for concise descriptions of each control coupled with links to an external informational page. Prior work has studied the challenges users face understanding lengthy technical jargon in privacy policies [33], [34]. We hypothesize that users will face similar difficulties with the verbose explanations on eBay and other platforms, though it is unclear if concise interfaces like Pinterest’s can convey enough information to users (Section 6.3.6). While nearly all platforms describe which user data factors into personalization, a large amount of that information is **ambiguous** and/or lengthy, echoing deficiencies identified with privacy policies [33], [34].

4.2. Inferences and Exploration tools

Ad explanations (Figure 5), uniquely, do not vary in design or location across platforms and are rare among ATSs (6 of 22). Platforms list one to three reasons, ranging from demographics, a user’s activity on the platform, to a user being a part of a “similar audience” (a thorough characterization of the content of ad explanations can be found in Wilkinson et al. [22]). Additionally, they can generally all be found by clicking a three-dot menu next to a post or link. In Section 6.3.1 we discuss how this design affects users.

Even though eBay’s ATS, by far, contains the longest descriptions of all the platforms we observe, eBay and 14 other platforms do not show any concrete data about

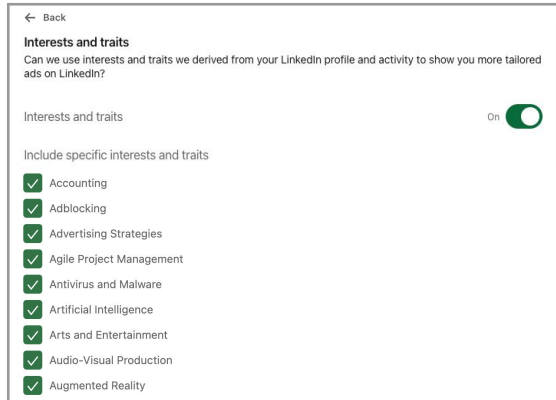


Figure 3: LinkedIn’s inferred interests page showing topics alphabetically.

topics that the platform infers that a user is interested in (*Inferred topics associated with user*; Table 1). While many platforms may suggest “*certain characteristics (e.g., your possible interests, your purchase intentions, your consumer profile)*” (eBay) are used, on these 15 platforms, users have no way to see or change their inferred interests. Their only option is to opt in or out entirely via all-encompassing controls like those in Figure 2.

Exploration tools are affordances like search bars and filters that allow users to navigate through data shown to them. Figure 3 shows an example of the inferred topic page from LinkedIn. When platforms show users inferred topics (7), they often contain a large amount of information, as suggested by Figure 3, which shows nine pages of inferred interests for a user. We observe that five of the seven platforms that show inferred topics or advertisers associated with the user have at least one feature available to help explore that data. For the other platforms, users must manually scroll through those pages of data to identify items of interest (Table 1), which can be overwhelming. Specifically, we observe **filtering**, **sorting**, and **searching** as the full range of exploration tools (Table 1). These features *should* help greatly reduce the number of options a user initially has *or* immediately surface relevant data; however, we question their utility in practice (Section 6.3.3).

Google, for example, has sorting options that include most popular, most recent, most seen, and alphabetically; the filtering options include trending, recent, and seeing more/less. From the perspective of a first-time user, two of these options (trending, most popular) are undefined and seeing more/less would be empty. While the other options could be useful, in Section 6.3.3 we analyze participants’ experiences with inferred topics pages and exploration tools.

Searching, however, may be a useful exploration tool for users, especially as a complement to useful sorting schemes. On Facebook, Instagram, and Stack Overflow, users can keyword search for interests or advertisers and control the frequency in which they see ads about those topics or from those advertisers. Notably, on Facebook and Instagram, the search is accompanied by invariant suggestions of “popular




Figure 4: The only control in Amazon’s ATS: its interests-based ads control.

search topics”: alcohol, parenting, pets, politics, gambling, and body weight control. Features like this could be useful in spurring ideas of what to explore, but could raise engineering challenges in matching search queries to the advertising topics (e.g., a user may search “fountain pen” and see no matches, but a search for “writing” shows the desired topic).

Under half of platforms (9) expose **inferred topics**, **associated advertisers**, or **ad explanations** to users. The absence of these may suggest to users that such data does not exist when that is likely not the case. While **exploration tools** appear among most platforms with inferred topics (4), we hypothesize that some of them may not be useful. Without useful exploration tools, user data is presented indiscriminately, creating barriers to meaningful analysis.

4.3. Personalization and data controls

Nearly every platform has a **personalization** or **data control**. Personalization controls affect the use of personalization in the system and allow users to control the frequency with which they see ads about certain topics, or certain advertisers as well as specific ads. The most common control was a **personalization toggle** (15, Table 1):

“No, don’t make my ads more relevant by using this information” (Facebook, Instagram).

“If you do not wish to receive more accurate ads, Please turn off this function.” (AliExpress).

“(1) Personalized ads (On/Off). (2) Use Web & App Activity to personalize ads” (Google).

Nine platforms have a personalization toggle that disables the whole interface, seemingly disabling all forms of ad personalization (Figure 4). This type of design seems to align with GDPR’s stipulation that “It shall be as easy to withdraw consent as to give consent” [13].

However, on eight platforms, the behavior of the toggle **diverges** from this expectation, potentially giving users the impression that all personalization has been disabled, while personalization remains in a complex, ambiguous state (Table 1). Some of the diverging designs are because we identify **multiple** personalization toggles (e.g., Google); otherwise, it is because the personalization toggle appears to have **no effect** on the other controls within the system (e.g., Google, Twitter, Facebook).

Less commonly, users are offered a relatively fine-grained level of control over their advertising experience. These are controls like hiding ads (6), removing ads about

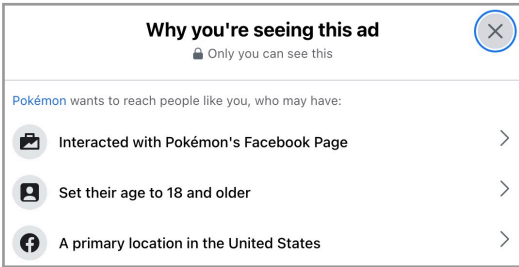


Figure 5: Facebook’s ad explanation interface.

certain topics (2), or blocking advertisers (4). Thus, for the majority of ATSS, advertising is largely unable to be customized to any great extent by the user explicitly, despite the prevalence of user desire for this type of customization (as found in our study (Section 6.2.2) and previous work [26]).

Data controls, a subset of **personalization controls**, indicate which **types** or **sources of data** will be used for ad personalization. Most controls in ATSS are data controls. Pinterest, for example, has controls for web activity outside the platform, audience lists, and web activity from within the platform (Figure 2). Notably, across platforms similar data controls often use different terminology for third parties:

“Can we use your information that you’ve given to businesses to show you more relevant ads?” (LinkedIn).

“Use sites you visit to improve which recommendations and ads you see on Pinterest” (Pinterest).

We note that many of these controls also neglect to explicitly define who “business partners” are or what “information” means. This not only made it difficult for us to reliably describe these features, but may also contribute to the insufficient descriptions of how data about a user is tracked and used across the internet observed by Farke et al. [16].

Further, platforms rarely enable control over the usage of familiar and unambiguous data types like location, gender, or age, suggesting a lack of transparency in control for users. While these controls are often the most straightforward to understand, only five platforms have this type of control (Table 1). While location and demographic information are almost certainly factors used in targeting on some platforms like Facebook and Instagram, as suggested by Figure 5, users have **no ability to control whether they are targeted by these fields on these platforms**. This raises the question for Facebook, Instagram, and the other 15 platforms without these controls: why are users unable to control whether these factors are used? While it is unrealistic to enumerate all possible types of data that are used or not used in these systems, the absence of these features is surprising, especially for large platforms.

Personalization and **data controls** present numerous potential challenges for users. In particular, platforms with *divergent* personalization toggles present potential legal and usability concerns. Disparate designs in these personalization and data controls may require additional effort for users to learn and understand for each platform they use.

5. User Study Overview and Methods

While our taxonomization of 22 current ATSS revealed differences across platforms, it inherently could not reveal how those differences impact users. Thus, we conducted an online user study to gauge the usability and subjective perceptions of these different design choices and affordances. To increase our study’s ecological validity, we had participants explore the ATS using their own account on a platform they regularly use. Here, we describe our methods.

5.1. Platform Selection

To ensure that a sufficient number of participants could provide feedback on each interface, we chose to test eight representative ATSS, rather than the full set of 22. To capture the breadth of features and affordances observed in the full set, members of the research team clustered the 22 platforms based on the codes in our taxonomy, categorizing platforms’ ATSS by the number and types of affordances. As the ATS interfaces vary from simple, single-page systems to larger, multi-page hierarchical systems, this clustering process also considered the design of the user experience. Because we expected most users to take the study with a mobile phone, we also prioritized platforms whose mobile ATSS were most similar to their desktop ATSS. The process resulted in three clusters, and we chose representative ATSS for each, capturing correlations with Tranco rankings and the amount of public discussion about advertising on each. Our eight representative ATSS were from eBay (*EB*), Facebook (*FB*), Google (*G*), Instagram (*IG*), LinkedIn (*LI*), Reddit (*R*), TikTok (*TTK*), and X (*X*), formerly known as Twitter. We used these italicized abbreviations to indicate a participant’s assigned ATS when presenting quotes. To date, the privacy literature has primarily studied the Google, Facebook, and Twitter ATSS. Thus, we analyze three previously studied platforms alongside five understudied platforms.

5.2. Recruitment and Screening Process

We recruited participants using the Prolific crowdsourcing service, using Prolific’s filters to restrict the study to U.S.-based participants. Because we needed to recruit only active users of at least one of the eight platforms we were studying, we first conducted a screening survey. This survey asked participants whether they had active accounts on a range of different platforms, including both distractor and nonexistent platforms so that participants would be unlikely to identify our inclusion criteria. We also asked respondents about their mobile device(s) and use of ad blockers. Our screening survey took two minutes to complete and we compensated respondents \$0.50 USD. In total, 553 people completed the screening survey.

Based on participants’ usage of at least one of the platforms we were studying, we invited them to complete the main study, assigning them round-robin to one of the eight ATSS. We recruited a total of 200 participants, 25 per ATS. We compensated participants \$6.00 USD for completing the

main study, which took roughly 29 minutes on average. We collected data between May and September 2024.

5.3. Ethics

We received approval for our study from the lead institution’s IRB. As part of this approved protocol, we followed an online consent process in which we showed participants a consent form at the beginning of the screening survey. Participants had to click a button to indicate that they “read the consent form, are at least 18 years old, and agree to participate in the research.” This consent form also described the study procedure and possible risks. Following the screening process, we invited qualified participants to the full survey. Since some time had passed since they originally saw the consent form, we again showed participants the consent form at the start of the full survey.

Only authors from the lead institution interacted with participants or had access to any raw survey data, which was de-identified by the first author (e.g., by removing Prolific IDs) and stored on machines with disk encryption enabled. In the study, participants did not share with us any personal information the platform has about them, such as any usernames or inferred interests. This user data is not visible to us, and we chose not to ask participants to share anything specific about their personal data other than the state of advertising controls (i.e., on or off).

5.4. Main Survey Structure

In this part of our study, participants completed a survey that directed them to explore specific parts of their assigned ATS using their mobile phone while logged into that platform. Before having participants explore their assigned ATS, however, we wanted to gauge their expectations and desires for the ATS. Thus, we asked them to list one to three questions they had “about how ads are targeted to [them]” and an additional one to three questions “about the data used for the ads targeted to [them]” that they would want the ATS to answer. We also asked them to list one to three “advertising preferences that [they] might want to verify or change.” At the end of the survey, participants reflected on the degree to which those questions were answered.

To analyze the usability and usefulness of the features of each ATS, we then asked the participant to complete the four tasks detailed below. We designed these tasks to cover a breadth of affordances offered by ATSs while also being general enough to be applicable across platforms. For each task, we asked the participant to navigate to the specified page of their ATS and respond to a series of questions corresponding to what they were seeing. Because we worried that some participants would not be able to find certain parts of the ATS, we configured the survey to enable participants to note that they could not find the requested page, at which point the survey would provide step-by-step instructions. In the case of ATS pages that included controls that could be changed, the survey asked the participant to report the state of each key setting both when they arrived

at that page and when they left that page. For each, the participant described why they changed or kept the setting, explained what that setting does, and rated their satisfaction.

Task 1: Navigate to the ATS landing page. To gauge whether participants know where to look for ad transparency information, we first described the concept of an ATS and asked participants to use their mobile device to navigate to their assigned platform’s ATS. We provided a partial screenshot of what they should expect to see upon finding the ATS so that they could confirm they had succeeded. Finally, we asked participants to describe their experience locating the ATS.

Task 2: Explore third-party data usage. Next, we asked participants to find where they could examine “third-party data usage.” We asked them to report the initial state of the applicable controls (i.e., on or off) and gave them the opportunity to change any controls if they wanted. We asked participants to explain any decisions to change settings and to explain in text what they thought the control would do, rating their confidence in their answers on a Likert scale. Finally, we asked participants to describe any difficulties they had in finding or understanding the page.

Task 3: Review inferred topics. Next, we directed participants to find where they could learn what interest topics the platform has assigned to them and to respond to similar questions as in Task 2 around adjusting and understanding controls. We also asked participants to describe in text how they think the platform assigns interest topics to them.

Task 4: View ad explanations. Finally, we asked participants to find an ad on the platform and try to find information about why they were shown that ad. This aspect of the study investigates ad explanations, which were a focus of prior work [3], [17], [22], [27]. We asked participants to rate the perceived completeness and usefulness of the explanation, justifying their answers in text.

5.5. Data Analysis

We use both quantitative and qualitative methods to analyze our data. For Likert-scale data, we first compare across all groups with a Kruskal-Wallis test. When this omnibus test is significant, we conduct post-hoc pairwise analysis using the Wilcoxon rank-sum test. We correct p-values for multiple testing using the Holm method. To analyze which factors influenced participants’ perceived difficulty in navigating the ATS (on a scale from extremely easy to extremely difficult), we fit ordinal logistic regression models.

To analyze qualitative data, we engaged in several rounds of exploratory open coding to build and refine a codebook. Two authors coded responses individually before meeting to resolve all disagreements. Because some parts of the qualitative data have dependencies across the study (e.g., the degree to which a participant’s pre-task question was answered of course depends on the question), we grouped a given participant’s responses to different questions. This holistic approach provides more fine-grained insights into the relevant mechanisms, data types, and concerns. We also analyzed participants’ explanations of why they did or did

not change ATS settings and their explanations of various ATS sections' effectiveness, usability, and usefulness. To improve readability, we fix minor mistakes in grammar, punctuation, and capitalization in these quotes. For clarity in reporting, we respectively bin all positive and all negative Likert-scale responses (e.g., "strongly agree" and "agree"). When the participant marked "neither agree nor disagree" and the coders did not feel the participant's free-text response clearly implied either agreement or disagreement, we marked those responses as inconclusive.

5.6. Limitations

Our methods have limitations common to most online user studies, including the potential for satisficing and other response biases. To control for potential differences across countries in ATSs due to differing privacy laws (e.g., GDPR or DSA), we only studied American participants using the American versions of ATSs, limiting generalizability. Due to the Tranco list's focus on domains, our method of selecting ATSs does not necessarily represent the popularity of mobile apps. That said, we conducted an informal assessment of the top free mobile apps, finding that only Snapchat and Uber were not in the Tranco Top 150 yet offered an ATS as of April 20th, 2024. Further, our participants represent a convenience sample and are not demographically representative of any larger population. In addition, nearly half of participants reported using some type of ad blocker on at least one of their devices. We had considered excluding users of ad blockers from the study. However, because ad blockers on desktop devices are quite popular, doing so would have excluded the viewpoints of a sizable fraction of users. Further, because participants viewed their assigned ATS during the user study on their mobile device, where ad blockers are less common, we do not expect that the use of ad blockers meaningfully impacted participants' ability to view personalized content in their ATS.

The tasks we selected for our user study represent a widely adopted set of ATS affordances across all eight platforms we tested. However, keeping these high-level tasks consistent across platforms excludes unique and nuanced features only a few platforms offered, such as turning off all personalization, restricting ads about sensitive topics, or controlling the number of ads from certain brands. To avoid influencing participants' perceptions of any platform, we did not direct participants to any other specific affordances. Instead, we suggested that participants explore other advertising-related settings after completing all tasks.

6. User Study Results

6.1. Participants

Of our 200 participants (25 per ATS), two participants provided responses that appeared to be AI-generated (including disclaimer sentences, using perfect punctuation). We removed those two participants from further analysis. Our

remaining 198 participants were primarily ages 25–34 or 35–44 (67% combined), white (66%) and women (59%). More than half of all participants held a bachelor's degree or higher (50%). Finally, only 27% of participants reported having seen their assigned platform's ATS before the study. Table 6 in the appendix details participants' demographics.

6.2. What do Participants Want from ATSs?

Participants produced 1,047 questions and 503 action items related to their assigned platform's advertising system. These actions and questions (summarized in Table 2) were primarily related to data collection and use (379 questions), personalization (210 questions, 15 actions), and ad categories (130 actions). For over half of these questions and items, participants did not find answers or relevant controls.

6.2.1. Understanding and limiting data collection and use. A majority of participants (168) wanted to learn about the data collection practices of the platforms they used. Some participants asked about specific data types (e.g., web activity, cookies, personally identifiable information), but others asked about the platforms' data collection practices generally. When they were specific, participants were most curious about **whether platforms tracked their web activity** (*Collection/use: web activity*, Table 2). Participants were curious about activity both within the platform (e.g., "content I react to," P-21-FB; "professional profile, connections, or engagement with content," P-13-LI) and outside it (e.g., "websites I visit outside of Twitter," P-6-X; "data from other apps," P-3-FB). While 49% of these questions were answered overall, Reddit users only found answers to one third of them, the least among platforms for this question category. P-12-R attributes this to Reddit's "lack of transparency," since there seemed to be "no solid answer" (P-21-R) as to whether their web activity was collected. This fact is reflected in our taxonomy, as Reddit did not define their data collection practices (*Data shared by users, Data inferred by platform, Data sharing ambiguous*, Table 1).

Participants also asked 210 questions about ad targeting (*Targeting*, Table 2); most of these were straightforward questions about how targeting works and how user data is used to select ads, in keeping with participants' overall tendency to want to better understand data use. Questions were often broad and straight to the point, e.g., "How are ads targeted to me?" (P-8-IG), although some participants had theories on how the system worked, e.g., "Does AI or a person determine the ads I am shown?" (P-12-IG). 43% of these questions were answered; for those that were not, participants often attributed this to lack of "specifics of the algorithms" (P-17-TTK). While this rate is comparable to those for questions about data collection and use, it is low given that 19 out of 22 platforms implement the "Platform personalization" affordance (Table 1).

Participants also commonly asked questions about **who could access their personal data and to whom it was sold** (*Storage/security and Data sharing/selling*, Table 2). These questions often addressed retention ("How long is

TABLE 2: Question and action item themes listed by participants.

Theme	#	(n)	✓	✗	?	Sample Quote	Most Relevant Affordance
Questions	1,047	(197)	32%	53%	15%		
Collection/use:	379	(168)	36%	47%	17%	<i>Do you target me with ads that people in my area show interest in?</i> P-12-R	
↳ web activity	116	(74)	49%	33%	18%	<i>How much do they know about my shopping habits?</i> P-2-G	Web activity
↳ PII	58	(37)	36%	52%	12%	<i>Are ads frequently targeted to me based on my race?</i> P-21-LI	Ad explanation
↳ location	15	(13)	33%	40%	27%	<i>Does my location influence the ads I see?</i> P-14-R	Location
Targeting	210	(135)	43%	48%	9%	<i>Are [ads] targeted to me by the content I react to on Facebook?</i> P-21-FB	Ad explanation
Storage/security	136	(88)	17%	75%	8%	<i>Do you keep this information forever?</i> P-22-IG	Data within platform
Data sharing/selling	115	(84)	21%	70%	9%	<i>How much money do companies actually make off my data?</i> P-16-IG	Advertisers associated with user
Looking for Y	68	(46)	41%	44%	15%	<i>What are default settings for a given users ad selection?</i> P-19-X	Searching
Spying	25	(22)	4%	80%	16%	<i>Do they use my phone's microphone to listen?</i> P-34-TTK	Data sharing ambiguous
Inferences	18	(17)	44%	44%	12%	<i>What type of person does Google think I am?</i> P-20-G	Data inferred by platform
Actions	503	(186)	43%	34%	23%		
Delete/disable:	243	(129)	39%	37%	25%	<i>Completely opt out of all data sharing/storing between companies.</i> P-17-FB	
↳ ad category	85	(63)	43%	32%	25%	<i>Can I turn off certain types of ads that I find upsetting?</i> P-2-R	Topic frequency
↳ all ads	23	(21)	13%	61%	26%	<i>I would like to change ads to simply not appear ever again.</i> P-1-IG	Ad frequency
↳ data collection	20	(17)	25%	45%	30%	<i>I want to prevent data gathering all together.</i> P-4-EB	Data outside/within platform
Unspecified control:	128	(89)	52%	27%	21%	<i>Do I have any control over what they do with my data?</i> P-5-FB	
↳ ad category	45	(39)	55%	27%	18%	<i>No food ads.</i> P-39-TTK	Topic frequency
↳ targeting	15	(15)	67%	13%	20%	<i>Are there options for me to adjust or refine the ad targeting settings...?</i> P-23-X	Personalization toggle
Decrease	54	(47)	39%	48%	13%	<i>How do I reduce the number of ads that I see?</i> P-39-LI	Topic, Advertiser, Ad frequency
Increase	44	(32)	43%	23%	34%	<i>Can I make a request for ads about travel?</i> P-6-X	Topic, Advertiser, Ad frequency
Verify/learn	34	(31)	47%	29%	24%	<i>I would want to verify who is collecting my data.</i> P-43-TTK	Ad explanation

#: number of questions or actions fitting each theme, n: number of participants who asked a question or shared an action fitting each theme, ✓: percent of questions answered after task, ✗: percent of questions not answered after task, ?: percent of questions which were inconclusive, a sample quote, and the most relevant affordance from Table 1.

that data held for?” P-23-FB), as well as who, within and outside the company, “gets access to [...] data” (P-17-R). Some questions about *Data sharing/selling* were closed-ended (e.g., “Does eBay sell my personal info to third parties?” P-16-EB). Others were more nuanced, including, e.g., “What specifically is being sold?” (P-15-FB), “Who are the companies who get the data?” (P-9-R), “How much money do companies make off my data?” (P-16-IG).

In contrast to questions about collection and use, questions about storage/access and sharing/selling were less frequently answered (17% and 21%). It is surprising that even one fifth of these questions were answered, given there were no features nor affordances embedded in the advertising settings pages of these platforms that tackled topics such as selling data to third parties (Table 1). Participants reacted strongly to not finding these answers, saying they were “frustrated” (P-24-X) and “uneasy” (P-8-TTK):

“Nothing that I saw today explained how my data is protected within the app. If anything, I have even more questions about my data protection, as I learned some of my data is used outside of the app” (P-12-LI).

When questions about data collection and use were answered, participants nonetheless expressed mixed feelings. Some appreciated the transparency. P-21-FB, for example, ultimately felt “good that they showed me what they do because it makes me have a bit more trust” despite concluding that Meta does “track all movements on all pages.” Others, however, were **surprised and uncomfortable with how much data platforms collect**. P-20-TTK learned TikTok “definitely [has] access to more than” they thought. P-16-IG agreed, believing users “should receive some benefit from this other than just getting advertised to all the time.”

Many participants did, in fact, seek this greater control, listing **desired actions including disabling or more broadly controlling data collection and retention** (*Delete/disable: data collection* and *Unspecified control*,

respectively, in Table 2). Participants who wanted to disable data collection only succeeded 25% of the time, and even then, reported doubts as to whether their preferences would be respected. P-1-IG, for example, was “sure they’re still collecting it.” P-7-FB, similarly, did “not fully trust” Meta and believed that their data would “never stop being collected regardless of what [they] chose.” Our taxonomy shows that 13 of 22 platforms allow users to toggle at least some in-platform data collection, while 10 of 22 allow toggling at least some external collection (Table 1). Our participants’ **low success rates in finding data-collection toggles show that they are hard to find** (e.g., “Too much research needed to do all this,” P-11-FB) and/or **not entirely meeting user needs** (e.g., “I feel like I still don’t have control and the options didn’t really help much,” P-3-R). This may be because some controls only affect a subset of data collection, as we describe in Section 4.3, or may reflect general distrust of the platform (e.g., “eBay discloses very little but obscures the actual answer,” P-24-EB).

6.2.2. Nuanced personalization via user customization.

While some participants wanted to take action to disable all ads or prohibit all data collection, others were more interested in being served **better ads**. 47 participants wanted to see **fewer ads fitting some description** (*Decrease*, Table 2), such as “dating ads” (P-14-IG), “sports” (P-14-LI), “intimacy” (P-1-R), and “political advertisements” (P-5-X). In other cases, participants wanted “less targeted ads and more general ones” (P-23-FB) and even specified which categories of personal information they did not want the personalization system to use (e.g., “limit the use of my job title, industry, or connections for determining ad relevance,” P-13-LI). Participants were able to find features mapping to these desired actions 39% of the time.

Some participants also sought ways of **specifying types of ads they wanted to see and increasing the frequency of existing ad categories** (*Increase* and *Unspecified control*:

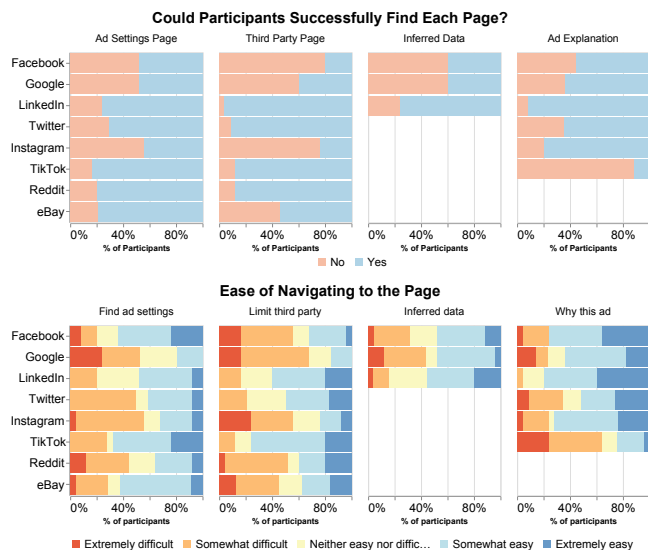


Figure 6: Whether participants found each ATS section on their own and their perception of how easy/difficult it was.

ad category, respectively, Table 2). These included “ads promoted by brands, not users,” such as online influencers (P-13-X), “precious metals” (P-22-EB), “kids and women related stuff” (P-13-FB), and “local businesses and entertainments” (P-11-G). Participants reported finding satisfying features or controls 43% of the time for *Increase* actions and 55% for *Unspecified control: ad category* actions, not dissimilar to *Decrease* actions or actions overall.

6.2.3. Summary: Desires for clarity and nuanced control often go unmet. At a high level, participants’ questions indicate they are searching for more transparency and clarity on how their data is collected and used. Participants’ top goals were not disabling targeting altogether, but rather adjusting the types of ads they are served, often to make the ads more relevant to them. As P-25-IG put it:

“I was pleased to see so many options to tweak my information shared or what is used to target me with ads. I feel it gives me quite a bit of control.”

While many platforms’ ATSs have features that address some of these transparency and control goals, we note that more than two-thirds of all questions and more than half of desired actions went unanswered or unaddressed, indicating that user needs are frequently unmet.

6.3. Experiences with Ad Transparency Systems

While our taxonomy describes the features of various ATSs, our user study allowed us to observe how participants interact with these features. We found that many affordances present challenges for navigation (Section 6.3.1), decision-making (Section 6.3.2), and understanding (Section 6.3.3).

6.3.1. Struggling to navigate. Figure 6 shows the number of participants who found the page in each task without

assistance from our guide, as well as the distribution of Likert difficulty ratings from participants across platforms. About half of participants on Facebook, Google, and Instagram were unable to find the ATS page; on other platforms, the majority of participants found it without assistance. However, 74 (37%) participants reported that navigation was difficult and affordances hard to find:

“I thought I would find ad settings under personal details because the ads are using my personal data. Also, some of the settings are buried really deep in the navigation” (P-21-IG).

We found a significant difference in difficulty across platforms (KW, $H = 22.7$, $p = 0.002$). To identify factors correlated with reported difficulty finding the ad settings page, we performed ordinal logistic regression with the Likert difficulty scale as the outcome variable (Table 3). We used binary variables to indicate if participants had seen the ad settings before, whether they used an ad blocker, whether they had obtained a Bachelor’s degree or higher, and whether the ad settings page had a clear title (Table 9). We also included categorical variables for age and a numerical variable representing the number of clicks and scrolls required to reach the ad settings page (Table 9). We found that participants who had *seen the settings before* were significantly more likely to view navigation as easier (OR = 3.26, $p < 0.001$). As one participant without experience noted, “I wish it was a lot easier to get to these settings. I don’t think a lot of people know that it exists” (P-4-IG). Participants who had a Bachelor’s degree or more perceived more difficulty in navigating to the ad settings page (OR = 0.46, $p = 0.005$). As might be expected, the more *clicks* required to find the ad settings page, the greater perceived difficulty (OR = 0.57, $p = 0.007$).

After struggling to find the initial ad transparency page, many participants found it difficult to find other controls needed to complete study tasks (41%, 31%, 29%, Figure 6). Across platforms, there were significant differences between participants’ difficulty rating for navigating to the third-party data usage page (KW, $H = 35.9$, $p < 0.0001$) and the Ad explanation page (KW, $H = 25.5$, $p = 0.0001$). We performed another ordinal logistic regression with similar factors as the previous model, now with the Likert difficulty responses to navigating to the third-party data usage page as the outcome variable. We added a binary indicator for whether the ATS is a *single page interface* and modified the *number of clicks* to be the *number of additional clicks and scrolls* it would take to reach the third-party data usage page from the primary ad settings page (which the participant had already reached). Participants on platforms like LinkedIn, Twitter, TikTok, and Reddit with a *single page interface* (OR = 3.59, $p = 0.0042$) were significantly more likely to view navigation as easier (Table 4). As before, participants who are *highly educated* (OR = 0.48, $p = 0.0063$) were significantly more likely to view navigation as more difficult. Interestingly, we find no significant effect from the *number of additional clicks and scrolls* required to reach the third-party data usage page, likely because this variable is partially

TABLE 3: Ordinal regression of participants’ perceived difficulty finding the **ad settings** page.

Factors	Est.	OR	Pr(> z)
Seen ad settings?	1.18	3.26	0.0002 ***
Education = Bachelor’s or more	-0.77	0.46	0.0045 **
Number of clicks	-0.57	0.57	0.0068 **
Clear title?	0.44	1.56	0.1307
Use ad block?	0.21	1.23	0.4477
Age = 25-34	-0.11	0.89	0.8220
Age = 35-44	0.41	1.51	0.4158
Age = 45-54	-0.34	0.71	0.5431
Age = 55-64	-0.70	0.50	0.3508
Age = 65+	0.37	1.45	0.7320

Signif. codes: *** < 0.001; ** < 0.01; * < 0.05

TABLE 4: Ordinal regression of participants’ perceived difficulty finding the **limit third-party data usage** page.

Factors	Est.	OR	Pr(> z)
Single page?	1.28	3.59	0.0042 **
Education = Bachelor’s or more	-0.74	0.48	0.0063 **
Seen ad settings?	0.32	1.38	0.3016
Clear title?	0.16	1.17	0.5862
Use ad block?	0.06	1.06	0.8389
Number of additional clicks	-0.07	0.93	0.6413
Age = 25-34	-0.37	0.69	0.4466
Age = 35-44	-0.13	0.88	0.7880
Age = 45-54	-0.65	0.52	0.2162
Age = 55-64	-1.04	0.35	0.1378
Age = 65+	-0.73	0.48	0.4584

Signif. codes: *** < 0.001; ** < 0.01; * < 0.05

associated with single-page design.

Further, when platforms deviate from common design patterns, users suffer; participants on TikTok have an extremely low rate of finding the *Ad explanation* because unlike every other platform, this affordance is under a “share” button instead of a three dot menu that every other platform uses (Figure 6). We propose that standardized design patterns across platforms could reduce these issues.

Overall, our findings indicate that **a meaningful subset of participants struggle to find the ATS at all, and then to find specific affordances within it.** While having previously seen the ad settings page seems to reduce the difficulty of finding the ATS, each platform has a different design and naming conventions for similar affordances. Users may struggle to transfer their experience across platforms.

6.3.2. Decision-making challenges. The third-party data usage section allowed participants to change how their data is *used*. More than half of participants (129, 65%) had at least one third-party data usage control “on” at the beginning of the study (Figure 7). Because only four participants changed a control from “off” to “on,” we focus on participants who started with at least one setting “on.” 67 (52%) of those participants changed at least one of those settings, while the remaining 62 (48%) did not change any settings. Surprisingly, a Kruskal-Wallis test found no significant difference between participant choices when their ATS had a **divergent** personalization toggle in the third-party data usage section compared to those without. Rather, participant decision-making seemed to be influenced by other factors.

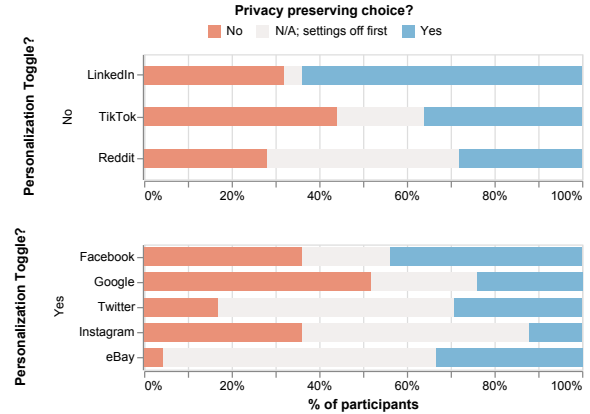


Figure 7: Participants whose third-party data usage was initially “on” and whether they made a privacy-preserving choice by turning it off.

Several participants reported that they wanted to keep the content they saw on the platform the way it was: “*I have not had any negative ads so far*” (P-14-IG). Other participants hesitated to change settings because, “*I would need to weigh the pros and cons and that would take some research and time*” (P-15-EB). Similarly, P-20-LI commented, “*I’m unsure how turning off the settings will affect my profile and experience*” and consequently kept their settings on. These responses suggest that **participants hesitate to change settings when trying to predict how their experience on the platform might be impacted.** A few participants, however, translated concerns about ambiguity into action. For example, P-5-X commented, “*I do not know who Twitter’s business partners are, therefore I cannot determine whether or not I trust them with my personal data.*” They chose to turn their settings off.

Participants also felt “*uneasy about apps or marketing agencies sharing information*” (P-17-LI). Others indicated that they “*did not even know these settings were enabled, and I’m glad to be able to turn them off*” (P-3-TTK). For participants on all platforms, lack of awareness, of both data collection practices and the ATS entirely, were major factors in why they had not already changed these control(s), suggesting discoverability and actionability are critical across all ATSs.

6.3.3. Understanding controls. Many participants did not fully understand the controls and associated descriptions of platform personalization, which could have impacted their inclination to make changes. We evaluated understanding by asking participants to select the options that best described what the controls do when set to “off.” For each platform there were 3-5 correct options. Participants selected more incorrect options than correct (Figure 8). TikTok and Instagram stand out in Figure 8 because participants were moderately or extremely confident in more than 50% of both correct *and* incorrect options. The Wilcoxon rank-sum test found significant differences between confidence

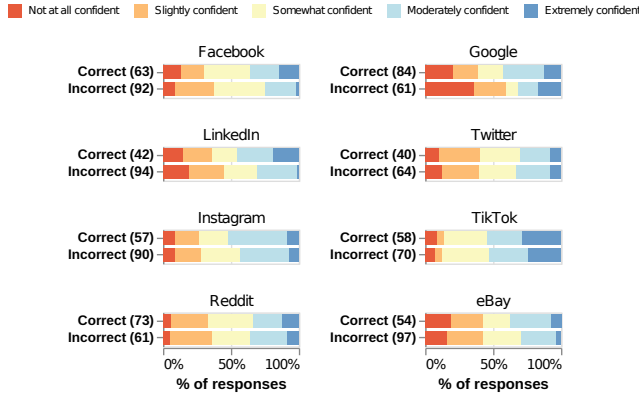


Figure 8: All confidence *responses* from participants grouped by platforms then correct and incorrect options.

in *incorrect* options reported by participants in TikTok and Facebook ($p < 0.001$), Google ($p = 0.001$), LinkedIn ($p < 0.001$), Twitter ($p = 0.02$), and eBay ($p < 0.001$). Similarly, the Wilcoxon rank-sum test found significant differences between confidence in *correct* options reported by participants in TikTok and eBay ($p = 0.04$).

To investigate the high levels of confidence we see from participants on TikTok and Instagram, we calculated the number of participants who are highly confident about correct and incorrect options (Table 7 in Appendix C). For all platforms except Google and Twitter, a majority of participants who selected correct options were highly confident about both correct and incorrect options, indicating that they did not have full understanding of the control’s behavior. This highlights ATs’ **failures to clearly communicate the effects of using personalization and data controls**. While some participants realized they did not understand the controls—“*I think it just needs more detail on what each of the settings actually do*” (P-12-R)—many others appeared unfazed by or unaware of potential complexity, “*It all seems straight forward to me*” (P-9-R). To address this, platforms should consider providing “*examples of how my interactions with businesses were used to target ads to me*” (P-7-LI) because “*the information provided is very basic and could be made a lot more transparent*” (P-2-R).

6.3.4. Understanding and usability of inferred data.

Next, we report on participants’ reaction to ATs presentation of inferred topics data. Our taxonomy suggests that a lack of *exploration tools* might limit participants’ ability to process this data. While Facebook lets users search through topics and Google has sorting/filtering capabilities, LinkedIn has no exploration affordances, instead requiring users to go through pages of topics. LinkedIn participants especially felt this was a barrier: “*I would have possibly removed some categories, but there were too many for it to make a difference. Not wasting time going through 100+ categories just to see *different* ads (vs. fewer)*” (P-2-LI).

Across all platforms, however, participants found topic-level personalization useful: “*I chose that [to see more about*

certain topics] because I’d rather see certain ads” (P-3-FB). For some participants, inferred topics affordances created additional, unanswered questions like “*Facebook seems to be tracking my Google searches and I’d like to know how much Google is sharing*” (P-6-FB).

LinkedIn participants demonstrated that *exploration tools* were essential to the utility of an inferred topics section. Inferred topics affordances did not address questions of how the inference was made (inferences were only explained in ad explanations); nor did they describe what seeing more or less of an ad topic would do, as P-2-FB suggested. Inferred topics affordances should be accompanied by information about *how* each inference was made and also clearer explanations about the tangible effect of choosing to see more or less of an ad topic.

6.3.5. Evaluating ad explanations.

Ad explanations across all platforms are designed similarly. About half of all participants assessed ad explanations as incomplete and not useful while the other half viewed them as complete and useful (Figure 9). Participants critiqued these when the inferences were inaccurate or did not fit to their expectations, “*The ad is for laser hair removal and the page says the activity that made this ad target me was that I had interacted with content about food & drinks*” (P-2-IG). Other participants felt that “*very generic phrases*” like “*your profile matches the intended audience*” (P-39-LI) were not at all useful.

When another segment of participants felt parts or all of the explanation fit, they tended to find the explanation more complete and more useful, “*I follow a movie company that’s involved with distributing this movie*” (P-1-FB).

We found no significant difference between participant evaluations, suggesting that the effect of different ad explanation implementations [2], [3], [17], [22] was negligible for participants on their perception of completeness and usefulness (Figure 9). These results imply that for many ads users see on all platforms, **ad explanations fail to give complete or useful reasons for why they were targeted**.

6.3.6. Final participant evaluations.

Overall, participants found systems difficult to use and not particularly useful. Participants on multi-page platforms (i.e., Facebook, Google, Instagram, eBay) found the structure/layout of the pages to be significantly more difficult than those on single-page platforms (KW, $H = 6.6$, $p = 0.01$), though the majority of participants found that the structure of ATs made it difficult to find information (Figure 9). While nearly half of participants found the terminology used by ATs to be simple, this may be to the user’s detriment. Simple language does not—and should not—rule out clear and accurate descriptions of a platform’s practices. Participants on Facebook, LinkedIn, Instagram, and TikTok seem particularly affected by the appearance of simplicity. On these platforms, 13, 11, 15, and 12 participants respectively were confident in correct and incorrect options to describe what controls would do when turned off (Table 7). This is communicated by participants feeling that ATs were not useful for informing them about how ads are targeted (Figure 9). Participants




Figure 9: Participants’ evaluations of the completeness, usefulness, and usability of the ATSS.

felt the language was “oversimplified,” “nebulous and incomplete,” “vague and not really informative,” and that it was “obscuring the true intentions” of the platform (P-8-X, P-24-EB, P-11-LI, P-11-R, respectively).

Furthermore, participants found that ATSS were more useful for controlling their personalization settings than for informing users at significantly different rates (37% vs 30%) (KW, $H = 4.2$, $p = 0.039$). While small, this difference indicates participants’ satisfaction with being able to “turn off the targeting of ads. And now they are more general. I know Facebook is still collecting my information but it’s not so obvious now” (P-23-FB). For a number of participants, a **personalization toggle was simple to understand and simple to use**, quickly satisfying their needs. However, that is predicated upon users being able to locate the ATS and understand the control(s).

7. Discussion

7.1. Issues unaddressed by existing systems

By comparing our findings from the taxonomy study with the findings from the user study, we argue that platforms are not meeting their users’ primary need: the ability to understand and have agency over how their personal data is collected, used, stored, and sold for advertising. 168 participants had questions about how their data was being collected and used, as we describe in Section 6.2.1. Data collection questions were mostly high-level (e.g., “How are you deciding what to show me?” P-16-R), but only 36% of them were satisfyingly answered (Table 2). Questions about personalization, similarly, were answered 43% of the time. This is surprising since, as we describe in Section 4, the “Platform personalization” affordance was implemented by the most platforms (19) out of any of the affordances. The reason for this discrepancy might lie in a tension we identified in many participants’ responses: platforms’ explanations were over-complicated, but at the same time too vague. We argue that platforms are devoting too much of their explanation on generic overviews of their ATSS instead of on the pieces of the system that directly touch user data.

Ultimately, **the details about how a platform shows ads to all of its users is less relevant to participants than the details about how a specific user’s personal data was used to show them a specific ad**. This claim is supported by the low frequency of implementation of the affordances that tackle these specific questions, namely “Ad explanation” (6 platforms) and “Specific data types” (4 platforms).

In addition to wanting to learn about and control data collection and use, nearly half of the participants wanted to know how their data was being stored, shared, and sold. However, **we did not identify any affordances that directly tackled these key issues of data storage and selling**. Given the lack of affordances dedicated to answering these questions, it is surprising that participants were even able to find answers to 17% of data storage and 21% of data selling questions. Participants’ concerns with the security and privacy of their data manifested in other scenarios as well, such as the 22 questions where participants suspected that their conversations were being recorded by the microphones on their devices (4% of these were answered). Much like in the case of data collection and use, we argue that platforms should be specific about how user data is being sold or shared and explain how user data is kept safe internally.

7.2. Recommendations

Interview industry experts to better understand platform priorities. While this and prior work largely focuses on user needs and priorities with transparency systems, they are not the only relevant stakeholders. Indeed, industry experts involved in the design of ATSS and ad transparency legislation have not been examined by the literature. While designing the study, we used blog posts about ATS updates from Facebook and Google to understand platform goals in designing these systems, though this granted us limited details. Future work should interview industry experts to better understand platforms’ priorities with ATSS and to understand how legislation can shape these priorities.

ATSS should describe the full provenance of data collection and inference pipelines. 37% of participants felt that

ATSS were not useful for informing them about for how ads are targeted. We contend that the vague language around data origins identified in the taxonomy plays a key role in this. Specifically, 168 (85%) participants asked questions about which data is used with many of those questions being about specific data types or sources. Many of those participants discovered that the information provided by the ATS did not meet the level of detail requested in their questions. ATSS could improve transparency and their informational capacity by explicitly stating the sources of inferences in ad explanations. Similarly, jargon like “*TikTok’s estimate of your interest*” should definitively describe *how* that interest was chosen and *what* data was used, not just gesture towards possible explanations. P-10-FB states, “*I would like to know exactly how the topics are determined.*”

Provide more options for user specification of interests and non-interests. Participants frequently wanted to specify the types of ads they want to see or not see (Table 2). However, few platforms offer the functionality to control specific ads or ad topics (Table 1), and these are usually limited to what a user can identify in the inferred topics affordance, which can be challenging, as shown in Section 6.3.4. We propose adding functionality for users to manually specify ad interests they have, focusing especially on in-situ controls to improve actionability, as prior work suggests [2], [26], [27]. The majority of personalization on these platforms comes from automated systems that users can only influence indirectly through their activity on platforms. Such functionality could increase the amount of personalization originating from intentional user action, rather than opaque automated systems and unspecified data flows. We view this as an opportunity for platforms to increase the amount of agency given to users while still showing relevant content.

Standardization could help users understand ATSS and make well-informed decisions. The systems we observed had few overlaps in features available, and fewer of those features were similar in verbiage and design. ATSS described similar third-party data usage controls as “*Manage your off-TikTok data,*” “*Create profiles for personalized advertising,*” “*Activity information from ad partners,*” “*Use partner info to improve which recommendations and ads you see on Pinterest,*” and more. Participants in the user study described this type of language as “*nebulous,*” “*vague,*” and “*oversimplified.*” Additionally, if a user wanted to modify their settings on multiple platforms, each platform’s new verbiage represents another obstacle to understanding and decision-making. To address these issues, we propose that platforms adopt industry standards to describe controls, data sources, and other aspects of ATSS. For example, the FTC proposes several understandable categories of data sources—e.g., data from users, data brokers, inferred information—that could be employed as standards [35]. Future research could explore which terms best describe different data sources and are most understandable to establish ATS standardization.

Acknowledgments

This material is based upon work supported by the National Science Foundation under Grants No. CNS-2047827, CNS-2149680, CNS-2151290, and CNS-2151837, as well as by the National Science Foundation Graduate Research Fellowship under Grant No. 2140001. We thank Miriam Metzger, Chris Geeng, Laurent Wang, and Gwen Petro for their assistance conceptualizing the study, as well as Nina Rebar and Catherine Golovin for their assistance with qualitative coding and data analysis.

References


- [1] M. Ali, P. Sapiezynski, A. Korolova, A. Mislove, and A. Rieke, “Ad delivery algorithms: The hidden arbiters of political messaging,” in *Proceedings of the 14th ACM International Conference on Web Search and Data Mining*, 2021.
- [2] J. Castleman and A. Korolova, “Why am I still seeing this: Measuring the effectiveness of ad controls and explanations in AI-mediated ad targeting systems,” in *Proceedings of the AAAI/ACM Conference on AI, Ethics, and Society*, 2024.
- [3] M. Wei, M. Stamos, S. Veys, N. Reitering, J. Goodman, M. Herman, D. Filipczuk, B. Weinschel, M. L. Mazurek, and B. Ur, “What Twitter knows: Characterizing ad targeting practices, user perceptions, and ad explanations through users’ own Twitter data,” in *Proceedings of the 29th USENIX Security Symposium*, 2020.
- [4] A. Andreou, M. Silva, F. Benevenuto, O. Goga, P. Loiseau, and A. Mislove, “Measuring the Facebook advertising ecosystem,” in *Proceedings 2019 Network and Distributed System Security Symposium*, 2019.
- [5] “Civil Rights Division | United States v. Meta Platforms, Inc., f/k/a Facebook, Inc. (S.D.N.Y.) | United States Department of Justice,” Jun. 2022. [Online]. Available: <https://www.justice.gov/crt/case/united-states-v-meta-platforms-inc-fka-facebook-inc-sdny>
- [6] V. Nagaraj Rao and A. Korolova, “Discrimination through image selection by job advertisers on Facebook,” in *Proceedings of the 2023 ACM Conference on Fairness, Accountability, and Transparency*, 2023.
- [7] B. Imana, A. Korolova, and J. Heidemann, “Auditing for discrimination in algorithms delivering job ads,” in *Proceedings of the Web Conference*, 2021.
- [8] D. Calvo, L. Cano-Orón, and T. Baviera, “Global spaces for local politics: An exploratory analysis of Facebook ads in Spanish election campaigns,” *Social Sciences*, vol. 10, no. 7, 2021.
- [9] J. Heawood, “Pseudo-public political speech: Democratic implications of the Cambridge Analytica scandal,” *Information Polity*, vol. 23, no. 4, pp. 429–434.
- [10] K. Steinmetz, “Lawmakers hint at regulating social media during hearing with Facebook and Twitter execs,” *Time*, 2018.
- [11] M. Wall, “Facebook, Google and Twitter in data regulators’ sights,” *BBC*, 2019.
- [12] European Parliament and Council, “Regulation (EU) 2022/2065 of 19 October 2022,” 2022. [Online]. Available: <https://eur-lex.europa.eu/eli/reg/2022/2065>
- [13] European Parliament and Council of the European Union, “Consolidated text: Regulation (EU) 2016/679 of the European Parliament and of the Council of 27 April 2016 on the protection of natural persons with regard to the processing of personal data and on the free movement of such data, and repealing Directive 95/46/EC (General Data Protection Regulation) (Text with EEA relevance),” May 2016. [Online]. Available: <https://eur-lex.europa.eu/eli/reg/2016/679/oj>

- [14] “California Consumer Privacy Act of 2018,” pp. 1798.100 – 1798.199.100, 2018.
- [15] F. M. Farke, D. G. Balash, M. Golla, M. Dürmuth, and A. J. Aviv, “Are privacy dashboards good for end users? Evaluating user perceptions and reactions to Google’s My Activity,” in *Proceedings of the 30th USENIX Security Symposium*, 2021.
- [16] F. M. Farke, D. G. Balash, M. Golla, and A. J. Aviv, “How does connecting online activities to advertising inferences impact privacy perceptions?” in *Proceedings on Privacy Enhancing Technologies*, 2024.
- [17] A. Andreou, G. Venkatadri, O. Goga, K. P. Gummedi, P. Loiseau, and A. Mislove, “Investigating ad transparency mechanisms in social media: A case study of Facebook’s explanations,” in *Proceedings of the 2018 Network and Distributed System Security Symposium*, 2018.
- [18] E. Gkiouzepe, A. Andreou, O. Goga, and P. Loiseau, “Collaborative ad transparency: Promises and limitations,” in *Proceedings of the 2023 IEEE Symposium on Security and Privacy*, 2023.
- [19] M. Ali, P. Sapiezynski, M. Bogen, A. Korolova, A. Mislove, and A. Rieke, “Discrimination through optimization: How Facebook’s ad delivery can lead to biased outcomes,” *Proceedings of the ACM on Human-Computer Interaction*, vol. 3, no. CSCW, Nov. 2019.
- [20] M. Ali, A. Goetzen, A. Mislove, E. M. Redmiles, and P. Sapiezynski, “Problematic advertising and its disparate exposure on Facebook,” in *Proceedings of the 32nd USENIX Security Symposium*, 2023.
- [21] T. Speicher, M. Ali, G. Venkatadri, F. N. Ribeiro, G. Arvanitakis, F. Benevenuto, K. P. Gummedi, P. Loiseau, and A. Mislove, “Potential for discrimination in online targeted advertising,” in *Proceedings of the 1st Conference on Fairness, Accountability and Transparency*, 2018.
- [22] D. Wilkinson, M. Namara, K. Patil, L. Guo, A. Manda, and B. Knijnenburg, “The pursuit of transparency and control: A classification of ad explanations in social media,” in *Proceedings of the Hawaii International Conference on System Sciences*, 2021.
- [23] T. Kim, K. Barasz, and L. K. John, “Why am I seeing this ad? The effect of ad transparency on ad effectiveness,” *Journal of Consumer Research*, vol. 45, no. 5, pp. 906–932, 2019.
- [24] A. Rao, F. Schaub, and N. Sadeh, “What do they know about me? Contents and concerns of online behavioral profiles,” arXiv:1506.01675, 2015.
- [25] P. Leon, B. Ur, R. Shay, Y. Wang, R. Balebako, and L. Cranor, “Why Johnny can’t opt out: A usability evaluation of tools to limit online behavioral advertising,” in *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems*, 2012.
- [26] H. Habib, S. Pearman, E. Young, I. Saxena, R. Zhang, and L. F. Cranor, “Identifying user needs for advertising controls on Facebook,” *Proceedings of the ACM on Human-Computer Interaction*, vol. 6, no. CSCW, Mar. 2022.
- [27] J. Im, R. Wang, W. Lyu, N. Cook, H. Habib, L. F. Cranor, N. Banovic, and F. Schaub, “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*, 2023.
- [28] H. Habib and L. F. Cranor, “Evaluating the usability of privacy choice mechanisms,” in *Proceedings of the Eighteenth Symposium on Usable Privacy and Security*, 2022.
- [29] H. Habib, M. Li, E. Young, and L. Cranor, ““Okay, whatever”: An evaluation of cookie consent interfaces,” in *Proceedings of the 2022 CHI Conference on Human Factors in Computing Systems*, 2022.
- [30] Y. Wu, S. Bice, W. K. Edwards, and S. Das, “The slow violence of surveillance capitalism: How online behavioral advertising harms people,” in *Proceedings of the 2023 ACM Conference on Fairness, Accountability, and Transparency*, 2023.
- [31] P. Sampson, R. Encarnacion, and D. Metaxa, “Representation, self-determination, and refusal: Queer people’s experiences with targeted advertising,” in *Proceedings of the 2023 ACM Conference on Fairness, Accountability, and Transparency*, 2023.
- [32] V. Le Pochat, T. Van Goethem, S. Tajalizadehkhooob, M. Korczyński, and W. Joosen, “Tranco: A research-oriented top sites ranking hardened against manipulation,” in *Proceedings of the 26th Annual Network and Distributed System Security Symposium*, 2019.
- [33] D. Ibdah, N. Lachtar, S. M. Raparathi, and A. Bacha, ““Why should I read the privacy policy, I just need the service”: A study on attitudes and perceptions toward privacy policies,” *IEEE Access*, vol. 9, 2021.
- [34] Y. Zou, S. Danino, K. Sun, and F. Schaub, “You ‘might’ be affected: An empirical analysis of readability and usability issues in data breach notifications,” in *Proceedings of the 2019 CHI Conference on Human Factors in Computing Systems*, 2019.
- [35] “FTC Staff Report Finds Large Social Media and Video Streaming Companies Have Engaged in Vast Surveillance of Users with Lax Privacy Controls and Inadequate Safeguards for Kids and Teens,” <https://www.ftc.gov/news-events/news/press-releases/2024/09/ftc-staff-report-finds-large-social-media-video-streaming-companies-have-engaged-vast-surveillance>, Sep. 2024.

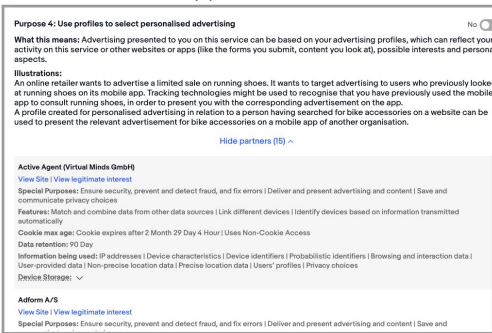
Appendix A. Codebook and Full Survey Instrument

The codebook and full survey instrument from our user study are available at the following link: https://osf.io/qgr94/?view_only=de2e5a845e544727a2cb5ffc19915d2b

Appendix B. Additional Screenshots of Interfaces



(a) LinkedIn



(b) eBay

Figure 10: Examples of **units** from LinkedIn and eBay.

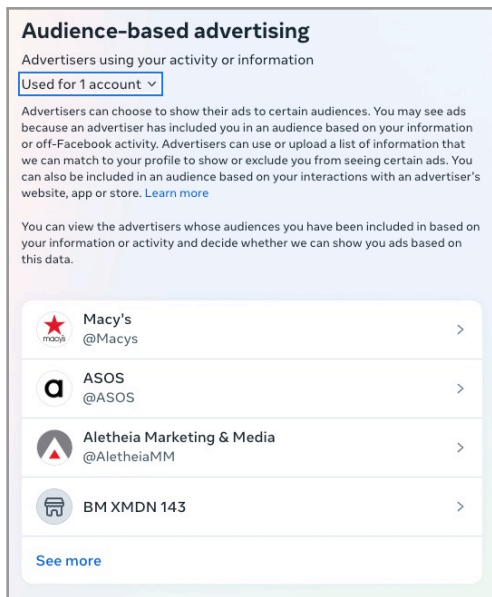


Figure 11: How Facebook displays the data from audience-based advertising.

Appendix C. Additional Tables

TABLE 5: Characteristics of the ATSeS when accessed via the Chrome browser on a desktop (🖥️), via the platform's Android app (📱), and via Chrome on Android (📲). “-” indicates the ATS cannot be accessed in that manner.

Platform	Tranco	# pages			# units		
		🖥️	📱	📲	🖥️	📱	📲
Google	1	9	8	9	43	30	43
Facebook	4	8	8	8	87	88	85
Microsoft	5	4	4	4	21	21	21
Twitter	8	10	11	11	29	30	29
Instagram	9	5	5	5	38	38	38
Apple	12	2	-	-	12	-	-
LinkedIn	14	3	3	3	43	34	43
Yahoo	19	17	17	17	47	47	47
Amazon	22	1	3	1	4	6	5
Warner Bros	24	1	-	1	7	-	7
Pinterest	34	1	1	1	9	14	9
Reddit	44	2	2	2	5	4	5
TikTok	74	-	4	-	-	13	-
Tumblr	85	5	5	5	31	26	30
Spotify	89	2	-	2	12	-	1
eBay	113	1	1	1	17	17	17
Canva	118	1	1	1	1	1	1
SoundCloud	123	2	1	1	4	3	3
Fandom	128	1	-	1	1	-	1
The Guardian	136	3	4	3	7	39	7
AliExpress	140	3	3	3	5	5	5
StackOverflow	142	2	-	-	12	-	-

TABLE 6: Participants' demographics and experiences.

Category	n	%	Category	n	%
Gender			Education		
Woman	118	59	No high school diploma	2	1
Man	72	36	GED	4	2
Non-binary	7	4	High school diploma	21	11
Other	1	0	<1 yr of college	10	5
Race/Ethnicity			>1 yrs of college	35	18
White	131	66	Associate's degree	26	13
Black or Af. Am.	23	12	Bachelor's degree	72	36
Asian	22	11	Master's degree	22	11
Hisp. or Lat.	13	6	Professional degree	4	2
Am. Ind. or Ala. Nat.	5	2	Doctorate	2	1
Mid. East. or N. Afr.	2	1	Use Ad Blocker?		
No answer	1	0	No	107	54
Nat. Hawai. or Pac. Isl.	1	0	Yes	91	46
Age Range			Previously Viewed Assigned ATSeS?		
18-24	17	9	No	145	73
25-34	70	35	Yes	53	27
35-44	63	32			
45-54	34	17			
55-64	11	6			
65+	3	2			

TABLE 7: The number of participants who were moderately or extremely confident in their answers and answered those questions correctly, incorrectly, or a mix of both.

Platform	Confident & Correct	Confident & Incorrect	Overlap
Facebook	12	13	9
Google	14	6	6
LinkedIn	12	11	8
Twitter	9	6	4
Instagram	17	15	12
TikTok	13	12	8
Reddit	12	8	7
eBay	10	9	8

TABLE 8: Comprehensive taxonomy of the components of ad transparency systems from the 22 systems we studied. # indicates the number of systems this feature or tool is present in.

Feature of system	#	Definition
Platform and user information	19	
Platform personalization	19	Mention and explanation of how ad targeting and personalization works for the platform.
Ad explanation	6	Explanation of techniques used to target user for specific ad.
Ad related activity logs	3	Historical log of ad activity (e.g., ads seen or clicked) with the date when this activity happened.
Ad, topic, and advertiser controls activity logs	5	Historical log of changes to ad settings (e.g., disabling personalization), timestamped.
Inferred topics associated with user	7	List of inferred topics or ad categories the user is associated with.
Advertisers associated with user	5	List of advertisers the user is associated with (e.g., an advertiser a user has seen ads from recently or an advertiser that the user has previously chosen to see more or less ads from).
Advertiser information	7	Information about the advertiser (e.g., location, tenure in platform, homepage).
Data shared by user	9	Explicit mention that information was shared by the user themselves.
Data inferred by platform	10	Explicit mention that information was inferred by the platform, not shared by user.
Data sharing ambiguous	9	Unclear if information was shared by user or inferred by platform.
Exploration tools	6	
Sorting	3	Sort the order in which you see ads a user has interacted with, brands, or topics.
Filtering	3	Not using keyword search, via a dropdown or boxes that can be checked instead.
Searching	3	Allows users to search using a text field.
Filter by advertiser	4	Search/explore by advertiser.
Filter by topic	2	Filter by the subject of an ad.
Personalization and data controls	18/19/20	<i># of platforms with personalization controls, data controls, and either type.</i>
Personalized ad toggle	15	Control states, "This will turn off personalized ads" {either entirely or from a specific data source}.
Report ad	8	Users can report an ad.
Ad frequency	1/6	Users can increase or decrease the frequency with which they see a particular ad or entirely remove an ad.
Topic frequency	4/4/2	Users can increase, decrease/limit, or entirely remove ads shown about a topic.
Advertiser frequency	1/2/4	Users can increase, decrease, or entirely remove ads shown by an advertiser.
Specific data types	4	E.g., gender, age, family and relationships, income, etc.
Web activity	13	Website or app activity (e.g., clicks, comments) tracked and logged.
Audience lists	10	Information sold/shared to the ad platform by another company.
Location	4	User's precise or approximate location and location history.
Data within platform	13	Any web activity used to show ads to users is only collected within the platform where the ad is shown.
Data outside platform	10	Platform uses data outside of platform to target ads to users.

TABLE 9: This table characterizes the user experience of the home page of each system. So that this information is usable for comparison in our study, we focus on updated mobile app interfaces of the eight platforms in the user study. **Single Page?** indicates whether the ATS interface is a single page. **Clicks & Scrolls** counts the number of clicks and scrolls required to reach the ATS, while **Additional Clicks & Scrolls** references the clicks required to reach the *Limit third-party data usage* page. **Page Title Clear?** states whether the page title indicates the page's purpose. For example, the Facebook ATS is on a page titled "Accounts Center", whereas the corresponding page on LinkedIn is titled "Advertising Data". Of the platforms with *all affordances*, only LinkedIn has a single page interface, while overall single page interfaces are relatively common as they correlate with smaller interfaces in general. While most platforms require a similar number of clicks and scrolls (tested on an iPhone XS and a Pixel 3a with the updated version of each app) to reach, the amount of time spent looking may change based on whether the title of the page they're looking for clearly indicates what it is they're looking for. Despite 5 of these platforms having a clear page title, 3 of these platforms require 5 clicks and scrolls to reach the interface, which could create more opportunities for a user to get lost while looking for the page.

Affordance Group	Platform	Single Page?	Clicks & Scrolls	Additional Clicks & Scrolls	Page Title Clear?
All Affordances	Facebook	No	4	4	No
	Google	No	5	2	Yes
	LinkedIn	Yes	3	1	Yes
Most Affordances	Twitter	Yes	5	1	Yes
	Instagram	No	4	4	No
	TikTok	Yes	4	0	Yes
Some Affordances	Reddit	Yes	4	0	No
	eBay	No	5	1	Yes

Appendix D. Meta-Review

The following meta-review was prepared by the program committee for the 2025 IEEE Symposium on Security and Privacy (S&P) as part of the review process as detailed in the call for papers.

D.1. Summary

This paper discusses ad-transparency systems (ATSs) that are essential in informing users about data collection for displaying personalized ads. It motivates the study by highlighting that existing state-of-the-art studies focus on major ATSs in isolation. Thus, an overall understanding is missing. To address this gap, this paper presents a new taxonomy of 22 popular ATSs selected based on popularity, aiming to study I) the features of different ATSs, II) what participants want from ATSs (expectation), and what the users can do with current ATSs (fulfillment). To answer these, the paper presents the findings of an online user study across eight platforms and reports several insights that can be used to improve the overall state of ATSs. The paper makes usability recommendations for the ATSs.

D.2. Scientific Contributions

- Provides a Valuable Step Forward in an Established Field

D.3. Reasons for Acceptance

- 1) Provides a Valuable Step Forward in an Established Field. Creating and studying the taxonomy of ATSs for 22 platforms is a welcome contribution to the existing research on the privacy and usability of ATSs. Prior work has investigated individual platforms' ad transparency systems and information provided to users about how their data is used to deliver ads (e.g., through privacy policies and privacy labels). This paper expands the field by conducting a broad comparative analysis of ATS from 22 platforms. It highlights inconsistencies in ATS design and how users engage with these systems. Since targeted ads raise privacy concerns, users want access to information about why they are being shown a particular ad and what data are being collected from them. ATS platforms help users by addressing these concerns. In this context, this paper provides a valuable step forward in understanding the current gaps and suggesting future directions for ATSs.