

Misalignments and Demographic Differences in Expected and Actual Privacy Settings on Facebook

Byron M. Lowens
University of Michigan
bmlowens@umich.edu

Sean Scarnecchia
University of Michigan
sscarnec@umich.edu

Jane Im
University of Michigan
imjane@umich.edu

Tanisha Afnan
University of Michigan
tafnan@umich.edu

Annie Chen
University of Michigan
anniechn@umich.edu

Yixin Zou
Max Planck Institute for Security and
Privacy
yixin.zou@mpi-sp.org

Florian Schaub
University of Michigan
fschaub@umich.edu

Abstract

Social media platforms pose privacy risks when data is used in unexpected ways (e.g., for advertising or data sharing with partners). Using a custom browser extension and an online survey with 195 Facebook users, we investigated (1) whether participants' expected values of their Facebook privacy settings were (mis)aligned with their actual settings; (2) demographic differences in privacy expectation-setting mismatches; and (3) participants' privacy concerns and trust towards Facebook. Our study presents a current and comprehensive analysis of Facebook users' privacy settings. We find that expectation-setting mismatches are prevalent: all participants had at least one mismatch; many had multiple, often expecting their settings to be more restrictive than they were. We also found that Facebook's default values are not aligned with people's expectations and/or actual settings, which suggests that those defaults are ineffective. Furthermore, mismatches differed along certain demographic variables. Participants' trust in Facebook decreased after they became aware of mismatches and their actual settings. Our empirical findings indicate that, despite increased public awareness, media scrutiny, and regulatory attention regarding privacy issues, there is still a substantial and concerning disconnect between how private people perceive their social media data to be and how exposed their data actually is, opening them up to both interpersonal and institutional privacy risks. We discuss design and public policy implications of our findings.

Keywords

Privacy, Social Media, Privacy Expectations.

1 Introduction

Social media companies claim to provide users adequate control over their data and prioritize user safety [112]. However, they have

also received public scrutiny due to privacy concerns [11, 67, 89]. Facebook, in particular, has received widespread media attention for its invasive personal data collection and targeted ad practices [9, 17, 87], especially after the Cambridge Analytica incident [21]. In response to regulatory scrutiny and emerging privacy laws, platforms have expanded their privacy settings. At the same time, the notice and choice paradigm has been heavily criticized [22, 90]: Consumers are often unaware of their data exposure and struggle to anticipate privacy risks and secondary data uses, such as online behavioral advertising [78, 79, 107]. If platforms' privacy settings were effective, users would know about and make use of them, and their expectations for data use would match the actual privacy settings in their accounts.

In our study, we investigated the latter, specifically, to what extent people's expectations of their privacy settings on Facebook are (mis)matched with their actual Facebook privacy settings collected from the platform using a browser extension. We focused on Facebook given that it is still the most popular social media platform [106], its large and diverse user base, prior privacy scandals, and constantly evolving privacy setting interfaces. Compared to prior work on people's Facebook privacy settings [44, 50, 58, 61], we cover a more comprehensive set of privacy settings, spanning general visibility, timeline, and ad settings. Furthermore, Facebook's user base and privacy settings have evolved significantly since some of these earlier studies were conducted. There has also been limited research on how exposure on Facebook (i.e., how open or private settings are) varies among different demographic groups.

We conducted a mixed-methods study with 195 participants. Whereas most prior work has relied on survey data alone, we achieve high ecological validity by developing a browser extension that—with a participant's consent—gathered their current values for 18 Facebook privacy settings (including general visibility settings, timeline settings, and ad settings) and information about their Facebook ad profiles. We then live embedded a participant's retrieved settings into an online survey. For each setting, we asked participants what they expected it to be, then showed them the actual setting collected from the platform, and elicited their level of concern with this setting. Specifically, our study investigated the following research questions:

This work is licensed under the Creative Commons Attribution 4.0 International License. To view a copy of this license visit <https://creativecommons.org/licenses/by/4.0/> or send a letter to Creative Commons, PO Box 1866, Mountain View, CA 94042, USA.



Proceedings on Privacy Enhancing Technologies 2025(1), 456–471
© 2025 Copyright held by the owner/author(s).
<https://doi.org/10.56553/popets-2025-0025>

RQ1: To what extent are participants' expected values of their Facebook privacy settings aligned or misaligned with their actual settings collected directly from their Facebook accounts?

RQ2: To what extent are there demographic differences in privacy expectation-setting misalignments?

RQ3: To what extent are participants concerned about privacy expectation-setting misalignments on Facebook?

Our study makes multiple contributions: (1) We identified substantial mismatches between participants' expected values and actual Facebook privacy settings. Many participants had multiple mismatches among the settings we analyzed, often expecting their settings to be more restrictive, revealing widespread misconception regarding how private or exposed their personal information is on Facebook. (2) Participants had, on average, 474 ad topics assigned to them and had been targeted with ads by 441 companies—much higher than what the majority of participants expected. (3) Our exploratory analysis indicates demographic differences likely exist in Facebook users' expectation-settings mismatches, with participants' age, race/ethnicity, gender, and political affiliations potentially being relevant factors. (4) While participants were neutral on average about their settings, they expressed heightened concern about certain settings (e.g., being shown ads off Facebook based on Facebook activity). (5) Participants' awareness of their actual settings and mismatches led to decreased trust in Facebook.

Many of our findings reaffirm prior work, revealing a persistent and continued gap between users' expectations and actual privacy settings [50, 58, 61], including prevalent and substantial settings mismatches for post and timeline visibility, demographic influences on settings behaviors, and users' unexpected overreliance on Facebook's privacy defaults. While our findings align with previous studies, it's important to recognize that Facebook's platform has evolved significantly since these earlier works. Key changes include Facebook's 2011 overhaul of privacy settings [74], the introduction of the Timeline feature [93], the Privacy Checkup Tool [46], adjustments following the Cambridge Analytica scandal in 2018 [52], the Off-Facebook Activity tool in 2019 [69], and the introduction of Facebook's Privacy Center in 2021 [96]. Our empirical findings demonstrate that a substantial disconnect persists between how private people perceive their social media data to be and how exposed their data actually is, which opens them up to both interpersonal privacy risks (e.g., from other users) and to institutional privacy risks (e.g., unexpected types of monetization of their data by the platform). This discrepancy is especially concerning given the increased attention privacy issues have received in recent years by the media, the public, and policymakers, as well as platform's claims to value and protect user privacy. We discuss our findings' implications for privacy design and public policy.

2 Related Work

Privacy settings on social networking sites have been widely studied. Our work measures mismatches between Facebook users' expected privacy settings and actual settings, while also exploring potential demographic differences. Unlike prior survey-based research, we analyzed participants' privacy expectations alongside their actual settings, retrieved from their Facebook account with our custom browser extension. Compared to earlier similar studies,

our work captures Facebook's increasingly diverse user base and a larger set of privacy settings.

2.1 Privacy Preferences and Behaviors

Prior work has investigated people's privacy attitudes, preferences, concerns, expectations, decisions, and behaviors [20, 86]. U.S. adults have increased concern over how their personal information is handled, with two-thirds of Americans saying "they understand little to nothing about what companies are doing with their personal data" [76]. The relationship between privacy behaviors and one's privacy preferences has been studied intensively, yielding the notion of a privacy paradox, which refers to a mismatch between the two [53]. Prior work has found some evidence of the privacy paradox in that people's privacy intentions do not always align with their behaviors on social media platforms [45, 50, 58, 61]. However, the privacy paradox has been contested [2, 33], as other work finds that this mismatch may be perpetuated by the diversity of scenarios and constructs used across different studies [24, 88, 91]. For example, some studies measure general attitudes, whereas others measure specific concerns or intentions [3].

The focus of our study—potential misalignments between Facebook users' expected and their actual settings from their accounts—could be viewed as a comparison of mental states and behaviors. While prior studies have examined the privacy paradox in the context of Facebook, many have employed qualitative methods [54, 80, 109], e.g., finding that college students adopt various strategies to address privacy concerns, but those concerns and strategies primarily revolve around social/interpersonal privacy [109] rather than institutional privacy with respect to a platform's data practices. Other studies have quantified risks and consequences of this behavior-preference gap on Facebook [5, 37, 92]. For instance, in 2005, Gross and Acquisti analyzed over 4k Facebook profiles and found that the majority shared a large amount of personal information, with only 0.06% of all users in their sample at that time choosing to limit access to their profile to just friends [37]. We take a novel approach by developing a custom browser extension to gather participants' actual privacy settings from their Facebook account, providing reliable data with high ecological validity, about both users' expectations of and their actual privacy settings.

2.2 Privacy and Identity Characteristics

There has been increasing interest in understanding at-risk populations' privacy concerns and experiences [67, 83, 104]. Most prior research on Facebook users' privacy concerns has been limited to college students or younger users.

Some prior work has studied differences in privacy concerns across demographics, primarily across age, gender, and education [31, 50, 87, 99], with varying results.

Older adults were found to be more likely to post publicly [31], have more mismatches between expected and actual privacy settings than any other age group [50], and are less concerned than younger adults about their privacy on social media [99]. More recent work examining Facebook's targeted advertising practices finds that older adults, along with Black and Hispanic participants, were more likely to be shown problematic Facebook ads, such as clickbait ads and scams [6]. Race and ethnicity influence privacy and

ad experiences on Facebook as well [68, 108]. Facebook has been critiqued for promoting ‘technological redlining,’ by displaying targeted job and housing ads to users of some racial groups while excluding others [7, 10]. Regarding gender differences, some studies find that women may be more likely to post non-publicly [31], express more concern about their privacy on Facebook [87], and are also more likely to be exposed to sensitive ad categories [17]. Privacy concerns on Facebook also seem to differ based on political, sexual, and religious affiliation [1, 13, 27, 64].

Thus, prior research indicates that there are differences in privacy concerns and behaviors among different Facebook users [83, 104], but studies often looked at specific demographic or identity characteristics in isolation. In our study, we conduct exploratory analysis on how expectation-setting mismatches might relate to multiple demographic factors.

2.3 Privacy Settings on Social Media

One way for individuals to address their privacy concerns is to make use of platforms’ privacy settings. Privacy settings on social media, including Facebook, comprise *interpersonal privacy settings* that allow users to regulate privacy exposure towards other users (e.g., post or profile visibility) and *institutional privacy settings* to configure how data is collected or used by the platform and/or third parties (e.g., opting out of personalized advertising) [80, 109]. Below, we review past work that has evaluated social platforms’ interpersonal privacy settings and institutional privacy settings.

Interpersonal privacy settings. Social media users often struggle to balance disclosing personal information with shielding their information from unwanted access by others [4, 28, 42, 51]. Users adopt various strategies to protect their personal information from interpersonal harm, such as declining friend requests from strangers, deleting old content, and managing audiences [55, 109]. Users configure interpersonal privacy settings according to their individual priorities. Wisniewski et al. identified six privacy management strategies: privacy maximizers, selective sharers, privacy balancers, self-censors, time savers/consumers, and privacy minimalists [105].

Institutional privacy settings. Additionally, many users now also hold institutional privacy concerns while navigating social platforms. These concerns arise from mistrust in entities rather than individuals and influence how users share personal information with corporations, advertisers, and data intermediaries online [76]. Institutional privacy settings often relate to advertising and specifically online behavioral advertising (OBA) [30, 39, 40]. Many platforms rely on OBA for revenue [9, 29, 32], i.e., tracking individuals’ online and offline activities [62, 101] to target them with ads customized to their inferred interests [14]. Users tend to be aware of social platforms’ online advertising, but do not fully understand OBA practices [34, 39]. A Pew survey found that most Facebook users did not know ad tracking categories existed and were uncomfortable with Facebook creating a list of categories about them [43]. Users are largely unaware of the existence of Facebook’s OBA settings [39, 44], suggesting issues of discoverability. Additionally, there are usability issues regarding the setting’s location, layout, and explanations [38, 39]. Recent work has contributed to design guidelines for more usable OBA settings [86], such as placing ad

settings at the top of one’s feed with links to more granular controls [47]. However, even users who have adjusted OBA settings report feeling indifferent about their adjustments’ effect—likely stemming from an overall lack of trust in Facebook’s claims [41].

Misalignments in Facebook privacy settings. A small number of studies have investigated mismatches between users’ expectations of and their currently enacted privacy settings [44, 50, 58, 61]. Liu et al.’s 2011 study with 200 Facebook users found that settings match users’ expectations only 37% of the time, and when incorrect, almost always exposed content to more users than expected [58]. Similarly, Madejski et al.’s 2012 study found mismatches between users’ sharing intentions and reality, with almost every participant in their study having at least one mismatch [61].

Our study also investigates potential disconnect between users’ privacy concerns and behaviors, building on closely related studies in several ways. (1) *Methodological differences:* compared to previous survey-based studies that relied on participants’ recall to hypothetical scenarios [50, 61], we collected participants’ actual privacy settings directly from their Facebook accounts and their expectations and level of concern for those settings with a custom browser extension. (2) *Different focus for mismatches:* for example, compared to Hsu et al.’s study [44], in which misalignments focused on a setting’s existence (“On Facebook, do you think a setting or group of settings exists to...?”), our study investigated misalignments regarding a setting’s collected value, providing new insights on participants’ true exposure, not just on their (un)awareness of settings. (3) *Broader coverage of settings:* Most prior work like Liu et al. [58] focused on interpersonal privacy settings, while our study investigates current Facebook settings for both interpersonal and institutional privacy settings, including some previously not investigated settings that relate to ad targeting and other OBA practices. (4) *Sample diversity:* We recruited a diverse sample of participants compared to smaller, convenience samples in some prior work [58, 61]. This allowed us to investigate differences in expectations and mismatches across demographic factors. (5) *Contemporary assessment:* Most closely related prior studies have been conducted over a decade ago [45, 58, 61]. Facebook’s privacy settings and public awareness of privacy issues have changed significantly since then. Our study provides a contemporary assessment of expectation-setting mismatches for Facebook’s current privacy settings and ad profiling practices in light of the increased public awareness, media attention, and regulatory scrutiny regarding privacy issues in recent years.

3 Methods

To answer our research questions, we conducted a mixed-method study in which participants were first asked to install a browser extension that collected their Facebook privacy setting values (with their consent) from their Facebook account, then we live embedded the retrieved values in an online survey, in which we asked participants for their expected values of 18 Facebook privacy settings, before showing them their actual settings and asking them to rate their level of concern with these settings.

Our study was approved by our Institutional Review Board (IRB). We first describe our browser extension and a qualitative pre-study to evaluate and refine the browser extension and study protocol; we

then describe the studied privacy settings, our main survey protocol and how it integrated with the browser extension, data analysis, and study limitations. All our study materials are available in an OSF repository.¹

3.1 Browser Extension System

To gather participants' Facebook privacy settings from their Facebook account, we developed a browser extension for Chrome and Firefox, and a backend server, using Django, Gunicorn, and Nginx. The extension initially checks if the user is logged into Facebook and asks them to log in if not. Once the extension detects the user is logged in, it displays a button that the user clicks to start data gathering. The extension gathers only the user's Facebook privacy settings, ad settings, and ad profile data; no other data from the user's Facebook account is collected. To do so, the extension opens browser tabs of the respective Facebook settings pages, identifies setting values using HTML classes and text descriptions, and stores the user's setting values. Once all setting values and ad profile data are collected, the extension opens a final tab with an aggregated table of the extracted settings. Participants could inspect their data before clicking a button to submit it to our server.

We tested the browser extension's reliability internally on a variety of different Facebook accounts and further confirmed that it functioned properly as part of our pre-study (see Section 3.2).

Ethical considerations. We developed the browser extension to only extract specific Facebook settings and ad profile information. The browser extension extracted no other personal information or behavioral data. We also let participants inspect their data before sending it to our server.

3.2 Qualitative Pre-study

Facebook has extensive visibility, privacy, and ad settings [47]. To strike a balance between comprehensive coverage of Facebook's settings and not overwhelming participants, we first conducted semi-structured interviews with our browser extension, allowing open exploration of different settings and participants' respective expectations. Based on these pre-study findings we selected settings to investigate in our main study (see Section 3.3) and refined our survey protocol (see Section 3.4).

Interview protocol. We interviewed 15 adult U.S. Facebook users, recruited through university research pools, in May to July 2022. We stratified our pre-study sample based on screening survey responses in terms of gender and ethnicity. Participants received \$15 compensation. The remote interviews lasted one hour. The pre-study materials are available in our OSF repository.¹

We first asked about participants' Facebook use, before asking about their perceptions and trust in Facebook, and perceived privacy issues. Next, participants ran our browser extension. Once the extension generated the final tab with all extracted settings, we asked participants to explore their settings and share their reactions (think aloud) regarding what, if any, settings or values were surprising to them or did not match their expectations. For inferred ad topics and companies who had advertised to them (ad companies), we asked participants to explore the lists and discuss any entries

that seemed surprising, invasive, or inaccurate; and those that were accurate or matched expectations. We concluded by asking about their overall reactions to their Facebook settings and ad profile.

Pre-study findings. Pre-study findings indicated that asking about all Facebook privacy and ad settings in our main survey would overwhelm and fatigue participants. Thus, we used the pre-study findings to select relevant settings. Our findings showed that the timeline and advertising settings often caused surprise and setting values frequently did not match participants' expectations. In contrast, general settings were often set to what was expected and even when misaligned elicited little concern. Because of this, we included the timeline and ads settings in the main survey to quantitatively investigate potential (mis)alignments (see Section 3.3). For the general settings, we included those that caused the most surprise to participants. These were related to people, Pages, or lists a user follows, friend requests/list, email address, and phone number (see Table 1).

Furthermore, pre-study interview participants were consistently concerned or surprised by the often extensive lists of ad topics and ad companies, as well as many of the specific companies and ad topics listed. However, the pre-study also made it clear that asking participants about all their ad topics and companies in a survey would be infeasible due to the length of these lists. Subsampling also did not appear promising as there was little variance in pre-study participants' reactions across different ad companies and topics (surprise and concern were generally high). Hence, we focused on whether their total number of assigned ad topics and ad companies was lower, higher, or about as they expected.

3.3 Studied Facebook Privacy Settings

As described in Section 3.2, based on the pre-study interviews, we selected 18 Facebook privacy settings from three settings categories (general, timeline, ad settings) for our main survey (see Table 1).

Under *general settings*, we looked at visibility settings for what or who a user follows, the user's friends list, and who can look up a user based on their email address or phone number, as well as who can send friend requests. All these settings relate to the general exposure of one's Facebook account to other people and entities. For instance, access to friends lists was a key aspect of the Cambridge Analytica incident [48]. Notably, all of these are set to public/everyone by default.

Timeline settings focus more specifically on visibility of user activity. We selected settings for visibility of a user's future posts, who can post on the user's profile, who can see what others post on a user's profile, who can see posts a user was tagged in and whether the user has to review them before tagged posts appear on the user's profile, and whether others' tags on the user's post have to be reviewed before they appear. These are primarily settings to control how others can see and interact with a user's posts. We were interested to see whether participants' knowledge of their exposure, as well as their actual exposure, has improved in comparison to much older studies on this aspect [37, 58, 61].

General and timeline settings relate to interpersonal privacy, for which related work suggests that people have a better understanding of respective risks and settings [50, 80, 109]. Our study

¹OSF repository with all study materials: https://osf.io/bt4cr/?view_only=2ac6d36899a24de59b99c86bae195e5e

Table 1: Overview of the Facebook privacy settings investigated in the main survey.

General Settings	Short name	Default	Description
Who can see the people, Pages, and/or lists you follow?	<i>visibility_following</i>	Public	Visibility of people, pages, and lists a user follows on FB.
Who can send you friend requests?	<i>send_friend_request</i>	Everyone	Others' ability to connect with a user.
Who can see your friends list?	<i>visibility_friends</i>	Public	Visibility of a user's list of friends on the platform.
Who can look you up using the email address you provided?	<i>visibility_email</i>	Everyone	Ability to find a user's profile based on their email address.
Who can look you up using the phone number you provided?	<i>visibility_phone</i>	Everyone	Ability to find a user's profile based on their phone number.
Timeline Settings			
Who can see your future posts?	<i>visibility_future_post</i>	Friends	Visibility of the user's posts to others.
Who can post on your profile?	<i>post_your_profile</i>	Friends	Others' ability to make posts on a user's profile.
Who can see what others post on your profile?	<i>visibility_others_post</i>	Friends	Visibility of posts made by others on a user's profile.
Who can see posts you're tagged in on your profile?	<i>visibility_tagged_post</i>	Fr. of friends	Visibility of posts a user is tagged in on their profile.
Review posts you're tagged in before they appear on your profile?	<i>review_tagged_post</i>	Off	Require a user's approval before posts they are tagged in appear on a user's profile.
Review tags people add to your posts before they appear on FB?	<i>review_tags_on_post</i>	Off	Require a user's approval of others' tags added to a user's post.
Ad Settings			
Show advertisements based on your job title?	<i>ad_job</i>	Yes	Use of a user's job information for ad targeting.
Show advertisements based on your employer?	<i>ad_employer</i>	Yes	Use of a user's employment information for ad targeting.
Show advertisements based on your relationship status?	<i>ad_relation</i>	Yes	Use of a user's relationship status for ad targeting.
Show advertisements based on your education?	<i>ad_education</i>	Yes	Use of a user's education information for ad targeting.
Who can see your social interactions alongside ads?	<i>ad_social_interactions</i>	Friends	Visibility of a user's interactions with Pages and events in ads.
Allow FB to show you personalized ads based on data from their third-party partners?	<i>ad_partners</i>	Yes	Use of data about a user provided by third parties for ad targeting.
Allow advertisers to show you ads on other platforms based on information FB has collected on your interests?	<i>ad_off_fb</i>	Yes	Use of FB's data about a user for ad targeting outside of FB.

investigated the extent to which participants' expectations regarding these settings matched their actual exposure.

There is limited prior work on institutional privacy issues and settings, e.g., regarding OBA, suggesting indifference or lack of understanding of such practices by users [41, 43, 78]. An open question is whether, when informed about the existence of such settings, people's expectations are accurate or misaligned. Thus, we included multiple potentially sensitive *ad settings*, specifically, whether a user's job, employer, relationship status, and education information can be used for ad targeting; and who can see a user's social interactions with ads. We also included settings on ad targeting based on data about the user provided by external parties to Facebook and ad targeting on other websites based on a user's Facebook data. These settings are relevant as they pertain to data practices that may be perceived as violations of contextual integrity [70]. Additionally, we asked participants about their perception of the number of companies that have targeted a participant based on their Facebook activity or information (*ad companies*) and the *ad topics* assigned to a user by Facebook.

3.4 Main Survey Protocol

The main survey protocol was informed by the pre-study interviews (conducted May to July 2022) and further refined through multiple rounds of informal pilot testing to remove ambiguities or comprehension issues. The main survey study was then conducted in January and February 2023. The final survey script is provided in our OSF repository.¹ In the survey, participants were first asked about their trust in Facebook (questions based on [19]) and their level of concern and perceived control regarding privacy on Facebook (based on [19, 63, 97, 100]). Next, participants were instructed to install and run our browser extension, which gathered their Facebook privacy settings and ad profile data from their Facebook account (see Section 3.1). Then, for each setting in Table 1, participants were shown the setting as a question (e.g., "Who can see your future posts?") with the possible setting values as response options. We further included an "I don't know" option

to let participants express uncertainty about their current setting rather than forcing them to guess. After selecting a response option, participants were then shown their response (expected setting) and their actual setting collected from their Facebook account together (e.g., participant chose "friends," actual setting is "public") and asked to rate their level of concern regarding their actual setting value on a 5-point scale ("not at all concerned" to "strongly concerned"). Throughout, we emphasized that these questions were not to test their knowledge, that responses would not affect compensation, and that participants should answer based on what they think their current setting is without looking it up. We also inserted two attention checks to confirm whether participants were paying attention to instructions and their actual settings shown.

For companies that had advertised to them (ad companies) and ad topics, we showed participants the total number for each and asked them to rate how the number compared to their expectations on a 5-point scale ("significantly lower than I expected" to "significantly higher than I expected"). At the end, participants were asked the same questions about trust, privacy concern, and perceived control regarding Facebook as at the beginning. This was done to assess whether interacting with the privacy settings data collected from their actual accounts had impacted their perceptions of Facebook.

3.5 Recruitment and Participants

We recruited U.S. participants through Prolific, which has shown to be reliable for studying privacy experiences [75, 81, 94]. Participants first completed a screening survey (same as the one used in the pre-study, see OSF repository¹), for which they were compensated \$1.50. The screening survey asked about participants' demographics, general technology use, and Facebook use. We collected a range of demographic variables to be able to explore potential correlations between these variables and participants' privacy settings and expectations. For sensitive demographic questions, such as those related to income and sexuality, we provided a 'prefer not to say' option to respect participants' privacy.

To qualify, participants needed to be at least 18 years old and physically located in the United States. They were required to have an active Facebook account set up in their name, used for at least a week, and to use Facebook at least once a week. Those without an active Facebook account or who used Facebook rarely or never were excluded. Additionally, participants had to be able to use the Google Chrome or Mozilla Firefox web browser on a desktop or laptop computer and log in to their Facebook account through these browsers. Participants who could not meet these criteria were excluded. Based on screening survey responses, we invited 400 participants to complete our main survey, with \$13.50 as compensation. 201 participants completed the main survey. After removing those who failed attention checks, our sample had 195 participants. Average completion time was 24.14 min. (median: 21.22 min.).

Table 2 shows our sample’s demographics. Our sample was gender balanced with slightly more women (53%). The average age was 37.43 years (Median=34, range=18–75); for our analysis we grouped participants into generations (22% Gen Z, 47% Millennials, 22% Gen X, 8% Baby Boomer / Silent Generation), following generational boundaries suggested by Pew Research Center [25]. Our sample was ethnically diverse: 39% identified as Caucasian, 19% as Asian (including South Asian), 11% as Black or African American, and 11% as Hispanic, Latine, or Spanish origin. Mixed-race participants, e.g., those identifying with two or more races, constituted 18% of the sample. Regarding education, 41% had a Bachelor’s degree, 32% had attended some college (without degree) or held an Associate’s degree, 14% had a postgraduate degree, and 12% had a high school diploma or equivalent (12%).

We also asked about political affiliation (57% Democrat, 11% Republican, 28% Independent) and religion (majority identified as non-religious/non-believers (43%) or Christian (38%); remaining participants identified as Buddhist (3%), Jewish (3%), Muslim (3%), or Hindu (2%), which we grouped for statistical analysis as Non-Christian Religions).²

In terms of Facebook usage, most participants (89%) had their account for over five years; 9% had an account for less than five years. Most reported using Facebook daily (48%) or multiple times per week (36%), suggesting a high level of regular interaction with Facebook. Most (54%) reported rarely posting (28% few times per month, 12% multiple times per week; 6% daily).

Some demographic variables collected, such as income and sexual orientation, were ultimately not included in our statistical analysis due to insufficient responses for these categories, limiting our ability to draw meaningful conclusions for them.

3.6 Data Analysis

For RQ1 (mismatches between expected and actual settings), we report descriptive statistics and present comparisons between expected and actual settings using bubble charts.

For RQ2 (demographic differences), we conducted a series of regression analysis to explore potential relationships between demographic factors, Facebook usage patterns, and mismatched expectations across various privacy settings. Specifically, we ran logistic

²Note that grouping non-Christian religions was done to enable statistical analysis without having to exclude those participants due to small sample sizes for these religious groups. We do not mean to suggest that these religions are the same/similar or that their adherents share the same experiences.

Table 2: Demographics of participant sample (n=195).

Category	Total
Age	
Gen Z (18–26)	22%
Millennials (27–42)	47%
Gen X (43–58)	22%
Baby Boomer / Silent Generation (59+)	8%
Not Specified	1%
Gender	
Man	46%
Woman	53%
Non-binary	0.5%
Prefer not to answer	0.5%
Race	
Caucasian	39%
Asian (incl. South Asian)	19%
Mixed Race	18%
Black or African American	11%
Hispanic, Latine, or Spanish origin	11%
Prefer not to answer	2%
Education	
High School or equivalent or less	13%
Some College or Associate’s degree	32%
Bachelor’s degree	41%
Postgraduate degree	14%
Political Affiliation	
Democrat	57%
Republican	11%
Independent	28%
Other	3%
Prefer not to answer	2%
Religion	
Christian (includes Protestant, Catholic, etc.)	38%
No religion, not a believer (includes atheist, agnostic)	43%
Non-Christian Religions (Buddhist, Hindu, Jewish, Muslim, other)	15%
Prefer not to answer	4%

regression models for each of the 18 expectation-setting misalignments identified in Section 3.3 (See Table 1).

Our analyses incorporated data from 178 participants, selected from an initial pool of 195, after excluding responses with missing values to ensure the integrity of our regression models. We employed listwise deletion, based on the assumption that missingness occurs completely at random (MCAR) to preserve robustness [57]. Categorical independent variables were coded as dummy variables, with the most prevalent category as the reference (e.g., for age, Millennials (27–42 years old) was our reference category). For the number of ad topics and ad companies, we ran linear regressions with the same independent variables.

For RQ3 (participants’ concerns about expectation-setting misalignments), we report descriptive statistics on participants’ levels of concern for individual settings. We used Wilcoxon Signed-Rank tests to evaluate differences among participants with matched and mismatched expectations, as well as changes in trust and privacy concerns at the beginning and end of the study.

3.7 Limitations

Our study design has some limitations. There may be self-selection bias, however, our recruitment message was vague about the study’s focus (“Experiences and Attitudes Towards the Facebook Platform”) and did not mention privacy. Our study also required installing a browser extension, which may have affected who chose to participate, as installing unknown software may be perceived as risky.

We went through extensive pilot testing to reduce participant concerns and provide transparency about the extension. We asked about settings in the same order for all participants, but did not observe ordering effects in the data. We investigated 18 Facebook settings, which is more than prior studies, but still less than all privacy-related Facebook settings. Expectation (mis)matches may look different for other Facebook settings. Through our pre-study interviews we identified settings of interest, these settings could possibly exhibit more mismatches, yet they are also settings with high consequence for people’s privacy and therefore relevant. In contrast to most prior work, we gathered participants’ actual Facebook settings from their Facebook accounts through our custom browser extension, however, we have no data on whether participants have previously interacted with or changed a specific setting. Asking participants to recollect whether they had changed a specific setting seemed too unreliable given recall bias and Facebook’s changing settings interfaces. Similarly, we do not know if a participant’s setting value may reflect Facebook’s default value for a given setting at the time the participant created their account, as Facebook has repeatedly changed the defaults for privacy settings. Nevertheless, our study provides comprehensive and contemporary insights on the extent to which participants’ current expectations of their Facebook privacy settings match their actual settings and how that relates to the default values at the time of the study. Future work could consider conducting longitudinal studies to investigate how individuals interact with and change their Facebook settings over time. Our sample consisted only of U.S. Facebook users—results in other regions and for other social media platforms may differ. However, in 2024, Facebook is still the most popular social media platform in the U.S. [106]. While we attempted to recruit a diverse sample, compared to Pew’s data on the characteristics of Facebook users [77], our sample was more educated, likely due to Prolific’s participant pool [95].

Our sample size had to strike a balance between overall cost (\$15 per participant) and test sensitivity (specifically for RQ2). Thus, our sample size is sufficient to detect medium effects in linear regressions, but may be underpowered to detect medium effects in logistic regressions. We still see value in reporting our logistic regression results for RQ2 as they indicate interesting/concerning demographic differences, but we consider them exploratory. Further research is needed to confirm the identified demographic differences.

4 Results

Next, we present our findings organized by research questions.

4.1 Expectation-Setting Mismatches (RQ1)

We examined participants’ expected settings versus their actual settings across Facebook settings in three categories: *General*, *Timeline*, and *Ad* settings (see Table 1 for all settings). Note that we classified “I don’t know” responses as mismatches because they indicate a lack of awareness or understanding of the actual privacy settings in use. We consider this lack of awareness in itself a privacy concern, as it suggests users may be unknowingly exposed to privacy risks.

4.1.1 Frequency of mismatches. Across the 18 settings we analyzed, all participants had at least one mismatch (100%). Participants had up to 11 mismatches, and 5.74 mismatches on average (SD=2.3,

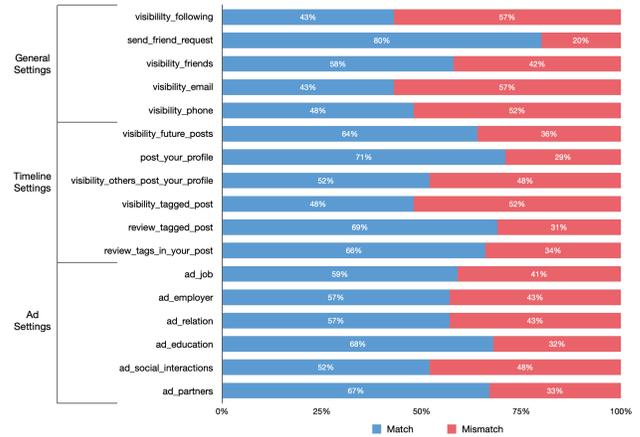


Figure 1: Expectation-setting mismatch ratio per analyzed setting, grouped by settings category.

median: 6); 13% (25) had nine or more mismatches, meaning that their expectation of current settings was inaccurate (i.e., either more or less private than expected) for at least half of the settings we analyzed.

The most mismatches occurred in the *Timeline* settings category: 61% of participants had a mismatch for at least one out of the six *Timeline* settings (median: 2; max: 5). 45% had at least one mismatch out of the six *General* settings (median: 2; max: 5). For *Ad* settings, 39% of participants had mismatches (median: 1; max 4).

Figure 1 shows the match-mismatch ratio per setting.³ “Who can send you friend requests” (*send_friend_request*) stands out with an 80% match rate. For all other settings, the mismatch rate was 30–57%, indicating substantial discrepancies between expectations and the actual settings collected from participants’ accounts.

4.1.2 Direction of mismatches. Consistent with Liu et al.’s findings [58], our participants often assumed their settings to be more restrictive than they were. For instance, for *visibility_email* and *phone*, participants often assumed “only me” when the actual settings allowed ‘Friends’ or even ‘Everyone’ to look them up. The same applied to most timeline settings, such as *visibility_tagged_post*, for which participants generally expected visibility to be restricted to ‘Friends,’ but for 43 (22%) participants it was actually set to ‘Friends of Friends.’

As a result, we can see that Facebook’s default settings often do not align with user expectations or the actual settings collected from their accounts. For example, the default at the time of our study for *visibility_email* and *visibility_phone* setting was ‘Everyone,’ whereas participants frequently expected these settings to be ‘only me.’ This discrepancy suggests that Facebook’s default settings are generally more open than what users expect or prefer, leading to potential privacy concerns.

³For some settings, we were unable to retrieve values for a small number of participants due to Facebook’s A/B testing, where those participants had a different version of Facebook’s settings where these specific settings no longer existed. The affected settings were *visibility_friends*, *ad_social_interactions*, *ad_partners*, and *ads_off_fb*. After excluding these participants, the total number of participants for each respective setting was 138 for *visibility_friends*, 189 for *ad_interactions*, 158 for *ad_partners*, and 97 for *ads_off_facebook*.

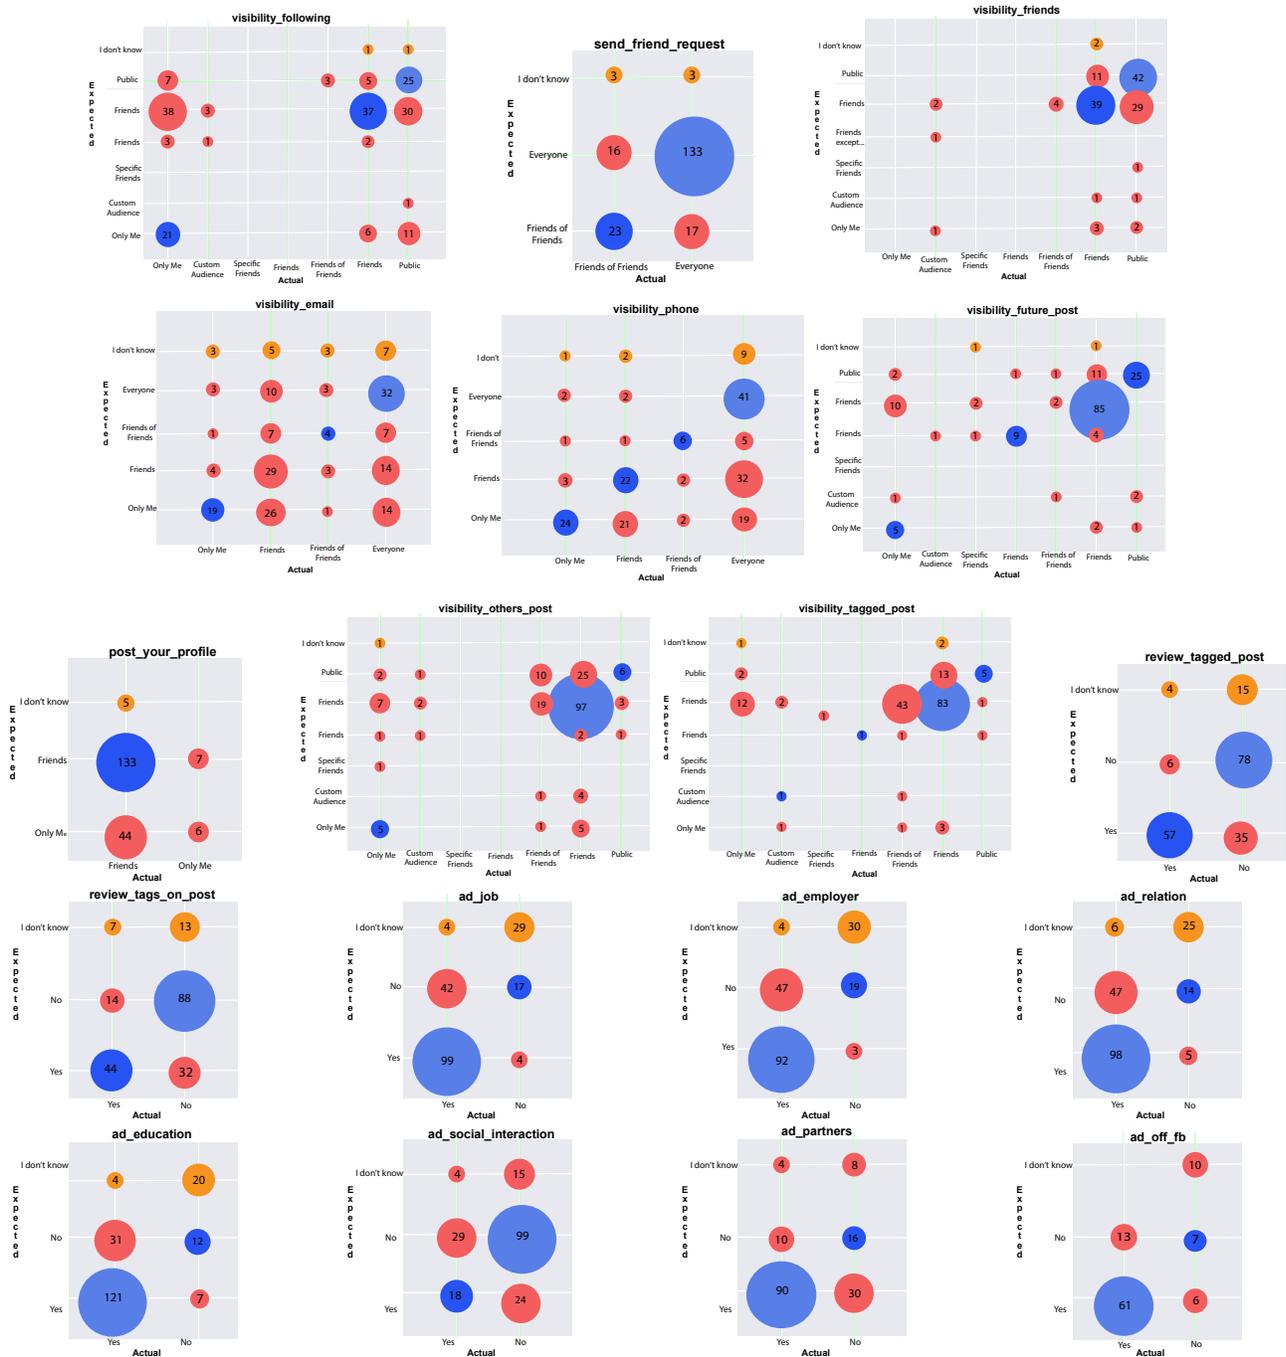


Figure 2: Mappings and frequency of participants’ expected and actual settings across General Settings, Timeline Settings, and Ad Settings (blue: expectation-setting match; light blue: default value; red: mismatch; orange: ‘I don’t know’ responses).

For ad settings—which were not covered by Liu et al.’s study [58]—we observed a similar trend, as participants frequently underestimated the extent to which their personal attributes are used for

ad targeting. For example, 47 (24%) participants erroneously believed their relationship status was not used for ad targeting. Similarly, 30 (15%) participants mistakenly believed that ads would not be targeted based on data about them from Facebook’s partners (*ad_partners*). Here again, Facebook’s default settings at the time of

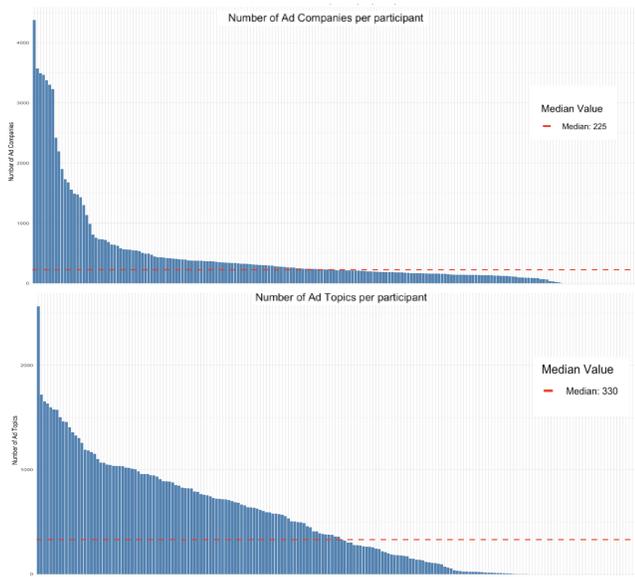


Figure 3: Number of ad companies per participant that had targeted them on Facebook (top) and the number of ad topics per participant (bottom).

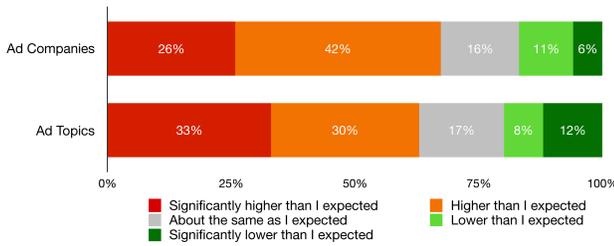


Figure 4: Participants’ perception of their number of ad companies (top) and ad topics (top).

our study were more permissive than many participants expected, highlighting a need for greater transparency and user awareness regarding ad targeting practices.

Expectation-setting mismatches can exist in both directions. For example, for *visibility_following*, most participants expected it to be ‘Friends,’ but for many, the actual setting was either the more restrictive ‘Only me’ (38, 20%) or the more liberal ‘Public’ (30, 15%). According to Altman’s boundary regulation theory both phenomena constitute privacy issues: undesired exposure or social isolation [8, 72].

4.1.3 Ad companies and ad topics. We find a high variance in the number of ad companies that have targeted participants and ad topics assigned to participants (see Figure 3). The number of ad companies on Facebook that had included participants in their audience ranged from 0 to 4,377 companies (mean: 441; median: 330). Most (68%) participants rated their number of ad companies to be “higher” or “significantly higher” than expected (see Figure 4).

Similarly, the number of ad topics Facebook had inferred per participant ranged from 0 to 2,565 (mean: 473.7; median: 330). Most (63%) indicated this to be “higher” or “significantly higher” than expected. These results suggest that participants were largely unaware of the breadth of ad-related inferences Facebook made about them and how many advertisers have targeted them.

Summary. Our findings show that the alignment between participants’ expected and actual collected Facebook settings is generally poor. Many participants believed their settings were more restrictive than they were, revealing a widespread misconception among Facebook users regarding how private their personal information is on Facebook. For ad settings, participants often underestimated how their information is utilized.

4.2 Demographic Differences (RQ2)

As described in Section 3.6, we ran a regression model for each of the 18 privacy settings and for the number of ad topics and ad companies to examine whether and how demographic factors and Facebook usage might explain variances in expectation-setting misalignments. We employed logistic regression models for binary outcomes related to expectation-setting mismatches, and linear regression models for the number of ad topics and ad companies associated with user profiles.

In the following, we summarize significant findings across those regression models for specific demographic variables to more clearly surface potential demographic effects, e.g., we discuss all significant findings for “age” together instead of going setting by setting.

Age: Older participants had more mismatches. Treating age generations as a categorical variable, we observed significant age differences for one *general setting*, two *timeline settings*, and most *ad settings* (See Table 1).

Gen X participants were more likely to have mismatched expectations for ad targeting based on job title (*ad_job*: $OR_{Mill}^{Gen X}=2.54, p=.04$), employer (*ad_employer*: $OR_{Mill}^{Gen X}=2.87, p=.021$), and third-party data (*ad_partners*: $OR_{Mill}^{Gen X}=3.42, p=.04$).

Gen X participants also had less exposure to ads, as they had much fewer ad companies (*ad_companies*: $\beta_{Mill}^{Gen X}=-413.19, p=.003$) and ad topics (*ad_topics*: $\beta_{Mill}^{Gen X}=-217.376, p=.02$). Gen X participants were also more likely to have mismatched expectations toward ads being shown off-Facebook (*ads_off_fb*: $OR_{Mill}^{Gen X}=5.66, p=.04$) and the audience of the posts they are tagged in (*visibility_tagged_post*: $OR_{Mill}^{Gen X}=2.74, p=.04$).

Baby boomers were also more likely to have mismatched expectations for the same two settings (*ads_off_fb*: $OR_{Mill}^{Boomers}=15.06, p=.01$; *visibility_tagged_post*: $OR_{Mill}^{Boomers}=4.82, p=.03$). Gen Z participants were less likely to have mismatched expectations for the visibility of people/pages/lists they follow than Gen Y participants (*visibility_following* ($OR_{Mill}^{Gen Z}=0.34, p=.03$)). However, Gen Z participants were more susceptible to mismatched expectations for tag-related settings (*visibility_tagged_post*: ($OR_{Mill}^{Gen Z}=6.27, p=.001$); *review_tags_in_your_post*: $OR_{Mill}^{Gen Z}=2.90, p=.04$).

Ethnicity: Asian, Mixed-Race, and Hispanic participants had more mismatches. With ethnicity as a categorical variable, we observed

significant differences for one general setting, two timeline settings, and three ad settings.

Hispanic participants were more likely to have a mismatch for the use of relationship status in ad targeting ($ad_relation: OR_{Cauc.}^{Hisp.} = 3.77, p=.032$) and for $ad_social_interactions$ ($OR_{Cauc.}^{Hisp.} = 0.23, p=.04$).

Participants in the Mixed Race group were more likely to have a mismatch regarding $visibility_tagged_post$ ($OR_{Cauc.}^{Mixed} = 3.66, p=.03$). Mixed Race and Asian participants were also more likely to have mismatches for ads_off_fb ($OR_{Cauc.}^{Mixed} = 9.73, p=.01$; $OR_{Cauc.}^{Asian} = 8.18, p=.04$).

Furthermore, Asian participants were more likely to experience mismatches regarding who can see their future posts ($visibility_future_post: OR_{Cauc.}^{Asian} = 3.64, p=.01$) and for $visibility_following$ ($OR_{Cauc.}^{Asian} = 3.88, p=.01$).

Political affiliation: Republicans had more mismatches. Treating political affiliation as a categorical variable, our results suggest that political affiliation could explain variances in mismatches across multiple timeline and ad settings, as well as the number of ad topics represented.

Republicans were more likely to have mismatches for $visibility_future_post$ ($OR_{Dem}^{Rep} = 7.68, p<.001$), $review_tagged_post$ ($OR_{Dem}^{Rep} = 4.55, p=.01$), $ad_relation$ ($OR_{Dem}^{Rep} = 3.91, p=.02$), and $ad_education$ ($OR_{Dem}^{Rep} = 4.53, p=.01$). Independents had a higher number of ad companies ($\beta_{Dem}^{Ind} = 398.23, p<.01$) and were more likely to have mismatches for $post_your_profile$ ($OR_{Dem}^{Ind} = 2.62, p=.02$). However, they were less likely to have mismatches for $visibility_email$ ($OR_{Dem}^{Ind} = 0.42, p=.02$).

Gender: Men had more mismatches and fewer ad topics. Treating gender as a binary variable,⁴ we find significant differences for two timeline settings and the number of ad topics. Men are significantly more likely to have mismatches for $visibility_tagged_post$ ($OR_{Dem}^{men} = 6.04, p<.01$) and $visibility_others_post$ ($OR_{Dem}^{men} = 2.59, p<.01$). However, men are likely to have fewer ad_topics ($\beta_{women}^{men} = -166.79, p=.02$).

Education: Limited impact on mismatches. With education as a categorical variable, we observed no significant differences for most settings. Compared to participants with a Bachelor's, those with some college or an Associate's degree are more likely to have mismatches for ads based on third-party data ($ad_partners: OR_{Bach.}^{Asso.} = 3.64, p=.02$). Participants with a postgraduate degree had been targeted by more $ad_companies$ ($\beta_{Bach.}^{PostGrad} = 678.38, p<.01$).

Religion: Not predictive of mismatches. Treating religion as a categorical variable, we did not find religion to be a significant predictor of any setting mismatches in our sample. However, these non-significant results should be interpreted with caution, as they may be influenced by our grouping of Non-Christian religions due to sample size limitations. This grouping could potentially obscure the distinct privacy concerns and behaviors associated with different religious affiliations.

⁴We received one non-binary response in our dataset, lacking statistical power to draw meaningful conclusions about non-binary Facebook users, we had to exclude them from the regression analysis.

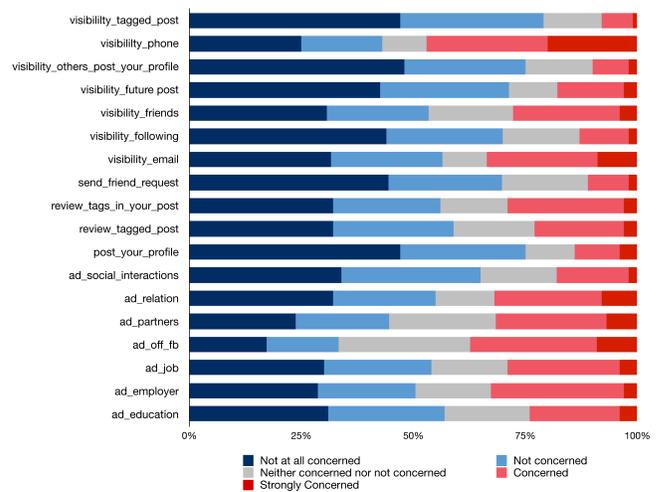


Figure 5: Participants’ level of concern (not at all concerned to strongly concerned) regarding their actual setting value for each of the 18 analyzed Facebook privacy settings.

Facebook usage: Mismatches linked to post frequency. We treated Facebook post frequency, account age, and usage frequency as categorical variables in our regression analysis.

Participants who post more frequently are more likely to experience mismatches in their expectations across a variety of settings. Those who post daily are more likely to have mismatches for $visibility_following$ ($OR_{Rarely}^{Everyday} = 6.52, p=.04$) and for $ad_relationship$ ($OR_{Rarely}^{Everyday} = 14.17, p=.02$). Those who post a few times a month experience more mismatches for $visibility_phone$ ($OR_{Rarely}^{Few\ times/month} = 0.36, p=.01$) and for $visibility_future_post$ ($OR_{Rarely}^{Few\ times/month} = 3.39, p=.03$). Interestingly, those who have never posted on Facebook, when contrasted against those who post rarely, are targeted by a larger number of $ad_companies$ ($\beta_{Rarely}^{Never} = 1203.509, p<.01$).

For account age, we find significant differences only for two settings: Those who had their account for less than five years had more mismatches for $visibility_others_post$ ($OR_{Over\ 5\ years}^{Less\ than\ five\ years} = 0.22, p=.02$) and for $ad_education$ ($OR_{Over\ 5\ years}^{Less\ than\ five\ years} = 3.93, p=.03$).

For account usage frequency, we only find significant differences for number of ad topics. Those who use Facebook a few times a month have significantly fewer ad_topics ($\beta_{Daily}^{Few\ Times/month} = -469.71, p<.001$). Those who use Facebook multiple times per week have fewer ad_topics ($\beta_{Everyday}^{Few\ times\ a\ week} = -351.97, p<.001$).

Summary. Our exploratory analysis indicates that demographic differences likely exist in Facebook users’ expectation-settings mismatches, with participants’ age, race, gender, and political affiliations potentially being relevant factors that should be investigated further. Education and religion appear to have limited impact on mismatches. Additionally, our results indicate that posting frequency is likely correlated with mismatches for some settings.

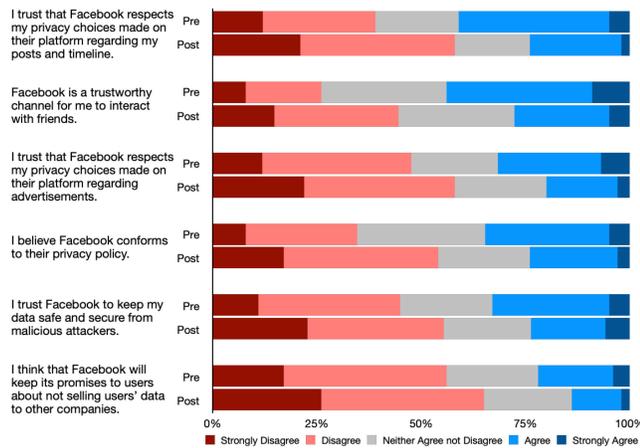


Figure 6: Trust in Facebook pre and post study.

4.3 Participant Concerns and Trust (RQ3)

We discuss participants’ concerns with their actual Facebook settings and effects on reported trust in Facebook.

4.3.1 Elevated concerns toward mismatched settings. Figure 5 shows participants’ concern ratings for each setting. The mean level of concern across settings was 2.35 (SD=1.26, median: 2), indicating that participants were generally neutral to not concerned about the value of their settings. Comparing matched and mismatched expectations, we see a subtle but not significant shift in concern. For mismatches the mean concern was slightly higher (mean: 2.37, SD=1.25, median: 2); for matches marginally lower (2.34, SD=1.27, median: 2).

Looking at individual settings, we find that *ads_off_fb* evoked the highest overall level of concern ($M=2.94$). In contrast, for *visibility_following* overall concern was the lowest ($M=1.88$), suggesting users are generally less concerned about the visibility of their social connections.

Looking at mismatched expectations, participants were most concerned about mismatches for *visibility_phone* ($M=3.58$) and *ads_off_fb* ($M=3.06$); and least concerned about mismatches for *visibility_following* ($M=1.95$) and *visibility_others_post* ($M=1.95$).

4.3.2 Seeing one’s actual settings decreased trust. We compared participants’ trust in Facebook at the beginning and end of the survey (see Figure 6). The questions in Figure 6 captured institutional trust in Facebook across six dimensions: trust in Facebook to keep its promises about data selling [19], to secure data from malicious attackers, to adhere to its privacy policy [65], to respect privacy choices for advertisements and timeline posts, and to serve as a trustworthy channel for interaction with friends [19].

After participants saw their actual privacy settings, we observed a notable decline in their trust across various dimensions of Facebook’s platform, from its data-selling promises to advertisement preferences. The mean level of trust in Facebook decreased from 2.88 (SD=0.94) to 2.48 (SD=0.97). Wilcoxon signed-rank tests showed that the decline in trust is significant for each of the six trust items (all $p < 0.001$).

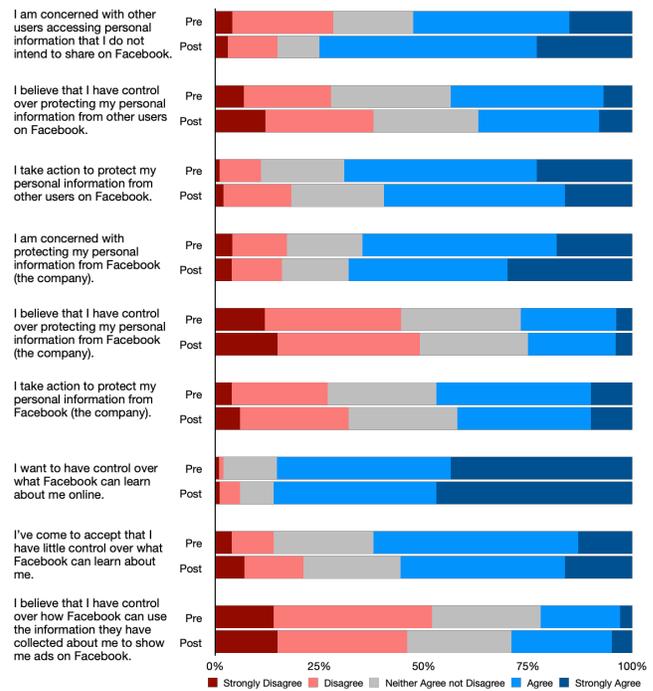


Figure 7: General privacy concerns regarding Facebook pre and post study.

While this decline in trust is concerning, it’s important to note that the survey also appeared to increase participants’ awareness of privacy settings. This increased awareness likely contributed to a more critical evaluation of Facebook’s practices, which, although leading to reduced trust, also indicates a positive shift in users taking more control over their privacy.

4.3.3 Increased interpersonal privacy concerns. We further compared participants’ privacy concerns (*interpersonal privacy, institutional privacy, resignation*) at the beginning and end of the survey (see Figure 7). After the survey, Wilcoxon signed-rank tests indicate that the increase in interpersonal privacy concerns is significant ($p < .001$ for all interpersonal privacy items). Changes for institutional privacy concerns were not significant.

Summary. Our findings reveal a clear relationship between perceptions of Facebook and expectation-setting mismatches. While concern with their actual settings was neutral or low on average, we observe a significant drop in trust towards Facebook and heightened interpersonal privacy concerns after participants’ learning about their actual Facebook settings.

5 Discussion

We first summarize the key findings of our study and then discuss our findings’ design and public policy implications.

5.1 Main Contributions and Implications

Compared to findings from prior studies discussed in Section 2.3, our study provides comparable findings, emphasizing a continued and persistent gap between users’ expectations and actual settings

on Facebook. Liu et al.'s 2011 study adopted the most similar approach to our work, focusing only on post visibility settings, using crawling techniques to gather Facebook privacy settings for 200 users. Over a decade later, our study confirms many of their findings for these settings, revealing that many users today still retain the more privacy-invasive default settings on Facebook. When mismatches occurred, they almost always exposed a user's content to a larger audience than expected. However, Facebook's settings and user demographics have changed drastically since 2011 [18]. Key updates include the 2011 overhaul of privacy settings [74], the introduction of the Timeline feature [93], the Privacy Checkup Tool [46], adjustments following the Cambridge Analytica scandal in 2018 [52], the Off-Facebook Activity tool in 2019 [69], and the Privacy Center in 2021 [96].

While these updates were intended to enhance user control and transparency, our findings suggest that they have not fully addressed the core issue of privacy expectation-setting mismatches. Users still face substantial gaps between their privacy expectations and the actual settings, leading to potential risks. For example, despite Facebook's introduction of tools like Privacy Checkup, users may still overlook or misunderstand critical settings, resulting in unintentional data exposure. Furthermore, the continued prevalence of default settings that are more permissive than users expect raises concerns about the effectiveness of these updates in genuinely protecting user privacy. This suggests that while Facebook's privacy settings have evolved, they still fall short in aligning with user expectations, necessitating further improvements in privacy design and user education.

Similar to our added focus on demographic characteristics, a 2018 study by Kanampiu and Anwar [50] also analyzed whether or not Facebook privacy setting behavior was related to user gender, age, and education, albeit in an artificial task. Their findings suggested a significant correlation between age and privacy settings behavior, which was reaffirmed by our own analysis, in which older participants had a higher number of mismatches. We expand on this work by additionally analyzing the influence of users' religious and political affiliations, and by comparing expectations against the actual privacy settings, which were collected through our custom browser extension, rather than inferring mismatches from an artificial task.

Our work also reaffirmed findings of other prior studies, including Madejski et al.'s 2012 study in which every participant also had at least one mismatch, and Hsu et al. [44]'s 2020 paper that found misalignments between expected values, particularly with regards to ad personalization. Overall, our work contributes comprehensive and updated insights on diverse users' actual Facebook privacy settings as well as their expectations of them, while emphasizing the continued need for better-designed privacy controls on Facebook. We discuss our specific key insights and their implications next.

5.1.1 Expectation-setting misalignments prevalent. Regarding RQ1, we find substantial and prevalent misalignments between participants' expected and actual Facebook privacy settings—for almost all of the 18 settings we studied, a third to over half of participants had a mismatch. We confirm Liu et al.'s respective findings for post visibility [58] and further identify substantial mismatches for general privacy settings, ad settings, and additional timeline settings. Most

participants expected their privacy settings to be more restrictive than they were (e.g., “friends only” versus “public”). The prevalence of mismatches in our study suggest that for many Facebook users, their data is much more exposed on Facebook to others than they expect, posing interpersonal privacy risks.

The misalignments for ad privacy settings and surprisingly high numbers of ad topics and companies highlight an important institutional privacy issue: participants' settings being more open than expected enables Facebook to monetize users' data for ad targeting in unnoticeable ways. Though we focused on Facebook, the discourse on surveillance capitalism [110, 111] and dark/deceptive design [35, 36] suggests that this is a systemic issue that likely extends beyond Meta to other companies. Platforms inundate users with a panoply of ineffective privacy settings under the guise of transparency and control, thereby overwhelming them to lose track of their actual settings and exposure. This deceptive design pattern is also known as “privacy Zuckering” in reference to Meta's CEO [15]. Our study provides empirical evidence for the detrimental effects of complex privacy settings on consumers' level of exposure and volume of data involved in ad targeting—on average, participants had over 400 assigned ad topics and over 400 companies targeted ads to them.

5.1.2 Potential demographic differences in mismatches. Regarding RQ2, our findings indicate a range of demographic differences in privacy expectation-setting misalignments in our sample: our findings suggest that participants who are older in age, Asian, Mixed-Race, and Hispanic/Latine (race/ethnicity), Republican (political affiliation), and men (gender) might experience more mismatches. Further research is needed to confirm and investigate these correlations. If our findings hold, people in these groups might be more likely to be more exposed and therefore to suffer disparate effects—both in terms of interpersonal privacy risks, e.g., online harassment or doxxing [56, 60], and institutional privacy risks, e.g., through discriminatory advertising, e.g., housing and employment ads that exclude certain racial and ethnic groups of users.

5.1.3 Actual settings decrease trust and raise concerns. For RQ3, we found that participants' concern with specific privacy settings was neutral to low overall. However, these findings should not be misconstrued as their current privacy settings reflecting participants' privacy preferences. Rather, most participants expressed desire for but resignation about their ability to control how Facebook uses their data (see Figure 7). Such resignation may be fueled by corporate practices [26] and, in the case of Facebook, by its history of changing privacy settings [16] and data leakage scandals [49].

The interpretation that participants' actual privacy settings did not reflect their desired privacy levels is further supported by the significant decrease in participants' trust in Facebook and increase in interpersonal privacy concerns by the end of the study. While there might be some priming effect here—we did make participants aware of expectation-setting mismatches, their actual settings, and the extent of Facebook's ad profiles in them—our study setup is not too different from someone reviewing their settings in their Facebook account. It appears that learning about these aspects had an aggregate effect on participants' trust in Facebook and interpersonal privacy concerns even though the per-setting concern was moderate.

Our findings regarding trust and interpersonal privacy concerns are not just alarming from a consumer perspective. Companies should also be concerned if users' trust decreases when learning about their actual privacy settings, as was the case for our participants, as it might indicate that the companies' privacy settings may not be effectively reaching users or may be causing misconceptions. Substantial mismatches between users' understanding and expectations of how their data is exposed or used and a company's actual practices may also lead to regulatory scrutiny and sanctions, e.g., by the FTC or European regulators regarding GDPR and DSA compliance.

5.2 Design and Policy Implications

Our findings provide further evidence that Facebook's approach to privacy settings is ineffective at giving users control over their privacy—with seemingly disproportionate impacts on certain demographic groups. We discuss considerations for design, education, and policy.

5.2.1 Embedding privacy controls into user experiences. Schaub and Cranor have advocated for not just usable but also useful privacy interfaces [84, 86]. They suggest that for privacy settings to be noticed and used, they must be a part of the user experience rather than be hidden away in difficult to find places [38]. Our study provides further evidence for the need to better surface privacy controls. Prior work already provides some directions for this. For instance, Habib et al. [39] gathered Facebook users' needs for ad controls and proposed design ideas; Im et al. [47] demonstrated how placing ad controls next to ads helps users find them; Farke et al. [30] showed that exposing users to data dashboards positively affects their trust in a platform and decreases privacy concerns; Schaub et al. [85] mapped the design space for privacy notice and control design. However, one challenge that our study surfaces is that the number of privacy settings platforms offer is overwhelming for users.

Instead of showing every setting to all users, an alternative could be to scaffold settings in a tailored way for user groups with different needs [39], e.g., marginalized populations [104]. More research is needed on understanding how to actually achieve this.

5.2.2 Improving digital literacy and online self-defense. A further aspect to consider is how consumers' digital literacy, especially privacy literacy, can be improved. Despite Facebook having been around for two decades and being scrutinized in the media and by regulators, our study shows there are still substantial gaps in how users understand and are aware of Facebook's data practices. This is neither necessarily consumers' fault nor should it be their responsibility. But, it is indicative of a dissonance between consumers' ease of engagement with platforms and their ability to reason critically about platforms' practices and associated risks.

A potential direction for addressing this is to rethink approaches to teaching digital literacy to both adults and children [23, 73]—including preschool-age children given their increasing technology use [59]. This could involve finding clearer metaphors to scaffold mental models of how interpersonal interactions on social platforms are enabled by institutional infrastructures that often aim to monetize user data [102]. It could also include teaching people

online self-defense, i.e., enabling them to take charge of their online privacy—at least as far as that is possible within the constraints of settings provided by platforms. While current privacy settings may not be ideal, and many privacy interfaces are fraught with usability issues, it is problematic when consumers are not even aware of how to leverage them.

5.2.3 Requiring privacy-friendly defaults. Our study adds further evidence to the failure of the notice and choice regime. Policymakers and regulators must take action to reign in and remediate the clear misalignment between consumers' expectations of how exposed or protected their data is and the realities of what platforms do with users' data. Many of our participants expected their settings to be more restrictive than they were in actuality, which suggests that many had a desire for more privacy-friendly defaults. Yet, platforms frequently set defaults to be more open and hide “opt-outs” deep in their settings or privacy policies [40]. Policymakers could more effectively require companies to practice privacy by default, as already mandated by GDPR.

Research on dark/deceptive design patterns has exposed widespread strategies to manipulate consumers into providing consent or sharing more data [15, 35, 66, 71, 98, 103]. Recent laws in California and elsewhere that prohibit deceptive and manipulative techniques are a step forward. However, we still lack clearer research-informed requirements on companies regarding how they must structure consent interfaces and interactions, as well as which privacy settings companies need to provide and how.

5.2.4 Setting meaningful limits for data processing. Additionally, policymakers and regulators should set more meaningful limits regarding what data processing practices are acceptable and which ones are not, given how surprised our participants were by the number of ad topics and companies associated with them. A promising avenue is the idea of making companies ‘information fiduciaries’ [12] and bestowing on them a ‘duty of loyalty’ to their users [82]. Under such a duty of loyalty, companies collecting and processing personal data would have to act in their users' best interests. Such an approach could move online privacy past the ruins of notice and choice to a more mutually respectful relationship between individuals and platforms, and reduce the unworkable over-reliance on meaningless consent interfaces and privacy settings.

6 Conclusion

Our study (n=195) gathered participants expectations and actual values for a range of Facebook privacy settings. We find that expectation-setting mismatches are prevalent—all participants had at least one mismatch; for many settings over half the participants had a mismatch. Many participants expected their settings to be more restrictive than they were. Our analysis suggests potential differences among some demographic groups in the number of mismatches. Participants expressed neutral to low concern for individual settings, but we saw a significant decrease in overall trust in Facebook and increase in interpersonal privacy concerns. Our findings provide further empirical evidence for the failure of the notice and choice approach to privacy, as the majority of our participants were unaware of the actual exposure of their Facebook data.

Acknowledgments

We are grateful to our participants. This research has been partially supported by the Defense Advanced Research Projects Agency (DARPA) under grant No. HR00112010010. Byron Lowens has been supported by a CRA/NSF CI Fellowship under grant No. 2030859. Jane Im has been supported by a Meta Graduate Research Student Fellowship. The content of the information does not necessarily reflect the position or the policy of the U.S. Government, and no official endorsement should be inferred. Approved for public release; distribution is unlimited.

References

- [1] Norah Abokhodair and Sarah Vieweg. 2016. Privacy & Social Media in the Context of the Arab Gulf. In *Proceedings of the 2016 ACM Conference on Designing Interactive Systems (DIS)*. ACM, Brisbane, Australia, 672–683. <https://doi.org/10.1145/2901790.2901873>
- [2] Alessandro Acquisti, Laura Brandimarte, and George Loewenstein. 2015. Privacy and human behavior in the age of information. *Science* 347, 6221 (2015), 509–514.
- [3] Alessandro Acquisti, Laura Brandimarte, and George Loewenstein. 2020. Secrets and likes: The drive for privacy and the difficulty of achieving it in the digital age. *Journal of Consumer Psychology* 30, 4 (2020), 736–758.
- [4] Alessandro Acquisti and Ralph Gross. 2006. Imagined communities: Awareness, information sharing, and privacy on the Facebook. In *Privacy Enhancing Technologies: 6th International Workshop, PET 2006, Cambridge, UK, June 28-30, 2006, Revised Selected Papers 6*. Springer, 36–58.
- [5] Hanan A Al-Asmari and Mohamed S Saleh. 2019. A conceptual framework for measuring personal privacy risks in Facebook online social network. In *2019 International Conference on Computer and Information Sciences (ICCIS)*. IEEE, 1–6.
- [6] Muhammad Ali, Angelica Goetzen, Alan Mislove, Elissa M Redmiles, and Piotr Sapiiezynski. 2023. Problematic Advertising and its Disparate Exposure on Facebook. In *arXiv preprint arXiv:2306.06052*.
- [7] Muhammad Ali, Piotr Sapiiezynski, Miranda Bogen, Aleksandra Korolova, Alan Mislove, and Aaron Rieke. 2019. Discrimination through optimization: How Facebook's Ad delivery can lead to biased outcomes. *Proceedings of the ACM on human-computer interaction* 3, CSCW (2019), 1–30.
- [8] Irwin Altman. 1975. The environment and social behavior: privacy, personal space, territory, and crowding. (1975).
- [9] Lori Andrews. 2012. Facebook is using you. *The New York Times* 4 (2012).
- [10] Julia Angwin, Ariana Tobin, and Madeleine Varner. 2017. Facebook (still) letting housing advertisers exclude users by race. *ProPublica*, November 21 (2017).
- [11] Brooke Auxier, Lee Rainie, Monica Anderson, Andrew Perrin, Madhu Kumar, and Erica Turner. 2019. Americans and privacy: Concerned, confused and feeling lack of control over their personal information. (2019).
- [12] Jack M Balkin. 2020. The Fiduciary Model of Privacy. In *Harvard Law Review Forum*, Vol. 134.
- [13] S Batool, S Sultana, and S Tariq. 2021. Social Media and Religious Minorities: Analyzing the Usage of Facebook Groups among Christian Minority to Highlight their Issues in Pakistan. *Global Mass Communication Studies Review*, VI (2021), 117–132.
- [14] Sophie C Boerman, Sanne Kruikemeier, and Frederik J Zuiderveen Borgesius. 2017. Online behavioral advertising: A literature review and research agenda. *Journal of advertising* 46, 3 (2017), 363–376.
- [15] Christoph Bösch, Benjamin Erb, Frank Kargl, Henning Kopp, and Stefan Pfatthicher. 2016. Tales from the dark side: privacy dark strategies and privacy dark patterns. *Proc. Priv. Enhancing Technol.* 2016, 4 (2016), 237–254.
- [16] danah boyd and Eszter Hargittai. 2010. Facebook privacy settings: Who cares? *First Monday* 15, 8 (Jul. 2010). <https://doi.org/10.5210/fm.v15i8.3086>
- [17] José González Cabañas, Ángel Cuevas, and Rubén Cuevas. 2018. Unveiling and quantifying facebook exploitation of sensitive personal data for advertising purposes. In *27th USENIX Security Symposium (USENIX Security 18)*. 479–495.
- [18] Pew Research Center. 2024. 5 facts about how Americans use Facebook two decades after its launch. <https://www.pewresearch.org/short-reads/2024/02/02/5-facts-about-how-americans-use-facebook-two-decades-after-its-launch/#:~:text=Today's%20teens%20tend%20to%20gravitate,those%20living%20in%20the%20suburbs> Accessed: 2024-07-28.
- [19] Shuchih Ernest Chang, Anne Yenching Liu, and Wei Cheng Shen. 2017. User trust in social networking services: A comparison of Facebook and LinkedIn. *Computers in Human Behavior* 69 (2017), 207–217.
- [20] Jessica Colnago, Lorrie Faith Cranor, Alessandro Acquisti, and Kate Hazel Stanton. 2022. Is it a concern or a preference? An investigation into the ability of privacy scales to capture and distinguish granular privacy constructs. In *Eighteenth Symposium on Usable Privacy and Security (SOUPS 2022)*. 331–346.
- [21] Nicholas Confessore. 2018. Cambridge Analytica and Facebook: The Scandal and the Fallout So Far. <https://www.nytimes.com/2018/04/04/us/politics/cambridge-analytica-scandal-fallout.html>. Online; accessed 11 September 2023.
- [22] Lorrie Faith Cranor. 2012. Necessary but not sufficient: Standardized mechanisms for privacy notice and choice. *J. on Telecomm. & High Tech. L.* 10 (2012), 273.
- [23] Laurien Desimpelaere, Liselot Hudders, and Dienneke Van de Sompel. 2020. Knowledge as a strategy for privacy protection: How a privacy literacy training affects children's online disclosure behavior. *Computers in human behavior* 110 (2020), 106382.
- [24] Tobias Dienlin and Sabine Trepte. 2015. Is the privacy paradox a relic of the past? An in-depth analysis of privacy attitudes and privacy behaviors. *European journal of social psychology* 45, 3 (2015), 285–297.
- [25] Michael Dimock. 2017. Defining generations: Where Millennials end and Generation Z begins. Pew Research Center. <https://www.pewresearch.org/short-reads/2019/01/17/where-millennials-end-and-generation-z-begins/>. (accessed 2023-09-13).
- [26] Nora A Draper and Joseph Turov. 2019. The corporate cultivation of digital resignation. *New media & society* 21, 8 (2019), 1824–1839.
- [27] Stefanie Duguay. 2016. "He has a way gayer Facebook than I do": Investigating sexual identity disclosure and context collapse on a social networking site. *New media & society* 18, 6 (2016), 891–907.
- [28] Nicole Ellison and Danah M Boyd. 2013. Sociality through social network sites. (2013).
- [29] Steven Englehardt and Arvind Narayanan. 2016. Online tracking: A 1-million-site measurement and analysis. In *Proceedings of the 2016 ACM SIGSAC conference on computer and communications security*. 1388–1401.
- [30] Florian M Farke, David G Balash, Maximilian Golla, Markus Dürmuth, and Adam J Aviv. 2021. Are Privacy Dashboards Good for End Users? Evaluating User Perceptions and Reactions to Google's My Activity. In *30th USENIX Security Symposium (USENIX Security 21)*. 483–500.
- [31] Casey Fiesler, Michaelanne Dye, Jessica L Feuston, Chaya Hiruncharoenvate, Clayton J Hutto, Shannon Morrison, Parisa Khanipour Roshan, Umashanthi Pavalanathan, Amy S Bruckman, Munmun De Choudhury, et al. 2017. What (or who) is public? Privacy settings and social media content sharing. In *Proceedings of the 2017 ACM conference on computer supported cooperative work and social computing*. 567–580.
- [32] Ariz Arrate Galán, José González Cabañas, Ángel Cuevas, María Calderón, and Rubén Cuevas Rumin. 2019. Large-scale analysis of user exposure to online advertising on facebook. *IEEE Access* 7 (2019), 11959–11971.
- [33] Nina Gerber, Paul Gerber, and Melanie Volkamer. 2018. Explaining the privacy paradox: A systematic review of literature investigating privacy attitude and behavior. *Computers & security* 77 (2018), 226–261.
- [34] Cami Goray and Sarita Schoenebeck. 2022. Youths' Perceptions of Data Collection in Online Advertising and Social Media. *Proc. ACM Hum.-Comput. Interact.* 6, CSCW2, Article 475 (nov 2022), 27 pages. <https://doi.org/10.1145/3555576>
- [35] Colin M Gray, Cristiana Santos, Natalia Bielova, Michael Toth, and Damian Clifford. 2021. Dark patterns and the legal requirements of consent banners: An interaction criticism perspective. In *Proceedings of the 2021 CHI Conference on Human Factors in Computing Systems*. 1–18.
- [36] Colin M Gray, Cristiana Teixeira Santos, Natalia Bielova, and Thomas Mildner. 2024. An ontology of dark patterns knowledge: Foundations, definitions, and a pathway for shared knowledge-building. In *Proceedings of the CHI Conference on Human Factors in Computing Systems*. 1–22.
- [37] Ralph Gross and Alessandro Acquisti. 2005. Information revelation and privacy in online social networks. In *Proceedings of the 2005 ACM workshop on Privacy in the electronic society*. 71–80.
- [38] Hana Habib, Sarah Pearman, Jiamin Wang, Yixin Zou, Alessandro Acquisti, Lorrie Faith Cranor, Norman Sadeh, and Florian Schaub. 2020. "It's a scavenger hunt": Usability of Websites' Opt-Out and Data Deletion Choices. In *Proceedings of the 2020 CHI Conference on Human Factors in Computing Systems*. 1–12.
- [39] Hana Habib, Sarah Pearman, Ellie Young, Ishika Saxena, Robert Zhang, and Lorrie Faith Cranor. 2022. Identifying user needs for advertising controls on Facebook. *Proceedings of the ACM on Human-Computer Interaction* 6, CSCW1 (2022), 1–42.
- [40] Hana Habib, Yixin Zou, Aditi Jannu, Neha Sridhar, Chelse Swoopes, Alessandro Acquisti, Lorrie Faith Cranor, Norman Sadeh, and Florian Schaub. 2019. An Empirical Analysis of Data Deletion and {Opt-Out} Choices on 150 Websites. In *Fifteenth Symposium on Usable Privacy and Security (SOUPS 2019)*. 387–406.
- [41] Rasia Haji and Wolfgang G Stock. 2021. User settings for advertising optimization on Facebook: Active customer participation or settings blindness? *Telematics and Informatics* 59 (2021), 101548.
- [42] Cory Hallam and Gianluca Zanella. 2017. Online self-disclosure: The privacy paradox explained as a temporally discounted balance between concerns and rewards. *Computers in Human Behavior* 68 (2017), 217–227.
- [43] Paul Hitlin and Lee Rainie. 2019. Facebook algorithms and personal data. Pew Research Center. <https://www.pewresearch.org/internet/2019/01/16/facebook>

- algorithms-and-personal-data/. (accessed 2023-09-13).
- [44] Silas Hsu, Kristen Vaccaro, Yin Yue, Aimee Rickman, and Karrie Karahalios. 2020. Awareness, navigation, and use of feed control settings online. In *Proceedings of the 2020 CHI Conference on Human Factors in Computing Systems*. 1–13.
- [45] Thomas Hughes-Roberts. 2013. Privacy and social networks: Is concern a valid indicator of intention and behaviour?. In *2013 International Conference on Social Computing*. IEEE, 909–912.
- [46] Andrew Hutchinson. 2020. Facebook updates its privacy check-up tool to clarify which elements you can control. *Social Media Today* (2020). <https://www.socialmediatoday.com/news/facebook-updates-its-privacy-check-up-tool-to-clarify-which-elements-you/569876/>
- [47] Jane Im, Ruiyi Wang, Weikun Lyu, Nick Cook, Hana Habib, Lorrie Faith Cranor, Nikola Banovic, and Florian Schaub. 2023. Less is Not More: Improving Findability and Actionability of Privacy Controls for Online Behavioral Advertising. In *Proceedings of the 2023 CHI Conference on Human Factors in Computing Systems*. 1–33.
- [48] Jim Isaak and Mina J. Hanna. 2018. User Data Privacy: Facebook, Cambridge Analytica, and Privacy Protection. *Computer* 51, 8 (2018), 56–59. <https://doi.org/10.1109/MC.2018.3191268>
- [49] Jim Isaak and Mina J. Hanna. 2018. User Data Privacy: Facebook, Cambridge Analytica, and Privacy Protection. *Computer* 51, 8 (2018), 56–59. <https://doi.org/10.1109/MC.2018.3191268>
- [50] Munene Kanampiu and Mohd Anwar. 2018. Privacy preferences vs. privacy settings: an exploratory Facebook study. In *International Conference on Applied Human Factors and Ergonomics*. Springer, 116–126.
- [51] Dilara Kekulluoglu, Kami Vaniea, and Walid Magdy. 2022. Understanding Privacy Switching Behaviour on Twitter. In *Proceedings of the 2022 CHI Conference on Human Factors in Computing Systems* (New Orleans, LA, USA) (CHI '22). Association for Computing Machinery, New York, NY, USA, Article 31, 14 pages. <https://doi.org/10.1145/3491102.3517675>
- [52] Matthew Keys. 2018. A Brief History of Facebook's Ever-Changing Privacy Settings. [https://medium.com/@matthewkeys/a-brief-history-of-facebooks-ever-changing-privacy-settings-8167dadd3bd0#:~:text=The%20overhaul%20was%20so%20complex,accounts\)%20watch%20a%20tutorial%20that](https://medium.com/@matthewkeys/a-brief-history-of-facebooks-ever-changing-privacy-settings-8167dadd3bd0#:~:text=The%20overhaul%20was%20so%20complex,accounts)%20watch%20a%20tutorial%20that) Accessed: 2024-06-28.
- [53] Spyros Kokolakis. 2017. Privacy attitudes and privacy behaviour: A review of current research on the privacy paradox phenomenon. *Computers & security* 64 (2017), 122–134.
- [54] Priya Kumar and Sarita Schoenebeck. 2015. The modern day baby book: Enacting good mothering and stewarding privacy on Facebook. In *Proceedings of the 18th ACM conference on computer supported cooperative work & social computing*. 1302–1312.
- [55] Caroline Lang and Hannah Barton. 2015. Just untag it: Exploring the management of undesirable Facebook photos. *Computers in Human Behavior* 43 (2015), 147–155.
- [56] Amanda Lenhart, Michele Ybarra, Kathryn Zickuhr, and Myeshia Price-Feeny. 2016. *Online harassment, digital abuse, and cyberstalking in America*. Data and Society Research Institute. https://www.datasociety.net/pubs/oh/Online_Harassment_2016.pdf.
- [57] Roderick JA Little and Donald B Rubin. 2019. *Statistical analysis with missing data*. Vol. 793. John Wiley & Sons.
- [58] Yabing Liu, Krishna P Gummadi, Balachander Krishnamurthy, and Alan Mislove. 2011. Analyzing facebook privacy settings: user expectations vs. reality. In *Proceedings of the 2011 ACM SIGCOMM conference on Internet measurement conference*. 61–70.
- [59] Sonia Livingstone, Mariya Stoilova, and Rishita Nandagiri. 2019. Children's data and privacy online: growing up in a digital age: an evidence review. (2019).
- [60] Mary Madden. 2017. *Privacy, Security, and Digital Inequality*. Data and Society Research Institute. <https://datasociety.net/library/privacy-security-and-digital-inequality/>.
- [61] Michelle Madejski, Maritza Johnson, and Steven M. Bellovin. 2012. A study of privacy settings errors in an online social network. In *2012 IEEE International Conference on Pervasive Computing and Communications Workshops*. 340–345. <https://doi.org/10.1109/PerComW.2012.6197507>
- [62] Miguel Malheiros, Charlene Jennett, Snehal Patel, Sacha Brostoff, and Martina Angela Sasse. 2012. Too Close for Comfort: A Study of the Effectiveness and Acceptability of Rich-Media Personalized Advertising. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems* (Austin, Texas, USA) (CHI '12). Association for Computing Machinery, New York, NY, USA, 579–588. <https://doi.org/10.1145/2207676.2207758>
- [63] Naresh K Malhotra, Sung S Kim, and James Agarwal. 2004. Internet users' information privacy concerns (UIPC): The construct, the scale, and a causal model. *Information systems research* 15, 4 (2004), 336–355.
- [64] Ben Marder, Emma Slade, David Houghton, and Chris Archer-Brown. 2016. "I like them, but won't 'like' them": An examination of impression management associated with visible political party affiliation on Facebook. *Computers in Human Behavior* 61 (2016), 280–287.
- [65] Kirsten Martin. 2015. Privacy notices as tabula rasa: An empirical investigation into how complying with a privacy notice is related to meeting privacy expectations online. *Journal of Public Policy & Marketing* 34, 2 (2015), 210–227.
- [66] Arunesh Mathur, Gunes Acar, Michael J Friedman, Eli Lucherini, Jonathan Mayer, Marshini Chetty, and Arvind Narayanan. 2019. Dark patterns at scale: Findings from a crawl of 11K shopping websites. *Proceedings of the ACM on Human-Computer Interaction* 3, CSCW (2019), 1–32.
- [67] Nora McDonald, Karla Badillo-Urquiola, Morgan G Ames, Nicola Dell, Elizabeth Keneski, Manya Sleeper, and Pamela J Wisniewski. 2020. Privacy and power: Acknowledging the importance of privacy research and design for vulnerable populations. In *Extended Abstracts of the 2020 CHI Conference on Human Factors in Computing Systems*. 1–8.
- [68] Tyler Musgrave, Alia Cummings, and Sarita Schoenebeck. 2022. Experiences of Harm, Healing, and Joy among Black Women and Femmes on Social Media. In *Proceedings of the 2022 CHI Conference on Human Factors in Computing Systems*. 1–17.
- [69] Facebook Newsroom. 2019. Off-Facebook Activity: More About Your Privacy. <https://about.fb.com/news/2019/08/off-facebook-activity/#:~:text=With%20Off%2DFacebook%20Activity%2C%20you,and%20websites%20share%20with%20Facebook>. Accessed: 2024-07-28.
- [70] Helen Nissenbaum. 2011. A contextual approach to privacy online. *Daedalus* 140, 4 (2011), 32–48.
- [71] Midas Nouwens, Ilaria Liccardi, Michael Veale, David Karger, and Lalana Kagal. 2020. Dark patterns after the GDPR: Scraping consent pop-ups and demonstrating their influence. In *Proceedings of the 2020 CHI conference on human factors in computing systems*. 1–13.
- [72] Leysia Palen and Paul Dourish. 2003. Unpacking "privacy" for a networked world. In *Proceedings of the SIGCHI conference on Human factors in computing systems*. 129–136.
- [73] Yong Jin Park. 2013. Digital literacy and privacy behavior online. *Communication research* 40, 2 (2013), 215–236.
- [74] Ben Parr. 2011. Facebook's Big Privacy Changes: An Overview [PICS]. *Mashable* (2011). <https://mashable.com/archive/facebook-privacy-changes-guide#TMdc0NovxZqX>
- [75] Eyal Peer, Laura Brandimarte, Sonam Samat, and Alessandro Acquisti. 2017. Beyond the Turk: Alternative platforms for crowdsourcing behavioral research. *Journal of experimental social psychology* 70 (2017), 153–163.
- [76] Pew Research Center. 2023. *How Americans View Data Privacy*. Technical Report. Washington, D.C. <https://www.pewresearch.org/internet/2023/10/18/how-americans-view-data-privacy/>
- [77] Pew Research Center. 2024. *Social Media Fact Sheet*. Technical Report. Washington, D.C. <https://www.pewresearch.org/internet/fact-sheet/social-media/>
- [78] Emilee Rader, Samantha Hautea, and Anjali Munasinghe. 2020. "I Have a Narrow Thought Process": Constraints on Explanations Connecting Inferences and {Self-Perceptions}. In *Sixteenth Symposium on Usable Privacy and Security (SOUPS 2020)*. 457–488.
- [79] Ashwini Rao, Florian Schaub, Norman Sadeh, Alessandro Acquisti, and Ruogu Kang. 2016. Expecting the unexpected: Understanding mismatched privacy expectations online. In *Twelfth Symposium on Usable Privacy and Security (SOUPS 2016)*. 77–96.
- [80] Kate Raynes-Goldie. 2010. Aliases, creeping, and wall cleaning: Understanding privacy in the age of Facebook. *First Monday* (2010).
- [81] Elissa M Redmiles, Sean Kross, and Michelle L Mazurek. 2019. How well do my results generalize? comparing security and privacy survey results from mturk, web, and telephone samples. In *2019 IEEE Symposium on Security and Privacy (SP)*. IEEE, 1326–1343.
- [82] Neil M Richards and Woodrow Hartzog. 2021. A Duty of Loyalty for Privacy Law. *Washington University Law Review* 99, 3 (2021), 961.
- [83] Shruti Sannon and Andrea Forte. 2022. Privacy Research with Marginalized Groups: What We Know, What's Needed, and What's Next. *Proc. ACM Hum.-Comput. Interact.* 6, CSCW2, Article 455 (nov 2022), 33 pages. <https://doi.org/10.1145/3555556>
- [84] Florian Schaub, Rebecca Balebako, and Lorrie Faith Cranor. 2017. Designing Effective Privacy Notices and Controls. *IEEE Internet Computing* 21, 3 (2017), 70–77. <https://doi.org/10.1109/MIC.2017.75>
- [85] Florian Schaub, Rebecca Balebako, Adam L Durity, and Lorrie Faith Cranor. 2015. A design space for effective privacy notices. In *Eleventh symposium on usable privacy and security (SOUPS 2015)*. 1–17.
- [86] Florian Schaub and Lorrie Faith Cranor. 2020. Usable and useful privacy interfaces. *An Introduction to Privacy for Technology Professionals* (2020), 176–299.
- [87] Sonya Scherini. 2020. Facebook: Where privacy concerns and social needs collide. (2020).
- [88] Hadas Schwartz-Chassidim, Oshrat Ayalon, Tamir Mendel, Ron Hirschprung, and Eran Toch. 2020. Selectivity in posting on social networks: the role of privacy concerns, social capital, and technical literacy. *Heliyon* 6, 2 (2020), e03298.
- [89] Matthew Smith, Christian Szongott, Benjamin Henne, and Gabriele Von Voigt. 2012. Big data privacy issues in public social media. In *2012 6th IEEE international*

- conference on digital ecosystems and technologies (DEST). IEEE, 1–6.
- [90] Daniel J Solove. 2012. Introduction: Privacy self-management and the consent dilemma. *Harv. L. Rev.* 126 (2012), 1880.
- [91] Daniel J Solove. 2021. The myth of the privacy paradox. *Geo. Wash. L. Rev.* 89 (2021), 1.
- [92] Agrima Srivastava and G Geethakumari. 2013. Measuring privacy leaks in online social networks. In *2013 International Conference on Advances in Computing, Communications and Informatics (ICACCI)*. IEEE, 2095–2100.
- [93] Donna Tam. 2012. Facebook resurrects old posts on Timeline, panic ensues. *CNET* (2012). <https://www.cnet.com/tech/services-and-software/facebook-resurrects-old-posts-on-timeline-panic-ensues/>
- [94] Jenny Tang, Eleanor Birrell, and Ada Lerner. 2022. Replication: How Well Do My Results Generalize Now? The External Validity of Online Privacy and Security Surveys. In *Eighteenth Symposium on Usable Privacy and Security (SOUPS 2022)*. USENIX Association, Boston, MA, 367–385. <https://www.usenix.org/conference/soups2022/presentation/tang>
- [95] Jenny Tang, Eleanor Birrell, and Ada Lerner. 2022. Replication: How well do my results generalize now? The external validity of online privacy and security surveys. In *Eighteenth symposium on usable privacy and security (SOUPS 2022)*. 367–385.
- [96] Social Media Today. 2021. Facebook Launches New Privacy Center to Facilitate More Control Over Privacy. <https://www.socialmediatoday.com/news/facebook-launches-new-privacy-center-to-facilitate-more-control-over-priv/616847/#:~:text= Accessed: 2024-07-28.>
- [97] Joseph Turow, Michael Hennessy, and Nora Draper. 2015. The tradeoff fallacy: How marketers are misrepresenting American consumers and opening them up to exploitation. Available at SSRN 2820060 (2015).
- [98] Christine Utz, Martin Degeling, Sascha Fahl, Florian Schaub, and Thorsten Holz. 2019. (Un) informed consent: Studying GDPR consent notices in the field. In *Proceedings of the 2019 acm sigsac conference on computer and communications security*. 973–990.
- [99] Evert Van den Broeck, Karolien Poels, and Michel Walrave. 2015. Older and wiser? Facebook use, privacy concern, and privacy protection in the life stages of emerging, young, and middle adulthood. *Social Media+ Society* 1, 2 (2015), 2056305115616149.
- [100] Evert Van den Broeck, Karolien Poels, and Michel Walrave. 2020. How do users evaluate personalized Facebook advertising? An analysis of consumer- and advertiser controlled factors. *Qualitative Market Research: An International Journal* (2020).
- [101] Fernando N Van der Vlist and Anne Helmond. 2021. How partners mediate platform power: Mapping business and data partnerships in the social media ecosystem. *Big Data & Society* 8, 1 (2021), 20539517211025061.
- [102] José Van Dijck, David Nieborg, and Thomas Poell. 2019. Reframing platform power. *Internet Policy Review* 8, 2 (2019), 1–18.
- [103] Ari Ezra Waldman. 2020. Cognitive biases, dark patterns, and the 'privacy paradox'. *Current Opinion in Psychology* 31 (2020), 105–109. <https://doi.org/10.1016/j.copsyc.2019.08.025> Privacy and Disclosure, Online and in Social Interactions.
- [104] Noel Warford, Tara Matthews, Kaitlyn Yang, Omer Akgul, Sunny Consolvo, Patrick Gage Kelley, Nathan Malkin, Michelle L. Mazurek, Manya Sleeper, and Kurt Thomas. 2022. SoK: A Framework for Unifying At-Risk User Research. In *2022 IEEE Symposium on Security and Privacy (SP)*. 2344–2360. <https://doi.org/10.1109/SP46214.2022.9833643>
- [105] Pamela J Wisniewski, Bart P Knijnenburg, and Heather Richter Lipford. 2017. Making privacy personal: Profiling social network users to inform privacy education and nudging. *International Journal of human-computer studies* 98 (2017), 95–108.
- [106] Belle Wong. 2024. Top Social Media Statistics And Trends. <https://www.forbes.com/advisor/in/business/social-media-statistics/>
- [107] Yaxing Yao, Davide Lo Re, and Yang Wang. 2017. Folk models of online behavioral advertising. In *Proceedings of the 2017 ACM Conference on Computer Supported Cooperative Work and Social Computing*. 1957–1969.
- [108] Chance York and Jason Turcotte. 2015. Vacationing from Facebook: Adoption, temporary discontinuance, and readoption of an innovation. *Communication Research Reports* 32, 1 (2015), 54–62.
- [109] Alyson Leigh Young and Anabel Quan-Haase. 2013. Privacy protection strategies on Facebook: The Internet privacy paradox revisited. *Information, Communication & Society* 16, 4 (2013), 479–500.
- [110] Shoshana Zuboff. 2015. Big other: surveillance capitalism and the prospects of an information civilization. *Journal of information technology* 30, 1 (2015), 75–89.
- [111] Shoshana Zuboff. 2019. *The age of surveillance capitalism: The fight for a human future at the new frontier of power*. Profile books.
- [112] Mark Zuckerberg. 2019. The Facts About Facebook. <https://www.wsj.com/articles/the-facts-about-facebook-11548374613>