

ATT vs. Personalized Ads: User’s Data Sharing Choices Under Apple’s Divergent Consent Strategies

Sagar Baviskar^{**} Iffat Chowdhury[†] Daniel Deisenroth[†]

Beibei Li^{*} Daniel Sokol[¶]

June 2024

Abstract

Several recent studies have found that Apple’s App Tracking Transparency (ATT) policy has shifted the advertising ecosystem, impacting millions of businesses reliant on digital advertising. ATT requires non-Apple apps to acquire user consent via a prompt that asks users if they will allow being “tracked”; meanwhile, Apple employs a different consent structure to deliver personalized ads (Apple Personalized Ads (PA) Prompt) on Apple apps. To explore the impact of these differential prompts, we conducted a large-scale survey experiment with 11,000 U.S. and U.K. online adults, randomly exposing them to the opt-in prompts Apple requires for third-party apps vs itself. Results indicate that ATT prompts reduce the likelihood that users opt-in to data sharing by 12.4pp—that is, opt-in rates with Apple’s PA prompt are approximately *double* opt-in rates with the ATT prompt (13% versus 25%). Individuals with a stated preference for personalized ads are even more discouraged to opt-in by an ATT prompt (a 15.1 percentage point difference). We also find that the influence of the ATT prompt is the same across a wide range of third-party applications, including social media, news, delivery apps, and consumer product goods. Finally, ATT increases the misconception that location data is being accessed by 9.2 percentage points. To our knowledge, this is the first study to experimentally test the impact of differential prompt language in the context of Apple’s ATT policy and contributes to the growing debate around the economic consequences of ATT.

^{*}Carnegie Mellon University.

[†]Meta.

[¶]University of Southern California.

^{*}Corresponding author. Email: sbaviska@andrew.cmu.edu

Acknowledgment: The authors acknowledge that they did not receive funding from Meta for this study.

Competing interests: S.B. and B.L. declare no competing interests. D.S. declares no competing interests, although he, in the past, has consulted for Meta. I.C. and D.D. are employees of Meta Platforms and hold a financial interest in Meta.

1 Introduction

In the contemporary landscape of mobile app ecosystems, users' demographic and behavioral data are utilized to deliver and measure the effectiveness of personalized ads. Personalized advertising enables developers to offer free or subsidized access to mobile applications while delivering relevant ads for products and services that users may be interested in (Ghose & Todri-Adamopoulos, 2016) (Bleier & Eisenbeiss, 2015) (Ghose, 2018) (Sun et al., 2023).

The unveiling of Apple's App Tracking Transparency (ATT) framework in June 2020 marked a shift in the landscape of online user data collection and strategies. ATT requires non-Apple apps to obtain active user consent for "tracking." If a user does not opt-in to "tracking" the app cannot access a unique anonymized identifier for their device, known as the Identifier for Advertisers (IDFA). The IDFA enables app developers, advertisers, and advertising platforms to measure the effectiveness of personalized ads, such as determining whether a user purchased a product from an advertiser after viewing an ad in a mobile application¹. Following its announcement, Apple formally implemented ATT as part of iOS 14.5 on April 20, 2021, with the update going live on April 26, 2021.

However, the emergence of ATT introduced challenges for developers in maintaining service quality and delivering relevant advertisements to users. ATT shifted user tracking practices for personalized advertising, moving from an opt-out to an opt-in paradigm. This marked a departure from previous industry norms and disrupted established revenue models (Kraft et al., 2023), possibly leading to a pivot towards alternative monetization strategies such as subscription fees or in-app purchases. Crucially, this policy change led to business exit, reduced business entry, market consolidation, and increased product prices (Deisenroth et al., 2024).

During the initial months of ATT's implementation, users were automatically opted-in to receive personalized advertisements from Apple in its app store and other applications². In September 2021, Apple began asking users for their consent to receive personalized advertising from Apple in an effort to bring parity with its policies with third-party developers. However, discrepancies emerged in the presentation of information to users across different consent prompts used for advertising through Apple, on its own applications, versus advertising on non-Apple applications (Competition and Markets Authority, 2021) (Sokol & Zhu, 2022). There is a contrast in the language and positioning of prompts Apple uses for itself versus that it requires third-party developers. While Apple's personalized ads prompt features positive language and prominently positions the "opt-in" option, the prompt required for non-Apple apps

¹ More generally, if a user does not opt-in to "tracking," Apple prohibits the app developer from "linking user or device data collected from [the] app with user or device data collected from other companies' apps, websites, or offline properties for targeted advertising or advertising measurement purposes." ([Apple Advertising & Privacy](#)).

² The information that Apple uses to personalize ads includes "music, movies, books, TV shows, and apps you download, as well as any in-app purchases and subscriptions," as well as "information stored on your device, such as the apps you frequently open." ([Legal - Apple Advertising & Privacy](#))

adopts potentially less favorable language and places the “opt-in” option less prominently. This disparity in presentation raises concerns that Apple’s privacy policy may bias user choices in favor of Apple’s products, potentially distorting competition in the app and advertising ecosystem.

While other papers have focused on the impact of ATT on businesses and advertisers, this paper is the first to examine the effects on end consumers resulting from the policy. To understand the potential impact on consumers, we conducted a large-scale survey experiment involving 11,000 self-identified iPhone users across the United States and the United Kingdom. Through this experiment, we sought to understand the nuanced ways in which the differences between the data-sharing prompts influence user responses. By analyzing user behavior and beliefs in response to varying prompt presentations, we aim to uncover insights into the potential implications of ATT on market dynamics and competition within the app ecosystem through the lens of the behavior of users.

Our survey experiment reveals that, across a wide variety of requestors including Apple and 10 other apps, opt-in rates, on average, double when users are presented with a PA-style prompt rather than an ATT-style prompt (25% v.s. 13%, or a -12.4pp difference). Furthermore, an ATT prompt decreases the likelihood that a respondent who says they like personalized ads will opt-into data sharing (a -15.1pp total effect of the ATT on this segment). In addition to affecting opt-in decisions, ATT prompts significantly increase the likelihood of saying that choosing to opt-in would allow the requestor to access location data (+9.2pp).

Our findings suggest that ATT prompts discourage data sharing with non-Apple app developers and increase certain misperceptions around data sharing that users may find particularly sensitive. These insights are important for policymakers, regulators, and industry stakeholders as they navigate the evolving landscape of company policies and their implications for consumers, businesses, and competition.

2 Related Literature

In the wake of the implementation of ATT, researchers have turned their attention to assessing its repercussions on various stakeholders within the digital ecosystem. Li and Tsai (2022) examine the consequences of apps' inability to utilize cross-app data for advertising, noting a decline in new downloads. Cheyre et al. (2023) observed a post-ATT initial dip and then rebound in Apple App Store apps, alongside a shift in Software Developer Kit (SDK³) usage, fewer Monetization and Ad Mediation SDKs, and more Authentication and Payments SDKs, indicating developers change their monetization strategy in response to ATT. Kesler (2022) investigates changes in app developers' monetization strategies post-ATT implementation, observing an uptick in the number of paid apps and those offering in-app purchases. This indicates the rise in price for the end users and potentially a gain for Apple through purchase commissions. Kollnig et al. (2022) explore the impact of ATT on data brokers and app developers, highlighting the framework's efficacy in curbing the collection of the IDFA cross-app tracking identifier. Consequently, data brokers encounter heightened difficulties in user data collection. Sokol and Zhu (2021) scrutinize the implementation of ATT as a potentially anti-competitive tactic, raising concerns about its implications for market fairness and end-user welfare. These studies collectively underscore the multifaceted impact of ATT on users' behavior, data privacy, app monetization strategies, and regulatory landscapes, highlighting the need for ongoing research and policy deliberations to not only navigate evolving privacy challenges in the digital age but also monitor the marketplaces anti-competitive practices (Cecere et al., 2023).

ATT has also sparked interest in exploring the effects of prompt design on user perceptions and choices. Apple asserts that it did not conduct A/B testing during the development of the prompts. However, insights from behavioral science literature suggest that variations in design and choice architecture could significantly influence opt-in decisions (Competition and Markets Authority, 2021).

A notable distinction between Apple's prompts lies in the portrayal of data sharing as "tracking" in the third-party prompt while refraining from such terminology in Apple's own prompt. This distinction may have some potential implications based on the extant literature regarding such prompts. Research indicates that users may associate 'tracking' with the collection and utilization of location data, impacting their comfort levels with data sharing (Hutton and Ellis, 2023) (Reidenberg et al., 2015). Users generally exhibit greater reluctance to share location data but may express more willingness to share offsite data for advertising purposes (Leon et al., 2015). The manner in which data sharing options are presented, including the order of opt-in and opt-out choices (Hogarth and Einhorn, 1992), can significantly influence

³ Advertisers use Software Development Kit (SDK), akin to an advertising pixel, to share information from apps to websites and other apps related to deliver and measure the effect of personalized ads.

user decisions and beliefs across various contexts, including data sharing (Thaler et al., 2010) (Acquisti et al., 2017). Adjerid et al. (2018) show how changes in choice frames implemented by service providers can influence consumers' choices, often in ways that they are unaware of and that may be destructive to them. Thus, the way choices are presented, could influence what a decision-maker chooses (Johnson et al., 2012) (Lin and Strulov-Shlain 2023), and the differentiation between ATT prompt for third parties' data use and Apple's data use is very evident.

Apple's Personalized Ads Prompt adopts positive framing and emphasizes the benefits of opting in, potentially enhancing users' perceptions of the value of data sharing (Levin et al., 1998) (Jin et al., 2017). Furthermore, the timing of presenting options following positive framing in Apple's prompt may positively influence user decisions (Turvey and Freeman, 2012). DeGiulio et al. (2021) delve into the presentation of tracking (ATT) requests by mobile apps to users and analyze how different design patterns influence decisions. Their findings suggest that ATT authorizations may reduce data sharing opt-in rates. Further, this study finds that there was no significant effect of a pre-prompt with information on request or framing on the opt-in decisions within the ATT structure. However, DeGiulio et al. (2021), does not evaluate Apple's consent prompt for its apps which we have noted is different from the ATT framework.

In broader terms, the visual design of prompts, including color schemes and screen layout, may impact users' cognitive processing and decision-making efficiency. The literature underscores the importance of effective information visualization in reducing cognitive burdens and facilitating quicker decision-making processes (Eberhard, 2023).

Our study adds a novel dimension to the existing literature by empirically examining the effects of the different prompts on user behavior and beliefs in the context of Apple's privacy policy that governs both its own and third-party apps. Through a large-scale survey experiment involving 11,000 self-identified iPhone users in the U.S. and U.K., we systematically varied prompt designs to explore their influence on user decision-making processes and subsequent effects. By providing empirical evidence on the impact of prompt design on user behavior and beliefs, our study contributes to a deeper understanding of the dynamics of platform privacy policies and informs the ongoing discourse on user consent, perceptions of privacy and market competition.

3 Research Design

We conducted a survey experiment aimed at understanding the effects of the ATT Tracking Prompt and Apple’s “Personalized Ads Prompt” on self-identified iPhone users in the United States (US) and the United Kingdom (UK). The experiment involved randomly exposing respondents to either an ATT Tracking styled-prompt or Apple’s Personalized Ads styled-prompt across various apps.

3.1 Sample Selection and Experimental Procedure

Our research sample consisted of iPhone users residing in both the United States and the United Kingdom. The survey was fielded by YouGov to its online panel of adults in the US and UK. To ensure comprehensive and representative data, we selected respondents from a wide range of demographic backgrounds. This selection process was designed to capture diversity in age, gender, socioeconomic status, and geographical location, providing a dataset reflective of the broader online adult population in these regions.

To enhance the accuracy and generalizability of our findings, we employed stratified sampling techniques. This approach allowed us to proportionately represent various demographic segments, ensuring that our sample was not skewed towards any particular group. This sampling strategy allows us to make reliable inferences about the broader population’s responses to the ATT Tracking Prompt and Apple’s Personalized Ads Prompt. The sampling frame is a smartphone-user representative frame of US adults, based upon the 2021 Current Population Survey (CPS) Voting and Registration supplements, and smartphone-user representative frame of UK adults, based on the 2018 Eurobarometer. The survey had post-hoc weights assigned by YouGov through that firm’s standard procedure; our results are robust to the inclusion of these weights.

The ATT and PA Prompts

Under ATT, users are prompted to grant or deny consent for “tracking” activities across diverse applications—this is an immediate decision-making scenario wherein users may not defer action. According to Apple, this represents a shift towards a more transparent and user-centric approach to data privacy (Apple Newsroom 2021). At the time of ATT implementation, iPhone users were opted in to receiving personalized advertising from Apple in its app store by default, as shown in Figure 1 below.

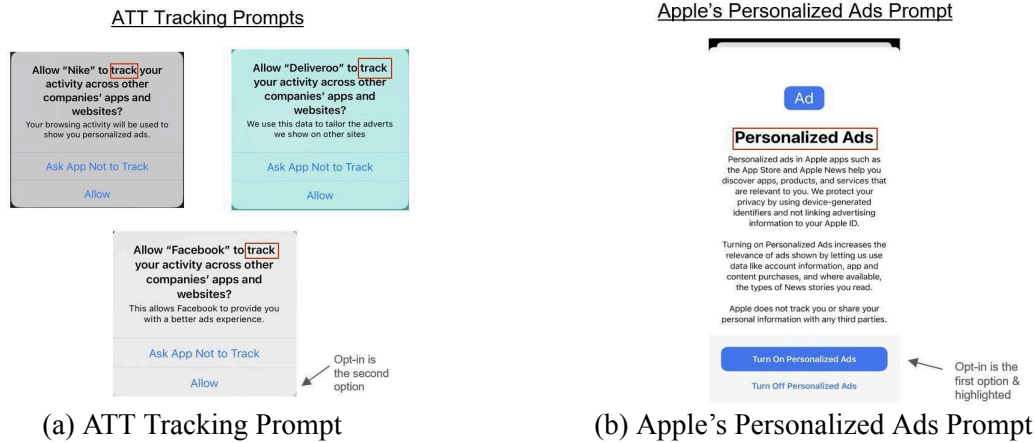


Figure 1: The ATT Tracking Prompt for third-party apps vs Apple's Personalized Ads Prompt for ad on its own apps

Commencing with iOS 14.5, developers that want to engage in personalized advertising are mandated to solicit explicit app-tracking authorization from users. In instances where authorization is withheld, the operating system restricts access to the device's IDFA, effectively curbing third-party access to user data.

Prompt Exposure Assignment:

In the initial stage of our experiment, we implemented a random assignment of respondents to one of two groups. The first group was exposed to the ATT-style Prompt; the second encountered Apple's Personalized Ads-style Prompt. By randomizing exposure, we mitigated potential confounding variables, ensuring that any observed differences in user behavior or attitudes could be attributed directly to the type of prompt encountered rather than to pre-existing differences among participants.

App-Specific Prompt Assignment:

In the subsequent stage of the experiment, respondents were further randomized to receive the request from one of seven different apps in the UK or one of ten apps in the US, including both third-party applications and Apple. This stage was designed to investigate whether the requester—whether it was Apple or a third-party app—had any influence on the respondents' perceptions and decisions regarding data sharing and privacy. A UK respondent could see a request from Apple, Facebook, Instagram, X, Nike, Sephora, or Deliveroo. A US respondent could see a request from Apple, Facebook, Instagram, X, Nike, Sephora, DoorDash, ESPN, MSNBC, Fox News, or Walmart.

The selection of the apps encompassed a variety of app categories and usage contexts. This diversity allowed us to examine how different types of apps, potentially associated with varying levels of user trust and engagement, might interact with the prompts to affect user behavior. The apps included high-usage social media platforms, commerce sites, and news apps, ensuring a comprehensive assessment of user responses across different digital environments. We chose apps from relatively well-known brands and only assigned respondents to apps they stated they have heard of. The prompt and app assignment randomizations were independent of each other.

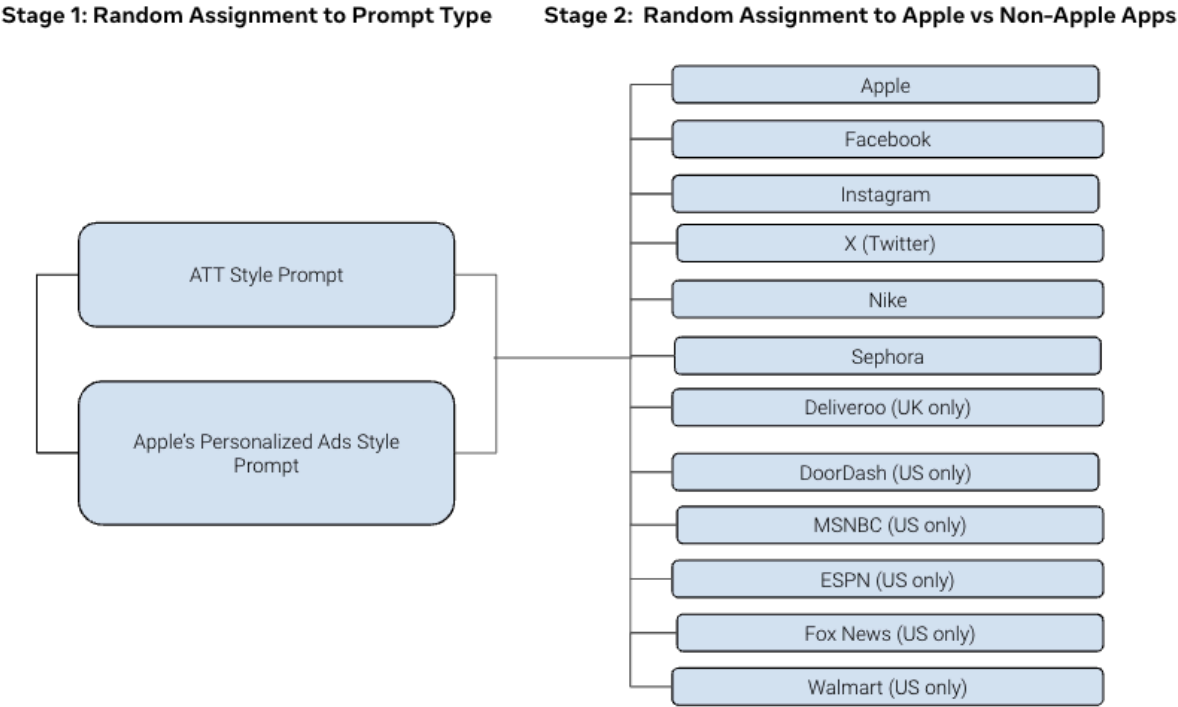


Figure 3: Research Design, respondents were randomly assigned to one of the requestors (apps) they stated they had heard of.

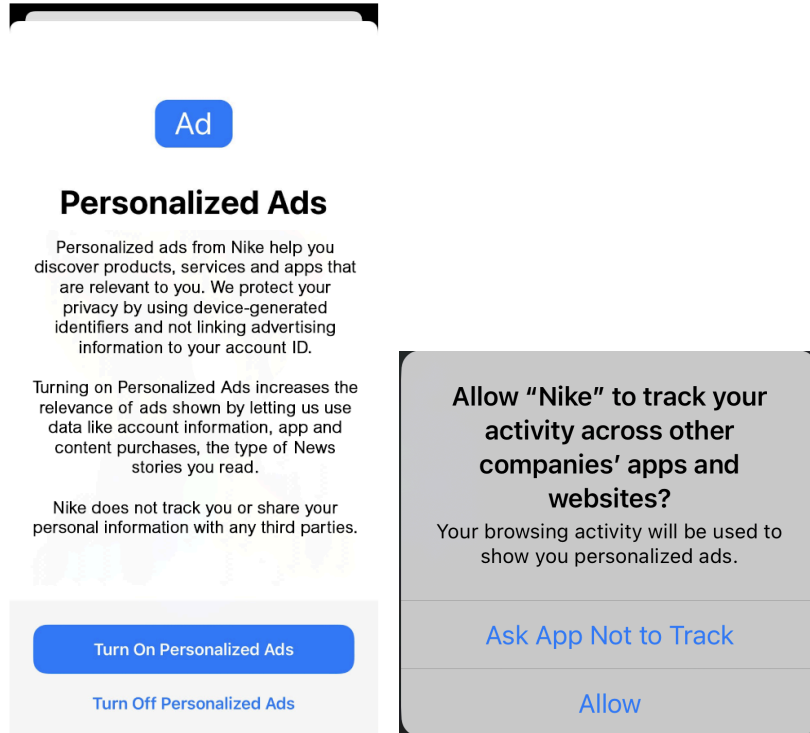


Figure 4 below illustrates two hypothetical prompts a survey respondent may have received. Left panel: the Nike Personalized Ads prompt, right panel: the ATT prompt used by Nike as of February 22, 2024.

3.2 Data Collection

We administered our survey through YouGov, a prominent survey vendor with a robust panel of online adults in both the US and the UK. Participants were presented with the designated prompt within the context of the assigned mobile application. Subsequently, participants engaged with a series of structured questions aimed at probing their perceptions, attitudes, and willingness to consent to tracking or personalized ads. These questions were crafted to capture nuanced insights into participants' reactions to the prompts.

1. Preliminary and Demographic Questions:

In the initial stages of our survey, we focused on gathering comprehensive demographic data and screening respondents for iPhone usage. Respondents were not informed that they needed to be iPhone users to qualify for the survey; instead, they were asked what type of device their primary personal smartphone was. Only respondents who did not select the iPhone were then told they did not qualify for this survey. In addition to demographic information, we also collected data on their awareness of non-Apple apps as well as their familiarity with Apple apps like the App Store, Apple Music, Apple News, Apple Fitness, and Apple TV.

2. Brand Sentiment and Preference for Personalized Ads:

We also included a survey “module” (block of questions) that asked about respondents' sentiments towards brands and their perceptions of privacy within the realm of digital advertising. Half of the respondents were randomized to answer these questions before seeing the prompt and answering questions about their opt-in decisions. While the other half were directed to this module after seeing the prompt and answering questions about their opt-in decisions. This randomization was independent of whether a respondent was assigned to an ATT prompt or a PA prompt and the assigned requesting app.

This module serves two purposes. First, it enables us to explore the extent to which brand sentiment, trust, and preference for personalized ads moderate the effects of the prompt—to do this, we use data from the respondents who answered these questions before exposure to the prompt. Second, it enables us to explore the relationships between exposure to the prompts and shifts in attitudes toward brands, as well as alterations in perceptions of privacy and data protection—for this, we use data from respondents who answered these questions after seeing the prompt.

3. Opt-in Outcome & Interpretations of Prompts:

Finally, we include questions that explore users' reactions to the prompt to which they were exposed. This includes our main outcome of interest: the reported decision on whether one would choose the opt-in option or not, as well as users' understanding of the potential consequences of allowing or denying tracking and enabling or disabling personalized ads to gain insights into their decision-making processes. Specifically, we explore whether users believe that opting in would allow the requesting app to access different types of data. This includes location data, emails, photos, microphone access, purchase information across apps and websites, and whether they clicked on an ad across apps and websites.

3.2.2 Data Management Processes

We adhered to established guidelines governing studies involving human subjects. Before participating in the survey experiment, respondents received a general overview of the study's nature. Respondents were not obligated to complete the survey and could withdraw at any time. Additionally, YouGov implemented standard measures to protect participants' privacy and maintain confidentiality throughout the data collection process. The researchers did not receive any personally identifiable information about the survey respondents. These measures included anonymizing responses, securely storing data, and limiting access to authorized individuals only. Table 1 provides summary statistics of demographic characteristics across prompt assignment.

Table 1: Summary Statistics

Variable	Personalized Ads (PA) Prompt			App Tracking Transparency (ATT) Prompt		
	N	Mean	SD	N	Mean	SD
App Assignment						
... Apple	5445	0.202	0.402	5555	0.204	0.403
... Deliveroo	5445	0.0494	0.217	5555	0.0499	0.218
... DoorDash	5445	0.0549	0.228	5555	0.0481	0.214
... ESPN	5445	0.0479	0.214	5555	0.0522	0.222
... Facebook	5445	0.0959	0.294	5555	0.103	0.304
... Fox News	5445	0.0575	0.233	5555	0.059	0.236
... Instagram	5445	0.106	0.307	5555	0.0986	0.298
... MSNBC	5445	0.0466	0.211	5555	0.047	0.212
... Nike	5445	0.0874	0.282	5555	0.0949	0.293
... Sephora	5445	0.0876	0.283	5555	0.0855	0.28
... Walmart	5445	0.0617	0.241	5555	0.0621	0.241
... X (formerly known as Twitter)	5445	0.103	0.304	5555	0.0956	0.294
Gender						
... Female	3984	0.547	0.498	4016	0.559	0.497
... Male	3984	0.453	0.498	4016	0.441	0.497
Race (US)						
... Black	3482	0.0942	0.292	3518	0.101	0.301
... Hispanic	3482	0.159	0.366	3518	0.159	0.366
... Other	3482	0.0939	0.292	3518	0.0995	0.299

... White	3482	0.653	0.476	3518	0.641	0.48
Race (UK)						
... Other race/nationality	1963	0.145	0.352	2037	0.154	0.361
...English / Welsh / Scottish / Northern Irish / British	1963	0.855	0.352	2037	0.846	0.361
Age (in years)	5445	45.1	17.2	5555	44.9	17.2
Age Group						
... 18-24	5445	0.134	0.341	5555	0.137	0.343
... 25-34	5445	0.2	0.4	5555	0.2	0.4
... 35-44	5445	0.193	0.395	5555	0.196	0.397
... 45-54	5445	0.154	0.361	5555	0.156	0.363
... 55+	5445	0.319	0.466	5555	0.311	0.463
Education (US)						
... College or more	3482	0.424	0.494	3518	0.441	0.497
... less than college	3482	0.576	0.494	3518	0.559	0.497
Education (UK)						
... College or more	1963	0.393	0.488	2037	0.399	0.49
... Less than College	1963	0.517	0.5	2037	0.514	0.5
... Student	1963	0.0902	0.286	2037	0.0869	0.282
Country						
... UK	5445	0.361	0.48	5555	0.367	0.482
... US	5445	0.639	0.48	5555	0.633	0.482

4 Empirical Analysis and Results

4.1 Econometric method

We have adopted regression-based techniques. Our approach involves utilizing a logit model to estimate Equation (1), allowing us to delve into the nuanced impacts of prompts on our key focal points—namely, opt-in decision and comprehension of data collection. Equation (2) will be leveraged to gauge the influence of moderator variables (X_i) on the primary effects, an aspect we expect to potentially alter both the magnitude and direction of the primary effects. We use the following logit models:

Estimation Equations:

$$Y_i = \beta_0 + \beta_1 \times \text{Prompt}_i + \epsilon_i \quad (1)$$

$$Y_i = \beta_0 + \beta_1 \times \text{Prompt}_i + \beta_2 \times X_i + \beta_3 \times \text{Prompt}_i \times X_i + \epsilon_i \quad (2)$$

4.2 Choice of Primary Estimands

4.2.1 User opt-in/Opt-out based on the prompt exposure

Our primary outcome is the average treatment effect of ATT prompt on user response to the question ‘Which option would you select in the featured prompt?’ The logit outcome of above mentioned equation 1, where the independent variable is survey respondents’ response on whether a respondent would select the opt-in option (coded as 1) or not (coded as 0), is presented below:

Table 2: ATT Prompt Opt-In Response

Coefficient Estimates				
Variable	Estimate	Std. Error	z value	Pr(> z)
Intercept	-1.0885	0.0313	-34.7724	0.0000***
ATT Prompt	-0.8351	0.0510	-16.3744	0.0000***

Marginal Effects				
Variable	dF/dx	Std. Err.	z value	P(> z)
ATT Prompt	-0.1244	0.0074	-16.8021	0.0000***

Note: The table reports coefficient estimates from a logit model. The dependent variable is the user opt-in decision in response to the exposed prompt. Consistent with Apple’s definitions, the opt-in decisions were ‘allow tracking’ and ‘turn on personalized ads’ for an ATT prompt and a PA prompt, respectively. The estimates indicate the effect of the ATT Prompt on the likelihood of user opt-in. Standard errors are provided in column 3. The sample size is **n=11,000**. Significance levels are indicated as follows: ***p < 0.01, **p < 0.05, *p < 0.1.

The marginal effects show the change in the probability of user opt-in due to the ATT Prompt. Standard errors for the marginal effects are provided in column 3. Significance levels for the marginal effects are indicated as follows: ***p < 0.01, **p < 0.05, *p < 0.1.

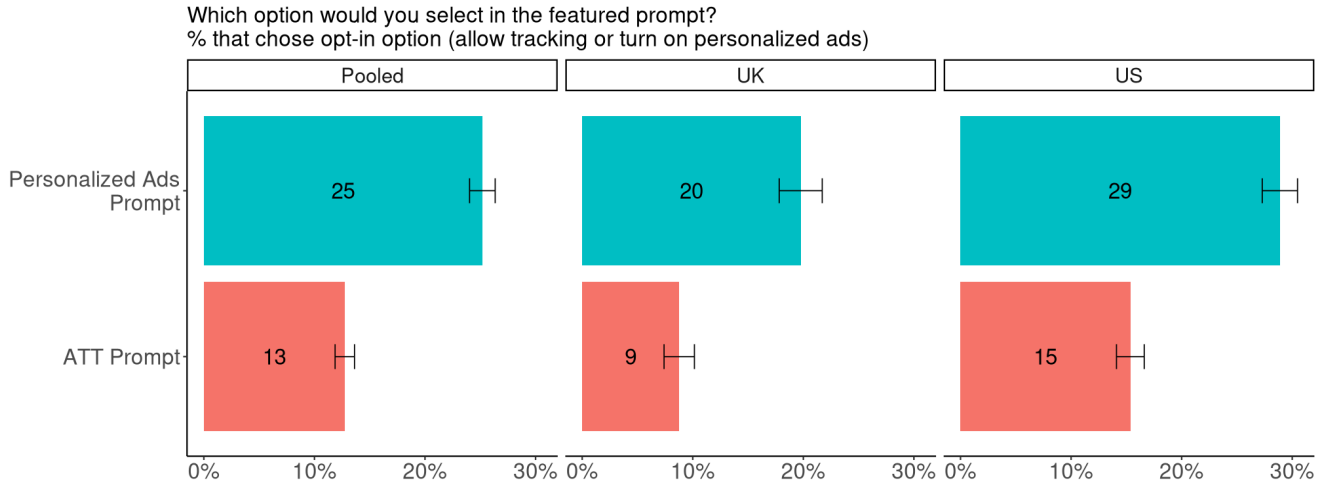


Figure 5: Average opt-in rates with 95% confidence intervals across prompt assignment

The logit model output reveals the impact of ATT prompts on user opt-in behavior. The coefficient estimates illustrate the baseline likelihood of users opting in, represented by the intercept. In this case, the intercept of -1.0885 signifies the log odds of users opting in when exposed to a PA prompt. Figure 5 shows the unadjusted average opt-in rates across the two prompt assignments for the entire sample as well as by country.

The coefficient for the ATT prompt, which is -0.8351, shows how exposure to the prompt influences user behavior. This negative coefficient suggests that the presence of the ATT prompt decreases the log odds of users opting in. The substantial magnitude of the coefficient, along with the high z value and statistically significant p-value ($\Pr(> |z|) = 0.0000$), underscores the robustness of this effect.

The marginal effect of the ATT prompt is calculated as -0.1244. This represents the change in the probability of opting in for each unit change in the ATT prompt variable. In more intuitive terms, for every user exposed to the ATT prompt, there is a 12.44 percentage point decrease in the likelihood of opting in; this suggests that the opt-in rate through the Personalize Ads prompt could be 92% higher. This significant decrease in the probability of opt-in, supported by the standard error and the z value, highlights the impact of the ATT prompt on shaping user decisions regarding data sharing. These results are robust to the inclusion of demographic controls, see appendix tables A1 and A2.

4.2.2 User opt-in/Opt-out - Preference for personalized advertisement

Furthermore, we explore if a user’s state opt-in decision interacted with their preference for personalized advertisement. We randomly asked half of the respondents whether they like personalized ads or not prior to the prompt and the other half, after seeing a prompt. This randomization was independent of both the prompt assignment randomization and the app assignment randomization. Therefore, we can explore whether preference for personalized ads moderates the effect of the prompts.

The impact of the ATT prompt on opt-in behavior varies depending on users’ preference for personalized advertisements. The interaction term (ATT Prompt:Like Pers. Ads) (-0.3828) sheds light on the combined effect of the ATT prompt and preference for personalized ads on user opt-in behavior. This coefficient suggests that the presence of the interaction term further decreases the log odds of users opting in, indicating a moderation effect is not in the direction that would be expected. In other words, those who like personalized ads are even more discouraged from opting in when presenting with an ATT prompt despite the fact that opting in would improve ad personalization and suggests that the ATT prompt, as it is presenting information on data sharing, could lead users to engage in behavior that is inconsistent with their preferences.

Table 3: ATT Prompt x Preference for personalized advertisement

Variable	Coefficient Estimates			
	Estimate	Std. Error	z value	Pr(> z)
Intercept	-1.8034	0.0625	-28.8467	0.0000***
ATT Prompt	-0.8546	0.1070	-7.9882	0.0000***
Like Pers.Ads	2.2765	0.1029	22.1242	0.0000***
ATT Prompt:Like Pers.Ads	-0.3828	0.1607	-2.3813	0.0173**

Note: The table reports coefficient estimates from a logit model. The dependent variable is the user opt-in decision, with an interaction term for user preference for personalized advertisements. A respondent is coded to ‘Like Personalized Ads,’(=1) if they select that they ‘Like personalized ads a lot/a little;’ they are coded to not like (=0) if they select ‘Neither like nor dislike personalized ads’ or ‘Dislike personalized ads a lot/a little.’ The estimates indicate the effect of the ATT Prompt, the preference for personalized ads, and their interaction on the likelihood of user opt-in. Standard errors are provided in column 3. The sample size is **n=5,506**. Significance levels are indicated as follows: ***p < 0.01, **p < 0.05, *p < 0.1.

Marginal Effects				
Variable	dF/dx	Std. Err.	z value	P(> z)
ATT Prompt	-0.1076	0.0130	-8.2584	0.0000***
Like Pers.Ads	0.4073	0.0219	18.5574	0.0000***
ATT Prompt:Like Pers.Ads	-0.0429	0.0164	-2.6122	0.0089***

Note: The marginal effects show the change in the probability of user opt-in due to the ATT Prompt, liking personalized ads, and their interaction. Standard errors for the marginal effects are provided in column 3. The sample size is **n=5,506**. Significance levels for the marginal effects are indicated as follows: ***p < 0.01, **p < 0.05, *p < 0.1.

Our main effects are robust to this specification. Firstly, the intercept (-1.8034) represents the log odds of users opting in when they do not prefer personalized advertisements and the PA prompt is presented. This baseline likelihood serves as a reference point for understanding the impact of subsequent variables.

The negative coefficient for the ATT prompt (-0.8546) indicates that exposure to the prompt decreases the log odds of users opting in, holding preference for personalized ads constant. This reinforces evidence that the ATT prompt acts as a deterrent, reducing the likelihood of users agreeing to data sharing. Conversely, the positive coefficient for preference for personalized advertisements (2.2765) signifies that users with a preference for personalized ads are more likely to opt-in, in the absence of the ATT prompt.

The marginal effect of the ATT prompt (-0.1076) highlights the change in the probability of opting in when exposed to an ATT prompt, holding preference for personalized ads constant. Similarly, the marginal effect for preference for personalized advertisements (0.4073) signifies the change in the probability of opting in for users who have a preference for personalized ads, in the absence of the ATT prompt. Finally, the marginal effect for the interaction term (-0.0429) reveals how the presence of the interaction term influences the likelihood of opting in.

4.2.3 User understanding and interpretation of prompts

With the similar estimation setup for previous estimands, we explore the user understanding and interpretation of the prompts too. Hutton and Ellis (2023) find that 43% of their participants held incorrect beliefs about what tracking does. The most common misconception was related to location data; nearly a quarter of respondents mistakenly believed that accepting a tracking request would share their location with the requesting app. This analysis delves into whether an ATT prompt itself leads to a greater misconception around location data.

ATT Prompt coefficient, which registers at 0.4395, suggests a significant positive increase in the likelihood in the misconception (by 9.1 p.p.) that opting in shares location relative to the PA

prompt. This finding suggests that users tend to link the act of allowing tracking with granting the assigned app permission to engage in location tracking activities; in the real world, this would affect non-Apple developers.

These findings highlight the impact of the ATT prompt on users' comprehension and interpretation of location data collection prompts. They underscore how certain cues and language choices influence users' perceptions of data collection that could differentially impact their understanding of how third-party developers and Apple handle their data for advertising.

Location:

Table 4: ATT Prompt x Understanding Location data collection

Coefficient Estimates				
Variable	Estimate	Std. Error	z value	Pr(> z)
Intercept	0.6416	0.0286	22.4697	0.0000***
ATT Prompt	0.4395	0.0421	10.4464	0.0000***

Note: The table reports coefficient estimates from a logit model. The dependent variable is the user interpretation of location data collection in response to the ATT Prompt. The estimates indicate the effect of the ATT Prompt on the likelihood of believing that opting in shares location with the assigned app. Standard errors are provided in column 3. The sample size is **n=10,967**. Significance levels are indicated as follows: ***p < 0.01, **p < 0.05, *p < 0.1.

Marginal Effects				
Variable	dF/dx	Std. Err.	z value	P(> z)
ATT Prompt	0.0916	0.0087	10.5212	0.0000

Note: The marginal effects show the change in the probability of interpreting whether opting in would share their location data due to the ATT Prompt. Standard errors for the marginal effects are provided in column 3. The sample size is **n=10,967**. Significance levels for the marginal effects are indicated as follows: ***p < 0.01, **p < 0.05, *p < 0.1.

4.2.4 User opt-in/Opt-out Across distinct applications (categories)

Table 5: ATT Prompt X App assignment

Coefficient Estimates				
Variable	Estimate	Std. Error	z value	Pr(> z)
Intercept	-0.978	0.068	-14.415	0.000***
ATT Prompt	-0.814	0.109	-7.491	0.000***
Deliveroo	-0.705	0.181	-3.890	0.000***
DoorDash	-0.116	0.150	-0.777	0.437
ESPN	-0.001	0.155	-0.008	0.994
Facebook	-0.270	0.125	-2.154	0.031**
Fox News	0.037	0.143	0.255	0.798
Instagram	-0.168	0.119	-1.411	0.158
MSNBC	-0.298	0.167	-1.785	0.074*
Nike	-0.019	0.124	-0.157	0.875
Sephora	0.057	0.122	0.467	0.640
Walmart	0.093	0.138	0.672	0.502
X (formerly known as Twitter)	-0.309	0.123	-2.507	0.012**
ATT Prompt:Deliveroo	-0.108	0.311	-0.347	0.728
ATT Prompt:DoorDash	0.017	0.250	0.067	0.947
ATT Prompt:ESPN	0.385	0.230	1.672	0.095*
ATT Prompt:Facebook	-0.177	0.207	-0.854	0.393
ATT Prompt:Fox News	0.175	0.222	0.787	0.432
ATT Prompt:Instagram	-0.192	0.202	-0.947	0.343
ATT Prompt:MSNBC	0.351	0.255	1.373	0.170
ATT Prompt:Nike	0.083	0.193	0.427	0.670
ATT Prompt:Sephora	-0.276	0.206	-1.338	0.181
ATT Prompt:Walmart	0.036	0.219	0.165	0.869
ATT Prompt:X (formerly known as Twitter)	-0.489	0.227	-2.158	0.031**

Note: The table reports coefficient estimates from a logit model. The dependent variable is the user opt-in/opt-out decision across distinct applications (categories), reference category is Apple. The estimates indicate the effect of the ATT Prompt and its interaction with various applications on the likelihood of user opt-in. Standard errors are provided in column 3. The sample size is $n=11,000$. Significance levels are indicated as follows: *** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$.

The analysis extends to the examination of user opt-in and opt-out behavior across different mobile applications. This includes apps from social media companies, e-commerce platforms, news media apps, and consumer goods like beauty products and sports products. The output of the model, based on Equation 2 above, unfolds crucial insights:

ATT Prompt coefficient, registered at -0.8139, reveals a significant negative effect between exposure to the ATT prompt and users' likelihood of opting in across various application assignments. The magnitude and statistical significance are similar to our main results. As discussed earlier, this implies that the presentation of the ATT prompt tends to discourage users from sharing their data with third-party companies. The coefficients corresponding to specific application assignments exhibit varying degrees of influence on user decisions, as evidenced by the magnitude and direction of their respective coefficients. The interaction terms between the ATT prompt and application assignments show non-moderation effects. With the exception of X (formerly Twitter), the interaction effects are not statistically different from 0. These interactions provide insights that the influence of the ATT prompt is applicable across different applications, meaning the app context and/or categories (such as sports, news, social media, retail, etc.) do not really moderate the ATT impact. ATT framing is affecting all. This evidence suggests that Apple's opt-in rates for personalized ads would drop by nearly half if it were to use the prompt it requires for third-party apps, see Figure 5.

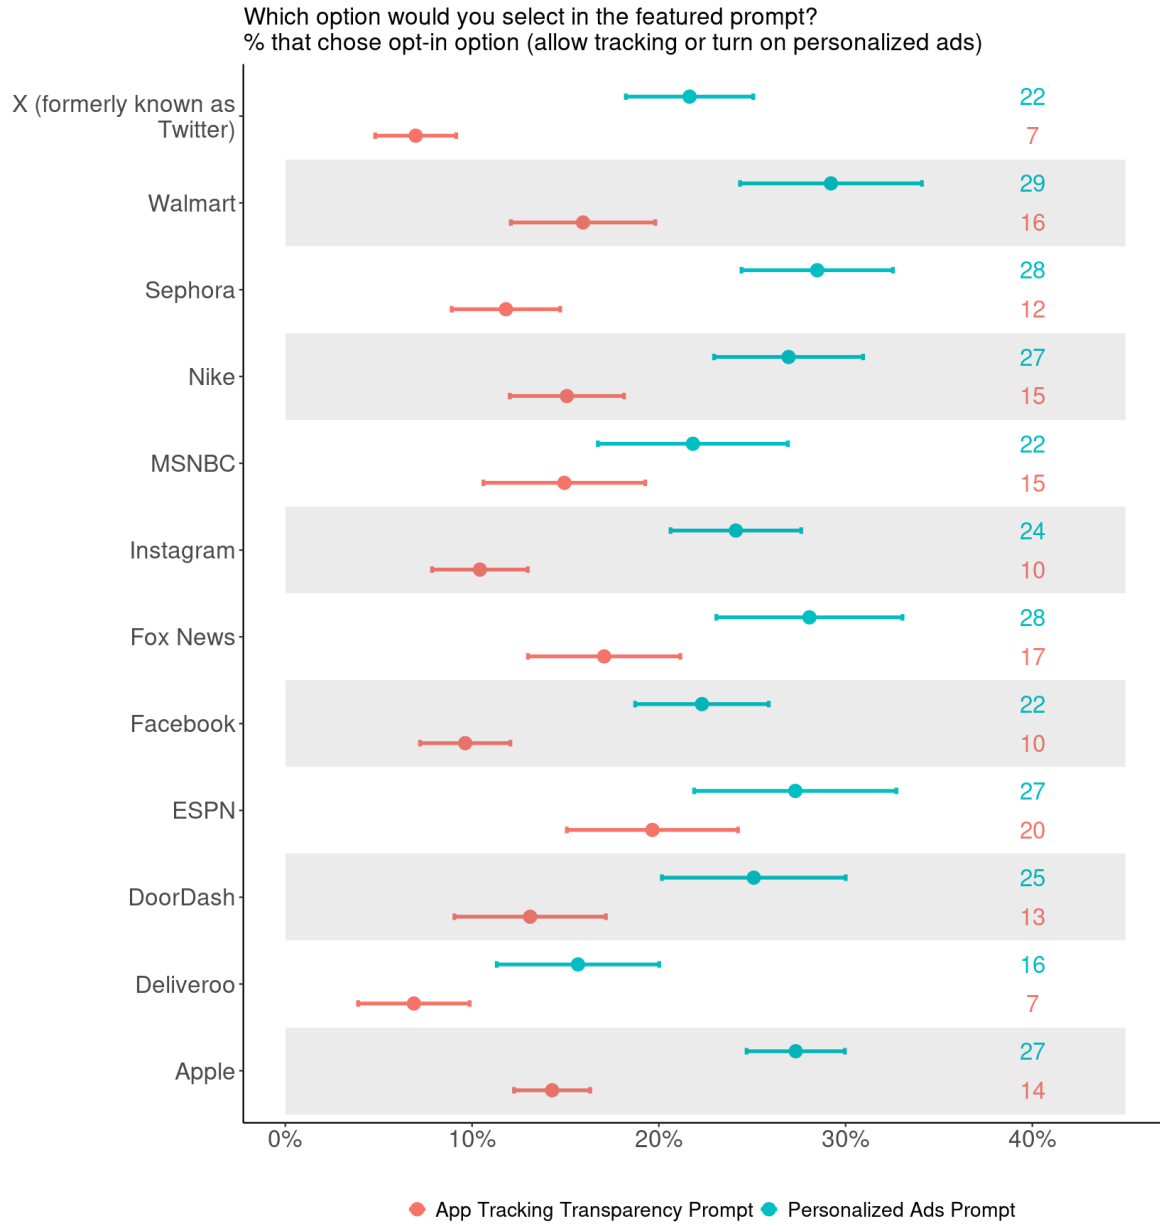


Figure 5: Average opt-in rates with 95% confidence intervals across app assignments.

The findings underscore the nuanced dynamics of user opt-in and opt-out behavior across various apps; the ATT prompt exerts a significant discouraging effect on data sharing with other companies. These insights emphasize the importance of impartial and equitable platform policies.

5 Conclusion

In this study, we conducted an empirical analysis to investigate the impacts of Apple’s App Tracking Transparency (ATT) prompt and Personalized Ads (PA) prompt on user behavior and perception within the digital landscape. Our findings reveal substantial disparities in user reactions depending on the specific prompt encountered, highlighting the pivotal role of choice architecture in shaping user decision-making processes regarding data sharing and ad personalization. Our research introduces a new dimension to the existing literature (e.g., Adjerid et al. 2018, Lin and Strulov-Shlain 2023) by empirically investigating how specific data consent prompts influence user behavior and beliefs within Apple’s ATT and PA prompt frameworks that potentially have wide-reaching effects across the entire mobile application and advertising ecosystem. We find that the ATT prompt can lead users to believe that their location is being tracked and can discourage users from opting into data sharing that would have real world effects on third-party companies that create iOS apps.

The results from this research bear several important policy implications. Specifically, while regulatory bodies and the general public express well-meaning concerns about privacy and data sharing, the actual implementation of interventions carries great meaning in terms of outcomes for users. Here, we find that nuances in the wording of data sharing opt-out prompts carry great implications in terms of user beliefs as well as user behavior. Crucially, in the present matter, we find that a large company (Apple), through changes in the nuanced framing of these opt-out prompts, may be able to influence users to opt out of sharing data with third-party entities (and negatively influence their beliefs about those companies’ data sharing behavior) while at the same time avoiding this outcome for its own data use. In other words, for policies that purportedly promote consumer privacy, the *implementation* of the policy may dictate outcomes that run counter to their stated objectives.

References

- Acquisti, A., Adjerid, I., Balebako, R., Brandimarte, L., Cranor, L. F., Komanduri, S., Leon, P. G., Sadeh, N., Schaub, F., Sleeper, M., Wang, Y., and Wilson, S. (2017). Nudges for privacy and security: Understanding and assisting users' choices online. *ACM Computing Surveys*, 50(3):1–41.
- Adjerid, I., Acquisti, A., and Loewenstein, G. (2018). Choice architecture, framing, and cascaded privacy choices. *Management Science*, 65(5):2267–2290.
- Apple Inc. (July 2024.). Apple advertising & privacy. Retrieved from <https://www.apple.com/legal/privacy/data/en/apple-advertising/>
- Bleier, A., Eisenbeiss, M. (2015) Personalized Online Advertising Effectiveness: The Interplay of What, When, and Where. *Marketing Science* 34(5):669-668
- Cecere, Grazia and Lemaire, Sarah, Have I Seen you Before? Measuring the Value of Tracking for Digital Advertising (December 1, 2023). SSRN: <https://ssrn.com/abstract=4659963>
- Competition and Markets Authority. (2021). *Mobile Ecosystems Market Study: Appendix J*. Source: <https://www.gov.uk/cma-cases/mobile-ecosystems-market-study>.
- DeGiulio, A., Lee, H., and Birrell, E. (2021). “ask app not to track”: The effect of opt in tracking authorization on mobile privacy. In Saracino, A. and Mori, P., editors, *Emerging Technologies for Authorization and Authentication*, pages 152–167, Cham. Springer International Publishing.
- Deisenroth, D., Manjeer, U., Sohail, Z., Tadelis, S., & Wernerfelt, N. (2024). The Economic Consequences of Data Sharing Restrictions. [Working paper]
- Eberhard, K. (2023). The effects of visualization on judgment and decision-making: a systematic literature review. *Management Review Quarterly*, 73:167–214.
- Ghose, A., & Todri-Adamopoulos, V. (2016). Toward a Digital Attribution Model: Measuring the Impact of Display Advertising on Online Consumer Behavior. *MIS Quarterly*, 40(4), 889–910.
- Ghose, A. (2018). *Tap - Unlocking the Mobile Economy*. The MIT Press, ISBN: 9780262536059
- Ghose A, Li B, Liu S (2019) Mobile targeting using customer trajectory patterns. *Management Science* 65(11):5027–5049.
- Hogarth, R. M. and Einhorn, H. J. (1992). Order effects in belief updating: The belief-adjustment model. *Cognitive Psychology*, 24(1):1–55.
- Hutton, H. J. and Ellis, D. A. (2023). Exploring user motivations behind ios app tracking transparency decisions. In *CHI '23: Proceedings of the 2023 CHI Conference on Human Factors in Computing Systems*, pages 1–12. ACM.
- Jin, J., Wuke, Z., and Chen, M. (2017). How consumers are affected by product descriptions in online shopping: Event-related potentials evidence of the attribute framing effect.

Neuroscience Research, 125.

Johnson, E. J., Shu, S. B., Dellaert, B. G. C., Fox, C., Goldstein, D. G., H'auhl, G., Larrick, R. P., Payne, J. W., Peters, E., Schkade, D., Wansink, B., and Weber, E. U. (2012). Beyond nudges: Tools of a choice architecture. *Marketing Letters: A Journal of Research in Marketing*, 23(2):487–504.

Kraft, Lennart and Skiera, Bernd and Koschella, Tim, Economic Impact of opt in versus Opt-out Requirements for Personal Data Usage: The Case of Apple's App Tracking Transparency (ATT) (October 10, 2023).

Kesler, R. (2022). The impact of Apple's app tracking transparency on app monetization. Available at SSRN 4090786.

Kollnig, K., Shuba, A., Kleek, M. V., Binns, R., and Shadbolt, N. (2022). Goodbye tracking? impact of iOS app tracking transparency and privacy labels. In *2022 ACM Conference on Fairness, Accountability, and Transparency*. ACM.

Levin, I. P., Schneider, S. L., and Gaeth, G. J. (1998). All frames are not created equal: A typology and critical analysis of framing effects. *Organizational Behavior and Human Decision Processes*, 76(2):149–188.

Leon, Pedro & Rao, Ashwini & Schaub, Florian & Marsh, Abigail & Cranor, Lorrie & Sadeh, Norman. (2015). Why People Are (Un)willing to Share Information with Online Advertisers.

Li, D. and Tsai, H.-T. (2022). Mobile apps and targeted advertising: Competitive effects of data exchange. Available at SSRN 4088166.

Lin, Tesary & Strulov-Shlain, Avner. (2023). Choice Architecture, Privacy Valuations, and Selection Bias in Consumer Data.

Meghanath Macha, Natasha Zhang Foutz, Beibei Li, Anindya Ghose (2023) Personalized Privacy Preservation in Consumer Mobile Trajectories. *Information Systems Research* 35(1):249-271.

Reidenberg, J. R., Breaux, T., Cranor, L. F., French, B. M., Grannis, A., Graves, J. T., Liu, F., McDonald, A., Norton, T. B., Ramanath, R., Russell, N. C., Sadeh, N., and Schaub, F. (2015). Disagreeable privacy policies: Mismatches between meaning and users' understanding. *Berkeley Technology Law Journal*, 30:39.

Sokol, D. D. and Zhu, F. (2021). Harming competition and consumers under the guise of protecting privacy: An analysis of Apple's iOS 14 policy updates.

Thaler, R. H., Sunstein, C. R., and Balz, J. P. (2010). Choice architecture.

Sun, T., Yuan, Z., Li, C., Zhang, K., Xu, J. (2023) The Value of Personal Data in Internet Commerce: A High-Stakes Field Experiment on Data Regulation Policy. *Management Science* 70(4):2645-2660.

Turvey, B. and Freeman, J. (2012). Jury psychology. In Ramachandran, V., editor, *Encyclopedia of Human Behavior (Second Edition)*, pages 495–502. Academic Press, San Diego, second

edition edition.

Appendix

1. ATT and Apple's Personalized Ads Prompts



Figure A1: ATT prompts used by the apps between February 2024 and April 2024.

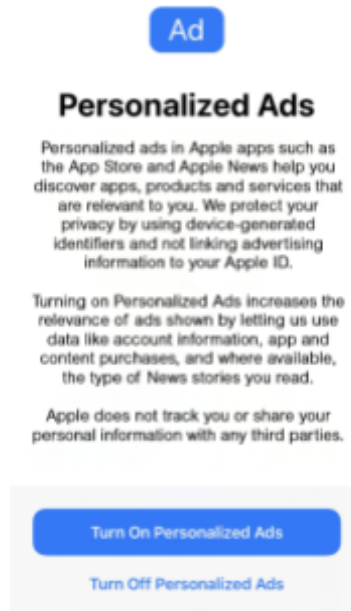


Figure A2: Apple’s personalized ads prompts used in February 2024.

2. Opt-in Responses with Demographic Controls

Table A1: ATT prompts effect on Opt-in with demographic controls

Coefficient Estimates				
Variable	Estimate	Std. Error	z value	Pr(> z)
(Intercept)	-1.194	0.087	-13.645	0.000***
ATT Prompt	-0.947	0.063	-14.926	0.000***
Gender=Male	0.329	0.061	5.381	0.000***
Age: 25-34	-0.171	0.095	-1.794	0.073*
Age:35-44	-0.336	0.098	-3.419	0.001***
Age:45-54	-0.644	0.108	-5.940	0.000***
Age:55+	-1.164	0.102	-11.455	0.000***
Country=US	0.743	0.063	11.733	0.000***

Note: The table reports coefficient estimates from a logit model. The dependent variable is the user opt-in/opt-out decision. The reference categories are as follows: Gender=Female, Age:18-24, Country=UK. Standard errors are provided in column 3. The sample size is **n=7,976**. Significance levels are indicated as follows: ***p < 0.01, **p < 0.05, *p < 0.1.

Table A2: ATT prompts effect on Opt-in with demographic controls (US)

Coefficient Estimates				
Variable	Estimate	Std. Error	z value	Pr(> z)
(Intercept)	-0.067	0.226	-0.296	0.768
ATT Prompt	-0.949	0.091	-10.476	0.000***
Gender=Male	0.509	0.089	5.718	0.000***
Age: 25-34	-0.029	0.163	-0.178	0.859
Age:35-44	-0.071	0.166	-0.428	0.668
Age:45-54	-0.690	0.180	-3.838	0.000***
Age:55+	-1.277	0.163	-7.825	0.000
Education=Less than college	0.082	0.098	0.828	0.408
Race=White	-0.534	0.145	-3.672	0.000***
Race= Hispanic	-0.420	0.169	-2.488	0.013**
Race=Other	-1.119	0.209	-5.345	0.000***
Income=Less than \$30,000	0.127	0.147	0.866	0.386
Income= \$30,000-\$69,999	-0.038	0.129	-0.296	0.767
Income= \$70,000-\$119,999	0.027	0.124	0.221	0.825

Note: The table reports coefficient estimates from a logit model. The dependent variable is the user opt-in/opt-out decision. The reference categories are as follows: Gender=Female, Age:18-24, Education=College or above, Race= Black, Income= \$120,000+. Standard errors are provided in column 3. The sample size is **n=3,326**. Significance levels are indicated as follows: ***p < 0.01, **p < 0.05, *p < 0.1.

Table A3: ATT prompts effect on Opt-in with demographic controls (UK)

Coefficient Estimates				
Variable	Estimate	Std. Error	z value	Pr(> z)
(Intercept)	-0.982	0.309	-3.178	0.001***
ATT Prompt	-0.919	0.139	-6.611	0.000***
Gender=Male	0.329	0.135	2.427	0.015**
Age: 25-34	-0.292	0.228	-1.279	0.201
Age:35-44	-0.602	0.242	-2.487	0.013**
Age:45-54	-0.649	0.260	-2.501	0.012**
Age:55+	-0.602	0.256	-2.346	0.019**
Education= Less than College	-0.041	0.147	-0.281	0.779
Education= Student	0.407	0.287	1.420	0.156
Race= English / Welsh / Scottish / Northern Irish / British	-0.007	0.187	-0.036	0.971
Income= Less than £30,000	-0.070	0.225	-0.309	0.757
Income= £30,000- £59,999	-0.195	0.210	-0.930	0.352
Income= £60,000 - £99,999	-0.143	0.221	-0.647	0.518

Note: The table reports coefficient estimates from a logit model. The dependent variable is the user opt-in/opt-out decision. The reference categories are as follows: Gender=Female, Age:18-24, Education=College or above, Race= Other ethnic category, Income= £100,000+. Standard errors are provided in column 3. The sample size is **n=1,879**. Significance levels are indicated as follows: ***p < 0.01, **p < 0.05, *p < 0.1.