"Ask App Not to Track": The Effect of Opt-in Tracking Authorization on Mobile Privacy*

Anzo DeGiulio
Pomona College
abdb2018@mymail.pomona.edu

Hanoom Lee
Pomona College
hlaa2020@mymail.pomona.edu

Eleanor Birrell
Pomona College
eleanor.birrell@pomona.edu

Abstract

App Tracking Transparency (ATT) introduces opt-in tracking authorization for iOS apps. In this work, we investigate how mobile apps present tracking requests to users, and we evaluate how the observed design patterns impact users' privacy. We perform a manual observational study of the Top 200 free iOS apps, and we classify each app by whether it requests permission to track, the purpose of the request, how the request was framed, whether the request was preceded or followed by additional ATT-related pages, and whether the request was preceded or followed by other permission requests. We then perform a user study with 950 participants to evaluate the impact of the observed UI elements. We find that opt-in authorizations are effective at enhancing data privacy in this context, and that the effect of ATT requests is robust to most implementation choices.

1 Introduction

App Tracking Transparency (ATT)—introduced into iOS 14.5 in April 2021—is a new authorization model that requires opt-in user consent for tracking by mobile apps. Any mobile app that collects and shares user data for tracking purposes is required to use the ATT framework to request tracking permission from the user. Users are then presented with standardized permission dialogue (Figure 1) and asked to choose between allowing tracking and asking the app not to track; Apple's User Privacy and Data Use policy prohibits apps from using identifiers, fingerprinting, or other techniques to track users who ask the app not to track. This work investigates how iOS mobile apps present tracking requests to users and evaluates how observed design patterns impact privacy.

We began by qualitatively coding ATT requests by the Top 200 free iOS apps on June 1, 2021. We found that almost all ($n = 197$) of the apps had been updated since the release of iOS 14.5 and that approximately half of the updated apps ($n = 91$) requested permission to track their users. Using a coding book developed based on experience on personal devices, we also classified each request by the purpose of the request, how the request was framed, whether the request was preceded or followed by additional ATT-related pages, and whether the request was preceded or followed by other permission requests.

To understand the effect of observed design patterns on data privacy, we conducted a user study with 950 participants recruited through Amazon Mechanical Turk. We found that tracking requests for non-advertising purposes were significantly more likely to be granted than tracking requests for advertising purposes ($p = .05$). The framing of the study (positive, neutral, or negative, and with or without the threat of future required payment) had no significant effect on opt-out rate. Surprisingly, the presence of a priming page appeared to reduce, rather than increase, the number of users who authorized tracking, although the difference was not statistically significant.

In a follow-up survey, 67.9% of respondents reported being somewhat or very uncomfortable with tracking, but 60.5% of respondents reported being very or somewhat satisfied the ATT opt-out provided; there were no significant differences in satisfaction between conditions. Moreover, 96.1% of respondents currently running iOS 14.5 or higher reported opting-out of tracking on their personal device at least a few times.

These results show that opt-in permissions are highly effective at enhancing data privacy in the context of tracking by mobile apps. Moreover, opt-in rates were relatively consistent across most conditions, which suggests that ATT is less impacted by dark patterns and other privacy-diminishing UI elements than other preference-setting mechanisms. This suggests that ATT—which requires opt-in consent with clearly defined options presented through a standardized interface—might prove an effective model for managing data privacy in other contexts.

*This work appeared in the 4th International Workshop on Emerging Technologies for Authorization and Authentication (ETAA), 2021.

Figure 1: Example ATT permission dialogue. The format of the pop-up, including the prompt and the opt-in/opt-out buttons, is standardized. The app developer specifies the app-defined text displayed between the prompt and the buttons.

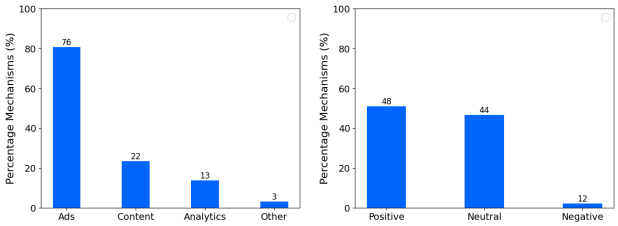

(a) Purpose (b) Framing

Figure 2: Frequency of different types of text in ATT requests among the Top 200 free iOS apps that request permission to tracking.

2 ATT Pop-ups in the Wild

To understand how apps present ATT requests, we conducted a manual user study of the Top 200 free iOS apps.

2.1 Methodology

We developed a coding book based on observed ATT requests during daily use of personal devices. Our coding book included seven distinct features:

1. Does the app request permission to track? (Yes / No)
2. Why does the app request tracking? (First party ads / Third party ads / Ads (unspecified) / Content / Analytics / Other)
3. How is tracking framed? (Positive / Neutral / Negative)
4. Does the app threaten payment if not allowed to track? (Yes / No)
5. Is there a pre-request ATT priming page? (Yes / No)
6. Is there a post-request follow-up page? (Yes / No)
7. How many other pop-up permissions does the app request? (n)

We manually examined and classified the Top 200 free iOS apps (as listed by Sensor Tower [36] on June 1, 2021) according to this coding book. Observations were conducted on iPhone 7s running iOS 14.6 on June 1, 2021.

For each app, one author installed and ran the app. To determine whether the app requested permission to track (1), we then ran the app for a minimum of 30 seconds: until we observed an ATT request, until it was clear that no request would be made, or until we ran into a roadblock that prevented us from proceeding.[1] If the app required an account, we created a new account for the purpose of this study.

For each app that requested permission to track, we qualitatively coded the app-defined text displayed in the ATT permission dialogue (Figure 1) for three features: the purpose for which the app requested permission to track (2), whether the request was framed positively (i.e., authorizing tracing would improve the user experience) negatively (i.e., denying permission to track would negatively impact the user experience) or using netural language (3), and whether the app threatened to start charging for service or features if users denied permission to track (4).

Finally, we classified the context in which the ATT request occurred: whether there was a pre-request priming page (5), whether there was a post-request follow-up page (6), and how many other permission requests were made by the app at the same time (7).

2.2 Results

We found that almost all ($n = 197$) of the apps had been updated since the release of iOS 14.5 and that approximately half of the updated apps ($n = 91$) requested permission to track their users.

The majority of apps that requested permission to track (81.9%) did so in order to support some form of behavioral advertising. Some of these apps ($n = 33$) specified that tracking was needed to support targeted ads on their app (i.e., first party ads) and some ($n = 21$) stated that tracking was used to support ads on other websites and apps (i.e., third party ads); other apps used vague language that did not clearly specify whether tracking data would be used for first party ads, third party ads, or both. The remaining 18.1% of apps requested permission for tracking only for non-advertising purposes; these purposed included improving or personalizing app content, analytics, and other purposes. These results are summarized in Figure 2a.

Most apps that requested permission used a positive framing that emphasized the benefits of allowing tracking: 51.1% claimed that allowing tracking would allow the app to serve better ads or offer a better experience. Only 2.1% of apps used negative framing (i.e., emphasizing the downsides of opting-out of tracking by claiming it would would result in

[1] Roadblocks included requests for payment or social security numbers (SSNs).

Cond.	App-defined Text	Purpose	Framing	Payment	Priming
0	Your data will be used to show you better and more relevant ads in this app.	Ads (1st)	Positive	No	No
1	Your data will be shared with our partners to show you better and more relevant ads outside of All News.	Ads (3rd)	Positive	No	No
2	Your data will be used to provide you with a better and more relevant ad experience.	Ads (vague)	Positive	No	No
3	Your data will be used to show you better and more relevant articles.	Content	Positive	No	No
4	This identifier will be used to deliver personalized ads to you.	Ads (vague)	Neutral	No	No
5	Selecting "Ask App Not to Track" will result in less relevant ads.	Ads (vague)	Negative	No	No
6	Your data will be used to provide you with a better and more relevant ad experience and keep All News free to use.	Ads (vague)	Positive	Yes	No
7	Selecting "Ask App Not to Track" will result in less relevant ads and may require us to start charging you to use All News.	Ads (vague)	Negative	Yes	No
8	Your data will be used to provide you with a better and more relevant ad experience.	Ads (vague)	Positive	No	Yes

Table 1: The nine conditions included in our user study.

less relevant ads or content); the remaining apps used neutral language. These results are depicted in Figure 2b.

Other factors we studied were relatively uncommon in the wild. Only 7.7% of tracking requests threatened to start requiring payment if users declined tracking. Less than a quarter of apps that requested permission to track (23.1%) primed users with an ATT-related page before requesting permission to track, and just one app presented a follow-up page when users asked the app not to track.

3 User Study Methodology

To understand the effect of the observed implementation choices for ATT requests on data privacy—and on user's likelihood of opting-out of tracking—we conducted a user study with 950 users.

In the user study, participants installed an aggregated news app called All News on their personal iOS device; they interacted with the app for a few minutes, and then answered a series of follow-up questions.

3.1 App Design and Conditions

All News is an aggregated news app developed for the purpose of this study. It fetches news articles from major sources using the News API and displays them to users. A screenshot of the All News home page is shown in Figure 3a. When opening the app for the first time, a pop-up identical to Apple's official ATT request appears; it asks the user for permission to track. We varied the app-defined language used in the ATT request between conditions, and each user was pseudorandomly assigned to a condition based on a hash of their IP address. An example ATT request is depicted in Figure 3b.

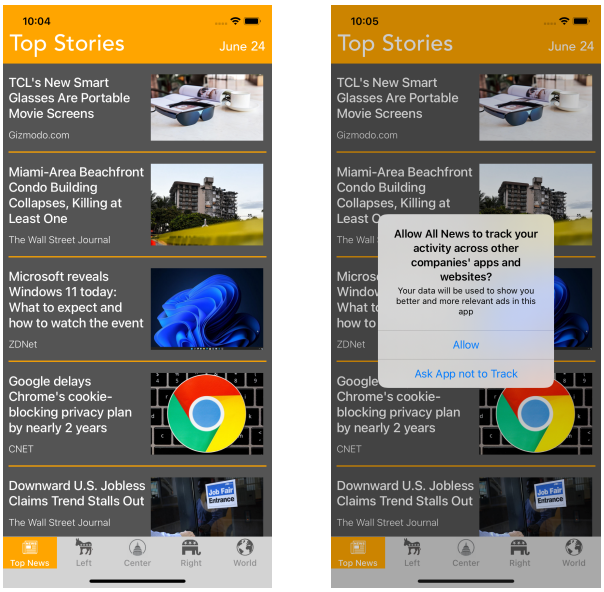

(a) Home screen (b) ATT permission request

Figure 3: Screenshots from the All News app used in our user study. The app-defined text in the ATT permission request dialogue varied between condition.

The first four conditions correspond to different purposes: first party ads, third party ads, ads (unspecified), and content. The next two conditions vary the framing of the request (from positive to neutral or negative). The next two conditions are variants of the positive and negative conditions that threaten to introduce payments if users ask not to track. Purpose and framing were controlled by adjusting the language of the app-defined text that appears in the ATT permission dialogue

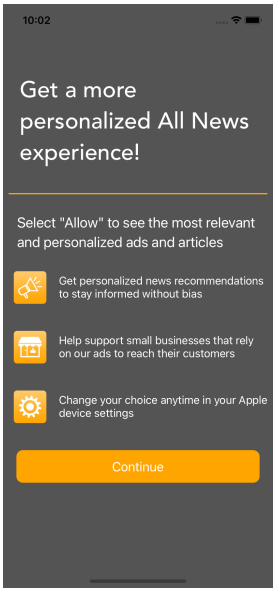

Figure 4: ATT priming page used in Condition 8.

between the standardized prompt and the standardized opt-in and opt-out buttons. The final condition introduces a priming page (shown in Figure 4) prior to the ATT permission request. These conditions are detailed in Table 1.

Data logged by the application included a unique, randomly generated, 32 character identifier that was assigned to a user upon opening the app. We also recorded which condition the participant was assigned to and whether or not the participant authorized tracking. No other information was collected by the app; the app did not actually receive or record any ATT tracking identifiers.

3.2 Participant Recruitment

In order to include the follow-up questions—and in order to vary the text that appeared in ATT requests, a feature that is not supported by the iOS App Store—we published our app through Expo Go and recruited participants through Amazon Mechanical Turk. The task was advertised as beta-testing an aggregated news app. Participation was restricted to iPhone users who had previously completed at least 50 HITs with an acceptance rate of at least 95%.

Before beginning the study, we informed users about our date collection practices and obtained participants' consent. We then guided them through installing and running the All News app and asked each participant to use the app for a few minutes. After collecting their user identifier in the survey, we asked them a set of follow-up questions about their experience with ATT pop-ups and mobile tracking in general. We also collected basic demographic information. The complete survey is provided in Appendix A.

Responses that did not include a valid app confirmation

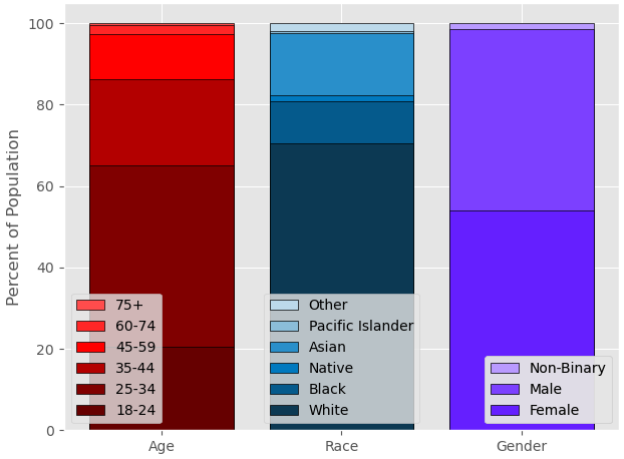

(a) Participant demographics.

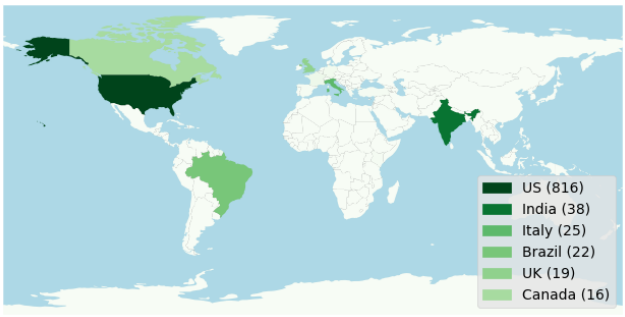

(b) Participant residency (for countries with 3+ participants)

Figure 5: Demographic summary of our 950 users study participants.

code (received after the user downloaded All News) or a valid qualtrics code (received after the user completed the follow-up survey) were rejected. Participants were compensated $1.20 upon successful completion of the survey.

We received 2298 preliminary responses; 1348 responses were rejected because they did not include a valid app confirmation code (received after the user downloaded All News) or they did not include a valid qualtrics code (received after the user completed the follow-up survey); we analyzed data from the remaining 950 participants. The median completion time for the user study was 5.4 minutes.

This study received an IRB exemption from the Institutional Review Board at our institution.

3.3 Participant Demographics

The 54.1% of our study participants identified as woman and 44.7% identified as men. The majority of study participants identified as white (70.5%), 15.2% identified as Asian, and 10.3% identified as Black. Less than 2% identified as Native American, as Pacific Islander, or as "other". 20.6% were 18-

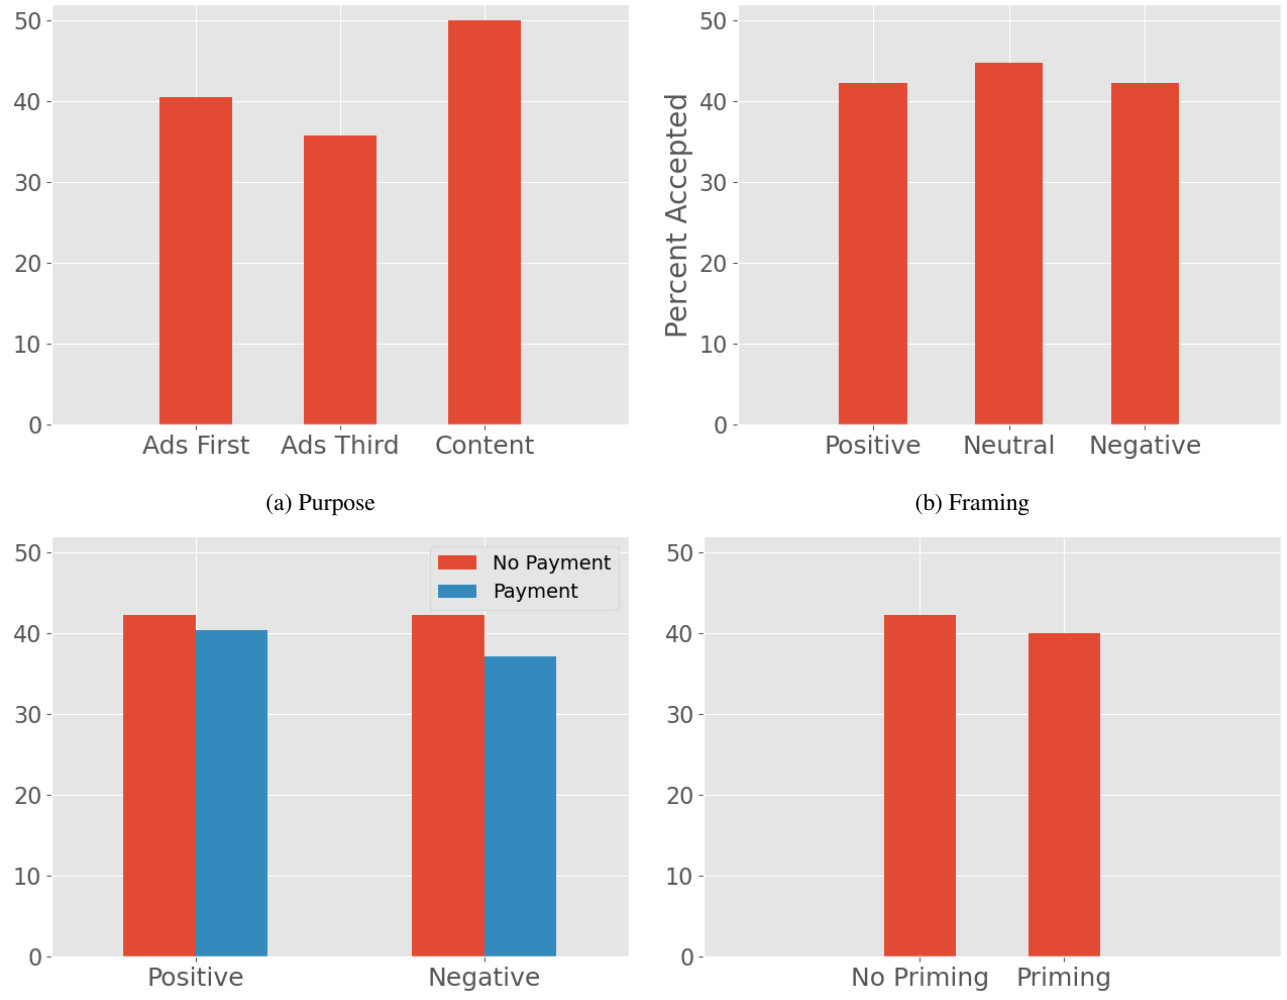

(a) Purpose

(b) Framing

(c) Payment

(d) Priming

Figure 6: Effect of various design patterns on tracking opt-in rates.

24, 44.5% were 25-34, and 21.2% were 35-44; the remaining participants were 45 or older. These demographics of the participants are shown in Figure 5a. Participants' self-reported country of residence is shown in Figure 5b.

We also collected data on participants' smartphone use. 99.3% of respondents used an iPhone as their primary smartphone, and 78.3% of them were running an iOS version 14.5 or higher (ATT enabled).

4 Results

Our user study evaluated the impact of four of the seven factors included in our observational study: (1) purpose, (2) framing, (3) payment, and (4) priming. Due to the infrequency of observed examples in the wild, we did not study the impact of follow-up pages or additional (non-ATT) permission requests. All conditions requested permission to track.

Purpose. We found that users were significantly more likely to authorize tracking for the purpose of personalizing content than for advertising purposes ($p = .05$). This effect was particularly strong for third-party ads: just 35.7% of users authorized tracking for the purpose of third-party advertising compared to 50.0% of users for content purposes($p = .04$). These results, depicted in Figure 6a, are consistent with prior work that has found users are more likely to accept permissions if they were essential to app functionality [29] [10].

Framing. Research into behavioral economics and decision theory has consistently found that the framing of a decision affects user choices [25, 37], results that have subsequently been extended to privacy interfaces [22] and trust in mobile apps [13]. We therefore expected to find that a negative framing—one that emphasized the potential negative impacts of asking an app not to track—would result in higher opt-out

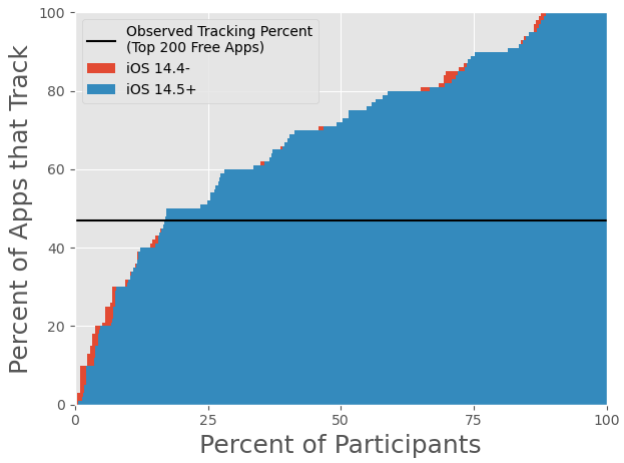

Figure 7: Percent of apps on a respondent's phone they believe are tracking them

rates than neutral or positive framings. However, we instead found no statistically significant differences between different framings. We saw that positive and negative framings had nearly identical opt-in rates (42.3% vs. 42.2%), both slightly lower than the neutral condition (Figure 6b). These results suggest that the standardization imposed on ATT requests is sufficient to negate the impact of framing observed in other contexts.

Payment. Prior work has found that most users put a low price on privacy [35], an effect that is amplified by framing effects [3]. We therefore expected users to opt-in to tracking at higher rates when told they might have to pay otherwise. However, we found no statistically significant effect due to payment (Figure 6c); in fact, the opt-in rates were slightly lower when users were threatened with payment if they opted-out of tracking. This effect might be due to a decrease in perceived trustworthiness when an app threatens to start charging for services if users ask it not to track, or it might be an artifact of the experimental design, in which users interacted with an app that they (presumably) did not intend to continue using in the future.

Priming. Prior work has found that priming can impact users decisions and privacy assessments in the context of mobile apps [5, 14]; we therefore expected the presence of a priming page to increase the rate at which users opt-in for tracking. Contrary to our expectations, there was no statistically significant effect due to priming; in fact, respondents allowed tracking 2.3% less often when presented with a priming page before the ATT permission request (Figure 6d).

User Experience. After interacting with the All News app, we asked each study participant a series of follow-up ques-

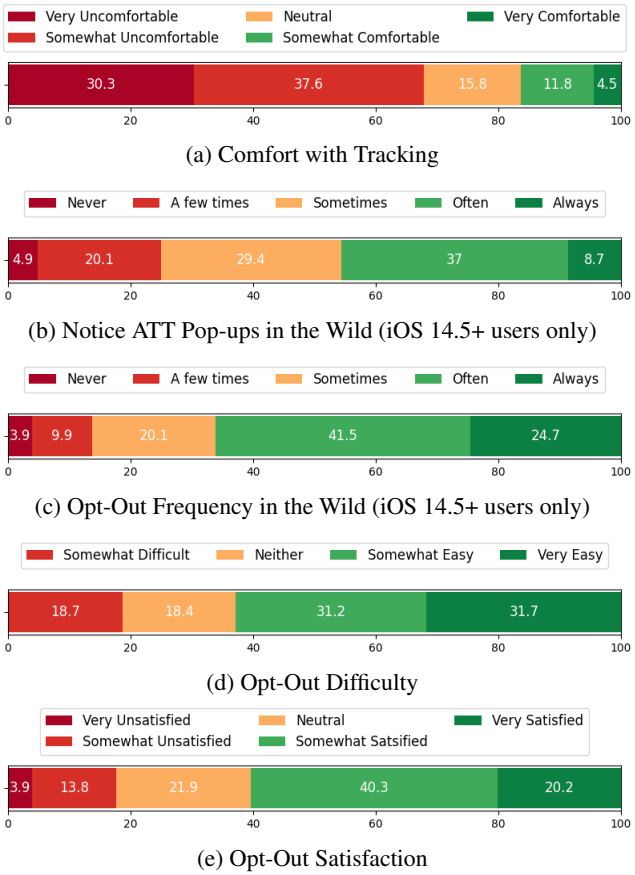

(a) Comfort with Tracking

(b) Notice ATT Pop-ups in the Wild (iOS 14.5+ users only)

(c) Opt-Out Frequency in the Wild (iOS 14.5+ users only)

(d) Opt-Out Difficulty

(e) Opt-Out Satisfaction

Figure 8: User experience with tracking and ATT tracking requests.

tions about their experience with mobile tracking. Overall, users reported being uncomfortable with the practice of tracking user behavior, with 67.9% saying that they were very uncomfortable or somewhat uncomfortable the the practice (Figure 8a). However, respondents thought that the majority of apps they had installed on their phone tracked them (Figure 7), with most users reporting higher rates of tracking than we observed in the Top 200 free apps. We saw no significant difference between users running iOS 14.5 and above versus 14.4 and below for this question.

In a promising sign for privacy, however, we found that 95.1% of users running iOS 14.5+ had noticed at least a few ATT tracking pop-ups prior to completing our study and 45.7% of respondents reporting seen them often or always, (Figure 8b). We also found that 66.2% of those running iOS 14.5+ opted-out of tracking often or always (Figure 8c).

After interacting with the All News app, most participants reported that the mechanisms was somewhat or very easy to use (Figure 8d), and most participants indicated that they were satisfied with the opt-out mechanism provided (Figure 8e). There were no statistically significant differences in difficulty or satisfaction between different conditions.

5 Related Work

To the best of our knowledge, this is the first work to examine Apple's ATT permission system. However, mobile permissions in general, the concept of nudging, and the impact of third-party tracking have been explored in many contexts.

5.1 Tracking

Previous research has highlighted the prevalence of third party trackers in mobile applications. Binns et al. decoded APKs for 959000 Google Play apps to map permissions to domain hosts, finding that 90% apps on the Google Play store had at least one tracker embedded [9]. Liccardi et al. found that of 528000 apps surveyed on Google Play, 46% collected personal data, while Vallina-Rodriguez et al. used ICSI Haystack to examine 1700 apps, finding that 60% connected to at least one ad tracking service (ATS) [28] [38]. Razaghpanah et al. developed the app Lumen to analyze device traffic and identify ad tracking domains, finding that the majority of data was shared both within and among organizations [34].

Less research has been conducted regarding third party trackers in applications on Apple's app store. Kurtz et al. found that approximately 40% of 1100 apps they analyzed connected to at least one ATS domain [27]. We found similar results from our data collection on the Top 200 free apps, where 47% of apps requested ATT permission, indicating their use of ATS domains. With such a high rate of third party tracking and low opt-in rates, Apple's ATT policy has a large effect on the mobile application economy.

5.2 Permissions and User Preferences

There is a large body of work focused on how users respond to permission requests. Mohamed and Patel outlined the differences in permission systems between Android and iOS, where Apple is more restrictive of developers access to sensitive subsystems [32]. The introduction of the ATT permission is another addition to these restrictions.

Before Android transitioned from pre-install permissions to ask-on-first-use (AOFU) with the upgrade to 6.0, it was established that users did not pay attention to, or understand the language of permissions requests [19] [18] [8] [26]. AOFU permissions offered some context about how a resource would be used, but were still seen as ineffective [10] [39]. We examine the impact of context in our second experiment by stating the purpose of tracking. This allows us to contribute data towards the impact of context, and expand on it by determining if one type of purpose is more readily accepted than another.

Similarly, other permissions work focused on defining the concept of privacy as expectations, where permissions were found to be accepted when they followed user expectations of an app's function. Lin et al. found that users were more comfortable when presented with a purpose for a requested permission, and felt least comfortable when any resource was used for advertising purposes [29]. Similarly, Bonné et al. noted a common reason for denying a permission was that the app shouldn't need it to function [10]. As Apple does not allow developers to remove functionality for users who reject tracking, this work may serve as an explanation for the documented low acceptance rates [23]. Our finding that users accept tracking more often when it improves app content contributes to this area of work.

5.3 Nudging

Previous research has examined persuasive design in a privacy context from a variety of perspectives. Research on framing specifically has returned mixed results. Gluck et al. found that neither positive nor negative framing had an effect on users' awareness of privacy notices, while Adjerid et al. found that the framing of a privacy notice affected how much personal information participants disclosed [20] [4].

Johnson et al. observed a significant framing effect when asking users to opt-in or out of a health survey with varying language [24]. They found that users were more likely to participate when presented with a positive frame. We build on this by including neutral language, allowing us to establish a baseline against which we compare different framings.

Other work on nudging has examined soft paternalism that leads users to better privacy decisions [7] [1]. Several researchers have built tools to help guide users through permission decisions, with nudges towards restricting permissions [30] [6] [39]. Apple's ATT policy is similar, expanding upon their existing permission system to give users more control over their privacy.

While research has focused on helping users improve their privacy, UI/UX is often designed with the opposite intentions in mind. Referred to as dark patterns, these design elements nudge users towards less privacy [12]. Researchers have categorized dark patterns in several ways, most recently splitting them into five categories: nagging, obstruction, sneaking, interface interference, forced action [15] [11] [21]. Dark patterns have been heavily studied on the web, with previous work examining major platforms including Windows 10, Google, and Facebook [16]. Other studies have cast a wider net, including one that found dark patterns in over 10% of 11000 shopping websites observed [31].

In the mobile context, researchers examined Android apps to determine that 95% of mobile apps contained at least one dark pattern, while 49% contained 7 or more [17]. We observed some dark patterns in the ATT priming pages we documented, falling under the category of interface interference. While we observed some examples, implementing dark patterns is forbidden in Apple's developer guidelines, explaining its low occurrence rate [23]. Similar policies limiting the use of dark patterns have been suggested by researchers previously [2] [33].

6 Conclusion

App Tracking Transparency (ATT) introduces opt-in tracking authorization for iOS apps. In this work, we investigate how mobile apps present tracking requests to users, and we evaluate how observed design patterns impact users' privacy.

This work conducts the first observational study to investigate how apps implement ATT requests in the wild. We perform a manual observational study of the Top 200 free iOS apps, and we report on our findings. We note, however, that the behavior of these apps may not be representative of the full app ecosystem; apps that are less popular and apps that charge for installation are likely to exhibit different behavior than the apps we examined. Further work will be required to determine to what extend our findings extend to mobile apps on the whole.

We also perform a user study with 950 to evaluate the impact of the observed designs. Our results show that opt-in authorization is highly effective at enhancing data privacy in the context of tracking by mobile apps. Moreover, opt-in rates (and thus privacy) were relatively consistent across most conditions, which indicates that ATT is less subject to dark patterns and other privacy-diminishing effect than other types of preference settings. Further work will be required to determine to what extend these findings generalize to non-iPhone users and to other types of apps, but our results suggests that ATT—which requires opt-in consent with clearly defined options presented in a standardized format—might prove an effective model for managing data privacy in other contexts.

References

[1] Alessandro Acquisti. Nudging privacy: The behavioral economics of personal information. *IEEE security & privacy*, 7(6):82–85, 2009.

[2] Alessandro Acquisti, Idris Adjerid, and Laura Brandimarte. Gone in 15 seconds: The limits of privacy transparency and control. *IEEE Security & Privacy*, 11(4):72–74, 2013.

[3] Alessandro Acquisti, Leslie K John, and George Loewenstein. What is privacy worth? *The Journal of Legal Studies*, 42(2):249–274, 2013.

[4] Idris Adjerid, Alessandro Acquisti, Laura Brandimarte, and George Loewenstein. Sleights of privacy: Framing, disclosures, and the limits of transparency. In *Proceedings of the ninth symposium on usable privacy and security*, pages 1–11, 2013.

[5] Tawfiq Alashoor, Grace Fox, and H Jeff Smith. The priming effect of prominent is privacy concerns scales on disclosure outcomes: An empirical examination. In *Pre-ICIS Workshop on Information Security and Privacy*, 2017.

[6] Hazim Almuhimedi, Florian Schaub, Norman Sadeh, Idris Adjerid, Alessandro Acquisti, Joshua Gluck, Lorrie Faith Cranor, and Yuvraj Agarwal. Your location has been shared 5,398 times! a field study on mobile app privacy nudging. In *Proceedings of the 33rd annual ACM conference on human factors in computing systems*, pages 787–796, 2015.

[7] Rebecca Balebako, Pedro G Leon, Hazim Almuhimedi, Patrick Gage Kelley, Jonathan Mugan, Alessandro Acquisti, Lorrie Cranor, and Norman Sadeh-Koniecpol. Nudging users towards privacy on mobile devices. 2011.

[8] Kevin Benton, L Jean Camp, and Vaibhav Garg. Studying the effectiveness of android application permissions requests. In *2013 IEEE International Conference on Pervasive Computing and Communications Workshops (PERCOM Workshops)*, pages 291–296. IEEE, 2013.

[9] Reuben Binns, Ulrik Lyngs, Max Van Kleek, Jun Zhao, Timothy Libert, and Nigel Shadbolt. Third party tracking in the mobile ecosystem. In *Proceedings of the 10th ACM Conference on Web Science*, pages 23–31, 2018.

[10] Bram Bonné, Sai Teja Peddinti, Igor Bilogrevic, and Nina Taft. Exploring decision making with android's runtime permission dialogs using in-context surveys. In *Thirteenth Symposium on Usable Privacy and Security*, pages 195–210, 2017.

[11] Christoph Bösch, Benjamin Erb, Frank Kargl, Henning Kopp, and Stefan Pfattheicher. Tales from the dark side: Privacy dark strategies and privacy dark patterns. *Proc. Priv. Enhancing Technol.*, 2016(4):237–254, 2016.

[12] Harry Brignull. Dark patterns, 2019.

[13] Eun Kyoung Choe, Jaeyeon Jung, Bongshin Lee, and Kristie Fisher. Nudging people away from privacy-invasive mobile apps through visual framing. In *IFIP Conference on Human-Computer Interaction*, pages 74–91. Springer, 2013.

[14] Isis Chong, Huangyi Ge, Ninghui Li, and Robert W Proctor. Influence of privacy priming and security framing on mobile app selection. *Computers & Security*, 78:143–154, 2018.

[15] Gregory Conti and Edward Sobiesk. Malicious interface design: exploiting the user. In *Proceedings of the 19th international conference on World wide web*, pages 271–280, 2010.

[16] Norwegian Consumer Council. Deceived by design, how tech companies use dark patterns to discourage us from exercising our rights to privacy. *Norwegian Consumer Council Report*, 2018.

[17] Linda Di Geronimo, Larissa Braz, Enrico Fregnan, Fabio Palomba, and Alberto Bacchelli. Ui dark patterns and where to find them: a study on mobile applications and user perception. In *Proceedings of the 2020 CHI Conference on Human Factors in Computing Systems*, pages 1–14, 2020.

[18] Adrienne Porter Felt, Serge Egelman, and David Wagner. I've got 99 problems, but vibration ain't one: a survey of smartphone users' concerns. In *Proceedings of the second ACM workshop on Security and privacy in smartphones and mobile devices*, pages 33–44, 2012.

[19] Adrienne Porter Felt, Elizabeth Ha, Serge Egelman, Ariel Haney, Erika Chin, and David Wagner. Android permissions: User attention, comprehension, and behavior. In *Proceedings of the eighth symposium on usable privacy and security*, pages 1–14, 2012.

[20] Joshua Gluck, Florian Schaub, Amy Friedman, Hana Habib, Norman Sadeh, Lorrie Faith Cranor, and Yuvraj Agarwal. How short is too short? implications of length and framing on the effectiveness of privacy notices. In *Twelfth Symposium on Usable Privacy and Security*, pages 321–340, 2016.

[21] Colin M Gray, Yubo Kou, Bryan Battles, Joseph Hoggatt, and Austin L Toombs. The dark (patterns) side of ux design. In *Proceedings of the 2018 CHI Conference on Human Factors in Computing Systems*, pages 1–14, 2018.

[22] Jens Grossklags and Alessandro Acquisti. When 25 cents is too much: An experiment on willingness-to-sell and willingness-to-protect personal information. In *WEIS*, 2007.

[23] Apple Inc. Human interface guidelines, 2021.

[24] Eric J Johnson, Steven Bellman, and Gerald L Lohse. Defaults, framing and privacy: Why opting in-opting out 1. *Marketing letters*, 13(1):5–15, 2002.

[25] Daniel Kahneman and Amos Tversky. Prospect theory: An analysis of decision under risk. *Econometrica*, 47(2):263–292, 1979.

[26] Patrick Gage Kelley, Sunny Consolvo, Lorrie Faith Cranor, Jaeyeon Jung, Norman Sadeh, and David Wetherall. A conundrum of permissions: installing applications on an android smartphone. In *International conference on financial cryptography and data security*, pages 68–79. Springer, 2012.

[27] Andreas Kurtz, Andreas Weinlein, Christoph Settgast, and Felix Freiling. Dios: Dynamic privacy analysis of ios applications. 2014.

[28] Ilaria Liccardi, Joseph Pato, and Daniel J Weitzner. Improving mobile app selection through transparency and better permission analysis. *Journal of Privacy and Confidentiality*, 5(2):1–55, 2014.

[29] Jialiu Lin, Shahriyar Amini, Jason I Hong, Norman Sadeh, Janne Lindqvist, and Joy Zhang. Expectation and purpose: understanding users' mental models of mobile app privacy through crowdsourcing. In *Proceedings of the 2012 ACM conference on ubiquitous computing*, pages 501–510, 2012.

[30] Bin Liu, Mads Schaarup Andersen, Florian Schaub, Hazim Almuhimedi, Shikun Aerin Zhang, Norman Sadeh, Yuvraj Agarwal, and Alessandro Acquisti. Follow my recommendations: A personalized privacy assistant for mobile app permissions. In *Twelfth Symposium on Usable Privacy and Security*, pages 27–41, 2016.

[31] Arunesh Mathur, Gunes Acar, Michael J Friedman, Elena Lucherini, Jonathan Mayer, Marshini Chetty, and Arvind Narayanan. Dark patterns at scale: Findings from a crawl of 11k shopping websites. *Proceedings of the ACM on Human-Computer Interaction*, 3(CSCW):1–32, 2019.

[32] Ibtisam Mohamed and Dhiren Patel. Android vs ios security: A comparative study. In *2015 12th International Conference on Information Technology-New Generations*, pages 725–730. IEEE, 2015.

[33] Arvind Narayanan, Arunesh Mathur, Marshini Chetty, and Mihir Kshirsagar. Dark patterns: Past, present, and future: The evolution of tricky user interfaces. *Queue*, 18(2):67–92, 2020.

[34] Abbas Razaghpanah, Rishab Nithyanand, Narseo Vallina-Rodriguez, Srikanth Sundaresan, Mark Allman, Christian Kreibich, and Phillipa Gill. Apps, trackers, privacy, and regulators: A global study of the mobile tracking ecosystem. 2018.

[35] Michel Schreiner and Thomas Hess. On the willingness to pay for privacy as a freemium model: First empirical evidence. 2013.

[36] Sensor Tower. Top charts: iphone - us - all categories, Jun 2021.

[37] Amos Tversky and Daniel Kahneman. Loss aversion in riskless choice: A reference-dependent model. *The quarterly journal of economics*, 106(4):1039–1061, 1991.

[38] Narseo Vallina-Rodriguez, Srikanth Sundaresan, Abbas Razaghpanah, Rishab Nithyanand, Mark Allman, Christian Kreibich, and Phillipa Gill. Tracking the trackers: Towards understanding the mobile advertising and tracking ecosystem. *arXiv preprint arXiv:1609.07190*, 2016.

[39] Primal Wijesekera, Arjun Baokar, Lynn Tsai, Joel Reardon, Serge Egelman, David Wagner, and Konstantin Beznosov. The feasibility of dynamically granted permissions: Aligning mobile privacy with user preferences. In *2017 IEEE Symposium on Security and Privacy (SP)*, pages 1077–1093. IEEE, 2017.

A Follow-up Survey Questions

In this Appendix, we provide the complete set of questions asked in our user study.

1. "What percentage of the apps you have installed on your phone do you believe track you?" (Chosen on scale from 0-100)

2. "If the mobile apps you use employed a permanent identifier to track your behavior across multiple apps and/or to link you to your other behavior online, how comfortable would you be with it?" (Very Comfortable / Somewhat comfortable / Neutral / Somewhat uncomfortable / Very uncomfortable)

3. "How often have you noticed apps you use giving you an option to opt-in or opt-out of sharing a tracking identifier with the app?" (Never / A few times / Sometimes / Often / Always)

4. "How often do you opt-out of tracking on the apps you use?" (Never Have / Have a few times / Sometimes / Usually / Always)

5. (If did not respond "Never" to Question 4) "How difficult on average did you find it to opt-out of tracking on apps you use?" (Somewhat difficult / Neither difficult nor easy / Somewhat easy / Very easy)

6. (If did not respond "Never" to Question 4) "How satisfied are you with the opt-out mechanisms you have used to opt out of tracking by mobile apps?" (Very satisfied / Somewhat satisfied / Neutral / Somewhat unsatisfied / Very unsatisfied)

7. "What sort of smartphone do you primarily use?" (iPhone / Android device / Other / None)

8. (If responded "iPhone" to Question 7) "What version of iOS is currently installed on your device?" (14.5 or higher / 14.4 or lower / I don't know)

9. "What is your current age?" (18-24 / 25-34 / 35-44 / 45-59 / 60-74 / 75+)

10. "What is your gender?" (Man / Woman / Non-binary person / Other)

11. "Choose one or more races that you consider yourself to be:" (White / Black or African American / American Indian or Alaska Native / Asian / Pacific Islander or Native Hawaiian / Other)

12. "In which country do you currently reside?" (list of countries)