Stop Following Me! Evaluating the Malicious Uses of Personal Item Tracking Devices and Their Anti-Stalking Features

KIERON IVY TURK and ALICE HUTCHINGS, University Of Cambridge, United Kingdom

Personal item tracking devices are popular for locating lost items such as keys, wallets, and suitcases. Originally created to help users find personal items quickly, these devices are now being abused by stalkers and domestic abusers to track their victims' location over time. Some device manufacturers created 'anti-stalking features' in response, and later improved on them after criticism that they were insufficient. We analyse the use of the anti-stalking features with five brands of tracking devices through a gamified naturalistic quasi-experiment in collaboration with the Assassins' Guild student society. To understand how these devices are used for stalking and how people make use of provided anti-stalking features, we provide 40 of 91 participants with a tracker that can be used to follow any other participant, and observe both how trackers are used and how people avoid being tracked. Despite participants knowing they might be tracked and being incentivised to detect and remove the tracker, the anti-stalking features were not useful and were rarely used. We also identify additional issues with feature availability, usability, and effectiveness. These failures combined imply a need to greatly improve the presence of anti-stalking features to prevent trackers being abused.

CCS Concepts: • **Security and privacy → Privacy protections**; **Usability in security and privacy**.

Additional Key Words and Phrases: Stalking, Surveillance, Technology-Facilitated Abuse, Trackers, Privacy

ACM Reference Format:
Kieron Ivy Turk and Alice Hutchings. 2024. Stop Following Me! Evaluating the Malicious Uses of Personal Item Tracking Devices and Their Anti-Stalking Features. In *The 2024 European Symposium on Usable Security (EuroUSEC 2024), September 30-October 1, 2024, Karlstad, Sweden.* ACM, New York, NY, USA, 21 pages. https://doi.org/10.1145/3688459.3688477

1 INTRODUCTION

Personal item tracking devices such as the Apple AirTag[1] and Tile tracker[2] have recently grown in popularity. These coin-sized devices are intended to be attached to belongings such as keys and bags to allow the owner to quickly find them when they are nearby, or to allow them to be remotely tracked if lost or stolen. Many different brands have started creating their own item tracking devices following the success of the Apple AirTag.

While these devices are useful for their intended purpose, they have begun to be used maliciously. News outlets have reported many cases of AirTags and similar devices being used for stalking and domestic abuse [7]. In late 2022, Apple were sued for this abuse of their devices [15]. Other malicious uses of item tracking devices have also been reported, such as planting trackers on cars to track where they are parked later to steal them [6, 19].

Manufacturers are aware of these issues and have implemented a range of "anti-stalking features" in response. These commonly include scanning for unwanted trackers following a user and alerting the user of the unwanted tracker. Most

[1]https://www.apple.com/airtag/
[2]https://www.tile.com/

Authors' Contact Information: Kieron Ivy Turk, kieron.turk@cl.cam.ac.uk; Alice Hutchings, alice.hutchings@cl.cam.ac.uk, University Of Cambridge, Cambridge, Cambridgeshire, United Kingdom.

Manuscript submitted to ACM

implementations additionally provide means to locate the tracker after alerting the user, such as using Bluetooth to indicate the distance to the tracker or making it play a sound alert.

These features later came under criticism [10] for not being sufficient to detect unwanted trackers. Common issues included the sound alerts being too quiet, separated-from-owner alerts taking too long to trigger, and only providing anti-stalking features for certain platforms. Apple later improved on their anti-stalking features following this feedback [2], however, many believe there is still room for improvement.

We set out to evaluate the effectiveness of the anti-stalking features of tracking devices. We wanted our evaluation to be as naturalistic as possible, but there are also important ethical and data protection issues to consider. We wanted our participants to consent to being tracked, but not necessarily know if or when they were being tracked, by whom, or with what type of tracker.

We designed our evaluation around a long-established student society, the Assassins' Guild. The Guild runs a competition that spans the entire term. Students enroll in the game and are given a list of targets (other players) to 'assassinate'. They also become targets but do not know who their adversary is. The Guild agreed to run a sub-game of participants who agreed to be tracked as part of this evaluation. There were 91 participants who consented to take part in the study and be tracking targets, and 40 of these participants were provided with trackers.

Our research questions are as follows:

RQ1: How easy is it for people to learn to use tracking devices for stalking?
RQ2: What countermeasures are used by those who are tracked to find and remove the trackers?
RQ3: How usable are the provided anti-stalking features?
RQ4: How effective are the provided anti-stalking features?

In this paper, we provide an overview of the background and related work in §3. We then discuss the trackers chosen for this study, alongside their features for regular use and stalking prevention in §2. Our research was approved by the department's Research Ethics Committee. We provide a detailed overview of the ethical considerations that were central to our research design in §4.4. We provide an overview of our methods in §4. Our findings are detailed in §5. In our discussion (§6) we consider the implications of our study before concluding (§7).

We make a novel methodological contribution for evaluating the use and effectiveness of anti-stalking features. We find that available anti-stalking features do not appear sufficient to prevent many cases of unwanted stalking by tracking devices. We make a number of recommendations so that the features provided by all manufacturers can be improved upon in the future. Following our research, a joint draft RFC by Apple and Google [16] has been announced to establish a standardised protocol for trackers to allow unwanted trackers to be universally detectable. Furthermore, Google have implemented background scanning for unwanted trackers by default for all Android users, with an additional manual scanning feature available in settings.

2 ITEM TRACKING DEVICES AND THEIR FEATURES

Personal item tracking devices are intended to allow the owner to quickly locate personal possessions. By attaching these devices to keys or bags, or inserting them into wallets and purses, they can be quickly located by distance or through sound without having to manually search a location. These devices are often the size of a large coin — up to 5cm in diameter, and about 1cm thick. Some manufacturers provide card-style versions of the devices, such as the Tile Slim or Chipolo Card, which can be inserted into purses, wallets or passport holders as an alternative to attaching a disc to personal possessions.

Table 1. Availability of Anti-Stalking Features

Application		Apple: Find My	Apple: Tracker Detect	Samsung: Smart Things	Tile	AirGuard
Available on	iOS	Built-in	✗	✓	✓	✓
	Android	✗	✓	✓	✓	✓
Detects	AirTag	✓	✓	✗	✗	✓
	SmartTag	✗	✗	✓	✗	✗
	Tile	✗	✗	✗	✓	✓
	Chipolo	✓	✓	✗	✗	✓
Scan Type	Background	iPhone Only	✗	Galaxy Only	✗	✓
	Manual	✗	✓	✓	✓	✓

Many of these devices also feature alerts when a user leaves them behind, ensuring necessities are not left at home. The most common use of the devices are to allow users to locate items that have been lost by showing their most recently reported location, allowing them to be recovered.

In this section, we explain our choice of tracking devices used for this study, then detail the different features available for these specific trackers. This includes remote location updates, methods for finding nearby trackers, and an overview of the anti-stalking features provided by the manufacturers of these devices.

2.1 Selected Devices

A wide range of different item tracking devices are available; for our study, we focus on AirTags, Tile trackers, Galaxy SmartTags, Chipolo trackers, and the Kuxian Keyfinder. AirTags are the most popular tracker and have gained media coverage for their use in stalking. Tile is the largest existing brand of item tracker before the AirTag, and provides several models: the Pro, Mate, Slim, and Sticker.As we are focusing on the use in stalking, we chose the Sticker for our study, which is the smallest tracker and also has a sticky surface to attach the trackers to other's possessions. The Galaxy SmartTag functions similarly to AirTags: the associated brand of phones all provide location updates for trackers, making them more effective at finding lost items. Chipolo trackers are another popular brand of trackers not associated with a phone brand, comparable to the Tile trackers. For the fifth tracker, we selected the most popular (by ratings) off-brand tracker from Amazon as a control, which is the Kuxian Keyfinder. This tracker does not provide any anti-stalking features.

These item tracking devices provide several features to help the owner locate them, in addition to anti-stalking features that build off of these. The trackers provide updates on their location when in Bluetooth range of the owner and remotely through several methods, and allow the user to locate the device when in range through Bluetooth and by making the device play a sound. The anti-stalking features revolve around providing mechanisms to scan for unwanted tracking devices, and then allowing the user to use the aforementioned features to locate any detected devices. Additionally, Apple provides "Separated from Owner" sound alerts periodically when the tracker is unable to connect to the owner's device for an extended period.

2.2 Location Updates

All of the trackers in our study use Bluetooth to provide the location of the tracker; they do not have any on-board GPS or other mechanism to update the owner on their current location. When they are in Bluetooth range of the owner's device, they are able to connect to inform the owner that the tracker is near their current location. We measured the Bluetooth range of each tracker and found that most trackers have a range of 55-75m; Apple's AirTags are an outlier to this at 26m.

Most trackers provide a mechanism to remotely update the location of the tracker. When the tracker moves in range of a compatible device, the location of the tracker is encrypted and uploaded to the manufacturer's servers, who then forward it to the owner to decrypt. Devices compatible with Apple's Find My protocol (AirTags, Chipolos) will provide this update when in range of any Apple device which is part of the network[3], while SmartTags can report via any Samsung Galaxy phone. Tile trackers provide updates through other Tile app users, and work on both iOS and Android.

2.3 Finding Nearby Trackers

Nearby trackers, either the owner's or detected trackers following a user, can be located through two methods. Bluetooth provides the distance to a connected device, which users of any tracker in our study can use to get closer to the tracker once in Bluetooth range. Apple AirTags additionally allow the user to see the direction to a tracker through the Ultra-WideBand (UWB) support in the M1 chip, as long as the user's phone supports UWB (i.e. iPhone 11 and later).

The other primary method for locating nearby trackers are sound alerts. When a tracker is nearby, users can make it play a sound so that they can easily locate it. We measured the volume of the different trackers in several common scenarios, shown in Table 2. These include common hiding places such as coats and bags; we also measured the volume at varying distances for the bags to show the dropoff at a distance. While the trackers are relatively loud for their size, the sound alerts are easily muffled to be quieter than speech (approximately 60dB).

Another additional feature provided by Apple AirTags are "Separated from Owner" sound alerts. When a tracker has been separated from its owner for a random period between 8 and 24 hours, the AirTag will play a 15 second sound alert every 6 hours. It will additionally attempt to send a notification to nearby iPhones at the same time. These alerts are intended as a fallback in case scanning mechanisms fail, but the short sound and long delay between alerts limit the utility of this feature.

2.4 Finding Unwanted Tracking Devices

There are two types of scanning available for unwanted trackers: background scanning and manual scanning. The availability of these scans is shown in Table 1. Background scanning listens for nearby tracking devices and records when and where they have been detected near the user. When a tracker has been detected in multiple locations over an extended period and distance, a notification is shown to the user stating that an unknown tracking device is following them. Apple provides background scanning for devices compatible with Find My for iPhones, and Samsung provide background scanning for Galaxy SmartTags for Galaxy phone users with the SmartThings app installed.

The other type of scan is manual scanning. This requires the user to initiate a scan through the provided application and then move around during the scan. The feature identifies which trackers, which do not belong to the user, are moving with them, and provides a report on the unwanted trackers. This requires multiple scans from different locations to avoid false positives and privacy invasions of other tracker users who are nearby when the user runs the scans.

[3]All Apple devices are part of the Find My network by default. Users are able to opt-out in their device's settings menu.

Table 2. Volume of Each Type of Tracker in Decibels. Background Volume is 34dB.

Tracker	Advertised Volume (dB)	Maximum Measured Volume in (dB)			
		Open Area	Coat Pocket	Bag (at 10cm)	Bag (at 5m)
AirTag	60	86	76.3	40	36.5
Galaxy SmartTag	85-96	89	84	68	54
Chipolo One	120	96	80	64.5	47.9
Tile Sticker	85-114	92	76	61	40
Keyfinder	—	85	70	55	38

Apple provides manual scanning to Android users through the Tracker Detect app; Tile provides manual scanning for all users of the Tile app; and Samsung provides manual scanning for Galaxy users through the SmartThings app.

The time to inform users of unwanted trackers is not publically reported by any item tracking device manufacturer. Heinrich et al. [13] created test scenarios to measure the time it takes Apple to report an unwanted tracker when planted in a pocket, a backpack, or on a car, and compared it with the AirGuard app that they developed. They found that Apple's tracker scanning located the unwanted devices after 1 hour 45 minutes in the first scenario, 4 hours 14 minutes in the second scenario, and failed to detect them in the third scenario. Their own AirGuard app detected the trackers in 35 minutes, 30 minutes, and 4 hours 18 minutes respectively.

We created our own tests to measure the times taken to detect each type of tracker with background scanning. We provided an iPhone user, a Galaxy user, and a Google Pixel user with all five trackers and installed the SmartThings, Chipolo, Tile, and AirGuard applications. Two users are authors of this paper, while the third was selected due to the type of phone they had, and the range of activities they would be undertaking throughout the day. Each user then went about their regular activities and measured the time for each application to alert the user of each unwanted tracker, if at all. Note that the AirGuard application is not available on iPhones, and Find My scanning is not available on Android phones. These tests took place in February 2023. The reports received are in Appendix A.

We found that very few of the trackers were detected, as shown in Table 4 in the appendix. The iPhone alerted the owner to the unwanted AirTag after 5 hours 41 minutes, shortly after connecting to the home Wi-Fi, but failed to detect the Chipolo tracker. AirGuard detected the AirTag and Tile tracker simultaneously on the Pixel phone after 25 minutes, but missed the Chipolo tracker and did not find the tags on the Galaxy Phone. In all cases, the Samsung Smart Things application did not find the Galaxy SmartTag.

3 BACKGROUND AND RELATED WORK

In this section, we discuss the variety of existing research and media surrounding personal item tracking devices. As a notable use of these devices is in domestic abuse scenarios, we additionally explore prior research on technology-facilitated domestic abuse and technology-enabled surveillance. Finally, we detail literature on the usability of privacy features due to its application to the anti-stalking features used for these devices.

3.1 Prior Work on Trackers

Despite a wide range of news outlets reporting on different abuses of personal item tracking devices, there is very little prior academic work on this topic. Motherboard [7] reported on 150 police records of AirTags being used in Intimate Partner Abuse (IPA). Apple was sued later in 2022 by two women who were stalked by their ex-partners, as reported by

the New York Times [15]. The devices were also planted on cars to steal them later, as reported by Mac Rumors [6] and the New York Times [19].

Following these reports, several groups performed informal tests using the devices for stalking. Hill [14] used a range of tracking devices on her husband to follow his movements throughout the day, and he was unable to locate many of the planted trackers. Fowler [10] consented to be tracked by a colleague to put the anti-stalking features to the test and found them severely lacking. Scott [25, 26] ran a pair of games where one participant attempted to complete tasks while the other used an AirTag to try and track them down. In several cases, the tracking participant came very close to finding the other player.

Wired [4] reported specific failures of the Apple anti-stalking features that made them a “gift to stalkers”, including quiet and infrequent sound notifications, long periods before alerting users, and the absence (at time of publication in early 2021) of any Android anti-stalking features. Apple [2] responded by creating the “Tracker Detect” app for Android, increasing the tracker volume, and reducing the period to notify users of unwanted trackers.

Shafqat et al. [24] performed in-depth analysis of the existing implementation of anti-stalking features, which do not have detailed public documentation. They measured the times to alert users in an experimental setup, and additionally tested using a cloned AirTag that modified the announcements to attempt to get around the safety features. They concluded with a set of recommendations to improve the safety features of these devices.

Heinrich et al. [13] compared Apple’s iOS tracker detection with their AirGuard app. By reverse engineering the functionality of Apple’s iOS AirTag detection, they created AirGuard to scan for AirTags and later Tile trackers. Their application detects unwanted trackers significantly faster than official options.

Turk et al. [27] analysed the implementations of anti-stalking features of item-tracking devices through a series of experiments with trackers combined with the authors’ expert opinions on the topic. They identified a series of limitations in the current anti-stalking features, as well as limitations on their future design. They provided a series of potential improvements that could be made to anti-stalking features in the future.

The limitations identified by Turk et al. [27] include various issues around how devices scan for unwanted trackers. They found that the time taken to identify unwanted trackers with background scanning was too long, allowing the stalker to follow the victim for many hours before they will discover that they are being tracked. Manual scanning was observed to be too difficult to use effectively due to the extensive user effort required. There were also problems with the availability of anti-stalking features: as shown in Table 1, background scanning is not available for all operating systems and devices, and users are required to install multiple applications to be able to detect all unwanted tracker types. Additionally, they identified shortcomings in the methods to locate unwanted trackers. Using Bluetooth to locate trackers is not always provided as an option to users, nor is Apple’s precise location feature. The sound alerts were also too easily muffled in common hiding spots used by stalkers.

There are also fundamental limitations of the effectiveness of anti-stalking features. In addition to the malicious use for stalking, there is the legitimate use case of locating stolen items [27]. This appears as though the tracker is stalking the thief and the anti-stalking features will help them remove the tracker. There is no way to distinguish this from malicious use, and companies are therefore forced to "pick a side" — anti-stalking or anti-theft. Apple advertises that the AirTag is not intended to be used to locate stolen items and that it should instead be used to find lost items and ensure necessary items are not left behind. Tile created an anti-theft mode which disables their anti-stalking features in favour of anti-theft, provided users send photographs of ID to Tile and agree not to use this feature to enable stalking under threat of being sued by Tile. Avoiding other forms of false positive, such as using public transport with someone who has a tracker but whose Bluetooth is turned off, also limits the design of anti-stalking features.

3.2 Technology-Enabled Domestic Abuse

Technology has become commonplace in domestic abuse, now called technology-facilitated abuse or "tech-abuse". Refuge [22] found that 72% of abuse cases involved tech-abuse, while The Safety Net Project [21] found 97% of domestic violence victims reporting technology being misused for abuse.

Multiple studies have explored how technology is used within abusive relationships. Matthews et al. [20] provide a framework for analysing the phases, practices, and challenges faced by survivors of technology-enabled domestic abuse. They identified three phases of abusive relationships: physical control, where abusers have full access to the victim's systems; escape, where the victim must hide their actions while severing ties with the abuser; and life apart, where the survivor builds and maintains a new life while avoiding returning to the physical control phase. Tracking devices have a major impact on all three of these phases: they enable further control while in the relationship, and restrict the victim's agency and ability to separate themselves from the abuser while apart.

Freed et al. [11] interviewed survivors and domestic abuse professionals to identify four categories of tech-abuse: ownership-based access, account or device compromise, harmful messages and posts, and exposure of private information. They describe abusers as the "UI-bound adversary" who only use the provided features of a device rather than using any technically skilled attacks. Tracking devices fit this model through their simple setup and use, with almost no technical ability required to use them maliciously. Douglas et al. [8] interviewed 55 survivors of coercive control and found abusers using technology to isolate victims from their support networks, harass them online, monitor their activity, and stalk them. Dragiewicz et al. [9] found similar types of technology-facilitated abuse, adding that the use of technology for abuse either starts or escalates in the escape phase of abuse [20]. Leitão [17] analysed posts on online forums to explore how technology is used for domestic abuse. They find that the key issues are monitoring and stalking in addition to harassment and intimidation as identified by other studies. All three of these studies emphasise the prevalence of stalking and monitoring especially in later stages of abusive relationships, highlighting the persistent issue of Intimate Partner Surveillance (IPS).

Sambasivan [23] found cyberstalking to be the most prominent form of digital abuse in South Asia, with 66% of their participants reporting instances of cyberstalking. Woodlock [28] found technology used to impose a sense of the abuser's omnipresence, and that the extent of the stalking forced the victims to isolate themselves to be able to escape. Item tracking devices such as AirTags and Tile trackers further facilitate these harms against victims of tech-abuse and IPS.

Levy and Schneier [18] looked at broader "relationships" where privacy is affected by misuses of technology. This includes IPS in addition to parent-child relationships, adults and the elderly, caregivers, and friends. They identified several common features of the threats posed and relate these to a series of implications for the design of future systems, which aids systems designers who wish to mitigate some of these privacy problems. This includes recognising that there is often a balance between multiple interests and values. This is particularly relevant for location trackers, where tracking may be justified by users due to safety concerns.

Bellini et al. [3] explored the justifications for IPS on online infidelity forums. They found that prior instances or suspicions of cheating are common justifications, as well as "suspicious behaviour" by the victim. They presented a four stage cycle of IPS, where abusers often decide on IPS as a "necessary hardship". Posters often share their stories to find advice and guidance on how to perform IPS, as well as to share or boast about their surveillance.

3.3 Usability of Privacy Features

Privacy features such as the anti-stalking features discussed in this study have important usability choices in their design. These can be designed for the users' benefit, making it easier to improve their own privacy, or to their detriment but to organisations' benefit, allowing for easier data collection. Features which are created for users' benefit should be designed in a manner which guides users towards using them.

Frik et al. [12] explored user's expectations of privacy features on their phones. They found that most users have not configured the settings and are unaware of the defaults, although they therefore rely on the defaults in place. Furthermore, users were not aware of which settings are available, and there are some demographic impacts on users' awareness of the features such as older participants interacting with privacy settings less often.

A common feature which aims to aid user privacy are nudges: designing decisions which "nudge" users towards certain choices, without removing any options. Caraban et al. [5] perform a systematic review of literature on nudge designs, identifying 23 mechanisms to nudge users towards better decisions. Acquisti et al. [1] discuss how these nudge designs can be used to aid users towards beneficial decisions. This includes providing more information to users to base their decisions on, presenting the choices in a manner where the "optimal" decision is made easier for users, and providing better default options to establish better privacy norms.

4 METHODS

We aim to emulate the real-world use of tracking devices for stalking in a naturalistic environment through a gamified study, allowing us to observe the ease of misuse of these devices as well as the methods that participants use to protect themselves from this threat. Our study extends the game of "Assassins" by providing participants with trackers to place on other "tracker-enabled" players. Several incentives were provided to encourage use of trackers. Players are rewarded for completing our "post-mortem" survey, requesting information about how the tracker was used and how the tracked player (attempted to) locate it.

4.1 The Assassins Game

In a game of Assassins, players are provided with three targets to track down and "kill", by using fake weapons such as a "knife" (pen or another object with "knife" written on it) or a water pistol. We chose to integrate our study into this game as it creates a gamified, naturalistic scenario similar to how stalkers operate in the real world, while taking careful ethical precautions to ensure that the similarity to this scenario does not cause harm to participants.

There are restrictions on when assassins can make kills, such as not playing the game in defined out-of-bounds areas like lecture theatres or in spaces where the general public may become concerned. After a kill, players write reports on both the accurate list of events and a dramatic retelling, which are published on the Assassins' website. In addition to having three targets, players are the target of three other players. The target that "died" is replaced after each kill, and a new player will be assigned to track a target if one of the players hunting them dies. The "targets" do not know who is trying to "kill" them, although if they discover their assassin during an attempt they are permitted to retaliate in-game.

Players are rewarded for making attempts and getting kills through "competence". When a player's competence runs out, they become "incompetent" and their information is published on the Assassins' website, making them a target for all players. Making two attempts or getting a kill adds a week to the player's competence period, incentivising players to remain active. In addition to regular players, there are "Police" who may only "kill" anyone on the incompetents list.

Any members of the Police who break the rules will become wanted and are subject to similar rules to incompetent players.

The game ran over 8 weeks (October to early December 2022), with three special events occurring during this time. After 2 weeks (the default "competence" provided for players to prove that they are taking part in the game), there is an "incobash" where players team up to kill as many incompetents as possible. This removes most of the players who signed up for the game but are not actively partaking, or who no longer want to take part. The last week of the game is "open season", in which all remaining players' information is provided on the website and players can target any other player still alive. A scoreboard is also published, and players can attempt to get into the top 6 who will go on to duel to win the game on a set date after the main game has ended.

4.2 Study Setup

We selected 5 types of trackers to be used in our study. We selected devices based on popularity and the presence of anti-stalking features, with an additional "control" tracker from Amazon that does not provide them. We first ran various tests with the trackers to understand how the anti-stalking features work. We placed all trackers on three users to measure time to alert for different phones/OSes. We also provide 8 of each type of tracker to participants to track other players with a total of 40 devices. The trackers we selected for evaluation are the Apple AirTag, the Tile Sticker, the Galaxy SmartTag, the Chipolo One Spot, and the Kuxian Keyfinder.

Participants were recruited through the local Assassins' Guild student society. The study was advertised alongside the game, and all players who signed up for the game were able to sign up as tracker-enabled (provided they completed the consent form) or tracker-disabled. We coordinated with the umpires to ensure that completed consent forms matched the tracker study sign-ups. 40 tracker-enabled players were randomly selected to receive trackers which were compatible with their phone. To distribute the 40 trackers, we first allocated 8 AirTags to randomly selected iPhone users, then 8 Galaxy SmartTags to Galaxy users, then assigned all remaining trackers randomly amongst the remaining participants. The participant's type of phone was collected at the time they signed up to take part in the study. We allowed players who were assigned trackers to keep them after the game to incentivise people to join, in addition to providing twice the competence bonus and four times as many points if they planted a tracker on the target and left it for at least 4 hours before the kill.

Players are explicitly told which of their targets are tracker-enabled and tracker-disabled. Players on the incompetent and wanted lists are marked with this information. Tracker-enabled players therefore have one, multiple, or no players tracking them at any given time. Players reported any planted trackers to the umpires, and if a player found a tracker planted on them they contacted the umpires to let them know it was found and then followed provided instructions to disable the device. Tracker-enabled players who were not provided with a tracker were invited to remain tracker-enabled and were rewarded with additional points for each kill they made, increasing the number of possible tracking targets.

Players taking part in the study were intentionally not informed about the available anti-stalking features. This makes the study naturalistic and allows us to better understand how people will attempt to find trackers in the real world, rather than pushing the use of technical countermeasures.

4.3 Data Collection

There were three surveys in our study, intented to enable data collection at distinct stages of the study. The first was a consent and demographics form completed at the time of recruitment. A "post-mortem" survey gathered information during the game, after each kill. A post-game survey was completed once the game had finished for the term.

The initial survey was completed after obtaining informed consent from participants, who were provided with detailed information about the study and how the game works when they are part of the study. The survey asks for information on the participant's devices and operating systems, prior experience with the Assassins' game and any prior use of trackers.

The "post-mortem" survey was completed after one player "kills" another during the course of the game while using a tracker. It asks the person using the tracker how they planted it and how easy and effective it was for stalking. Both the tracking and tracked players were asked to provide information on how they searched for trackers that had been planted on them. They are asked to provide the current location of their own tracker so that it can be returned to the owner.

The post-game survey asks a series of questions initially overlapping with the post-mortem survey. The intention of this part is to gather the same information on how trackers were (intended to be) used, as well as how players were checking for trackers, even if they did not get the opportunity to use a tracker and/or were never tracked by other players. Additionally, the survey asks participants if they used each anti-stalking service, and if so how useful was it. Prior to this, participants had not been explicitly told what anti-stalking features were available to them, so this allows us to measure how often each tracker's anti-stalking features are used.

4.4 Ethics

We obtained ethics approval from the Cambridge University Department of Computer Science Research Ethics Committee before commencing our study. Our main ethical considerations included informed consent, possible misuse of trackers, recruitment and incentives, data collection and storage, and anonymisation.

All participants provided informed consent before any information was collected. The consent form explains the study, the data collected and how it is later anonymised, how to withdraw from the study, and who to contact with further questions. This immediately preceded the demographic data collection. Participants explicitly consented to be tracked and to only track other tracker-enabled players. They agreed not to use the tracker on people or any items that were not their own after the study. They also had to inform us of the tracker's location after it had been planted in case it needed to be retrieved.

While we ensure that participants are aware of whom they are allowed to plant trackers on, there is room for accidental errors when planting the tracker on a target's possessions. To mitigate this risk, all trackers were labelled with a sticker with a QR code and URL. Both of these pointed to a webpage[4] explaining the study with a clearly labelled link to a set of instructions to disable the tracker and contact the organisers, ensuring that the tracker does not cause unintended harms.

We collected participant contact information, demographic information, and prior experience with technology and the Assassins' Guild. We also collected participants' responses to the post-kill surveys during the 8-week game. Data were not anonymised during the course of the game as we needed to maintain contact to gather survey responses. At the end of the game, we pseudonymised the results by replacing player names with unique identifiers.

Participants are able to contact the researchers at any time to request a copy of their data or to withdraw from the study. We established a protocol for withdrawal if it were to occur while the game was underway: we would delete all data stored about the participant and contact the person tracking them to ensure that the tracker was returned. We would also obtain the current location of the participant's tracker so that we could remove it. In the game, the

[4]https://www.cl.cam.ac.uk/~kst36/tracker_study.html#found-tracker

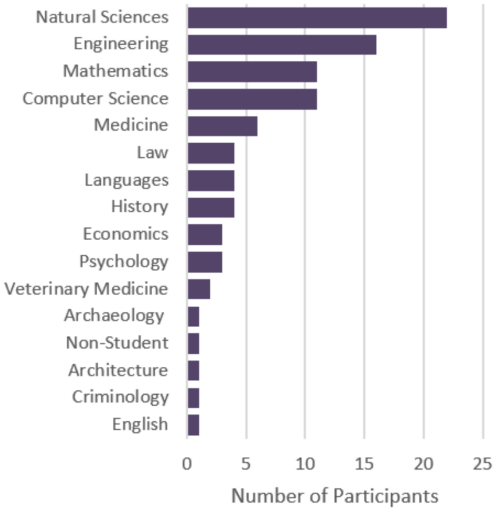

Fig. 1. Distribution of Degree Courses of Participants

withdrawing player is marked as tracker-disabled, and any players targeting them at the time were explicitly informed of this change to ensure they did not try to plant a tracker on them after withdrawing from the study.

5 RESULTS

In this section, we discuss the results of our study with respect to the usage of trackers within the game and the results of our surveys. §5.1 describes the background of the participants and §5.2 discusses how trackers were used in the game. §5.3 explores how effectively players used the trackers to locate other participants. §5.4 discusses how participants attempted to locate trackers planted on their person, and the use of anti-stalking features by participants.

5.1 Participant Demographics

372 players signed up to the overall game. Of these, 91 signed up as tracker-enabled and completed the consent form, and we provided 40 of them with trackers. The remaining 51 all opted to remain tracker-enabled despite not having a tracker (i.e. consented to remain as possible targets for tracking), in return for bonus competency points for their own kills.

Of the 91 tracker-enabled players, 45 (49.4%) had iPhones, 16 (17.6%) had Samsung Galaxy phones, and 30 (33.0%) had another brand of Android phone. 67% of participants were playing for the first time, 29% had played between 1 and 3 games before, and 4% had played 4 or more games before.

Most participants in our study are studying STEM subjects, as shown in Figure 1. Participants rated themselves as having above-average knowledge of technology, although most participants reported limited experience with tracking devices. A summary of participant's levels of experience with technology and tracking devices is shown in Table 3.

Table 3. Participants' Self-Reported Level of Technical Knowledge and Experience Using Item Tracking Devices

	Tracker Experience					
Tech Knowledge	None	A Little	Moderate	A Lot	A Great Deal	Total
Far Below Average	1	0	0	0	0	1
Somewhat Below Average	7	1	0	0	0	8
Average	23	6	2	0	0	31
Somewhat Above Average	22	18	5	0	0	45
Far Above Average	4	0	1	0	1	6
Total	57	25	8	0	1	91

5.2 Events During the Game

Trackers were assigned when signup closed and delivered to participants within 12 hours of the game starting. Participants were then notified if they were given trackers to ensure they could begin to use them as soon as possible. Participants were also reminded every 2 weeks to use the trackers and fill in the post-mortem survey to boost participation and avoid participants forgetting about the tracker component of the game.

One player successfully used the trackers twice to track down targets before achieving in-game kills, and two other players managed to plant their trackers although they were unsuccessful in using them before being killed.

In the first successful use of a tracker, the participant planted the tracker on the target when walking past them:

> *I placed the tracker in the open side pocket of their bag whilst passing them in a public place as they were walking somewhere from their lecture.*
>
> *I used the tracker to check which lecture theatre they would be in as I knew they must be in one of 2 [from the location provided by the app].—G7*

This allowed the participant to accurately locate where the target was, and to wait outside for them to appear. Later in the game, the same player managed to plant their tracker on another target:

> *When I arrived at their room, the key was left in the door. I attached the tracker to their keychain to find them later.*
>
> *I found out they were at a meeting, then found out they were riding a bike due to how fast they were. I headed to intercept them at the bike rack closest to their room. —G7*

In this case, the participant determined the mode of transport being used from the location updates of the tracker. From this, they were able to predict where they could go to intercept the person, before using other features of the tracker to identify exactly who it is:

> *I decide to turn on the tracker's alarm. I go to stab the person I believed was [the target], but moments before, the alarm rings on a blonde-haired girl nearby, saving the innocent and myself. —G7*

Despite not having seen the person they were targeting beforehand, the tracker's sound alert allowed the participant to identify who they were tracking and avoid assassinating the wrong person.

5.3 Using Trackers for Stalking

Many of the participants in our study found it difficult to plant trackers on their targets. Of 19 participants with a tracker who responded to the end of study survey, 7 attempted to plant their tracker at least once, and three were successful. Participants explained that the difficulty of planting trackers stemmed from attempting to plant them while

in public spaces — it was too easy to be spotted by their target or by the general public while planting the tracker, and there was the possibility of either the in-game or real-world police spotting the tracker being planted.

One participant was additionally concerned by an auxilliary feature of the Chipolo tracker (also provided by Tile and the Kuxian Keyfinder): when the tracker is clicked twice, the owner's phone will ring out, ignoring if the phone is on silent.

> *The tracker I was given to plant, a Chipolo, when double clicked rings the phone of the tracker regardless of whether their phone is on silent or not, which made it difficult for me to place my own tracker. —C8*

The possibility of being discovered this way was deemed too high of a risk by the participant, who then did not plant their tracker. The participant did not discover that this can be disabled in the Chipolo app.

Several participants attempted to plant trackers but found it difficult and were unable to plant them successfully. One attempted to plant an AirTag on a bicycle but were unsure how to attach it. Two others had difficulty finding a situation where they would be able to avoid raising suspicion from the public while also avoiding interaction with their target.

> *It was difficult in many situations to do so without arousing suspicion from other people. —G4*
> *Planting trackers is also way too difficult, and I needed multiple attempts before I could actually find my target unaware and for me to actually feel like the police wouldn't be called if seen. —G7*

There were some successful uses of trackers in our study. Two participants planted trackers in others' backpacks, one player attached it to the underside of a bike seat, and another attached it to their target's key chain. Two of these cases were with a Galaxy SmartTag, one was an AirTag, and one was a Chipolo One Spot. Two of these players were unable to use the trackers on their target before being killed: one player complained that the Chipolo tracker was not providing any remote location updates:

> *The tracker was absolutely useless. As soon as the target went further away from my Bluetooth range, the app had no idea where the tracker is. —C6*

The participant was unable to locate the tracked player before being attacked by in-game police. The other was attacked during the 4 hour waiting period after planting the tracker. The locations used to plant trackers are largely a reflection of the participants of our study being students, who are in constant possession of backpacks and commonly use bikes to travel.

The most effective trackers for stalking were the AirTag and Galaxy SmartTags. Both of these trackers are tied to a common phone brand (half of the participants had an iPhone, while 17.8% had a Galaxy phone). This made the trackers provide more frequent updates, in one case often enough to allow the owner to determine what mode of transport the tracked target was using and predict where they were heading to intercept them. Chipolo trackers were notably harder to use, with one participant planting their tracker and then complaining to the study organisers about the lack of updates on the tracker's location.

5.4 Use of Anti-Stalking Features

Most players in our study either did not search for trackers planted on them or relied on periodic manual searches of their possessions. Of 19 respondents to the post-game survey, 9 reported checking inside their bags, coat pockets, under their bike seats, and otherwise manually looking for trackers. Only one player out of 19 respondents attempted to use any of the anti-stalking features provided by manufacturers, although several noticed that their iPhones automatically

scan for AirTags. None of the players reported that any anti-stalking feature alerted them to the presence of an unwanted tracker.

Ten participants did not search for trackers that may have been planted on them, even though they knew that they could be tracked at any time and had an incentive to prevent tracking. Commonly reported reasons for this were that they had realised how hard it was to plant trackers and had little concern about trackers being planted on them, or that they felt manual searching was sufficient. The participant who used some anti-stalking features discovered Apple's Tracker Detect app as well as the third-party AirGuard app and used the background and manual scans available in the respective applications, although they did not know about the Tile, Chipolo, or Galaxy anti-stalking features. Although they used the two applications to search for unwanted trackers, they did not locate any trackers planted on their person. This participant rated themselves as having "somewhat above average" technical knowledge, although they had never used a tracker or participated in an Assassins' game before this study. This could suggest that only the more technical users will use these applications.

Four participants were aware of anti-stalking features but did not use them. Of 19 respondents to our post-game survey, 4 knew of at least 1 anti-stalking feature but none of them used the features. None of these respondents explained why they did not use features they knew of; one possible explanation is they remembered about these features when completing the post-game survey, but did not think to use them during the game. None of the respondents knew about the Samsung SmartThings application.

The most common anti-stalking feature used was the iPhone background scanning for Airtags. Even though participants did not look for technical means to detect trackers planted on them, Apple has unwanted AirTag detection built in to the OS of their phones and enabled by default, which means that users of these devices already had some anti-stalking protection by default, requiring no manual intervention.

6 DISCUSSION

There are a number of implications of the results of our study. In this section, we discuss the scenarios in which trackers are more easily abused for stalking, in addition to which devices are easiest to abuse and why. We then describe the key limitation of anti-stalking features identified in this study, as well as how default options impact this. We finish by discussing the industry response to our work.

6.1 Misuse Scenarios

Many participants in our study reported that it was difficult to plant trackers due to having to attempt to plant them in public spaces. This is analogous to how stalkers may attempt to use trackers, and reflects issues they would also face: being spotted while planting a tracker, or being unable to access the victim's possessions to place a tracker as the victim is in close proximity to them. The difficulty may vary depending on if the stalker is targeting a specific person, or opportunistically stalking a victim when presented with a moment to plant the tracker surreptitiously.

On the other hand, it is much easier to plant these trackers in domestic abuse scenarios, where the abuser has access to all of the victim's possessions and can easily plant a tracker without the victim knowing (such as when they are out of the house or occupied with chores or downtime activities). This is reflected by real-word reports of tracking devices being used for stalking: the vast majority of stories of their misuse are in domestic abuse cases, while few cases see trackers used on strangers.

6.2 Ease of Abusing Trackers

Our first research question explores the usability of tracking devices for malicious purposes. We find that the effectiveness of item finding devices for stalking depends primarily on the remote location update mechanism used. Trackers which report their location when near to a phone of the same brand (AirTags and Galaxy SmartTags) are most effective for tracking, due to the frequent updates of the device's location. Other trackers are less effective as it is rare for them to update, which made them useless for participants in our study.

While improved remote tracker location enables malicious use, it also has legitimate benefits. If a personal possession is left at an unknown location, the owner can see where it has been left; if it is lost by airport or mail staff, the owner can inform the company of the current location; and if an item is stolen, it can be tracked down through remote tracking features.

6.3 Almost No-one uses Anti-Stalking — Even When They Know They're Being Stalked

Our most striking finding is that of the 19 respondents to the post-game survey who knew they could be stalked as part of this study, only one participant found available anti-stalking features and attempted to use them. Other participants either searched their possessions manually or did not search for trackers at all. This provides an unexpected response to our second research question concerning the countermeasures used to locate and remove trackers, as users appear not to put much if any effort into removing tracking devices planted on their person.

In a real-world scenario, people will not know if they are being tracked, however in our study all participants consented to be tracked beforehand and were explicitly made aware that they were likely to be tracked during the study. While iPhone users have background scanning by default, this only detects trackers that use the Find My network. In our evaluation, the participants were incentivised to detect and remove any trackers. Furthermore, our scenario in which skilled "assassins" are hunting each other while knowing they are being hunted by skilled players is one of the optimal scenarios for ensuring participants will use all available mechanisms to defend themselves and find unwanted trackers. However, they are rarely aware of or use the anti-stalking features provided by manufacturers — which means that in the real world, even fewer people are likely to use anti-stalking features.

It is worth noting that in our study, no participant reported using the anti-stalking features successfully to locate a tracker. This is primarily due to the lack of successfully planted trackers in combination with very few players using anti-stalking features. This limits our evaluation of anti-stalking features to their availability and accessibility, rather than allowing evaluation of the features' usability and effectiveness discussed in RQ3 and RQ4.

6.4 Defaults

One limitation of the Samsung SmartThings app is that background scanning is disabled by default. While it is great that background scanning is available, the easier option presented to users is to manually scan for trackers; they have to go into settings to enable background scanning for unwanted trackers. This is an unnecessary barrier to accessing better anti-stalking features and should be enabled by default. In contrast, background scanning for AirTags is provided by default for all iPhone users, even if they are not aware of anti-stalking features.

6.5 Reflections On Study Design

Our study uses a novel methodology to attempt to overcome the significant difficulties associated with this challenging area of research. Using a gamified study allows ethically difficult scenarios to be observed in the real world. Using a

naturalistic study, in that participants did not know when they were being tracked (the study ran for 8 weeks), who by, or by what type of tracker, is consistent with how trackers may be used for stalking and leads to more realistic observations than a laboratory study. Rewarding participants for using and detecting trackers encourages controlled misuse of the devices, which strengthens our findings about their lack of uptake of anti-stalking features.

However, a number of limitations remain. Despite having 91 tracker-enabled players and 40 trackers, we only saw 4 trackers successfully planted on targets during our study. Future work may implement stronger incentives to use the trackers provided, or design the study such that tracking other players is a requirement to participate. Additionally, we note that the one-life system of the game limits the number of opportunities to track other players and prevents repeated use of trackers if caught using them during an attempt, which likely contributes to the small sample size. A future study may choose to allow "respawning" a few hours after a player is "killed", enabling players with trackers to try again and avoiding a shrinking pool of targets for tracking. While the small sample limits the conclusions that can be drawn about the effectiveness of each anti-stalking feature, we note the purpose of the study was to explore users' experience with anti-stalking features, rather than evaluate the usefulness of the trackers themselves.

In our study, participants were explicitly made aware that they would be tracked by other players to ensure that they can properly consent to this study. Although this makes the study slightly less naturalistic, in that people in the real world will not know if and when they are being tracked, this does imply that even fewer people in the real world will use the anti-stalking features.

The participants in our study are university students primarily from a STEM background with a self-reported higher than average technical knowledge. This implies that the participants are more likely to be able to use the trackers, and anti-stalking features, more effectively than the average user. This suggests that in the real world, even fewer people would make use of the anti-stalking features than in our study — however, almost no participants made use of them. The demographic of our participants, combined with them explicitly being made aware of possible tracking and not using anti-stalking in spite of this, suggests that the anti-stalking features are extremely unlikely to be discovered and effectively used by real world users.

6.6 Reporting and Response

The results and recommendations from this paper were communicated to the four companies which provide anti-stalking features in early March 2023. The summary of our recommendations are as follows:

- Anti-stalking should be enabled by default on all devices.
- Anti-stalking mechanisms should provide background scanning for unwanted trackers as a primary mechanism, with manual scanning available as a secondary option.
- All scanning mechanisms must provide methods to locate unwanted trackers after they are detected, such as providing an estimate of distance through Bluetooth and triggering sound alerts on the devices.
- Trackers should be designed such that all anti-stalking software is able to detect any brand of tracking device.
- Anti-stalking features should be build into the operating system where possible, instead of requiring users to download many apps to detect unwanted trackers.

Apple responded with a list of improvements they made in June 2021 and February 2022 despite the study running at the end of 2022, but did not directly respond to our results and suggestions. We additionally reached out to Google with our results, aiming to integrate anti-stalking into the Android OS to allow for universal anti-stalking features. We met

with several people on Google's safety teams to discuss our work and recommendations, in addition to possible issues that may arise (such as false positives) and how to design around them.

In May 2023, Apple and Google announced a joint draft RFC [16] to standardise anti-stalking features and how they are implemented. This includes standardising the background scanning in operating systems, ensuring all trackers provide methods to trigger sound alerts, suggesting other mechanisms for locating identified unwanted trackers in future designs, and ensuring trackers provide a feature similar to Apple's away-from-owner alerts. The standard includes all of our suggestions for improving existing implementations. Other companies including Samsung, Chipolo, Tile, Pebblebee and eufy Security have expressed interest in this draft, suggesting they will incorporate the changes once the standard is complete.

7 CONCLUSIONS

In this study we have analysed the use of personal item trackers for stalking and analysed the use of anti-stalking features. We see that it is difficult to plant these devices on strangers and their possessions, making them more of a threat in domestic-abuse scenarios than stalking of strangers, although both are risks. We found that even in the ideal case of technically skilled people knowing that they will be stalked and what by, very few people use or are even aware of anti-stalking features available to them, implying that the feature sees little to no use in the real world. This, combined with other failures of the anti-stalking features, implies that they are little more than security theatre and are intended to quell concerns rather than to improve user safety. We provide recommendations for improving the availability of these features so that they are more accessible to the users who need them.

ACKNOWLEDGMENTS

We thank our colleagues at the Cambridge Cybercrime Center for their feedback, and the Cambridge University Assassins' Guild for collaborating on this project. We would also like to thank Henry Caushi for inspiring the project.

This work was supported by the UK Engineering and Physical Sciences Research Council (EPSRC) grant EP/T517847/1 and the European Research Council (ERC) under the European Union's Horizon 2020 research and innovation programme (grant agreement No 949127).

REFERENCES

[1] 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. Nudges for privacy and security: Understanding and assisting users' choices online. *ACM Comput. Surv. 50*, 3 (aug 2017).

[2] Apple. An update on airtag and unwanted tracking. https://www.apple.com/uk/newsroom/2022/02/an-update-on-airtag-and-unwanted-tracking/, 2022.

[3] Bellini, R., Tseng, E., McDonald, N., Greenstadt, R., McCoy, D., Ristenpart, T., and Dell, N. "So-Called Privacy Breeds Evil": Narrative justifications for intimate partner surveillance in online forums. *Proceedings of ACM on Human-Computer Interaction 4*, CSCW3 (January 2021).

[4] Cahn, A. F. Apple's AirTags are a gift to stalkers. https://www.wired.com/story/opinion-apples-air-tags-are-a-gift-to-stalkers/, 2021.

[5] Caraban, A., Karapanos, E., Gonçalves, D., and Campos, P. 23 ways to nudge: A review of technology-mediated nudging in human-computer interaction. In *Proceedings of the 2019 CHI Conference on Human Factors in Computing Systems* (New York, NY, USA, 2019), CHI '19, Association for Computing Machinery, p. 1–15.

[6] Charlton, H. Apple's AirTag item trackers increasingly linked to criminal activity. https://www.macrumors.com/2021/12/31/airtag-increasinglylinked-to-crime/, 2021.

[7] Cole, S. Police records show women are being stalked with Apple AirTags across the country. https://www.vice.com/en/article/y3vj3y/apple-airtags-police-reports-stalking-harassment, 2022.

[8] Douglas, H., Harris, B., and Dragiewicz, M. Technology-facilitated domestic and family violence: Women's experiences. *British Journal of Criminology 59*, 4 (2019), 551–570.

[9] Dragiewicz, M., Harris, B., Woodlock, D., Salter, M., Easton, H., Lynch, A., Campbell, H., Leach, J., and Milne, L. *Domestic violence and communication technology: Survivor experiences of intrusion, surveillance, and identity crime.* The Australian Communications Consumer Action Network (ACCAN), Australia, 2019.

[10] Fowler, G. A. Apple's AirTag trackers made it frighteningly easy to 'stalk' me in a test. https://www.washingtonpost.com/technology/2021/05/05/apple-airtags-stalking/, 2021.

[11] Freed, D., Palmer, J., Minchala, D., Levy, K., Ristenpart, T., and Dell, N. "A Stalker's Paradise": How Intimate Partner Abusers Exploit Technology. In *Proceedings of the 2018 CHI Conference on Human Factors in Computing Systems* (New York, NY, USA, 2018), CHI '18, Association for Computing Machinery, p. 1–13.

[12] Frik, A., Kim, J., Sanchez, J. R., and Ma, J. Users' expectations about and use of smartphone privacy and security settings. In *Proceedings of the 2022 CHI Conference on Human Factors in Computing Systems* (New York, NY, USA, 2022), CHI '22, Association for Computing Machinery.

[13] Heinrich, A., Bittner, N., and Hollick, M. AirGuard – Protecting Android users from stalking attacks by Apple Find My devices. https://arxiv.org/abs/2202.11813, 2022.

[14] Hill, K. I used Apple AirTags, Tiles and a GPS tracker to watch my husband's every move. https://www.nytimes.com/2022/02/11/technology/airtags-gps-surveillance.html, 2022.

[15] Holpuch, A. Two women sue Apple over AirTag stalking. https://www.nytimes.com/2022/12/06/business/apple-airtag-lawsuit.html, 2022.

[16] Ledvina, B., Eddinger, Z., Detwiler, B., and Polatkan, S. P. Detecting Unwanted Location Trackers. Internet-Draft draft-detecting-unwanted-location-trackers-01, Internet Engineering Task Force, Dec. 2023. Work in Progress.

[17] Leitão, R. Technology-facilitated intimate partner abuse: a qualitative analysis of data from online domestic abuse forums. *Human–Computer Interaction 36*, 3 (2021), 203–242.

[18] Levy, K., and Schneier, B. Privacy threats in intimate relationships. *Journal of Cybersecurity 6*, 1 (2020), 1–13.

[19] Mac, R., and Hill, K. Are Apple AirTags being used to track people and steal cars? https://www.nytimes.com/2021/12/30/technology/apple-airtags-tracking-stalking.html, 2021.

[20] Matthews, T., O'Leary, K., Turner, A., Sleeper, M., Woelfer, J. P., Shelton, M., Manthorne, C., Churchill, E. F., and Consolvo, S. Stories from survivors: Privacy & security practices when coping with intimate partner abuse. In *Proceedings of the 2017 CHI Conference on Human Factors in Computing Systems* (New York, NY, USA, 2017), CHI '17, Association for Computing Machinery, p. 2189–2201.

[21] National Network To End Domestic Violence. A glimpse from the field: How abusers are misusing technology. https://www.techsafety.org/blog/2015/2/17/a-glimpse-from-the-field-how-abusers-are-misusing-technology, 2015.

[22] Refuge. Refuge, the UK's largest domestic violence charity says 72% of service users experience abuse through technology and launches chatbot to provide real-time information to survivors. https://www.refuge.org.uk/72-of-refuge-service-users-identify-experiencing-tech-abuse/, 2019.

[23] Sambasivan, N., Batool, A., Ahmed, N., Matthews, T., Thomas, K., Gaytán-Lugo, L. S., Nemer, D., Bursztein, E., Churchill, E., and Consolvo, S. "They Don't Leave Us Alone Anywhere We Go": Gender and digital abuse in South Asia. In *Proceedings of the 2019 CHI Conference on Human Factors in Computing Systems* (New York, NY, USA, 2019), CHI '19, Association for Computing Machinery, p. 1–14.

[24] Shafqat, N., Gerzon, N., Van Nortwick, M., Sun, V., Mislove, A., and Ranganathan, A. Track you: A deep dive into safety alerts for apple airtags. *Proceedings on Privacy Enhancing Technologies 4* (2023), 132–148.

[25] Tom Scott plus. Can you stalk someone with an Apple AirTag? https://www.youtube.com/watch?v=GmC05wOc5Dw, 2022.

[26] Tom Scott plus. He tracked me with an AirTag. now it's my turn. https://www.youtube.com/watch?v=NuEgjAMfdIY, 2022.

[27] Turk, K. I., Hutchings, A., and Beresford, A. R. Can't keep them away: The failures of anti-stalking protocols in personal item tracking devices. In *Security Protocols XXVIII* (Cham, 2023), F. Stajano, V. Matyáš, B. Christianson, and J. Anderson, Eds., Springer Nature Switzerland, pp. 78–88.

[28] Woodlock, D. The abuse of technology in domestic violence and stalking. *Violence Against Women 23*, 5 (2017), 584–602.

A ANTI-STALKING ALERT TIMING TESTS

Table 4. Detection Time for each Tracker by each Device using Background Scanning

Tracker	AirTag		Tile	Galaxy	Chipolo	
Detected By	Find My	AirGuard	AirGuard	SmartThings	Find My	AirGuard
iPhone	5h 41m	—	—	Not Found	Not Found	—
Galaxy	—	Not Found	Not Found	Not Found	—	Not Found
Pixel	—	25m	25m	Not Found	—	Not Found

A.1 User 1: Galaxy Phone

I have a Samsung Galaxy A40. My phone's Bluetooth and location were switched on all day. The trackers were planted at 9:45am, in the department next door to mine. I then walked to my department, and the trackers remained in my bag all morning, with my phone less than a metre away most of the time (the only exceptions being three trips to get coffee from the departmental machine taking a few minutes each). The phone automatically connected to the local University Eduroam WiFi and remained connected from 9:45-1pm. At 1pm, I took my bag (and trackers) to the gym to get changed for a run. The walk to the gym took 5 minutes with my phone's network data being switched on. At the gym, the phone automatically connected to the local Eduroam wifi. At 1:15pm I took the trackers with me on the entire 10.71 mile run (which took 1.5 hours). The trackers were stored in a pocket a few centimetres away from my phone at all times. My phone was connected to my phone's network data at all times during the run (and bluetooth & location were switched on). I didn't stop at all during the 1.5 hours and ran across fields and through local villages during the run. I returned to the gym at 2:45pm at my phone automatically connected to the university Eduroam wifi. At 3pm I walked back to my University department and my phone again connected to the university Eduroam wifi automatically at 3:10pm. The trackers then stayed within 1 metre of my phone all afternoon (with the exception of two trips of less than 5 minutes duration each for coffee from the local departmental machine). At 5.10pm I walked home. My phone automatically disconnected from university Eduroam wifi and my phone's network data automatically took over. I arrived home at 17:39pm and my phone automatically connected to my home wifi. My phone remained connected to the home wifi all evening. I didn't receive any notifications regarding the tags at all. The Chipolo app briefly gave one notification at around 1pm that the Chipolo app was "running in the background" (at approximately the time my phone would have disconnected from Eduroam and connected to my phone's network data when I started walking to the gym) but no notifications regarding tags detected. None of the other apps gave any notifications.

A.2 User 2: iPhone

I have an iPhone X. I carried the trackers in my pocket all day. I cycled to the end of the street (so as to no longer be on my home Wi-Fi) before turning my Bluetooth on at 10:15am. I then continued cycling to my destination (2.1 miles). I was at this destination for around 40 minutes, before cycling to my next destination (2.3 miles). I left after approximately one hour and wandered around town on foot. I then cycled 1.3 miles to the grocery store and did a weekly shop, before cycling .3 miles home.

I returned home at 3:50pm. At 3:56pm I received a notification on my iPhone that it had detected an AirTag moving with me. I tapped on the notification, and it opened the Find My app. It first showed me a 'what's new in Find Me' page. After I hit continue it showed me a map of my travels, and informed me it first detected the AirTag at 10:20am. One option presented was to learn more about the AirTag. I selected this and was instructed to bring the phone near the device, then tap on the on-screen notification. I tried to do this for several minutes, but no notification was delivered.

None of the other apps I installed detected any of the trackers. My iPhone only detected the AirTag, not the other trackers. It only notified me of this over 5.5 hours after I started being tracked, soon after I returned home and connected to the home Wi-Fi.

A.3 User 3: Google Pixel

I have a Google Pixel 4. I obtained the trackers at my department and put them in the lower pocket of my trousers, below the pocket containing my phone. I kept my Bluetooth off until 4:01pm when I left the department. I then walked 1.4 miles to the

city centre and began grocery shopping, and shortly after arriving I received 2 simultaneous notifications from the AirGuard app at 4:26pm. The notifications alerted me to an unknown AirTag and Tile Sticker respectively, and opening the notifications showed the path I had taken into the city centre. I then continued shopping around the city centre, walking approximately 1.6 miles total, before walking 0.8 miles home. I stayed at home for 2 hours, then walked the 0.8 miles back to the city centre for a social event. After several hours, I walked back home. Aside from the two early notifications, I was not alerted to any unwanted trackers present on me.

B SURVEY QUESTIONS

B.1 Demographics

- What course are you studying?
- How would you rate your knowledge of technology? (5 point scale)
- How much prior experience do you have using tracking devices? (5 point scale)
- How many previous games of assassins have you played? (0, 1-3, 4-6, 7-9, 10+)
- What type of phone do you have?
- What other devices do you own which can detect a tracker?

B.2 Post-Mortem Survey

- Please enter your name
- Were you the assassin or assassinated?
- *Only if Assassin:*
  - Which tracker did you use to track your target?
  - Where did you plant the tracker on your target? (bag, pocket, ...)
  - How did you manage to plant the tracker on your target? (placed when they were away, planted on their person, ...)
  - How did you use the tracker to locate your target?
  - How useful did you find the tracker for locating your target? (5 point scale)
  - How easy was it to plant the tracker on your target? (5 point scale)
- *Only if Assassinated:*
  - Which tracker was used against you?
  - Did you locate the tracker planted on you?
  - Did you use any anti-stalking features of the tracking device? Which ones? e.g. Notifications of a device following you, manual scan, sound alerts, locating via Bluetooth...
  - How easy was it to discover the anti-stalking features of the device? (5 point scale)
  - How useful were the anti-stalking features provided by the device? (5 point scale)
- If you were killed and currently have a tracker on another person that has not been recovered, or if you killed another player and the tracker is still on their person, we need to return your tracker to you for their privacy. Please provide as much information as possible as to the current location of your tracker.

B.3 End of Study Survey

- **Participant Information**

  - Please enter your name
  - Were you provided with a tracker?
  - Would you be happy to be contacted at a later date to take part in a publicity video about this study?
- **Planting Trackers**
  - How did you plant trackers on your targets? (If used multiple times, separate answers with newlines; if not used, write N/A)
  - How easy was it to plant trackers on your target(s)? (5 point scale)
  - How useful was the tracker for locating your target(s)? (5 point scale)
  - Do you have any further comments about your experience using the tracker?
- **Locating Trackers**
  - How did you search for trackers that could have been planted on you?
  - Did you use any of Apple's anti-stalking features?
  - *(If yes)* How useful were Apple's anti-stalking features? (5 point scale)
  - Did you use any of Tile's anti-stalking features?
  - *(If yes)* How useful were Tile's anti-stalking features? (5 point scale)
  - Did you use any of Samsung's anti-stalking features?
  - *(If yes)* How useful were Samsung's anti-stalking features? (5 point scale)
  - Did you use any of Chipolo's anti-stalking features?
  - *(If yes)* How useful were Chipolo's anti-stalking features? (5 point scale)
  - Did you use the third-party AirGuard app??
  - *(If yes)* How useful was AirGuard? (5 point scale)
  - Did you use any other electronic means to find trackers placed on you?
  - Do you have any other comments about finding the trackers?