Cart-ology: Intercepting Targeted Advertising via Ad Network Identity Entanglement

ChangSeok Oh
Georgia Institute of Technology
changseok@gatech.edu

Chris Kanich
University of Illinois Chicago
ckanich@uic.edu

Damon McCoy
NYU
mccoy@nyu.edu

Paul Pearce
Georgia Institute of Technology
pearce@gatech.edu

ABSTRACT

Targeted advertising is a pervasive practice in the advertising ecosystem, with complex representations of user *identity* central to targeting. Ad networks are incentivized to tie ephemeral cookies across devices to lasting durable identifiers such as email addresses in order to develop comprehensive cross-device user profiles. Third-party ad networks typically do not have relationships with users and must rely on external parties such as merchant websites for durable identity information, introducing intricate trust relationships. We find attackers can exploit these trust relationships to confuse an ad network into linking an unprivileged attacker's browser to a victim's identity, thus "impersonating" the victim to the ad network.

We present *Advertising Identity Entanglement*, a vulnerability to extract specific user browsing behavior from ad networks remotely, knowing only a victim's email address, with no access to the victim, ad network, or websites. This new fundamental flaw in cross-device tracking allows attackers to pass erroneous identity information to third-party ad networks, causing the networks to confuse attacker and victim. Once entangled, the attacker receives advertisements intended for the victim across the entire ad network. We find identity entanglement is a significant user privacy vulnerability where attackers can learn detailed victim browsing activity such as retail websites, products, and even specific apartments or hotels the victim has interacted with. The vulnerability is also bi-directional, with the attacker able to cause specific ads to be shown to the victim, introducing the possibility of embarrassment attacks and blackmail. We have disclosed the vulnerability; Criteo, one of the largest third-party ad networks, acknowledges the attack.

CCS CONCEPTS

• **Security and privacy** → *Web application security*; Domain-specific security and privacy architectures.

KEYWORDS

Web Privacy; Targeted Advertising; Ad Networks; Tracking

CCS '22, November 7–11, 2022, Los Angeles, CA, USA

ACM ISBN 978-1-4503-9450-5/22/11.
https://doi.org/10.1145/3548606.3560641

ACM Reference Format:
ChangSeok Oh, Chris Kanich, Damon McCoy, and Paul Pearce. 2022. Cart-ology: Intercepting Targeted Advertising via Ad Network Identity Entanglement . In *Proceedings of the 2022 ACM SIGSAC Conference on Computer and Communications Security (CCS '22), November 7–11, 2022, Los Angeles, CA, USA.* ACM, New York, NY, USA, 14 pages. https://doi.org/10.1145/3548606.3560641

1 INTRODUCTION

Targeted online advertisements reveal sensitive private information about a person. Ad networks invest significant resources to develop comprehensive profiles of user activity in order to more effectively target advertisements [1]. To create these profiles, ad networks leverage durable user information (e.g., user accounts or email addresses) to link together various devices operated by a user (e.g., a desktop, a laptop, and a phone) [2, 31], and then aggregate detailed browsing and shopping behavior by leveraging a litany of first and third-party trackers embedded into websites across the Internet [4].

A particularly privacy invasive form of targeted advertising is *retargeted ads*, where a specific product that a person has previously viewed is included in an ad delivered to their device [35, 41]. Retargeted ads, by design, can "follow" users to their other devices via complex cross-device tracking. Cross-device tracking methods enable persistent tracking by linking together third-party ephemeral tracking cookies to a user's *durable identity* (e.g., email address) [10]. These ads are sometimes mistakenly shown to the devices of family members and friends that often share network connections [32]. Such "leaked" retargeted advertisements can cause a range of harms from ruining a surprise gift to revealing sensitive personal information such as sexual orientation, religious affiliation, or pregnancy status [11, 28]. The cause of some of these anecdotally reported privacy leaks are likely errors in the cross-device tracking methods deployed by third-party ad networks. These reports of privacy leaks motivated us to explore the security of cross-device tracking methods, seeking to understand how secure these methods are against an adversary that is intentionally attempting to subvert them.

Ad networks such as Google's DoubleClick have, in essence, first-party relationships with users given Google's user-facing products such as Gmail or YouTube. Direct user relationships enable such ad networks to have durable and verifiable notions of user identity that can be used to definitively link together a user's various browsers and devices (i.e., cross-device tracking). Third-party ad networks, however, lack direct relationships with users, making it challenging to link devices. One way to compete with first-party ad networks and develop similarly effective targeting based on comprehensive

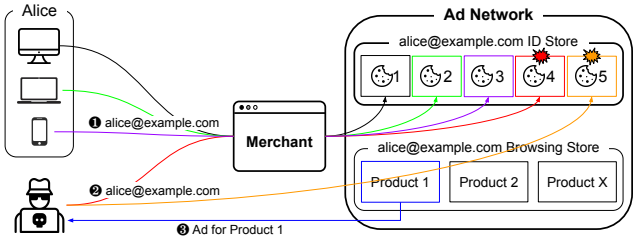

Figure 1: *Identity Entanglement*: (1) Ad networks link Alice's browsers, devices, and product views into a single profile using her email address. (2) Attacker inserts their browser's cookies into Alice's profile, via either editing requests to the ad network or confusing a merchant website, resulting in an *entangled* profile comprising both attacker and Alice. (3) Attacker receives ads for Alice's products.

user profiles is for third-party ad networks to find a source of both identifier information and user activity to link to their own tracking cookies. These requirements can be fulfilled by merchant websites; these sites *do* have a first-party relationship with users and form a symbiotic relationship with third-party ad networks wherein the ad networks learn durable user identity information and behavioral activity (via embedding trackers into the merchant websites), and the merchants are then able to more effectively target their own products back to users via ads shown across the Internet. This level of identity indirection introduces a problem of federated trust, where the third-party ad network must now rely on outside parties to verify and supply durable identity information used by the ad network to perform cross-device tracking and generate retargeted ads based on user behavior across all devices. Worse still, this identity information flows through a user's web browser without integrity or authenticity, putting it under the attacker's control.

This paper explores security at the intersection of retargeted advertising, cross-device tracking, and identity. We discover a significant new class of privacy vulnerability wherein **an attacker extracts specific victim browsing behavior knowing only the victim's email address, with no prior knowledge of the victim's behavior, ads they see, or access to their computer or accounts**. Attackers achieve this by abusing the intricate relationship between merchant websites and third-party ad networks and entangling their own browser's tracking cookie with an ad network's cross-device representation of a user's identity.

To exploit this *Identity Entanglement* vulnerability, an attacker identifies a merchant website, dubbed a "patsy," which is expected to send durable identity information (e.g., an email address) to a third-party ad network. The ecosystem and structure of third-party ad networks makes the exchange of durable identifiers without integrity within the web browser commonplace. We show these limitations allow an attacker to entangle themselves into the ad network's cross-device representation of the victim in one of two ways. First, the attacker can rewrite HTTP requests sent to the ad network but initiated by the patsy website, replacing their own email address with the victim's. Second, alternatively, the attacker can trick a patsy website without email verification (which we show is common) into generating HTTP tracking requests to the ad network which embeds the victim's email address with the attacker's cookies. In both cases, the third-party ad network has no ability to disambiguate a new attacker device from a new user device. Once entangled, an attacker is considered an additional device of the victim, and begins receiving retargeted ads intended for the victim. This is a severe and harmful privacy violation that reveals sensitive information to an arbitrary attacker with minimal assumptions. Figure 1 provides an overview of the attack.

In this work we discover the problem of Identity Entanglement and develop an attack method. However, even after an attacker's tracking cookie is entangled with the ad network's representation of a user, it is still challenging to extract specific user activity given the complexity of ad network algorithms. We therefore develop a measurement technique that allows attackers to extract user behavior from background advertising. We then explore the scope of the problem using our measurement technique in two contexts, discovering problems with both. First, we examine Criteo, the largest third-party ad retargeting network [15], which serves more than 4 billion advertisements a day [18], covering 75% of the world's shoppers [20], and has been shown to be the largest source of retargeted ads by an order of magnitude [6]. We find an attacker can entangle their device with Criteo's cross-device representation of a user, extracting detailed user activity such as retail items viewed, and apartments or hotels the user searched. Second, we explore Yahoo Ad's Analytics and Ads network which reaches tens of millions of users daily [45]. We find that attackers can similarly rewrite email addresses in Yahoo's Analytics interface, leading to the same identity entanglement with Criteo. We have responsibly disclosed the vulnerability, and Criteo has acknowledged the attack.

Specifically, our contributions include:

- Developing and describing a new type of vulnerability where an attacker receives retargeted ads intended for a specific user, knowing only the user's durable identifier, such as email address. The attacker requires no access to the user's devices, accounts, prior browsing behavior, or other information.
- Demonstrating how an attacker can extract victim browsing information from Criteo, one of the largest third-party ad networks. We also show how an attacker can utilize the Yahoo Analytics interface to cause Criteo to leak potentially sensitive information. Combined, we find these networks cover 35% of all tracker-containing websites.
- Developing an attack method to determine if entanglement is successful and then (with no prior knowledge of user behavior or ads) extracting private information about the user from the ad network including: 1) what online merchants they visit, 2) what *specific* items they interact with, and 3) in some cases the user's location and travel plans (e.g., viewed apartments and hotels).
- Demonstrating that entanglement is bi-directional, allowing an attacker to influence which retargeted ads are shown to the victim. This bi-directionality allows blackmail and embarrassment.
- Discussing near and long-term mitigations, as well as how the problem may evolve with changes to ad and browser ecosystems.

Ad networks currently have no way of confirming that the durable identity provided by merchants is correct [16]. Worse, much of the tracking occurs inside the user's browser (without integrity), under the control of the attacker. As such, we believe this vulnerability to be fundamental, requiring significant changes technically in how third-party ad networks perform cross-device tracking, and to better align incentives for ad networks to protect privacy.

2 BACKGROUND

The majority of online advertising is targeted [5] in order to improve marketing Return on Investment (ROI) [27]. Targeted advertising is a way for marketers to present users with customized ads based on their specific features such as demographics (i.e., gender, age), location, interests, browser history, and shopping behavior. We now provide background on the technology that enables the tracking of users and the collection of information that is used for ad targeting.

2.1 Tracking Cookies

Third-party tracking is the process by which an entity other than the website a user is visiting tracks activity on that website [37]. For example, if a user visits a merchant such as macys.com, a third-party tracker such as facebook.com placed on the webpage by Macy's could track the user's Macy's activity. Such tracking ultimately seeks to link user activity on macys.com with their activity across all websites that embed the tracker [37].

There are numerous forms of client state that enable tracking, the most pervasive and well-known of which are *cookies*. Conceptually, cookies are an *ephemeral identifier* represented as a tuple of (domain, key, value) that is stored inside the browser and are accessible for reading and writing whenever a user visits that domain. Cookies set by the specific domain the user visits (e.g., macys.com) are *first-party* cookies, and cookies set by different domains or scripts embedded in that page (e.g., facebook.com) are *third-party cookies* [37].

Cookies that facilitate third-party tracking are sent via JavaScript that runs on the website the user visits or via HTTP headers ("Set-Cookie") [37]. Once a cookie is set, it is automatically sent out with HTTP responses via the header or programmatically via JavaScript APIs. Cookie usage is governed by the *Same-Origin* policy that ensures that cookies cannot be shared between unrelated domains [37]. Users ultimately have control over if cookies are stored and or sent via advanced settings in their browser, but disabling cookies is uncommon and renders the web unusable [37].

Mozilla's Firefox was the first major browser to block third-party cookies by default in 2019, Apple's Safari also started blocking them in 2020, and Google has stated that Chrome will block third-party cookies by 2023 [8, 13, 43]. However, blocking of third-party cookies does not prevent browser fingerprint-based tracking [24, 34], nor the visited website from providing information about a user to a tracking company. Tracking cookies are also limited to a single application or browser which provides an incomplete view when users use multiple applications, browsers, or devices.

The tracking cookie identifiers used by the ad networks we explore are random UUID-like opaque values. They contain no durable identifier information such as email address. All connections between the cookie identifiers and victim identity are maintained by the ad networks outside the scope of the victim or attacker.

2.2 Cross-Device Tracking

Online advertising and analytics companies have developed several methods to track users across devices. Ad networks that have a direct login relationship with a user, such as Facebook or Google, leverage the user logging into their account from all of their devices to enable cross-device tracking. After a user logs into their account, their tracking cookies are synced cross-devices, or the device tracking cookies are linked in the ad network's backend database so a more complete profile of user activity can be utilized for ad targeting [10, 47]. Conceptually, this can be thought of as a graph of numerous ephemeral tracking cookies that are all linked together with a single durable identity. This constellation of ephemeral identities constitutes an ad network's "profile" of a user. In the context of our identity entanglement attack, we are conceptually inserting a new erroneous edge into this graph.

Third-party tracking companies do not have login relationships with consumers directly, but do have relationships with merchants that do have a login relationship with users. These merchants can provide identifying information such as email addresses or hashed versions of these identifiers during login to a third-party tracking company. If the same email address or hash is transmitted to the same third-party from different devices, the company can then match the strings together and either sync the tracking cookies on the user's devices or link the two tracking cookies in their backend database to enable cross-device tracking. However, in this case, the third-party tracking company must trust that the merchant: 1) has verified that the user actually controls the email address, and 2) is not malicious. We will show that most merchants do not verify that the user controls the email address which enables an attacker to insert their device into another user's identity profile. This can then result in information leaking from targeted advertisements.

2.3 Retargeting Advertisements

A retargeting ad is an advertising technique that often includes a product a person engaged with previously. Retargeting ads are commonly targeted based on Personally Identifying Information (PII) (i.e., email, phone number). The information for retargeted ads is often provided by the merchant to the ad network and is based on searching, viewing or adding a product to the shopping cart. This highly personalized advertising is intended to remind customers to revisit their shopping websites and purchase products that they previously expressed interest in and that are included in the ads.

Retargeting ads have raised several privacy concerns with end users. First, advertisements that accurately reflect a consumer's interest (i.e., contain a product they previously viewed) explicitly indicates that a user's shopping activity is being tracked and provided to ad networks. Second, retargeted ads often contain privacy sensitive information, so the ad network could unintentionally become a privacy leaking channel for adversaries. Anecdotally, there are stories of information about purchases such as gifts leaking to another user's device in the same household potently when someone logs into a family member's device or through other mistakes in cross-device linking techniques [9]. However, to the best of our knowledge, there has not been an analysis of cross-device tracking techniques to understand if an attacker can intentionally entangle their device with another user's identity to learn potentially sensitive private information from retargeted ads.

2.4 The Centrality of Criteo

The majority of our experiments center on Criteo, the largest third-party tracking company [15, 26]. Per whotracks.me Criteo exists across 31% of all sites containing trackers. The next largest third-party ad network is Yahoo, which we also explore. We, however, find that Yahoo only marginally adds 4% coverage to websites. SimilarTech, another industry tracking data source also indicates that

Criteo is the dominant player in the space [39], with other entities being Facebook or non-ad network identity providers. They indicate Criteo's reach is 27x that of the next third-party ad network (Yahoo) [39]. Criteo's own marketing purports to serve more than 4 billion ads a day and has data on 75% of Internet shoppers. Criteo's prevalence in this space, combined with their third-party relationship with web users, indicates their centrality to this ecosystem and the severity of identity entanglement on their platform.

3 THREAT MODEL AND ETHICS

We make minimal assumptions across the victim, advertiser, ad network, and merchant websites. We assume merchants share information with the third-party ad network that we studied for the purposes of targeted and retargeted advertising, victims use merchant sites with JavaScript and cookies enabled, an attacker knows some piece of information about a user (e.g., email address).

We assume that the victim: 1) Visits merchant websites with JavaScript and cookies enabled. 2) Does not block ads. 3) Has an account with a durable identifier (e.g., email address) at one or more merchant websites partnered with the third-party ad network.

We assume that the attacker: 1) Knows the durable identifier (e.g., email address) of a specific user they wish to target, which that user used for at least one merchant website account. 2) Is able to identify one merchant in the third-party ad network that sends the user's identity information. We call such a website a *"patsy" website*. We note that such behavior is common [14, 22, 26], and that the victim **does not** need to have ever visited or interacted with the patsy website. 3) We **do not** assume the attacker has access to the victim's computer, any of the victim's accounts, or is in any way co-located with the victim on the Internet. We also **do not** assume the attacker has any knowledge of the victim's activity, behavior, or ads they see.

We assume that merchant websites the victim visits: 1) Use the third-party ad network that we studied. 2) Pass user activity and identity information to the ad network. We note such behavior is common [14, 22]. 3) Sends user identity information to the ad network via the user's browser, with no integrity or authenticity (this is true for both Criteo and Yahoo's APIs), or that the website itself does not properly verify email addresses when creating accounts (which we find is true for 84% of merchant websites we examine).

We assume that the third-party ad network: 1) Builds comprehensive profiles of user activity across devices. This assumption is reasonable and well documented [17]. 2) Solicits identity information from websites on user activity. This assumption is reasonable and well documented [19]. 3) Has no direct relationship with users. e.g., users do not have an account with the ad network.

The above set of assumptions is reflective of the current state of merchant websites, retargeting, and third-party ad networks. Both Criteo and Yahoo's Ad networks fit into these sets of assumptions, as do the merchant websites we explore.

Implicit in these assumptions is the existence of third-party cookies. While there are discussions about eliminating such cookies in the future [8], in Section 7, we discuss that such cookies are not strictly necessary if ad networks and merchant websites collude.

Trust Model. Third-party tracking relies upon comprehensive profiles created by combining and accumulating user behavioral data. Accumulating this data across different devices, websites, and after cookies expire or are deleted requires linkage through a *durable* identifier, that is, one which does not change frequently like a session ID. Because the user only maintains a first-party relationship with a few or even one participating merchant site, the advertising network must trust merchant sites to convey that information to them, and cannot independently verify that information. Furthermore, due to the competitive nature of the online advertising ecosystem, ad networks are not incentivized to perform strict quality control or integrity checks on this durable identity linkage, and the networks we observed accept a report of the identity directly from the browser on sites where the library is loaded.

Roughly speaking, there are two core issues enabling this vulnerability: 1) Ad networks are not able to verify if the email addresses given to them by merchants are correct. 2) Often, a person's email addresses are public information.

Ethics. We exclusively used synthetic user and attacker profiles we created. At no time did we attack or attempt to attack any real user. We responsibly disclosed the vulnerability, with Criteo having acknowledged the problem.

4 IDENTITY ENTANGLEMENT

We now introduce *Identity Entanglement*, a new vulnerability that harms users' privacy. Identity entanglement occurs when an attacker tricks an ad network into linking the attacker's tracking cookie to a chosen victim's durable identity, allowing the attacker's browser to receive ads as if it were the victim. Entanglement is bi-directional, allowing the attacker to also influence ads shown to the victim. Figure 2 shows an overview of the problem.

The attack proceeds in three phases, which we describe briefly here and in more detail below. First, the victim must use their durable identity to authenticate with any merchant website participating in the ad network, thus causing their device's tracking cookie to be linked to their durable identity.

Second, the attacker must perform *identity entanglement* to associate their own tracking cookie with the victim's durable identity. The attacker can leverage either the lack of integrity on information being sent from merchant to ad network through the attacker's browser to rewrite HTTP requests, replacing their own email address with the victim's, or a lack of identity verification by merchants to trick the merchant into sending the victim's email directly.

Finally, once the cookies are entangled, the attacker periodically polls sites that display ads from the ad network to infer products and merchants that the victim visits. Similarly, the attacker can inject ads into the victim's browsing experience.

4.1 Victim Behavior

In our identity entanglement attack, a victim authenticates with at least one website that performs cookie syncing with a vulnerable ad network. In this instance, the syncing is based on the victim's email address as the durable identifier. Cookie syncing is performed when the merchant website invokes various third-party ad network's library functions, e.g., sending an event that indicates that this tracking cookie has a specific email address and has interacted with a specific product. We note that this event is sent through the user's web browser without integrity or authenticity, and is subject to attacker manipulation. Figure 2, steps 1-10 outline this process and show the victim interacting with a merchant website.

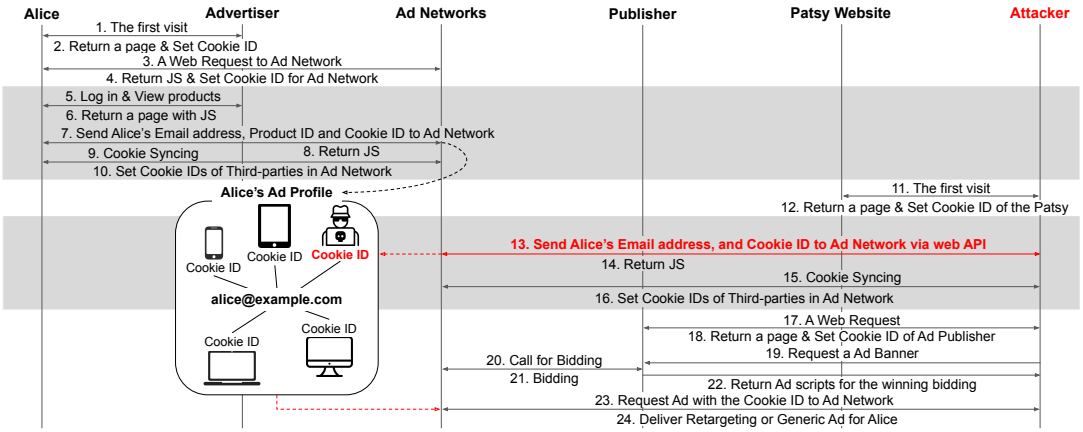

Figure 2: Step-by-step interactions between the victim, merchant, ad network, and attacker. Step 13 (color red) is the critical step leading to the vulnerability. In steps 1–10 a victim interacts with a merchant website that is member of a third-party ad network. This interaction may occur across multiple devices. As a result of this interaction the ad network creates a profile of the user's ephemeral IDs (Cookie IDs). At some point later in steps 11–16 an attacker interacts with a (potentially) unrelated patsy website, and either edits the email address (replacing the attacker's own email address with the victims) in the HTTP request, or erroneously convinces the patsy website they are the victim by giving them the victim's email address. In steps 17–24 the attacker probes unrelated publisher pages to observe targeted and retargeted advertisements.

After the API is invoked, the ad network can tie *all* user activity across *any* website in the ad network together, as the ad network has successfully linked the user's durable identifiers to the ad network's tracking cookies. This cookie is sent to the ad network for visits across all websites that utilize the ad network, making all of the victim's activity vulnerable to attack, not only the single website they authenticated with. At this point, the victim continues their normal browsing activity, and may or may not interact with specific products on this website or any other website. Websites that use the ad network for retargeted ads will use the API to report products being viewed and/or added to the shopping basket.

4.2 The Attack

The first step the attacker takes is to entangle their browser's tracking cookie with the victim's durable identifier inside the ad network. Entanglement is done via interaction with a **patsy website**. A patsy website is any merchant website in the ad network that utilizes the ad network's tracking API to send users' email addresses (or similar) to the ad network as part of login. This practice is commonplace [14, 22] and documented by ad networks [3, 16, 46]. As part of the login process, the merchant generates an HTTP request to be sent to the ad network, originating from the user's web browser, which contains the tracking information. Email addresses are potentially hashed, but lack salt and are thus deterministic. This information is sent without integrity or authenticity, *through* the attacker's web browser, allowing the attacker to manipulate it.

Figure 2, steps 11–16 shows entanglement, with step 13 being the critical step that triggers entanglement, as the victim's email address is communicated to the ad network along with the attacker's cookie. Entanglement is performed in one of two ways. Either 1) the attacker intercepts the HTTP request generated by visiting the merchant website (step 13) containing the attacker's own email address, and rewrites the email address to be the victim's address. Or 2) the attacker can find a patsy website that does not verify a user's identity when creating an account. For example, an attacker goes to a website and enters the victim's email address to create an account, and the account is created *without* verifying that the attacker controls the email address. Our results (Section 6.8) show that, surprisingly, not only is it trivial to find such sites, but *most* retail websites demonstrate this behavior. We formulate this lack of verification method as even if some form of integrity over the email field was deployed (which would denote a significant departure from current practices, potentially involving large-scale key management), the vulnerability would still exist as the ad network still cannot verify the address sent to them by the merchants is *correct*.

We emphasize the victim does not need to have any prior relationship to the patsy website, and does not need to engage with or visit the patsy during the attack. The role of the patsy is only to provide an entry point for the attacker's profile into the ad network.

4.2.1 *Entanglement Case Studies*

We now explore the specifics of the entanglement in both the Criteo and Yahoo ad network APIs, denoted by step 13 in Figure 2. In both instances, merchant websites pass identity information directly to the ad network, unauthenticated, without integrity controls. As the merchant website passes this information within the web browser, it is under attacker control.

```
https://sslwidget.criteo.com/event?a=18015&v=5.6.2&p0=e%3Dce%26m%
3D%255Bvictim%252540email.com%255D&p1=e%3Dexd%26site_type%3Dd&p2=
e%3Dvh&p3=e%3Ddis&adce=1&tld=[Merchant]&dtycbr=75575
```

Figure 3: HTTP API request to Criteo from a merchant website. The merchant conveys a user's identity (e.g., email address) to Criteo. We did not identify any integrity checks on this request.

Figure 3 shows a sample API call between the merchant website and the Criteo ad network (with the merchant and email address redacted), relayed through the user's browser. The critical field, highlighted in bolded in blue, is a user's email address. Other fields include boilerplate parameters, the merchant's ID at the ad network, and the merchant URL. As this request is sent in the browser, a user's cookies (including their Criteo ID) are sent along with the request. The attacker is able to manipulate this email field directly, writing the victim's email address into the query which is then sent to the

ad network along with the attacker's cookie, causing entanglement between the victim's identity and the attacker's cookie. We did not identify any integrity checks on the request.

```
https://sp.analytics.yahoo.com/sp.pl?a=1000210819854&d=[Current
Date]&n=4d&b=[Product]&.yp=427149&he=89bead80173b00cdf78015157df3
76ad8b569ecf2d5979ed48213fd723a7f916&f=[Merchant]&enc=UTF-8&yv=1.
12.0&et=custom&ea=ViewProduct&product_id=12489441&tagmgr=gtm
```

Figure 4: HTTP API request to Yahoo from a merchant website. The merchant conveys a user's identity (e.g., email address) to Yahoo. The user's email address is SHA-256 hashed without salt. We did not identify any integrity checks on this request.

Figure 4 shows a sample API call between the merchant website and the *Yahoo* ad network (with the merchant and product redacted), relayed through the user's browser. As before, the user's email address is passed in the bolded blue field along with various boilerplate and merchant information. Unlike before, the user's email address is SHA-256 hashed, without a salt. Also as before, an attacker is able to manipulate this field to cause the victim's identity to become entangled with the attacker's cookie. We did not identify any integrity checks on the request.

4.3 Privacy Vulnerability

After an attacker successfully entangles their ephemeral tracking cookie with the victim's durable identifier, the ad network incorporates the attacker's browser into the constellation of identifiers that make up the victim's profile in their network—identity entanglement is now complete. This is shown in Figure 1. Both attacker and victim browsers will receive personalized ads from the ad network as if the two devices belong to the same user. Similarly, any subsequent activity by either user will contribute to future changes in the ad network profile. The attacker can now 1) observe ads intended for the victim, and/or 2) cause ads to be shown to the victim.

Causing ads to be shown to the victim is straightforward; the attacker visits and interacts with specific items, which will then trigger retargeting of the item to all devices associated with the profile, including the victim's browser. Inferring specific victim activity from the stream of ads shown to the attacker is challenging; ads appear for many reasons, and ads resulting from the victim interacting with merchants or products must be differentiated from all other advertisements. Inferring victim behavior requires extensive methodological development and is discussed in detail in Section 5.

4.4 Attack Scope

Identity entanglement is designed to compromise the privacy of specific individuals, knowing only their email addresses or other identifiers. The primary challenges of scaling the attack to large numbers of users simultaneously falls across two dimensions: resources and impact on the ad ecosystem.

Resources. An attacker seeking to extract private details from a large group of users at-scale would need to create, manage, and orchestrate browser profiles for each victim. Similarly, as the attacker scales the attack, they may encounter automated ad network fraud detection systems designed to prevent large-scale crawling and advertising abuse. Attackers may be able to re-use baseline profiles across victims, but the specifics of baseline reuse are unclear as such behavior may influence the distribution of ads themselves, leading to the next challenge.

Ad Ecosystem Impact. Ad networks operate on budgets and campaigns. An attacker attempting to extract private information at scale may, by virtue of the process of measuring ads, deplete ad budgets and shift the distribution of ads they seek to observe. At small scale, these effects are likely negligible, but as the attacks scale so do the risk of measurement error and behavioral shifts.

We believe both challenges are surmountable with further research. However, for the purpose of this work, we limit ourselves to targeted attacks on specific individuals.

5 BROWSING BEHAVIOR INFERENCE

Once the attacker has entangled their browser's ephemeral tracking cookie with the ad network's profile of the victim, they must then develop a method to extract user browsing activity from the advertisements being shown to them. Conceptually this problem can be broken down into three steps: 1) executing the identity entanglement attack, 2) collecting advertisements served to the attacker's entangled profile, and 3) determining which subset of these advertisements has been retargeted from another browser linked to the same durable identity (instead of being general advertisements). Section 4 provides an overview of the identity entanglement attack. In the remainder of this section, we provide an overview of how an attacker extracts victim browsing behavior from that stream of advertisements served to the attacker after entanglement.

Core to accomplishing step 3 above is a metric we call *normalized difference*, which compares the baseline ad distribution of an unentangled profile with the ads served to an attacker profile which is entangled with the victim profile. The victim profile, in this case, has performed the behavior that leads to retargeted ads—adding items to the shopping cart on various merchant sites.

5.1 Measurement Website (Publisher Page)

Once an attacker's profile has been entangled, the attacker must view ads that have been customized according to the victim's durable identity. To do so, the attacker collects advertisements using entanglement and baseline profiles (discussed subsequently) at a publisher page (i.e., a webpage that shows ads). We denote the specific publisher page an attacker uses as the *measurement website*. An appropriate measurement website must: 1) show display ads (including retargeted ads) from the ad network the attacker wishes to entangle, 2) not itself be a merchant website, otherwise the act of measurement could influence the composition of ads shown in the victim profile, 3) show ads without needing to create an account. Many news websites meet these requirements, including yahoo.com, which we use as our measurement website.

5.2 Browser Profile Setup and Baselines

Advertising campaigns and algorithms constantly change. To account for this, we introduce the notion of attacker created "baseline" profiles that capture general ad campaigns that are not reflective of victim behavior. The attacker records ad activity via the measurement website in parallel for the attacker and baseline profiles.

The attacker creates the various browser profiles: baseline and attack. In addition, we create a simulated victim profile for our experiments. For all experimentation, each profile is created with separation between attacker, victim, and baseline profiles on different machines and originate from different IP addresses, all starting from new browser profiles.

Attack Profiles. The attacker creates an **attack profile** and then performs the identity entanglement attack by registering an account at the patsy website and then either editing the HTTP requests replacing their own email address with the victim's, or by using the victim's email address if the merchant doesn't verify email addresses. After the attacker logs into the patsy website, they visit the landing page to ensure the email address is reported to the targeted ad network.

Baseline Profiles. The attacker also creates two baseline profiles and primes them with tracking cookies for the target ad network by visiting the patsy website. (We assume that using the patsy website in this way precludes identifying victim activity related to the patsy website, discussed further in Section 7.2.) The attacker creates two profiles for redundancy in understanding the distribution of general advertisements, which is especially helpful in identifying if the entanglement was successful (see Section 6.1). In addition, to control for ad customization that occurs as a result of the creation of an account at the patsy website, the attacker also creates two account baseline profiles that create accounts using attacker controlled email addresses but they do not perform identity entanglement.

Victim Profiles. A necessary prerequisite for understanding the effectiveness of our attack is that an attacker is able to infer the behavior of a victim. Ethically we cannot perform this attack against a real user. We are also ethically precluded from performing this attack against ourselves, as the identity entanglement attack could reveal detailed personal information about the researchers to each other. We therefore simulate victim profiles via several mechanisms. First, we create "basic" profiles by registering a free email account at a large webmail provider which is then used to create an account at a target merchant website. We then prime each profile by adding randomly selected products to the shopping cart.

Profile Locations. The *attacker-controlled* profiles (i.e., attack and baseline) should be geographically near each other. This is necessary to ensure that the attacker's baseline and attack profiles will receive roughly the same set of advertisements, so that the baseline profiles can identify general ads.

5.3 Normalized Difference

After the attacker has collected the distribution of observed ads on the measurement webpage across all of the profiles, they need a mechanism to extract unrelated ad campaigns and ad targeting from the victim's actual merchant activity. To achieve this goal, we introduce the notion of *Normalized Difference*. Normalized Difference attempts to understand the distribution of advertisers (or ads) across the baseline profiles and then subtract those occurrences from the entangled attacker profile's ad distributions. Equation 1 defines the formula we use for Normalized Difference (per advertiser, or per specific item, depending on context).

$$\text{Normalized Difference} = \frac{max(A - B, 0)}{\text{M}} \tag{1}$$

A := attacker's observed count, B := baseline observed count
M := the max value in attacker's observed counts.

Normalized difference assumes background advertising noise derived from advertising campaigns that appear to the attacker profile will also appear to unrelated baseline profiles.

5.4 Experimental Implementation

To understand the effectiveness of Identity Entanglement, we create the previously described set of attacker and victim profiles. Once we have created this collection of browser profiles, we then use them to extract information from the victim's advertising profile by viewing displayed advertisements.

After executing the entanglement attack and identifying a measurement website, an attacker repeatedly loads the measurement website at regular intervals across all profiles (besides the victim profile) and records all advertisements. For each visit across each profile, we record timestamps and full HTTP requests and responses associated with all advertisements. Merchant and item identifier information is embedded in the recorded HTTP request parameters.

We perform these measurements across 6 browser profiles for three days by leveraging the `Puppeteer` browser automation library to load the measurement website and record all necessary page information. The system is built on top of Puppeteer 10.4.0 executing `Google Chrome` 94.0.4606.54 and is run concurrently across different physical systems. We continually reload the measurement website every 30 seconds. The total quantity of ads served varies from one profile to another; to address this imbalance, we normalize all distribution comparisons rather than give raw counts.

While we conduct the priming of the victim profile with synthetic shopping behavior separately from the measurement using entangled attacker profiles and analysis, an attacker could potentially perform such analysis in real-time by continually investigating the ad distribution received by the profiles.

5.5 Implementation Specifics

For our experiments and crawling we used four different systems from different IP addresses, with attacker and victim profiles split across different networks (although within the same US city) to limit the possibility of non-entanglement fingerprinting influencing our results. The victim device is an Apple MacBook Pro laptop, with an Intel Core i7-4960HQ and 16 GB of RAM. The attacker device is an AMD Opteron Processor 6328 server running Ubuntu 20.04 with 126GB of RAM. Baseline profiles are run from two different machines, an Intel Xeon Gold 6140 server running Ubuntu 18.04 with 126 GB memory and Intel Core 286 i7-8565U desktop running Fedora 34 with 16 GB of memory, respectively. All profiles and tests are run on "bare metal" with no use of virtualization or cloud services. For each experiment we created and used a brand new browser profile. Simulated victim and attacker machines used IP addresses were located in the same US city, but on different networks.

To identify merchants using Criteo for entanglement experimentation, we crawled the Alexa top 10,000 using our non-victim machines, looking for utilization of the Criteo retargeting API in the browser development tool logs. From this set we then selected merchants at random for our activity identification experiments. We note this method produces a lower bound on possible merchants using Criteo, as merchants may not utilize Criteo at all times, on all pages, or may employ cloaking that simple crawling does not evade. Such limitations produce a lower bound and are acceptable for identifying a set of merchants to explore, but do not accurately convey the full scope of Criteo. To better understand Criteo's scope we rely on whotracks.me, discussed previously and further in Section 6.8.

6 RESULTS

To demonstrate the feasibility of the identity entanglement attack, we conducted a series of "blind" (i.e., the attacker has no prior knowledge of the victim's behavior or ads) experiments leveraging the attack (Section 4) and behavior extraction technique (Section 5).

We are able to extract browsing behaviors of a victim profile through retargeted ads placed via a vulnerable third-party ad network, **with no information about the victim beyond their email address.** To explore the effectiveness of this attack, we perform experiments to answer a series of research questions (RQs):

- RQ1: Can attackers confirm successful identity entanglement?
- RQ2: Can attackers identify what merchants a victim visits?
- RQ3: Can attackers identify the items a victim interacts with?
- RQ4: How quickly can the attack succeed?
- RQ5: Can attackers influence what ads the victim sees?
- RQ6: Can attackers entangle solely rewriting requests?
- RQ7: Are other ad networks vulnerable to identity entanglement?
- RQ8: What is the potential scope of the problem across websites?

For RQs 1-5 we utilize the merchant websites not verifying the victim's email address entanglement method. In RQ6, we refine the attack to rewriting HTTP requests within the attacker's browser (demonstrating a lack of integrity over the transmitted information, and that the information is not transmitted via other means). We note that not verifying email addresses is a more challenging problem to solve than adding communication integrity or communicating out of band, as it involves restructuring the incentives of the ad ecosystem, requires perfect security across all merchants, and cannot survive a malicious merchant. For RQs 1-6, we focus on Criteo, and in RQ7 we extend email injection to Yahoo Analytics.

We find identity entanglement is feasible, victim browsing activity can be extracted at website and item granularity, and the problem exists across ad networks. While we investigate these RQs largely independently and from an empirical perspective, a hypothetical attacker with knowledge of how various targeting mechanisms work will likely have a much lower threshold for being confident in the success of their attack, and more tolerance for noise.

6.1 RQ1: Can attackers confirm successful identity entanglement?

Step one in understanding attack feasibility is identifying if attacker engagement with the patsy website is sufficient for successful identity entanglement. Our hypothesis is that performing entanglement results in a significant change in the distribution of ads—the number of advertisers and the distribution of ads across those advertisers—as the ad network begins retargeting the victim's ads to the attacker. To explore this, we compare the distribution of ads seen from Criteo across the five profiles: two unrelated unused "baseline" profiles, two unrelated unused "account baseline" profiles (Section 5.2), and the attacker's attack profile. All profiles are created and mechanistically measured in the same way, with separation between attacker, victim, and baselines on different machines with different IPs.

We explore RQ1 via the results in Figure 5 and Figure 6. These figures show a cumulative distribution function plot of the observed occurrence of advertisers and individual product ads respectively across all measurement profiles in our dataset.

Figure 5 explores the distribution of advertisers seen in our dataset. Advertisers correspond directly to merchant websites. Our

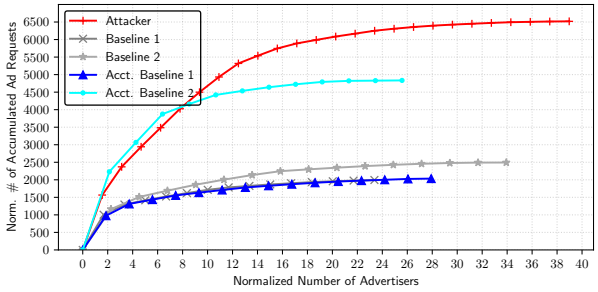

Figure 5: The accumulated appearance of advertisers across the baseline and attacker profiles, post entanglement, normalized per thousand views. We observe varied distributions across baseline profiles corresponding to non-deterministic ad network behavior. We note that the entangled attacker profile has significantly more advertisers and ad appearances, corresponding to retargeting of the victim.

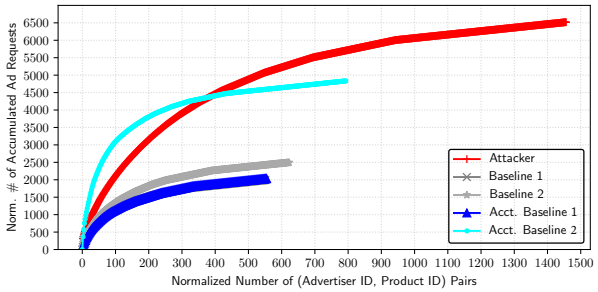

Figure 6: The accumulated appearance of specific item ads across the baseline and attacker profiles, post entanglement, normalized per thousand views. Similar to Figure 5, we observed varied ad distributions across the baseline profiles and a significant difference in the number of individual item ads seen by the entangled attacker profile. However, this difference is much more pronounced, with the attacker profile receiving nearly twice as many different ads than any baseline profile. This change corresponds to a significant increase in retargeted ads related to the victim.

hypothesis is a baseline attacker profile successfully entangled with a victim profile would have a substantially different distribution than the baseline profiles, as the ad network has specific retargeting information which can be leveraged to generate different ads.

We find that while three of the baseline profiles have similar distributions, one baseline profile ("Account Baseline 2") has a substantially different distribution. Manual investigation reveals this distribution difference is due almost entirely to a single merchant running an aggressive marketing campaign directed at only this profile. Such non-determinism is expected, validating our strategy of deploying numerous baseline profiles and the use of normalization.

While "Account Baseline 2" showed a very different distribution than the other baseline profiles, the attacker entangled profile showed both significantly more individual advertisers and more individual ad requests. We note Criteo is one of several ad networks utilized by the measurement website. As a result, ad networks compete via bidding for the right to show ads, and the ad network we measured will not necessarily win all of the placements. For this reason, it's possible that different profiles may receive different numbers of ads from the ad network even if each profile performed the same number of page loads [33].

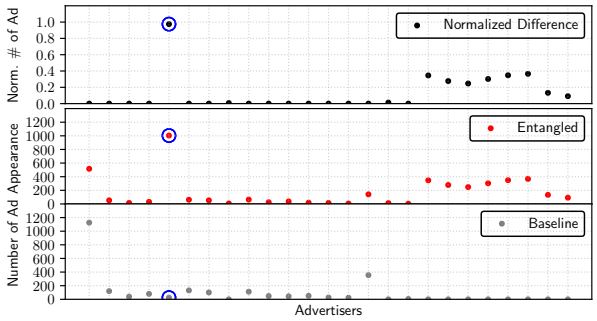

Figure 7: Count of observed ads per advertiser for all advertisers. Advertisers (x-axis) are ordered by appearance. The normalized difference is the normalization of attacker counts considering what the baseline profiles observe as background ad campaigns. The advertiser data points circled represent the ground-truth victim activity. Seven of the eight advertisers that appear on the right attacker graph are thematically similar to the target merchant, indicating targeting.

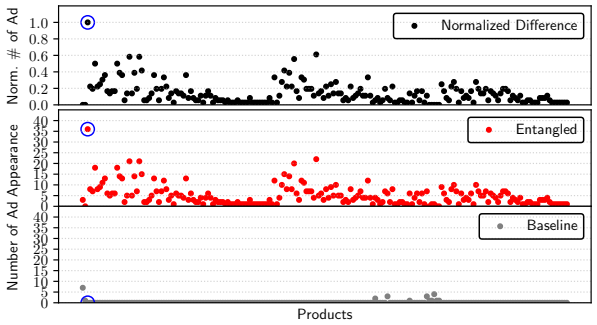

Figure 8: Count of observed ads per item for the previously discovered website. Items (x-axis) are ordered by appearance. The normalized difference is the normalization of attacker counts considering what the baseline profiles observe as background ad campaigns. The specific item the victim interacted with is circled in blue. We observe the victim-interacted item is an outlier in the dataset.

Based on these observations, we explore the distribution of ads for specific items (instead of the distribution of advertisers) in Figure 6. This CDF shows the distribution of ads for specific items over all of the baseline profiles. We find not only does the attacker entangled profile observe 35% more ad requests than "Account Baseline 2", it receives 203% more requests than the other three baselines on average. A stronger signal exists in the distribution of the number of specific items being advertised, with the entangled profile seeing 84% more items than "Account Baseline 2" and 154% more items than the other three baselines on average. This significant change in distribution is due to a large volume of retargeted products from the victim's profile at the target merchant.

Across both experiments, we observe a substantial shift in the distribution of advertisers and ads upon the execution of the entanglement attack. From this analysis, we demonstrate not only that identity entanglement is possible, but we can also confirm when successful identity entanglement has occurred. It is also important to note that, because the attacker has control over when they perform the entanglement of their profile, they do not need to detect when the distribution changes, but can rather perform the entanglement at a controlled moment in time and then inspect the subsequent distributions between the various profiles.

6.2 RQ2: Can attackers identify what merchants a victim visits?

After confirming that the identity entanglement attack is successful (RQ1), our next goal is to understand whether we can effectively identify which merchant websites the victim visits. We hypothesize that because retargeted advertisements are so lucrative, ad networks will aggressively retarget the victim's merchant activity across the ad network. As the attacker's account is entangled with that of the victim, the attacker should also observe that activity. Further, we hypothesize that the merchant website (or websites) the victim visits should represent an outlier in the observed advertisements.

Figure 7 shows the distribution of all ad appearances for each merchant in the entangled and baseline profiles. The top-most plot is the computed normalized difference (Equation 1) between the baseline and the attacker's entangled profile. The advertiser circled in blue is the merchant the victim interacted with in our experiment.

The website the victim visits is a clear outlier in the observed dataset, receiving 95% more advertisements than any other website in the dataset. When utilizing the normalized difference, that result becomes more pronounced, with the target website receiving 167% more ads in the normalized space. The x-axis on the figures is ordered by the occurrence of each advertiser. The cluster of advertisers that appear later are the result of retargeting toward the victim, and are thematically similar to the target website.

From this substantial outlier observation, we conclude that we are able to identify the website visited by the victim, knowing only their email address, and with no prior knowledge of their browsing activity. We explore the reliability of these observations further with RQ5. We next explore whether we can identify specific products viewed by the victim at a given merchant that has been identified using our approach.

6.3 RQ3: Can attackers identify the items a victim interacts with?

After identifying the merchant the victim interacted with, we aim to identify the specific items the victim has interacted with. We hypothesize that the specific items the victim interacted with will be aggressively advertised compared to other items from that merchant (regardless of any overall shift in items shown).

Figure 8 shows the distribution of ads shown for different items on the target website discovered via RQ2. In this figure, we compare the distribution of ads shown to the entangled attacker profile and the baseline profiles. We note that the baseline profiles received a small but positive number of advertisements for the merchant the victim interacted with. The top-most portion of the figure shows the normalized difference between the profile sets. The specific item the victim interacted with is a significant outlier in the dataset, having been served 64% more than any other item in the dataset.

This analysis allows us to determine the specific item that the victim interacts with on a specific merchant website. This represents a significant privacy concern, as **the attacker can extract specific user browsing activity knowing only the email of the victim.**

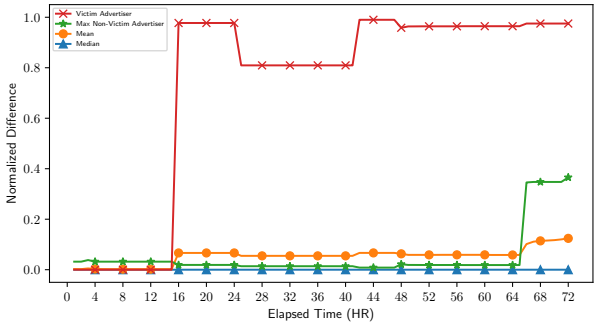

Figure 9: Normalized difference of attacker-observed merchants over time. This is a time-based view of Figure 7. The attacker executes the attack at time 0, and entanglement becomes visible roughly 15 hours later, as denoted by the set of merchants observed in the attacker's profile shifting significantly. Once the attacker profile deviates from the baseline at all, the victim's merchant dominates the dataset.

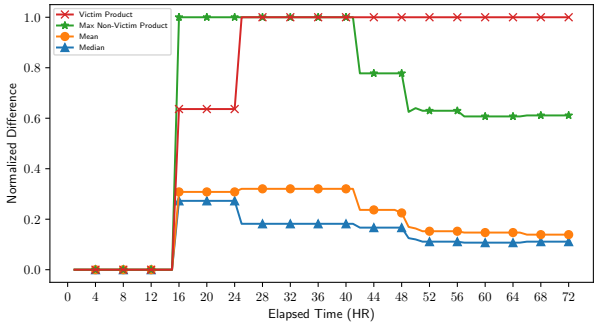

Figure 10: Normalized difference of attacker-observed products over time. This is a time-based view of Figure 8. The attacker executes the attack at time 0, and entanglement becomes visible roughly 15 hours later, as denoted by the set of products observed in the attacker's profile shifting significantly.

6.4 RQ4: How quickly can the attack succeed?

While RQ2 and RQ3 sought to explore *if* we can identify merchants and items that victims interact with, in RQ4 we seek to understand *how quickly* attackers can identify victim behavior. To explore this question we perform a time-based analysis of the same dataset explored in Figures 7 and 8.

For this analysis, we compute a rolling normalized difference over time. At one hour increments we snapshot all ads received across the attacker and baseline profiles, computing the normalized difference on all data observed up until that moment in time. We then plot the maximum, mean, median, and victim merchant/product normalized difference. We note that the attacker has no knowledge of the victim's merchant/product; we denote that line given our knowledge of the simulated victim.

Figure 9 shows the normalized difference over time of the *merchants* in our experiment from RQ2 (Figure 7). The merchant the victim interacted with dominates the dataset throughout the majority of the experiment, and rises as soon as the attacker profile begins seeing any advertisements that do not exist in our baseline profiles. Victim activity begins showing up roughly 15 hours into the experiment, with the attacker executing the attack at time 0. We note such behavior is common, with there being some delay between attack and changes in attacker profile. We speculate this delay corresponds to advertising campaigns, ad network back-end processes, or other opaque phenomena within the complex ecosystem of the ad network. While there is some delay, the dominance of the merchant as soon as the attacker profile deviates from baseline can be used as a signal of entanglement. Towards the end of the experiment, possibly as a result of non-engagement, other merchants begin being marketed to our attacker profile, which causes the uptick in another merchant ("Max Non-Victim Advertiser").

Figure 10 shows the normalized difference over time of the *products* in our experiment from RQ3 (Figure 8). As with Figure 9 entanglement occurs and the attacker's profile shifts after 15 hours. The victim's product is an outlier from other products immediately upon entanglement, and is the most prevalent product for a majority of the time post-entanglement. Of note is the victim's product not having the highest normalized difference early on in the experiment (but still being an outlier). Investigation reveals the

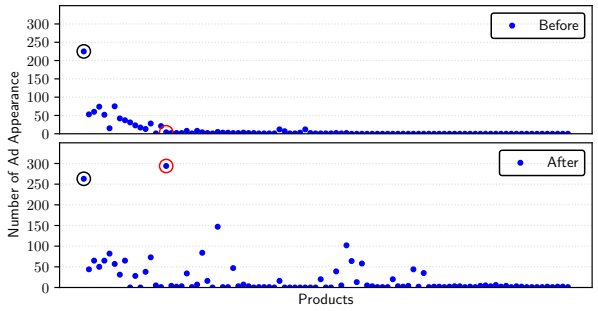

Figure 11: Count of observed ads per item seen by the *victim*. The top plot shows the distribution of ads before the attack, and the bottom plot shows the distribution of ads after. An item the victim interacted with is circled black, and the attacker-placed item circled red. We find that the attacker can influence ads shown to the victim.

items are extremely similar, are from the same category, and have similar properties and physical appearances. We speculate that this is remarketing from the advertiser attempting to sell similar items due to the victim's engagement, which conveys victim behavior albeit not the precise item. We do however note that over time the specific item the victim interacted with dominates the dataset.

6.5 RQ5: Can attackers influence what ads the victim sees?

We now explore if entanglement is bi-directional, i.e., we seek to determine if the attacker is able to influence what ads are shown to the victim after entanglement. Our hypothesis is that the influence is bi-directional and the attacker can influence what ads are shown to the victim. To answer this question, we designed a separate experiment that begins with the victim and attacker methodology from RQ3. After the attacker successfully entangles their profile with the victim's, the attacker then identifies a merchant the victim interacts with. Next, the attacker goes to that merchant and adds an item to their shopping cart. We note that the attacker *does not* have access to the victim's account at the merchant, and the shopping carts between the attacker and victim are separate.

Figure 11 shows an overview of the distribution of ads shown to the *victim* before and after the identity entanglement and product placement attack. We find that immediately after the attacker

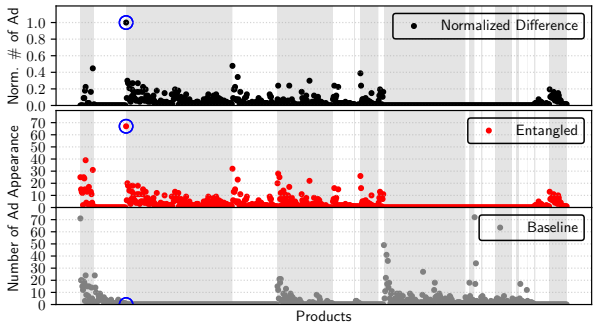

Figure 12: Count of observed ads per item for all advertisers. Items (x-axis) are grouped by the same advertiser. Different shades separate adjacent advertisers. We rewrite HTTP requests from the merchant to the ad network, replacing the attacker's email address with the victim's. The significant differences in ad volume allow the attacker to identify both the merchant and specific item from the victim.

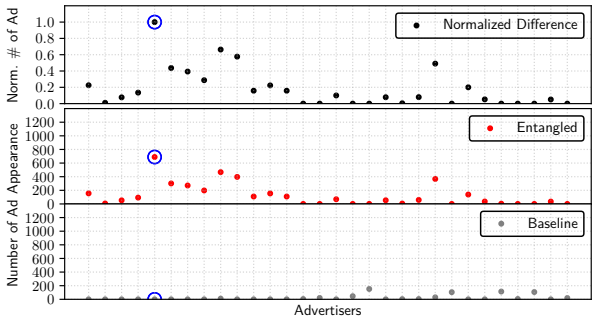

Figure 13: Count of observed ads per advertiser for all advertisers from Yahoo-driven Entanglement. The normalized difference is the normalization of attacker counts considering what the baseline profiles observe as background ad campaigns. The advertiser data points circled represent the ground-truth victim activity. The victim's activity is an outlier, 35% higher than any other advertiser.

interacts with an item at the merchant website, the victim begins receiving retargeted ads for that item. This result is extremely concerning, as it allows an attacker to control what is shown to the victim, allowing for embarrassment and blackmail attacks.

6.6 RQ6: Can attackers entangle solely rewriting requests?

Thus far, our exploration of entanglement relied upon a patsy website not verifying a user's email address before accepting it. This is the stronger of the two attack scenarios, as it controls for technical deficiencies such as integrity over communication between the merchant and ad network (either cryptographically or via out-of-band communication), but cannot solve the fundamental challenge of outsourcing identity verification (poorly implemented or malicious merchants). We now explore if an attacker can accomplish the same attack solely by editing HTTP requests that flow through their browser, replacing their email address with the victim's.

To explore this question, we conduct a new full identity entanglement experiment against a victim profile where we take an existing query from a merchant website and rewrite it such that that the user information associated with the query is the victim's (hashed) email address. This mimics the same API behavior generated by retail websites sent through the user's browser, as if the victim had logged into the website using their email.

Figure 12 explores the results, which bridge both RQ2 and RQ3 by simultaneously attempting to identify both the website the victim visited and their specific product interactions. We observe similar outlier behavior as prior experiments, and are able to identify the specific product the user interacted with, without the need of a patsy website that does not verify email addresses. We conclude that direct manipulation of the requests between the merchant website and ad network is possible, and sufficient to cause entanglement. This result again points to the severity of the attack, as the ad network is reliant not only on retail websites, but also on data sent through a user's browser with no integrity checks. We emphasize we did not identify any attempts by the ad networks to implement integrity checks over these API calls.

6.7 RQ7: Are other ad networks vulnerable to identity entanglement?

In the prior experiments, we explored Criteo, the largest retargeting third-party ad network. In this experiment we explore a different ad network, Yahoo Ads. More specifically, we focus on the Yahoo Ads Analytics interface, which has a very similar behavioral pattern to Criteo. When users log into merchant websites, an HTTP request to Yahoo's Analytics API relays the user's email address and browsing activity. Surprisingly, we find that Criteo is a dominant player within Yahoo's Supply-Side business [44], with Yahoo-driven entanglement leading to significant retargeting from Criteo.

For this experiment, we conduct the same methodology as in RQ2, but instead of rewriting HTTP requests to Criteo, we rewrite requests *only* to Yahoo's Analytics API, replacing the attacker's (hashed) email address with the victims. When performing the entanglement, we block all traffic to Criteo in order to demonstrate the passing of unverified email addresses that can be manipulated by an attacker is driven by interactions with Yahoo, not Criteo.

Figure 13 shows the results of our experiment. The merchant that the victim interacts with gets 35% more ads than the next most prevalent merchant, a significant outlier. From this experiment, we conclude that Yahoo's Ad service is passing unverified email addresses to their advertisers leading to entanglement, although the specifics of how information is exchanged within the complex ecosystem of advertising bidding remain opaque. We anecdotally note that numerous other ad networks delivered ads for the merchant website the victim interacted with, via Yahoo, post entanglement, again pointing to a broader ecosystem problem.

6.8 RQ8: What is the potential scope of the problem across websites?

To explore the potential scope of the attack both across the Internet, and in terms of severity, we now focus on three questions. 1) How many websites utilize the two ad networks we explore? 2) What are some specific examples of activity that is exposed by this attack?, and 3) How common is a lack of merchant email verification?

How many websites utilize the two ad networks we explore. To understand the potential reach of the vulnerability, we turn

whotracks.me [26], a longitudinal and continuously updated dataset of trackers across the Internet. whotracks.me collects data from the popular Ghostery [23] extension, and reports a comprehensive dataset of websites and their embedded trackers. We identify the trackers relating to Criteo and Yahoo's Analytics APIs, and compare those URLs with the whotracks.me dataset.

Using the March 2022 whotracks.me snapshot of the top 6,000 global sites that embed tracking technology, we find that 31% of sites utilize Criteo's tracking, 19% utilize Yahoo's, and 35% utilize either of the two. Given that the 6,000 sites documented by whotracks.me contain *any type* of tracking, not just advertising-related tracking, we expect 35% to be a lower bound on the advertising impact.

Specific activity examples. The list of merchant websites we have extracted the victim's retargeted ads from includes known retailers such as newegg.com, macys.com, zappos.com, and groupon.com. Perhaps most concerning, we find priceline.com and apartments.com were susceptible to this attack. The retargeted ads from these websites included specific details and links to properties that could allow an attacker to track information about the locations the victim is interested in. **The ads themselves show images of the apartment or hotel**. If the attacker clicks on the ad, they are taken to a landing page at the respective website containing the precise properties the victim viewed along with the search parameters (e.g., number of occupants, number of rooms) of their query. This finding highlights the severity of this vulnerability; we expect further exploration will result in other sensitive privacy leaks.

How Common is lack of email verification. To understand the prevalence of merchants not verifying email addresses, we crawled the Alexa Top 10,000, and identified 334 sites utilizing Criteo retargeting services. For each site, we manually created an account using a non-existent email address. We found that 84% (279) of all potential websites **did not** verify a user's email address before logging them in, indicating merchant websites that do not verify email addresses are both prevalent and easy to find. We note that we consider our measure of the number of sites using the Criteo retargeting API a lower bound; merchants may not utilize Criteo at all times, on all pages, or may employ cloaking to evade crawling.

7 DISCUSSION

Our results show issues stemming from third-party ad networks' use of unverified identity information from merchants. We demonstrated the potential for such attacks to substantially invade user privacy, motivating the need for systematic exploration of: the problem at other ad networks, the full scope of the vulnerability, and potential mitigations. Our experiments reveal significant leakage of private user information, knowing only a victim's email address. In some experiments, not all victim browsing activity could be recovered. While we expect additional work could further identify specific user browsing activity, we note that **any** leakage of user browsing behavior to an external attacker knowing only an email address **should be impossible**, and our findings point to a significant and impactful real-world vulnerability needing remediation.

Further investigation to encompass more complex victim behaviors, aimed at extracting more complete and detailed browsing profiles in the presence of more significant noise, across different platforms (e.g., mobile web and native apps), and with different durable identifiers (i.e., mobile identifiers), is important. Unfortunately, to the best of our knowledge, Criteo has not deployed any mitigation in the ensuing months since they acknowledged the attack. This highlights the misaligned incentives and the challenges of convincing ad networks to deploy mitigations.

7.1 Mitigations

Another future work component is exploring mitigations within the scope of existing system designs and assumptions. We break these down across three mitigation axes: client, server, and ecosystem.

7.1.1 Client Mitigations

Blocking third-party cookies. Industry leaders [7, 21] and browser vendors [8, 13, 43] have announced plans to limit or eliminate third-party cookies. However, given the revenue generated from targeted advertising by ad platforms and merchants, it seems likely alternative tracking forms will emerge. For example, rather than passing tracking information via tracking cookies within the user's web browser, merchant websites could pass user activity and identifiers (i.e., hashed email addresses) via direct communication to ad network APIs outside of the browser. Such a change would both allow identity entanglement attacks to still occur, and also worsen the landscape of user privacy, since the flow of information would not be visible or blockable by the user directly within their browser.

Email Address Aliases. Another mitigation strategy is for users to use distinct email addresses when creating an account with merchants in addition to blocking third-party cookies. There are many services that provide "disposable" email addresses that could be used for registering accounts [42]. However, this solution poses a usability challenge, requiring users to remember the email address used to register with each merchant similar to managing distinct passwords [36]. Apple has created a system for generating and managing unique email addresses but it is limited to their operating systems [25]. Future work could entail creating an open system for creating and managing unique email addresses that can be used for merchant account registration. There is also the challenge that merchants and ad networks might start using identifiers that are more difficult to compartmentalize, such as phone numbers, or increase their use of browser fingerprinting.

7.1.2 Server Mitigations

The core issue revealed by our identity entanglement attack is that ad networks depend on merchant websites to properly verify user identity for cross-device tracking. When a merchant website does not verify the users' identities or does not provide integrity over that data, then the ad network becomes vulnerable to our attack. Possible mitigations revolve around controlling verification, or trusting retailers and then security communication of that identity.

Identify Verification. Ad networks could require that all account information ultimately shared with them first be vetted through verification steps performed directly by the ad network, e.g., an email ownership verification step that is hosted by the ad network. Such a measure would be a radical departure from current relationship dynamics, and also would represent merchant websites giving up control of their account creation and verification process (unlikely), or requiring additional friction for account creation.

Trusted Retailers + Cryptographic Solutions. If ad networks could trust retailers to properly verify user identity via non-technical

means (e.g., audits, policies, incentives), the problem of solving identity entanglement then centers around securely communicating user identity information from retailer to ad network, via the user's web browser. This could be done with traditional cryptographic integrity mechanisms, if ad networks established cryptographic keys with retail websites. To our knowledge no such cryptographic measures are deployed in the ad ecosystem to secure data exchanged inside users browsers. So while a departure from current trends, we are not aware of any *technical* barriers to deployment.

Trusted Retailers + Out-of-Band Mechanisms. Similar to the aforementioned cryptographic solutions, in an ecosystem of trusted retailers, ad network could relay user identity information outside of the user's web browser, e.g., via direct communication between the retailers and ad network, thus removing identity information from the attacker browser. While technically possible, we note such defenses would be a significant departure from existing ad network norms and procedures.

7.1.3 *Ecosystem Mitigations*

Due to misaligned incentives, it is challenging to convince ad networks to deploy mitigations. There is also the potential for an arms race between people attempting to protect their privacy and tracking companies. This suggests the need for regulation to limit the ability of companies to collect information and track people. Future work crafting such legislation should take into account that information collected by tracking companies will likely be vulnerable to attacks similar to ours. Thus regulation should limit the data that can be collected about users and used for targeting.

7.2 Limitations

While we were able to successfully build an inference mechanism in many situations, our attack has limitations in specific scenarios. First, the effectiveness of our attack relies upon sufficient successful bids for retargeted ads, which itself relies upon the non-deterministic nature of advertisers' marketing strategy, budget, and competing bids. Second, a victim who places several items in their shopping cart could degrade the accuracy of our inference mechanism by splitting the merchant's retargeting budget amongst several products, thus driving down the normalized difference between retargeted items and all other items. Third, our measurement (not the attack itself) relies upon comparing ad display counts across multiple browser profiles. Unfortunately, this approach can be distorted by intermittent ad campaigns (heavily promoted items during e.g., Black Friday, etc.), or ad campaigns that randomly target a certain cohort (e.g., an A-B test). Additionally, we did not exhaustively explore the role of geolocation on the process of entanglement, but we have no evidence to support geolocation playing a role.

8 RELATED WORK

Papadopoulos et al. revealed how cookie synchronization enables third-parties to track a victim across websites by exchanging identifiers via HTTP cookies [35]. Chen et al. traced how a user identifier that third-parties set in a first-party cookie flowed to the third-parties through a JavaScript taint analysis [12]. Sanchez-Rola et al. explored the cookie ecosystem in detail, focusing on the roles of advertising entities and the relationship among them [38].

Researchers have focused on measuring and disrupting online tracking. XRay [29] is a personal data tracking system that aims to improve transparency in user data collection and when it is used for ad targeting. It diagnoses who has what data by comparing input (i.e., personal data) and output (i.e., personalized services). Sunlight [30] is able to reveal the causation of personalized web services including retargeting ads with a modular diagnosis system based on statistics and machine learning. Our attack methodology of normalization to reduce noise takes inspiration from these works. Bashir et al. studied information flows between ad exchanges by leveraging retargeted ads and categorized them based on several matching rules [6]. whotracks.me [26] is a longitudinal open-source database of online trackers and the websites that embed them, collected via telemetry collected from users of the Ghostery browser extension [23] across the Internet. Our work introduces how an attacker could weaponize the retargeted ads to extract private information from a specific ad network beyond analyzing the information flow or trackers in the ad network, in part, utilizing whotracks.me data to understand the scope of the problem.

Cross-device tracking has become increasingly studied by researchers with its deployment by most ad networks. Brookman et al. discovered personal information (e.g., email, username) played a crucial role in correlating different devices, and shared with third parties [10]. Zimmeck et al. discussed a probabilistic device correlation technique based on machine learning with IP address and browsing history [47]. Solomos et al. developed Talon [40], showing it could detect cross-device tracking with a data-driven methodology. They found cross-device trackers by observing if a retargeting ad triggered with a certain behavior on one device is delivered to another device. This work motivated us to evaluate the security of cross-device tracking implementations. Our work is distinct from this prior work in that we focus on evaluating the security of the underlying third-party cross-device tracking techniques and how they might be exploited in combination with retargeted ads. We show that insecurities in cross-device tracking can be exploited to exfiltrate sensitive user information from ad networks.

9 CONCLUSION

We identified a fundamental vulnerability in how the largest third-party ad networks perform cross-device tracking. Our analysis of retargeted ads served to an attacker after performing this attack demonstrates that there is data leakage occurring which could allow an attacker to determine merchant websites and products viewed by a victim, as well as control what ads are shown to the victim. We also present several defensive directions and associated tensions. This finding highlights yet another privacy implication of online tracking, which is currently a hotly contested topic in the popular conversation. It is unclear if a single "effective" solution exists. Our findings demonstrate the need for further security analysis of ad networks to identify additional vulnerabilities that could lead to exfiltrating users' private information. We also stress that this is not a purely technical issue and that the economic incentives must be taken into consideration when assessing the feasibility of solutions.

10 ACKNOWLEDGEMENTS

The authors thank Masood Ali and the anonymous reviewers whose thoughtful feedback and engagement during rebuttals improved the work significantly. This work was supported in part by funding from ONR award N00014-18-1-2662 and NSF CNS award 2151837.

REFERENCES

[1] AdRoll. 2015. STATE of the INDUSTRY: A close look at retargeting and the programmatic marketer. https://www.iab.com/wp-content/uploads/2015/07/US_AdRoll_State_of_the_Industry.pdf.

[2] AdRoll. 2016. Demystifying Cross-Device Marketing. https://pages.adroll.com/rs/964-WFU-818/images/Collision_Adam_Berke_Marketing_Stage.pdf.

[3] Google Ads. 2022. Tag your website for dynamic remarketing. https://support.google.com/google-ads/answer/3103357?hl=en&ref_topic=10070359#.

[4] AWS. 2019. Identity Graphs on AWS. https://aws.amazon.com/neptune/identity-graphs-on-aws/.

[5] Paul Barford, Igor Canadi, Darja Krushevskaja, Qiang Ma, and S. Muthukrishnan. 2014. Adscape: Harvesting and Analyzing Online Display Ads. In *Proceedings of the 23rd International Conference on World Wide Web* (Seoul, Korea) *(WWW '14)*. Association for Computing Machinery, New York, NY, USA, 597–608. https://doi.org/10.1145/2566486.2567992

[6] Muhammad Ahmad Bashir, Sajjad Arshad, William Robertson, and Christo Wilson. 2016. Tracing Information Flows Between Ad Exchanges Using Retargeted Ads. In *Proceedings of the 25th USENIX Security Symposium (Security)*. Austin, TX.

[7] Chetna Bindra. 2021. Building a privacy-first future for web advertising. https://blog.google/products/ads-commerce/2021-01-privacy-sandbox.

[8] Dieter Bohn. 2021. Google delays blocking third-party cookies in Chrome until 2023. https://www.theverge.com/2021/6/24/22547339/google-chrome-cookiepocalypse-delayed-2023.

[9] Adina Bresge. 2018. Online ads spoil Christmas surprises, raising privacy concerns. https://www.cbc.ca/news/science/online-ads-christmas-spoilers-1.4942461.

[10] Justin Brookman, Phoebe Rouge, Aaron Alva, and Christina Yeung. 2017. Cross-Device Tracking: Measurement and Disclosures. *Proc. Priv. Enhancing Technol.* 2017, 2 (2017), 133–148.

[11] Juan Miguel Carrascosa, Jakub Mikians, Ruben Cuevas, Vijay Erramilli, and Nikolaos Laoutaris. 2015. I Always Feel like Somebody's Watching Me: Measuring Online Behavioural Advertising. In *Proceedings of the 11th ACM Conference on Emerging Networking Experiments and Technologies* (Heidelberg, Germany) *(CoNEXT '15)*. Association for Computing Machinery, New York, NY, USA, Article 13, 13 pages. https://doi.org/10.1145/2716281.2836098

[12] Quan Chen, Panagiotis Ilia, Michalis Polychronakis, and Alexandros Kapravelos. 2021. Cookie Swap Party: Abusing First-Party Cookies for Web Tracking. In *Proceedings of the 30th International World Wide Web Conference (WWW)*. Virtual Event.

[13] Catalin Cimpanu. 2020. Apple blocks third-party cookies in Safari. https://www.zdnet.com/article/apple-blocks-third-party-cookies-in-safari/.

[14] Eliza Crawford. 2020. Website Tracking: Why and How Do Websites Track You? https://www.cookiepro.com/blog/website-tracking.

[15] Criteo. 2018. Criteo Ranked Number One in AdTech Worldwide Market Share According to Leading Analyst Firm Report. https://www.criteo.com/news/press-releases/2018/09/criteo-ranked-number-one-in-adtech-worldwide-market-share/.

[16] Criteo. 2018. OneTag for CSP. https://www.criteo.com/wp-content/uploads/2018/09/CSPOneTag_v1.1.pdf.

[17] Criteo. 2020. Criteo Ad Tech Explained - Shopper Graph. https://youtu.be/s3UVXOmCtmg.

[18] Criteo. 2022. About Us. https://labs.criteo.com/about-us.

[19] Criteo. 2022. Criteo OneTag advanced settings. https://help.criteo.com/kb/guide/en/criteo-onetag-advanced-settings-M2TiX6m90K/Steps/886908,887075.

[20] Criteo. 2022. Shopper Graph | Criteo. https://www.criteo.com/technology/shopper-graph/.

[21] The Trade Desk. 2021. How the advertising industry is preparing for life after cookies. https://www.thetradedesk.com/us/news/what-the-tech-is-unified-id-2-0.

[22] Steven Englehardt, Jeffrey Han, and Arvind Narayanan. 2018. I never signed up for this! Privacy implications of email tracking. *Proc. Priv. Enhancing Technol.* 2018, 1 (2018), 109–126.

[23] Ghostery GmbH. 2022. Ghostery. https://www.ghostery.com.

[24] Alejandro Gómez-Boix, Pierre Laperdrix, and Benoit Baudry. 2018. Hiding in the crowd: an analysis of the effectiveness of browser fingerprinting at large scale. In *Proceedings of the 27th International World Wide Web Conference (WWW)*.

[25] Apple Inc. 2021. What is Hide My Email? https://support.apple.com/en-us/HT210425.

[26] Arjaldo Karaj, Sam Macbeth, Rémi Berson, and Josep M. Pujol. 2018. WhoTracks.Me: Shedding light on the opaque world of online tracking. (2018). arXiv:1804.08959 [cs.CY]

[27] Pavel Kireyev, Koen Pauwels, and Sunil Gupta. 2016. Do display ads influence search? Attribution and dynamics in online advertising. *International Journal of Research in Marketing* 33, 3 (2016), 475–490. https://doi.org/10.1016/j.ijresmar.2015.09.007

[28] Steve Kroft. 2014. The Data Brokers: Selling your personal information. https://www.cbsnews.com/news/the-data-brokers-selling-your-personal-information.

[29] Mathias Lécuyer, Guillaume Ducoffe, Francis Lan, Andrei Papancea, Theofilos Petsios, Riley Spahn, Augustin Chaintreau, and Roxana Geambasu. 2014. Xray: Enhancing the web's transparency with differential correlation. In *Proceedings of the 23rd USENIX Security Symposium (Security)*. San Diego, CA.

[30] Mathias Lecuyer, Riley Spahn, Yannis Spiliopolous, Augustin Chaintreau, Roxana Geambasu, and Daniel Hsu. 2015. Sunlight: Sunlight: Fine-grained Targeting Detection at Scale with Statistical Confidence. In *Proceedings of the 22nd ACM Conference on Computer and Communications Security (CCS)*. Denver, Colorado.

[31] Evan Neufeld. 2016. Best practices in cross-device and cross-channel identity measurement. https://cimm-us.org/wp-content/uploads/2012/07/CIMM_Best-Practices-in-Cross-Device-and-Cross-Channel-Identity-Measurement.pdf.

[32] Oliver. 2018. Does YouTube Recommend Videos Watched by People on the Same Wi-Fi as You? https://weakwifisolutions.com/does-youtube-recommend-videos-watched-by-people-on-the-same-wifi-as-you/.

[33] Michalis Pachilakis, Panagiotis Papadopoulos, Evangelos P Markatos, and Nicolas Kourtellis. 2019. No More Chasing Waterfalls: A Measurement Study of the Header Bidding Ad-Ecosystem. In *Proceedings of the 19th ACM Internet Measurement Conference (IMC)*. Amsterdam, Netherlands.

[34] Emmanouil Papadogiannakis, Panagiotis Papadopoulos, Nicolas Kourtellis, and Evangelos P Markatos. 2021. User tracking in the post-cookie era: How websites bypass gdpr consent to track users. In *Proceedings of the Web Conference 2021*. 2130–2141.

[35] Panagiotis Papadopoulos, Nicolas Kourtellis, and Evangelos P. Markatos. 2019. Cookie Synchronization: Everything You Always Wanted to Know But Were Afraid to Ask. In *Proceedings of the 28th International World Wide Web Conference (WWW)*. San Francisco, CA, USA.

[36] Sarah Pearman, Shikun Aerin Zhang, Lujo Bauer, Nicolas Christin, and Lorrie Faith Cranor. 2019. Why people (don't) use password managers effectively. In *Proceedings of the ACM Symposium on Usable Privacy and Security (SOUPS)*. Santa Clara, CA.

[37] Franziska Roesner, Tadayoshi Kohno, and David Wetherall. 2012. Detecting and Defending Against Third-Party Tracking on the Web. In *USENIX Symposium on Networked Systems Design and Implementation (NSDI)*.

[38] Iskander Sanchez-Rola, Matteo Dell'Amico, Davide Balzarotti, Pierre-Antoine Vervier, and Leyla Bilge. 2021. Journey to the Center of the Cookie Ecosystem: Unraveling Actors' Roles and Relationships. In *Proceedings of the 42th IEEE Symposium on Security and Privacy (Oakland)*. Virtual Event.

[39] SimilarTech. 2022. Retargeting Technologies Market Share and Web Usage Statistics. https://www.similartech.com/categories/retargeting.

[40] Konstantinos Solomos, Panagiotis Ilia, Sotiris Ioannidis, and Nicolas Kourtellis. 2019. {TALON}: an automated framework for cross-device tracking detection. In *Proceedings of the 22th International Symposium on Research in Attacks, Intrusions and Defenses (RAID)*. Beijing, China.

[41] Catherine E. Tucker. 2012. The economics of advertising and privacy. *International Journal of Industrial Organization* 30, 3 (2012), 326–329. https://doi.org/10.1016/j.ijindorg.2011.11.004 Selected Papers, European Association for Research in Industrial Economics 38th Annual Conference, Stockholm, Sweden, September 1-3, 2011.

[42] Vishak. 2020. 12 Best Temporary Email Services To Protect Your Privacy In 2021. https://codeandhack.com/temporary-email-services-to-protect-privacy/.

[43] Marissa Wood. 2019. Today's Firefox Blocks Third-Party Tracking Cookies and Cryptomining by Default. https://blog.mozilla.org/en/products/firefox/todays-firefox-blocks-third-party-tracking-cookies-and-cryptomining-by-default/.

[44] Yahoo! 2022. Supply Side Platform (SSP) Advertising | Yahoo Ad Tech. https://www.adtech.yahooinc.com/advertising/publishers/solutions/ssp.

[45] Yahoo! 2022. Yahoo | Our Trusted Brands | Verizon Media. https://www.adtech.yahooinc.com/our-brands/yahoo.

[46] Yahoo! 2022. Yahoo Native Dot Tags. https://developer.yahooinc.com/native/guide/audience-management/dottags.

[47] Sebastian Zimmeck, Jie S Li, Hyungtae Kim, Steven M Bellovin, and Tony Jebara. 2017. A Privacy Analysis of Cross-device Tracking. In *Proceedings of the 26th USENIX Security Symposium (Security)*. Vancouver, BC, Canada.