

YourAdvalue: Measuring Advertising Price Dynamics without Bankrupting User Privacy

MICHALIS PACHILAKIS, University of Crete / FORTH, Greece
PANAGIOTIS PAPADOPOULOS, Telefonica Research, Spain
NIKOLAOS LAOUTARIS, IMDEA Networks Institute, Spain
EVANGELOS P. MARKATOS, University of Crete / FORTH, Greece
NICOLAS KOURTELLIS, Telefonica Research, Spain

The Real Time Bidding (RTB) protocol is by now more than a decade old. During this time, a handful of measurement papers have looked at bidding strategies, personal information flow, and cost of display advertising through RTB. In this paper, we present *YourAdvalue*, a privacy-preserving tool for displaying to end-users in a simple and intuitive manner their advertising value as seen through RTB. Using *YourAdvalue*, we measure *desktop* RTB prices in the wild, and compare them with desktop and *mobile* RTB prices reported by past work. We present how it estimates ad prices that are encrypted, and how it preserves user privacy while reporting results back to a data-server for analysis. We deployed our system, disseminated its browser extension, and collected data from 200 users, including 12000 ad impressions over 11 months.

By analyzing this dataset, we show that desktop RTB prices have grown 4.6× over desktop RTB prices measured in 2013, and 3.8× over mobile RTB prices measured in 2015. We also study how user demographics associate with the intensity of RTB ecosystem tracking, leading to higher ad prices. We find that exchanging data between advertisers and/or data brokers through *cookie-synchronization* increases the median value of displayed ads by 19%. We also find that female and younger users are more targeted, suffering more tracking (via cookie synchronization) than male or elder users. As a result of this targeting in our dataset, the advertising value (i) of women is 2.4× higher than that of men, (ii) of 25-34 year-olds is 2.5× higher than that of 35-44 year-olds, (iii) is most expensive on weekends and early mornings.

CCS Concepts: • **Information systems** → **Online advertising; Display advertising**; • **Networks** → *Network privacy and anonymity*.

Additional Key Words and Phrases: Advertising Transparency; Real Time Bidding Ad-auctions; User Privacy and Anonymity

ACM Reference Format:

Michalis Pachilakis, Panagiotis Papadopoulos, Nikolaos Laoutaris, Evangelos P. Markatos, and Nicolas Kourtellis. 2021. YourAdvalue: Measuring Advertising Price Dynamics without Bankrupting User Privacy. *Proc. ACM Meas. Anal. Comput. Syst.* 5, 3, Article 32 (December 2021), 26 pages. <https://doi.org/10.1145/3491044>

1 INTRODUCTION

In the last few years, around 300 billions dollars are spent per annum in digital advertising globally [1]. The dominant portion of this amount (84.9% for the case of US) go to *programmatic advertising* [2]. In programmatic advertising, real-time auctions that last less than 100 ms take place before a visitor finishes loading a web-page, to determine which advertisements they get to see. The most popular medium for conducting such real-time auctions is the Real-Time Bidding (RTB)

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than the author(s) must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from permissions@acm.org.

© 2021 Copyright held by the owner/author(s). Publication rights licensed to ACM.

2476-1249/2021/12-ART32 \$15.00

<https://doi.org/10.1145/3491044>

protocol, amounting for 90% of all programmatically purchased ads [3] (the remaining percentage is split between Header-Bidding [4, 5], Programmatic Direct [6] and Private Exchange buying (PMP) [7]).

The main promise of programmatic advertising and RTB is that visitors get to see only ads tailored to their actual interests, instead of generic ones that may be of little relevance. The key benefits of RTB include: (i) granular and targeted ad-buying, (ii) cost optimization, efficiency, speed, and (iii) enhanced transparency to both the advertiser and the publisher. But *what about the transparency to the factors that impacts the pricing dynamics and generates the data for these decision engines, i.e., the user?* Typically, users learn little, if any, regarding the prices that advertisers pay to buy the real estate of an ad-slot on their displays.

There are only two studies aiming to shed light on the pricing dynamics of the RTB ecosystem. [8] reports cleartext RTB prices from artificial personas, as well as a handful of prices (400) from real users, to assess how ad prices vary across different DSPs in three countries. [9] showed how to estimate the value of RTB winning bids, even when they are encrypted. This was done by running purpose-built probe campaigns to obtain ground truth data about encrypted prices, that were then used to train a classifier that predicts encrypted prices based on features of users, ads, and domains. To demonstrate the effectiveness of their classifier, the authors tested their approach on datasets from a proxy for mobile users.

In this paper, we present *YourAdvalue*: a first-of-its kind transparency software [10] for programmatic advertising over RTB. *YourAdvalue* allows end-users to learn in real-time while browsing, how much money the programmatic advertising ecosystem pays to capture their attention by rendering ads on their display. *YourAdvalue* takes into account both cleartext and encrypted winning bids. Users help in estimating the value of the latter by committing ad prices and auction metadata in a crowd-sourced database, in a privacy-preserving manner. Data gets first anonymized, and then feature aggregation is applied, prior to their transmission towards the database. This allows *YourAdvalue* to re-train its classifiers, and improve its accuracy on estimating the most up-to-date encrypted prices. By better understanding the value of their digital footprint, we expect that end-users will be able to make more informed choices about how they wish to participate in the emerging data economy. *YourAdvalue* can also help data protection authorities, regulators, and watch-dog groups for signs of bias and discrimination against particular demographic groups [11, 12].

Over a period of 11 months, we collected information from 12000 ads shown to 200 users of *YourAdvalue*. This is, to the best of our knowledge, the largest desktop RTB data collected thus far (30× more impressions shown to real users than [8]). We analyzed this dataset for factors affecting the prices (age, gender, geo-location, time, user tracking, etc.), studied the time-evolution of pricing with respect to past reports from [8, 9], and assessed the privacy protection offered to end-users. Overall, our findings are as follows:

- (i) We show that *YourAdvalue* achieves an efficient trade-off between user anonymity and estimation precision for encrypted RTB prices.
- (ii) We show that desktop RTB prices have grown 4.6× over desktop RTB prices measured in 2013.
- (iii) Compared to mobile RTB prices measured in 2015, desktop RTB prices have grown by a factor of 3.8×.
- (iv) The advertising value of women appears to be 2.4× higher than that of men.
- (v) In terms of age, 25-34 year old users appear to be 2.5× more valuable to advertisers than 35-44 year users.
- (vi) Different age and gender groups appear to receive varying attention (and targeting) by different DSPs.

Table 1. Real world example of a RTB price notification URL with the ADX of MoPub and the winning DSP of PocketMath. True IDs are omitted for space efficiency.**Example of Winning Price Notification URL**

```

cpp.imp.mpx.mopub.com/imp?ad_domain=mobileacademy.com&
auction_time=1420072833&bid_price=1.17&bidder_id=...&
bidder_name=PocketMath&campaign_id=...&charge_price=0.95&
request_id=...&response_id=1420072832890&units=0&adgroup_type=
marketplace&adgroup_priority=9

```

- (vii) Turning to time-of-day and day-of-week correlations with RTB prices, we see that early mornings and weekends, is when the most expensive ads get rendered.
- (viii) Finally, we see that if prior to the rendering of an ad, exchange of information has taken place between advertisers and/or data brokers via *cookie-syncing*, ads end-up having a 19% higher median value than when there's no cookie-sync taking place before.

2 BACKGROUND ON RTB

A Real-Time Bidding (RTB) auction is a programmatic, instantaneous type of auction, where a publisher's advertising inventory is bought and sold on a per ad slot basis. During such an auction, advertisers place bids for an ad-slot in a publisher's website and the one with the highest bid gets to render its ad on the user's display. RTB is more efficient in terms of revenue than the traditional static ad-buying for both the advertisers and the publishers.

A typical RTB auction: A typical transaction for an ad-slot begins with a user visiting a website. This triggers a bid request from the publisher (or Supply Side Platform (SSP)) to an Ad-Exchange (ADX), usually including various pieces of user's data (e.g., interests, demographics, location, cookie-related info, minimum acceptable price, etc.). Then, multiple Demand-Side Platforms (DSPs) programmatically submit their impressions and their bids in CPM (i.e., cost per thousand impressions [13]) to the ADX. All bids are sealed so every participant places only one bid for a particular ad-slot; this allows the RTB auction to finish within milliseconds (the entire RTB protocol usually runs in around 100 ms). The ad slot goes to the highest bidder and its impression is served in the user's display. Typically, the charge price for the ad slot is the second higher price following the Vickrey auctions [14].

Charge price notification: When an ADX selects the winning bid of an auction, the corresponding bidder must be notified about its win and the price to be paid to the ADX. This happens with a notification message (or nURL) conjoined with the price, piggybacked in the ad-response. This nURL passes through the user's browser and acts as a call-back to the DSP. This ensures the DSP that the winning impression was indeed delivered (the callback is fired soon after the impression is rendered on the user's device), and also gives the opportunity to drop a cookie on the user's device. The nURL includes the winning DSP's domain, the charge price, the impression ID, the auction ID and other relevant logistics (see example in Table 1). In this work, we monitor such nURLs and we study the prices embedded in them, as well as how they associate with the users' browsing behavior and other personal information.

Estimation of encrypted charge prices: Major ad companies like DoubleClick, OpenX, RubiconProject and PulsePoint tend to encrypt important information in the nURL (e.g., the charge price) to ensure the integrity of the reported values and avoid the prying eyes of competitors. Olejnik et al. in [8] assumed that the encrypted prices are following the same distribution with the cleartext ones. After observing an increase on the number of ad-companies that were using encryption, Papadopoulos et al. in [9] designed a methodology to estimate the value of encrypted prices showing that encrypted ad prices are (1.7×) more expensive than cleartext.

Header Bidding (HB): In the last few years, a new advertising standard has been proposed, called Header Bidding [4]. In this protocol, the ad-auction takes place in the user’s browser with the help of a javascript library), instead of a remote ADX as in RTB. Advertising partners receive bid requests from the browser library for an ad slot, and respond directly with their bid. The publisher (via the library in each user) is therefore informed of all bids and can choose the highest bid to sell the impression. In contrast to RTB, in HB the publisher has full control of the auction, and knowledge of the advertising partners and bids they place.

3 SYSTEM DESIGN

In this section, we elaborate on the guiding principles behind the design of the system, the various components it needs and the functionalities they need to execute to compute correctly this cost, while delivering a proper end-user experience.

3.1 Guiding Design Principles

We group the design principles into three classes: user-related, system-related and ad-related. The *user-related* principles are summarized as follows:

- (i) **Simple to understand.** The user interface (UI) of the tool should display metrics and visuals that are simple to understand by users who are neither tech- nor privacy-savvy.
- (ii) **Unhampered user experience.** The tool should be easy to install and not degrade the user’s existing Web experience. It should not change a webpage’s appearance, it should not slow down the loading of a webpage, should not affect the network traffic induced by a webpage, and should use a minimum of computing resources on the user’s device.
- (iii) **Preserve user anonymity.** The tool should not, in any case, leak sensitive information that could help an entity to re-identify the end-user on the Web.

The *system-related* principles are summarized as follows:

- (i) **Transparency.** The tool’s operations and functionality should be transparent, and an end-user or engineer should be able to audit the tool and its functionalities.
- (ii) **Scalability.** The tool’s architecture should be able to scale to accommodate activity delivered by thousands of users.
- (iii) **Fault tolerance.** The tool should continue to function even after encountering unexpected behaviour, without the need of user interaction, or causing issues to the user’s browser.
- (iv) **Privacy-by-design.** The tool should allow users to opt-in to reporting metadata from ads detected, but in a privacy-preserving fashion. That is, the data transmitted from the tool should not expose the user’s identity.

The *ad-related* principles are summarized as follows:

- (i) **Real-time operation.** The tool should operate in real time and be able to inform the user of updated measures as soon as new RTB ads are detected.
- (ii) **Generality of RTB detection.** The tool should detect RTB ads regardless if delivered in regular webpages (e.g., news) or closed-up publishers (e.g., social media).
- (iii) **Untampered RTB ad-protocol.** The detection mechanism should be simple and not tamper with the RTB protocol responsible for the impressions logistics, nor add on its expected latency to operate correctly.
- (iv) **RTB price detection.** The detection mechanism should detect and consider the value for both encrypted and cleartext prices from RTB nURLs.

YourAdvalue satisfies all aforementioned principles. It provides a simple UI to its end-users that reports in real time RTB information without altering the user experience or the RTB protocol, and

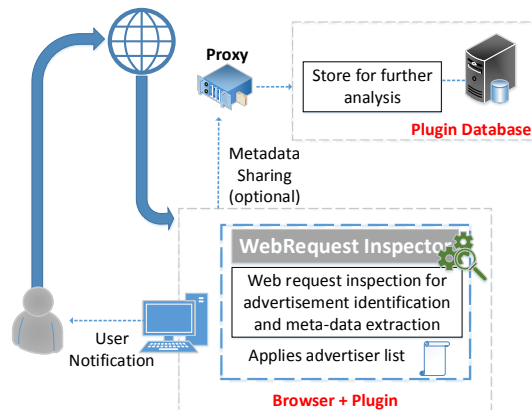


Fig. 1. Overview of the *YourAdvalue* architecture. The user visits the publisher’s website and the extension inspects all the webRequests to find any possible ads. When an ad is detected, the extension extracts metadata related to the ad, anonymizes them and, if the user chooses to do so, sends them to a proxy which forwards them to the server, for storage and historical analysis. At the same time, the UI incorporates the new price extracted from the ad, and displays an updated total value of the user to advertisers.

these two aspects were confirmed after extensive testing. At the same time, its development process was guided by security and privacy-by-design principles to guarantee security and anonymity of online users. *YourAdvalue* is an open source plugin and its code is available in our Github repository. More information can be found in the plugin’s webpage hosted at <https://youradvalue.tid.es:2222>.

3.2 System Components & Functionalities

Following these guidelines, we design the system and its various components. In particular, the system has two main components: 1) the end-user client (web browser extension), 2) the web server (running behind a web proxy). Figure 1 illustrates an overview of the system’s architecture. In the next paragraphs, we explain each of the components and how they fulfill the listed properties. The extension is written in a few thousand lines of JavaScript and is available for all major browsers. Additionally, the extension can be inspected at the user side while running on the browser.

3.3 User Interface & Input

For the end-user client to adhere with the user-related principles, we opt for building a browser extension which allows easy access to the user’s browsing sessions, can provide an easy to use UI, and can perform network traffic inspection without interfering with the normal operation of the website under visit. The extension executes a web request inspector, which analyses the incoming and outgoing network traffic of a website and detects RTB-related requests (nURLs). An example of this user interface (UI) can be seen in Figure 2.

This UI is simple and intuitive and was designed with the help of UX experts. In this UI, the user is prompted to indicate (if they want) their gender and age range. Furthermore, it displays to the end-user various simple metrics such as the total cost computed for the said user since the tool’s installation, as well as the cost of the user for the current session. The RTB ad detection and ad-related metadata extraction methods are described next. The design of the extension is light and its operation was tested with hundreds of different websites to confirm it does not hinder the normal loading or operation of each site. Finally, buttons allow the user to switch it on/off, and opt-in/out at any time from anonymous data reporting.




Fig. 2. The user interface in the popup that the user gets while using *YourAdvalue* browser extension. It includes: (i) gender and age fields which are user-specified and optional, (ii) pie chart which displays the type of ads a user encounters in a nice format, (iii) statistics field, where all time and session information about the ads are displayed, and (iv) menu, where the user can tune the extension.

3.4 Web Server & Database

At the backend, the server receives the anonymous metadata reported for further analysis. Upon arrival, the data are shuffled (reordering the table to remove contiguous rows from the same user) with the already collected data to break any relationship of reported data with their reporting users. The collected data are subsequently cleaned and analyzed for retraining and updating the decision tree (DT) model for the encrypted prices (as described in [9]). The updated DT is sent at frequent intervals to the user clients. For this transmission, the tree is serialized using a proper and agreed structure (e.g., XML format) and deserialized at the browser extension.

3.5 RTB ad detection & price extraction

nURL detection. The browser extension analyses all network traffic outgoing from the browser to detect all RTB-related requests. For each request leaving the browser, the extension collects metadata to check if it is a nURL about an ad. To detect such requests, we employ the *webRequest* API, which is common and available in all browsers we support. In particular, we create a listener using the event *onBeforeRequest*, an event which is triggered when a request is about to be made. The goal is to detect RTB nURLs; we don't want to block or redirect the nURLs and we don't want to tamper with the RTB protocol or affect the user experience.

We note that in RTB, the auction information is opaque to the client and the only information that can be inferred is through the parameters of the nURL. In contrast, in the HB protocol, browser DOM events are triggered that contain metadata directly available at the user browser. We use such events to clearly distinguish between RTB and HB-related prices, thus, focusing on the former.

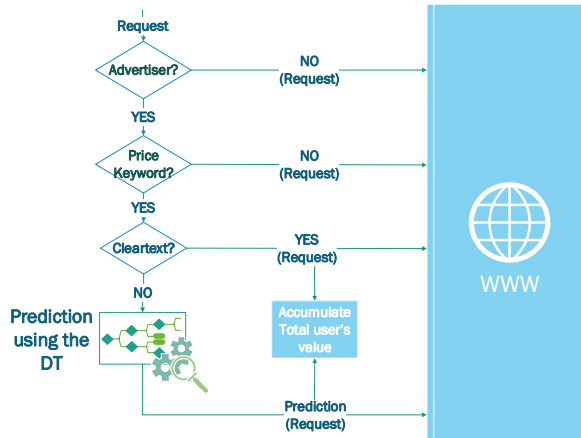


Fig. 3. The extension checks if the request corresponds to a known advertiser and if the price is encrypted. If it is not a known advertiser, it is let through. If it is cleartext, the tool extracts the value of the ad and adds it to the user's total ad-cost. If it is encrypted, the tool applies a decision tree model to infer its value before adding it to the user's total ad-cost.

DSP matching. After the tool catches an outgoing request, it checks if it is a nURL or not, by comparing the destination to a set of known DSPs (Figure 3). If it does not match with any of the DSPs, it is allowed to pass. If it is, we consider it a potential notification URL and extract the metadata related to it (i.e., all http variables and values accompanying them).

Price extraction. We then check if there is a possible price keyword in the nURL's metadata, and if this keyword is correlated to the nURL's advertiser (we utilize lists of keywords provided by [8, 9]). The argument here is that each advertiser or DSP is associated with a specific keyword used in the RTB. Therefore, if the price keyword does not match the typical keyword used by the detected DSP, it is considered a false positive and let through. However, if the keyword matches what the DSP uses, the tool extracts the associated price value.

Encrypted vs. cleartext prices. By leveraging the methodology of [9], the extension is able to detect both cleartext and encrypted RTB prices (i.e., if the price is numeric or alphanumeric). If the value is numeric (cleartext), it is normalized to CPM in US dollars and added to the user's total ad-cost. Otherwise, the tool applies the provided Decision Tree (DT) to estimate the value, and then add it to the total ad-cost.

Encrypted price inference with DT model. Classifying an encrypted price from a nURL is not a trivial task. First, before the tool even detects a nURL, it needs to load the DT model provided by the web server. The model is constructed at the server and exported as an XML file which is rolled out with the extension. The extension parses this XML file and creates an internal representation of the model that the server has created. We choose to embed (or preload) the DT model in an XML representation inside the extension for two reasons: (i) the DT does not change frequently and an existing version can be preloaded, and (ii) to reduce the number of requests from the extension to the server to minimum needed. Whenever the DT is updated from the server (normally every few months, when the server has enough new data to retrain the DT), the extension receives it in the form of an update of about 60-70KB. For the actual prediction of the encrypted price, we implemented a lightweight and efficient function that parses the DT and outputs the price prediction, so it does not interfere with the user experience, nor slow down the user's browsing (each price prediction takes less than 1ms). If the extracted feature values cannot be matched to the

Table 2. Metadata that the data-contributing users send to the server for further analysis and model updates.

	Feature	Feature Description
User Information	Gender	User's chosen gender
	Age	User's chosen age
	Location	User's current location
Browsing Information	Time Of Day	Time of the day an ad was detected
	Day Of Week	Day of the week an ad was detected
	Category	Topic category of 1st party domain
	DoNotTrack	DoNotTrack flag enabled/disabled
Advertisement Information	Ad Format	Impression's size in display
	Winner DSP	The DSP that won the auction
	Price Keyword	Keyword with the RTB ad-price
	Price Value	RTB price detected in the nURL
	Cookie Syncing	Cookie Synchronization detected

DT model, the extension uses a rolling time average of the cleartext prices as an estimation of the detected encrypted price that could not be inferred.

Feature extraction. When an encrypted price is detected, the tool extracts the features that the DT requires to make the price inference. For the feature extraction process (explained next), we collect metadata about the publisher and the detected ad. These features are provided to the prediction function to perform the inference.

Extension overhead. We tested our extension and did not notice any impact in the user experience. The analysis of a nURL takes less than 1ms, even when the price is encrypted. Also, when the user allows it, data sharing happens in an asynchronous manner which does not affect the page load time, or the total available bandwidth.

3.6 Ad-related feature extraction

Location extraction. For the location extraction, we use an online API. By making a GET request over a secure channel (HTTPS), the extension obtains and stores the user location on a country level. The location is stored in the client for the whole session (as long as the user keeps the browser open, does not restart the extension, etc.), reducing requests and other network resource needs.

Ad-format extraction. After analysing several nURLs, we created a list with keywords stating width and height of ads. Based on that list, we examine the nURLs' parameters for possible matches that are then stored in the price features.

Cookie Synchronization extraction. Except from detecting RTB, we are interested in detecting if a user is being tracked by a 3rd-party entity. To achieve that, we again employ the *webRequest* API and in particular, the *onBeforeRequest* event. We attempt to detect if user identifiers are sent from the publisher's webpage to other hostnames. We first load all the cookie identifiers stored in the user's browser, discarding session cookies and those with values less than 10 characters because most of times they have values unrelated to user identifiers. For each request from a tab, we check if a user identifier is present in the URL, and if that URL is sent to a tracker. Also, we make sure that the URL detected is not from the same domain that the user is currently visiting in that tab. Our method is able to detect a user identifier if it is included in the URL's parameter values, or as part of the URL path. When we detect such tracking from a 3rd-party, we store the information in a binary flag, indicating that user tracking with cookie synchronization took place in a specific tab to a specific domain.

IAB category extraction. Websites can be categorized based on taxonomies such as IAB's [15], which are well adopted by the ad-ecosystem. In order to extract the IAB category for each 1st-party, we created a list with a mapping between websites and their IABs. We choose to include only very popular websites to keep the list small and the extension lightweight. This list contains the top 500

Alexa sites, and can be updated from the server as frequently as the DT. When the user visits a website, we extract the 1st-party domain and check the list. In case we have a match, we report the IAB found, otherwise we report “undefined”.

3.7 User- & ad-related metadata reporting

As mentioned earlier, *YourAdvalue* allows users to contribute with their data like ad-prices and auction metadata. In Table 2, we list the metadata reported to the server. The list of features is carefully selected to reflect the needed features for the DT modeling to happen at the server. It also enables further analysis of RTB prices, how they evolve through time and if the advertisers target particular user categories based on their demographics (e.g., gender, age, etc.).

Being very cautious to preserve the user anonymity and privacy, the data reported do not include any user identifiers or other Personally Identifiable Information (PII). Also, they are anonymized at the client side before being sent to the server. Each feature can be reported at different granularities. However, one crucial question in this work is what is a sufficient granularity for a given feature to allow effective price modeling, but not compromise user anonymity? In Section 7, we elaborate on the aggregation or other obfuscation methods via noise addition that are possible for the given features, and the guarantees or limits on anonymity they can offer to the reporting users.

Also, in order to overcome the problem of an honest but curious server trying to re-identify the sender, each extension reports its data at random times throughout the day and week. Because of this induced randomness, reports may delay from several seconds to days after they are created, before leaving the user’s browser. Therefore, the server cannot identify a user based on the rate of its reports, since reporting from different users gets mixed-up while being received at the server side. In Section 7, we also investigate possible threats on anonymity of users reporting such data, and how the system’s design can protect them from such de-anonymization attacks.

4 DATASET COLLECTION & LIMITATIONS

YourAdvalue is designed with user anonymity and privacy-by-design in mind. Thus, we need to ensure that all contributing users cannot be (re)identified or correlated with their reported data, and that user profiles cannot be constructed. To keep that promise, we do not store user PII, and do not make attempts to correlate user data shared with the server. This comes with the limitation that we cannot know exactly how many users contribute data, and with what rate, i.e., when and what is reported, and by whom.

To attract users for our study, we advertised the *YourAdvalue* plugin across different social media networks and groups, and received interest from users across different demographics and communities. Since we made our plugin available for download from the Chrome store ¹, we were able to get an intuition of the active users from the information page that Google provides to the developers, which indicates the active users at each time. We observed that the rate of data sharing followed a similar trend to the active users of our plugin. That is, while the number of active users (i.e., that installed the plugin) increased, so did the rate of data sharing. Similarly, when users removed our plugin or deactivated it, the reporting rate dropped. Overall, a total of 200 users agreed to download and install our extension in their primary browser, providing us with real user data, but also imposing some limitations as well.

First, the users were free to use ad-blocking software. Disabling the ad-blocker is not obligatory for the tool to work, but we advised the users not to use ad-blocking, so that they see the tool in action and receive a better estimate of their ad-cost. This is because ad-blocking reduces the number of ads delivered successfully to the user browser. Therefore, RTB ad detection on the browser, as

¹<https://chrome.google.com/webstore/detail/youradvalue/apipekdgpjmboidaohhecfpacahaap>

Table 3. Summary of the datasets or values available from our work and past studies.

Property	Desktop (U) <i>YourAdvalue</i>	Mobile [9] (M) Papadopoulos et al.	Desktop [8] (D) Olejnik et al.
Time period	2019 (11 mon.)	2015 (12 mon.)	2013 (1 mon.)
Real users	~200	~810	~100
RTB Ads detected	11946	78561	~400
Gender	Yes	No	No
Age	Yes (5 bins)	No	No
Location	10 countries	Spanish cities	3 countries
Time Of Day	Yes (8 bins)	Yes (8 bins)	Yes (3 bins)
Day Of Week	Yes (7 days)	Yes (7 days)	No
Cookie Synching	Yes	No	No
DoNotTrack	Yes	No	No
adFormat	Yes	Yes	No
Winner DSP	Top 20 entities	Top 15 entities	Top 25 entities
IAB category	30 IAB categ.	30 IAB categ.	14 Web categ.

well as data reporting to the server, can become sparse or even non-existent, depending on the effectiveness of the ad-blocker in place.

Second, the number of users was not stable, since they were allowed to install the plugin, and then enable it and disable it at will, or even uninstall it completely. Therefore, the number of *YourAdvalue* users we report (200) is the distinct users that installed the extension at some point during the data collection, and likely contributed with some data.

Third, the reporting users are assumed to be real users, since they need to have a Google account, and proceed to several manual actions such to: search, find and install the specific plugin (or reach it via the tool’s website²), disable their ad-blocker if they can (for useful data to be collected), and enable the plugin to report anonymized ad-related data to our server, and optionally, declare their specific demographics (as explained earlier). We do not have control over any of the above actions the users choose to perform (or not), and therefore, cannot know how many of them are truthful in their data reporting. However, since *YourAdvalue* is a plugin related to a specific interest, and was advertised in different privacy-aware communities via social media, we believe our audience has some knowledge of the topic and no intention to be malicious.

Fourth, users of specific demographics may have preference to particular (categories of) websites. Indeed, we have no control over the websites visited by users, which introduces a limitation on comparing users of opposite-sided demographics (e.g., younger vs. older, or male vs. female users). In our analysis in the next sections, we assume that the data collected per demographic represent well each user category, and that we can cautiously generalize our findings as such.

Dataset breakdown: The data collected from said users are summarized in Table 3. As explained earlier, by design, we are not able to identify which users contribute and with how many impressions. However, we can provide a breakdown of the demographics vs. impressions reported. In particular, we received 4024 impressions from male and 6640 from female users. Also, we received 9826 impressions from the 25-34 age group, 1190 from the 35-44 age group. Finally, we received the following impressions per country: 6600 from US, 2685 from Spain, 1278 from Switzerland, 146 from Greece, 114 from Austria, 24 from Cyprus, 13 from Italy, 6 from Belgium, 4 from Germany and 3 from Israel. Differences from the totals reported in Table 3 are due to users choosing not to report all demographics (e.g., they report gender but not age and vice versa), the returned country having the value “unknown” because of issues with the location API, etc.

Comparison with past datasets (Table 3): First, we use statistics reported in a past study on RTB in desktop [8] in 2013 (no dataset released) for 100 users and various automated crawls (D). Second, we use a previously collected and shared dataset from [9] on RTB in mobile, for 2015-2016 (M). This is a larger dataset, containing 1 year-long activity of RTB ad metadata from 2015-16,

²<https://youradvalue.tid.es:2222>




Fig. 4. Portion of ads per DSP, per gender.




Fig. 5. Portion of ads per DSP, per age bin.

with about 80k ads from 810 mobile users, and each entry containing several features. The last dataset is the aforementioned U , with data from real users who installed the *YourAdvalue*. This dataset consists of about 12000 ad impressions detected from 200 users for a period of 11 months. We note that 64.7% of the ads collected have cleartext values and 35.3% encrypted. We also note that the M dataset has multiple times more ads reported than in the U dataset. This is because it is constructed from a traffic log of about $4 - 5 \times$ more users, and collected at a Web getaway proxy, instead of using a plugin, thus allowing them to detect all possible RTB-related ads delivered to these users. However, *YourAdvalue* collects similar metadata to the ones in the M dataset, so we can directly compare the price distributions for the two time periods, for mobile vs. desktop. Also, in contrast with previous datasets, we collect anonymized demographics from our users to study their association with RTB prices, something not done before in the past studies [8] and [9].

5 RTB TARGETING VS. DEMOGRAPHICS

In this section, we take a first look at data reported by ~ 200 real users who installed the *YourAdvalue* browser extension, in 2019 (U). Please see Section 9 for details on ethical collection of data. Table 3 summarizes the data and statistics available for this dataset.

In the next paragraphs, we focus on answering the following research question: How do the different demographics (i.e., gender, age and location) of online users affect their ad-cost? In particular, we focus on gender and age, and investigate if a specific age-gender group is being targeted more intensely by the RTB ecosystem (Section 5.1). Similarly, we study this question from the point of view of user location, at a country level (Section 5.2). Importantly, we proxy the intensity of targeting by studying 2 key elements of RTB: cookie synchronization (CS) during ad delivery, and the corresponding DSP delivering the ad.

5.1 Gender and Age Targeting

First, we assess how the different genders are targeted, by examining the portion of CS detected in RTB ads in conjunction to the gender of the user. We find that female users received ads with CS in 99.1% of the time, in comparison to 86.2% of ads with CS for male users. Thus, female users are targeted with 12.9% more CS than males. Similarly, we assess how different age groups are targeted, by examining the intensity of CS. We find that users of age 25-34 years received ads with CS in 96.2% of the time, compared to 79.3% of ads with CS for the older age group (35-44). Thus, younger users are targeted with 16.9% more CS than older users.

Second, we look into genders and the DSPs involved in ads. Figure 4 shows the CDF of 18 different DSPs involved, broken down by gender. We find that particular DSPs have preference for specific genders. For example, *adnxs.com* and *adsrvr.org* delivered many more ads to female than male users.




Fig. 6. Portion of ads with CS per country.

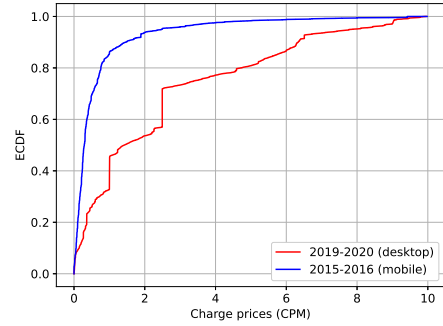


Fig. 7. Price distribution for M (2015-2016) and U (2019-2020).

In contrast, *casalemedia.com* and *criteo.com* had a preference for targeting male users. Similarly, we study age group and the DSPs involved in ads. Figure 5 shows the CDF of the DSPs involved, broken down by age. We find that particular DSPs have preference for specific age. For example, *adnxs.com* and *adsvr.org* delivered more ads to younger users (25-34). In contrast, *bluekai.com*, *criteo.com*, and *openx.com* had a preference for targeting older users (35-44). Interestingly, and as we will see in the next section, this variability or preference in targeting per gender and age, can potentially impact the RTB prices of ads delivered.

Finding: Female users are targeted more compared to male. Similarly, younger users are targeted more than older users. DSPs may have a preference for specific genders and age groups. This can typically be the result of the type of campaigns they were running at the time of data collection.

5.2 Country Targeting

We finish this investigation by comparing the ads delivered to users of different countries, with respect to the CS they received. Figure 6 shows for top 7 countries a breakdown of the percentage of ads delivered to users in these countries, that contained CS. We find that all countries had high portion of ads delivered with CS involved, i.e., 70.6% or more. Unsurprisingly, US as the most advanced ad-market, has one of the highest CS portions, along with smaller countries such as Belgium, Greece and Switzerland.

Finding: US users are targeted more than EU users. This could be explained, since US has a more mature and aggressive advertising system, and still a more relaxed online privacy regulation framework compared to EU.

6 RTB PRICE EVALUATION

In this section, following the analysis of demographics, we perform a price analysis and observe how the RTB ad-prices change over time, and how characteristics of users may be correlated to these prices. To do this analysis, we use all available results reported in literature on RTB prices. In particular, we aim to answer the following research questions:

- (1) How have RTB ad-prices changed over time? (Sec. 6.1)
- (2) Which day of the week, and hour of the day, do ads cost more? (Sec. 6.2)
- (3) For which IAB website category do ads cost more? (Sec. 6.3)
- (4) For which gender, age, and country groups do ads cost more? (Sec. 6.4, Sec. 6.5, and Sec. 6.6)
- (5) Does the use of cookie synchronization lead to higher costing ads? (Sec. 6.7)
- (6) Which DSPs pay more to reach users? (Sec. 6.8)




Fig. 8. RTB prices per day, for the M (2015-2016) and U (2019-2020) datasets.




Fig. 9. RTB prices per hour slot, for the M (2015-2016) and U (2019-2020) datasets.




Fig. 10. RTB prices per IAB category, for M (2015-2016) and U (2019-2020) datasets.

As a small suggestion, it might be nice to include some summary statistics to confirm these assertions, as well as a pointer to the section(s) that have further details.

6.1 General RTB price trends & changes

In Figure 7, we compare the CDFs of RTB cleartext prices for the two time periods we have detailed data: mobile M data [9] and desktop U data. We observe that the median price in U is 1.42 CPM (cost per mile, in USD), whereas in the older, mobile dataset it was 0.29 CPM, pointing to a $3.8\times$ increase between mobile and desktop RTB prices in a 4-year period. Also, the top 5% of RTB prices of the new dataset U are $\sim 2\times$ larger (7.85 vs. 3.7 CPM) in comparison to M dataset. Olejnik et al. in D [8] report an average price of \$0.43 CPM per user, with a min and max price of \$0.04 and \$1.98 CPM, respectively. As explained, we do not collect user identifiers, so we can not analyze prices per unique user. For this reason, we can only compute the total average price of ads collected across all potential users we had in the data collection period. We find an average price of \$2.44 CPM (st.dev=\$2.48) per user, which is more than $4.6\times$ higher than the average RTB price per user, 7 years ago in dataset D . Our results are verified by corporate studies as well, reporting a $4.3\times$ increase [16, 17]

Finding: During a period of 7 years, RTB prices have increased by $4.6\times$. The prices between desktop and mobile also increased by $3.6\times$ in the last 4 years alone.

6.2 Which day & hour costs more?

In Figure 8, we compare the prices found for each day of the week. For M , we observe a relatively stable distribution for each day with a median value around 0.3 CPM, and weekdays receiving ad-prices with a higher variability. In contrast, in the new dataset U , we find that during the weekends (Friday-Sunday) and Monday, the ad-prices tend to be higher than weekdays, with medians varying between 0.83 and 2.5 CPM. Interestingly, on Monday, the median price peaks to 2.5 CPM. Overall, and acknowledging M is on mobile whereas U is on desktop, we see an increase in prices advertisers pay in 2019 compared to the prices from 4 years ago.

Next, we study ad-prices paid with respect to the time an ad impression was displayed (Figure 9). To compare our results with the reported from Olejnik et al. [8] 7 years ago, we follow their format and we create 3 time slots: morning (0-8h), midday (8-16h) and night (16-24h) for the U dataset, and compare with D . As we see, their results confirm our findings: charge prices are higher in the morning slot than afternoon or evening. However, the prices detected in our dataset (U) are in general higher: morning (3.18 CPM vs. 0.57 CPM), midday (2.77 CPM vs. 0.53 CPM) and early night (1.61 CPM vs. 0.47 CPM).

When we compare our findings from U (2019), with the ones from M (2015), we note an increase in 2019 for the morning slot (3.18 CPM vs. 0.65 CPM), in midday (2.77 CPM vs. 0.64 CPM), and




Fig. 11. RTB prices per gender, for the U dataset (2019-2020).




Fig. 12. RTB prices by age group, for the U dataset (2019-2020).

night (1.61 CPM vs. 0.67 CPM). These differences (i.e., when the prices peak with respect to time of day and day of week) can be justified due to different usage of devices: users browse on mobile devices at different hours and days than on desktop [18], and with different intent [19], thus being delivered ads of different cost [20].

Finding: Ad prices are higher on weekends and during the beginning of the week. Also, prices are higher during the beginning of the day. Advertisers are more aggressive during the time users are more relaxed and out of office.

6.3 Which website category costs more?

In Figure 10, we set out to compare the distributions of ad-prices detected in U , and how they compare with M . In general, we detected prices across 11 different IAB categories in U , compared to 20 in M , with 9 overlapping. We find that the most popular IAB categories in U are News ($IAB12$), Technology and Computing ($IAB19$), and Unspecified. Similarly, in M , the most popular IAB categories are Sports ($IAB17$), Technology and Computing ($IAB19$), or Unspecified. We also compared the U 's IAB categories with the Web categories reported in D , if they loosely match, since no IAB taxonomy was used in [8]. In D , the most popular category is News ($IAB12$), followed by Entertainment ($IAB1$), Games ($IAB9$) and Technology and Computing ($IAB19$).

Using the general statistics provided in [8], and the distributions of data from [9], we compare U with the means of distributions of prices in 2013 (D), and 2015 (M). We observe that the average price paid per IAB category of the publisher is not the same across categories. This has also been reported in [8] and [9] and it is to be expected, since websites visited by users indicate behavioral preferences and economic status or buying power of the customer. RTB prices in 2019 are higher, on average, for all reported IABs (statistical significance at p -value < 0.001) compared to 2015 and 2013 data. Comparing D with U , the median for $IAB12$ (News) increased from 0.38 to 1.58 CPM, the median for $IAB19$ (Games) increased from 0.38 to 0.82 CPM, and the median for $IAB1$ (Entertainment) increased from 0.33 to 0.88 CPM. Similarly, when comparing M with U , we find the median for $IAB17$ (Sports) to be 0.34 CPM in M vs. 2.5 CPM in U , for $IAB19$ (Games) to be 0.24 CPM vs. 0.43, the median for $IAB12$ (News) to be 0.27 CPM vs. 1.0 CPM and in case of Unspecified, 0.28 CPM vs. 2.2 CPM. Consequently, some browsing interests are more highly targeted by advertisers and drive prices of ads to higher costs. However, the most popular IAB categories are not the most expensive with respect to ad-cost. In particular, $IAB17$ (Sports) and $IAB8$ (Food & Drink) are the most expensive, with median cost 2.5 CPM and 2.49 CPM in 2019, respectively.

Finding: Widely visited or popular website categories (based on the IAB taxonomy) tend to serve more expensive ads.




Fig. 13. RTB prices per country, for the *U* (2019-2020) dataset.




Fig. 14. RTB prices vs. CS, for the *U* (2019-2020) dataset.

6.4 Which gender costs more?

Next, we focus on some demographics of our real user-base from the 2019 *U* dataset (past studies did not capture these demographics). In Figure 11, we see that advertisers are willing to pay more to target women than men. In particular, for female users, the median price was 2.49 CPM, compared to 1.0 CPM for men, and 0.36 CPM for the unspecified. This can be explained by the earlier analysis on demographics and targeting, which showed that female are more targeted by specific DSPs and with more CS. This trend confirms past reports (e.g., [21]) which claim that women are generally more intensely targeted customers than men, thus, costing more to reach them.

Finding: Women are 2.5× more expensive to reach with ads compared to men. Advertisers employ more aggressive strategies to target women and are willing to spend more on ads to reach them.

6.5 Which age group costs more?

In Figure 12, we see the prices paid by advertisers based on the self-defined age of users. We see that advertisers are willing to spend higher prices for younger users. In particular, for users in the age group 25-34, the median price was 1.81 CPM, compared to 0.72 CPM for the group 35-44. Also, for the younger group, advertisers are willing to pay as high as 9.4 CPM, whereas in the older group the maximum reaches 5.21 CPM. Again, these results for younger users can be explained by higher CS and the tendency of specific DSPs to target this demographic.

Finding: Advertisers are willing to pay almost double when targeting younger people. Young people could be more receptive to ads, and thus a more profitable audience to target.

6.6 Which country costs more?

In general, we detect prices from 11 countries. In Figure 13, we plot the prices for the top 7 by volume. The top 3 are USA, Switzerland, and Spain, with median ad-prices 2.49 CPM, 1.47 CPM and 1.00 CPM, respectively. We note that the median price in Spain on the *M* dataset (mobile) was 0.29 CPM, whereas in *U* (desktop) is 1.0 CPM. That indicates an 2.4× increase in CPM in almost 4 years. Also, in *D* (2013), the average price in USA was found to be 0.69 CPM, indicating a 3 – 4× increase with respect to USA ad-prices in *U* (2019).

Finding: Users browsing from wealthier countries tend to be more expensive to target. Also, RTB prices increased around 4× in the past 4 years, globally. These results may show a hyperinflation of the specific ecosystem in comparison to the general growth of wealth in countries.

6.7 Does Cookie Synchronization (CS) indicate more costly ads?

CS [22, 23] is a very important technique for user data sharing, tightly connected with the RTB ad-auctions [8, 24, 25]. Interestingly, and as shown in Figure 14, the existence of CS is associated with higher RTB charge prices. In particular, when CS is detected, the median prices are 1.46 CPM,




Fig. 15. Major DSPs detected, for the M (2015-2016) and U (2019-2020) datasets.




Fig. 16. Portion of cumulative spending per DSP, in the U (2019-2020) dataset.

or 19% higher than when CS is not detected. In fact, the maximum prices with CS can reach 8.73 CPM, in comparison to 5.21 CPM without CS. These results support the intuition built by past studies [24], that when CS is performed before a RTB auction bid, it facilitates user ID syncing between the CS partners, thus enabling user re-identification and retargeting, and a more informed bidding for the DSP.

Finding: When CS is enabled, the median prices are 19% higher than when it is not. This verifies the hypothesis that targeted users are more valuable to advertisers compared to new users.

6.8 Which DSPs pay more to reach users?

Figure 15 presents a breakdown of the 16 major DSPs involved in RTB for the detected ad-prices. We find that 6 DSPs are involved with more than 75% of the ad prices detected in the U dataset. Interestingly, when compared to the DSPs found in M , there are differences which are to be expected, given that M was focused on mobile traffic. In particular, the top DSPs are *adnxs*, *adsvr*, and *critéo*, with median prices (U vs. M): 4.56 vs. 0.27 CPM, 1.0 vs. 0.34 CPM, and 1.0 vs. 0.31 CPM, respectively. When comparing U to the D dataset, we find that the top DSPs are *invoitemedia*, *turn.com*, and *mathtag*. From those DSPs we were able to detect advertisements only from *mathtag* with median price (U vs D): 1 CPM vs. 0.09 CPM. These results demonstrate that the top DSPs have increased the ad-prices they pay to reach desktop audiences by 3 – 17 \times through a period of 4-7 years. Also, in Figure 16, we show the portion of ad-cost that each DSP spends in RTB ads, in the U dataset. As expected, we find the top DSPs spend the most in order to reach online users. In particular, *adnxs* spends 55.6%, *critéo* spends 7.46%, and *adsvr* spends 12.58% of the total amount spent to reach users.

Finding: All DSPs increased the prices they are willing to pay for ads in the past 7 years, by a factor of 3-17 \times .

7 PRIVACY ANALYSIS

In order for the *YourAdvalue* system to accurately estimate encrypted charge prices, and keep its price modeling engine up-to-date with the current pricing dynamics in the ad-ecosystem, it requires data from the users. Given the system’s principle for privacy-by-design, *YourAdvalue* needs to protect users’ anonymity and privacy against re-identification attacks after a possible data-breach or a malicious server controller. In this Section, we first elaborate on possible threats that can expose a user’s identity, allowing an attacker to link a user to his reported data, and discuss how we solve these threats, thus maintaining user anonymity (Section 7.1). Then, we analyze limits on user privacy by features reported by *YourAdvalue* (Section 7.2). To do that, in Section 7.3 we focus on the

larger, M dataset of 810 users and 80k ads, used in [9] with the same features as our system's. This dataset allows us to study real users' distributions of features, and how user anonymity changes with feature aggregation. Finally, in Section 7.4, we study the inherent trade-off between feature aggregation and price estimation accuracy.

7.1 User de-anonymization threats

Below, we discuss possible de-anonymization threats that the system could face during data reporting, and how the system design prevents them from materializing.

Threat 1: De-anonymization using users' online behavior or demographics. Most users have a consistent online behavior, meaning they tend to visit the same websites through their week (e.g., a user may visit `cnn.com` to read daily news, vs. another user who does that via `boston.com`). This behavior makes the user identifiable, as already shown in [26]. In addition, the 1st-party domain that an ad was detected on, could be a private or sensitive domain or a domain that other users are very unlikely to visit. Also, browsing behavior could be linked to particular demographics like age and gender. We address these threats by not reporting the 1st-party domains visited by users but their corresponding IAB categories [27, 28]. In this way, multiple different sites can appear under the same IAB (e.g., all news sites will be under the same IAB). Furthermore, the users may choose to provide a self-defined gender (binary) and age (aggregated in coarse-grained 10-year bins).

Threat 2: De-anonymization using advertising-related activity. Users could be monitored by different trackers while browsing the web [29–31], with user-specific IDs and other identifiers associated with them. In order to reduce the probability of de-anonymization of a user, we do not report any identifiers or trackers, or even advertisers involved. Instead, we extract metadata from the nURL such as the winner DSP (i.e., the intermediary entity between advertiser and ADX), the price keyword of the ADX, the price value and ad size, if available. We also report if cookie syncing was detected.

Threat 3: De-anonymization using grouped records into user sessions. An attacker may attempt to de-anonymize a user, by assuming that consecutive records of incoming data could be reported by the same user. The system proceeds to the following actions to reduce or even eliminate such a threat: (1) As explained in Section 3.6, the reporting client does not send data for every detected ad, but instead collects a set of records which it then shuffles and sends to the server at random times. This strategy also breaks possible time linkage between records, as they are not reported as soon as they are created. (2) The server randomly shuffles incoming data with stored data frequently, depending on the amount of data stored. This enables the server to break any possible sessions or relation that records have.

Finding: The privacy-preserving design of *YourAdvalue* protects users from de-anonymization attacks in various typical or extreme scenarios.

7.2 User uniqueness via surprisal analysis

Overall, the system design can protect users' anonymity by removing identifiers or attributes that contain unique examples of users' activity, reducing the granularity of reported data (with respect to frequency of reporting or detail), and randomizing data reporting and storing to avoid possibility of linking records together into sessions. However, it is important to study the bounds of anonymity that can be achieved, when features reported to the server are considered together while attempting to de-anonymize a user.

In this section, we formally investigate this problem. In particular, we are interested to study what granularity of reporting is low enough to protect the anonymity of each user, while still maintaining

full or partial utility of the reported data. For this investigation, we employ two commonly known tools to understand this trade-off. First, we study how unique a user is, given a particular reporting setup (i.e., combination of features and granularity classes). and (similar to other works [32]) we estimate this by measuring the number of *surprisal bits* that such a setup achieves. Second, we compute the number of users who are expected to be found or collide in a given reporting setup. The second study is a classic k-anonymity analysis, in which we measure how many users have similar behavior and are likely to report the same combination of features and granularity classes. Next, we offer these analyses of surprisal and k-anonymity for the features considered in Section 3.6.

Background on Information Theory: In information theory, given an event E with probability $P(E)$, surprisal is $I(E) = -\log_2(P(E))$, and is measured in bits; the higher the surprisal the more unique an event is. Something that is certain has no surprisal (0 bits); flipping a fair coin is associated with 1 bit of surprisal; winning the lottery has 24 bits of surprisal. In our case, higher surprisal of a user's reporting data means they can be uniquely identified easier.

If the said event E is dependent on N independent attributes or features ($i \in N$), then the overall probability of $P(E)$ can be expressed as:

$$P(E) = p_1 \times p_2 \times \dots \times p_i \dots \times p_N. \quad (1)$$

If we assume that each feature i has different discrete classes, and they are equally likely to appear (i.e., they are governed by a uniform distribution), then every feature i has a probability $p_i = 1/\#$ of distinct classes of feature i .

Uniqueness using Uniform Distributions: Under the assumption of a uniform distribution of classes of features (Eq. 1), all the combinations of features have the same probability to occur. This assumption allows us to study the problem of data reporting and define baselines of features and their classes to be reported.

In a more realistic scenario, an event E , or in our case experimental setup E (that dictates specific class for each of the features), is dependent on N independent features that are governed by real distributions. Thus, the above probabilities are not equal for all classes of a given feature. Instead, for class j of feature i we have $p_{ij} = PDF_i(E_{ij})$, where $PDF_i()$ is the real (or observed) probability distribution function of feature i , and $PDF_i(E_{ij})$ is the probability returned by this function for the event E to occur at class j of this feature i , E_{ij} (i.e., the particular level or class that the experimental setup takes (j) for the feature i). For example, feature i could be the time of day, and the different classes of this feature could be the 24 hours available (j). Then, the probability of an ad to be detected at the 11am slot, is given by the $PDF_{i=\text{time of day}}(E_{ij=11am})$.

We can compute overall $P(E)$ for a specific combination of features $i \in N$ and classes j :

$$P(E) = PDF_1(E_{1j}) \times \dots \times PDF_N(E_{Nj}) = \prod_{i=1}^N PDF_i(E_{ij}) \quad (2)$$

where: E = the experimental setup of interest, comprised of a selected class j , for each feature E_i , $PDF_i(E_{ij})$ = probability of E to occur, as computed from the PDF of feature i for its class j , $P(E)$ = overall probability of E to occur.

User uniqueness via surprisal analysis assuming uniform distributions: Following the assumption of a uniform distribution of classes of features (Eq. 1), all the combinations of features have the same probability to occur in our dataset. This assumption allows us to study the problem of data reporting and define some baselines of features and their classes that we could report. In Table 4 we compute the theoretical surprisal bits of a system that has the features mentioned earlier and are shown in individual rows in the table. The granularity of reporting, i.e., the number of independent and uniformly distributed classes is reported in each column for each row.

Table 4. Number of surprisal bits for different distinct classes of each feature, assuming uniform distribution of appearance of each class, for each feature.

Feature	Number of classes									
Gender	2	2	2	2	2	2	2	2	2	2
Age	100	5	5	5	5	5	5	5	5	5
Country	240	240	240	240	240	120	120	25	25	25
Time Of Day	24	24	8	8	8	8	8	8	8	8
Day Of Week	7	7	7	7	7	7	7	7	7	7
Cookie Sync	200	200	200	2	2	2	2	2	2	2
DoNotTrack	2	2	2	2	2	2	2	2	2	2
adFormat	20	20	20	20	20	20	20	20	20	3
winnerDSP	200	200	200	200	200	200	200	200	200	200
category	500	500	500	500	30	30	30	30	30	30
Price Keyword	1	1	1	1	1	1	1	1	1	1
Price Value	200	200	200	200	200	200	50	50	50	3
Surprisal Bits	60.2	55.8	54.3	47.6	43.6	42.6	40.6	38.3	31.5	

We begin by assuming that the *YourAdvalue* reports back the features exactly as collected without any generalization. This scenario is hypothetical, but can give us a baseline understanding of what we are dealing with respect to surprisal bits. We assume that the age ranges from 0 to 99 years, the user could be from any of the 240 different countries of the world and that the reporting is in 1-hour bins. We also assume that the client reports back the 1st-party domains they browse through (e.g., 500 distinct values) and the trackers they have encountered and performed cookie synchronization (e.g., 200 distinct values). The price keyword can take only 1 distinct value because each DSP is using one keyword, but the price value can obviously vary. Thus, we assume this price value can take 200 distinct values. In this hypothetical scenario, there is no aggregation so the surprisal, as expected, is high (> 60 bits).

Next, we assume that the tool performs aggregation of the distinct classes of some of the features, to explore a more realistic setup. For example, by just grouping the age into 5 distinct ranges, we reduce the surprisal bits by 5. Grouping the reported time of day in 3-hour intervals seems to have only a small impact to the overall surprisal rate. In contrast, by eliminating the cookie synchronization DSP, and reporting back only if CS took place (2 distinct values) we greatly reduce the overall surprisal by almost 7 bits. Also, if instead of reporting the 1st-party domain, we report only 30 distinct IAB categories, we reduce the surprisal by over 4 bits.

In the far right case, in which all possible features are aggregated, we can effectively halve the initial surprisal rate. But this can impact greatly the utility of our data. In this extreme case, no actual value is sent. Instead, the client only reports if the price was low, medium or high. This is the same case with the ad format, where the client reports if the ad was small, medium or large. Finally, instead of the actual country, the client reports only a country zone (e.g., Central EU, North America, etc.). The combination of distinct features and their impact on the surprisal rate can be seen in Table 4. From the previous threat analysis and the computation of surprisal bits assuming uniform distribution of classes, we can extract lessons on what features can be reported and with how many classes each.

7.3 Real-world data privacy analysis

Building on the previous theoretical analysis, we conduct an anonymity analysis of the reported features on a real user dataset. However, in this case, we employ the formula of Eq. 2, where the probability of a class (or level) in a feature is no longer uniform, but is computed using the probability distribution function of the feature. The dataset used is M , with 80k ad impressions

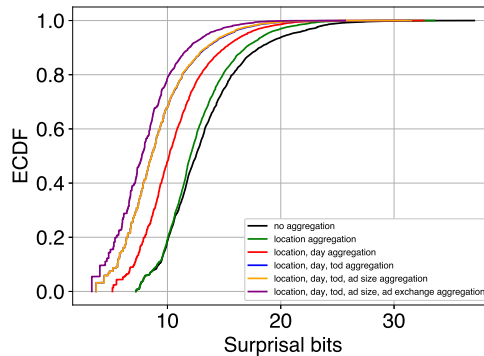


Fig. 17. Surprisal bits after incremental feature aggregation.

from 810 different users, each entry containing 27 different features. Out of these 27, we selected the ones reported back from *YourAdvalue* as well, ending up with a set of 8 features: user location, day of week, time of day, ad size, ADX involved, IAB category of 1st-party, price keyword and price value. Based on the classes in each feature, we calculate the probability of each class.

We start our analysis by examining a pessimistic scenario, which assumes an experimental setup where all features are reported at their classes with lowest probability observed. That is, we take the lowest probability class for each feature and compute the surprisal bits of this scenario. Indeed, the surprisal bits for a user who reports such unlikely event are about 60 bits. Interestingly, this number is very close to our previous theoretical worst case scenario, where no kind of anonymization was applied.

Next, we proceed to analyze the scenarios that did happen in the real dataset, and compute their surprisal bits. In Figure 17 (black line), we show the CDF of the surprisal bits for the observed setups. All setups have fewer than 37 bits, and 95% of the cases have fewer than 21 bits. Also, the median setup has 12.5 bits. Indeed, someone could argue that the high surprisal bit rate for a small portion of entries could expose some users. We remind the reader that in the absence of any sort of PII, the link to specific users can be difficult if not impossible inside this database.

Similarly with the analysis for uniform distribution, we can attempt to improve anonymity of users by performing aggregation of features to fewer classes. We demonstrate this effort in Table 5, where we apply incremental aggregation of each feature into larger groups or levels. For example, locations can be reported in higher abstraction such as districts instead of cities, days can be reported as weekdays vs. weekends, 4-hour slots can be grouped into larger bins, etc. In Figure 17, we show the CDF of the overall surprisal bits for the different setups included in the real dataset, and different aggregation efforts of Table 5. Interestingly, when we perform aggregation for the different features, we see a steady reduction of surprisal bits across all cases, besides the location aggregation which is effective for a portion of the ads. This is due to the fact that in such ads, the locations are popular and summarizing them does not have an impact in the surprisal rate. After applying aggregations in all classes, the surprisal rate of the median setup drops from 12.5 bits to 7.7 bits, and the 95% case drops from 21 bits to 13.7 bits.

Finding: With feature aggregation, the median surprisal bits under various distributions of classes (uniform or real) can be halved, in comparison to no aggregation scenarios. Also, location aggregation may not reduce user uniqueness as much as other features such as aggregation by time of day, or by day of week.

We continue this investigation by performing a k -anonymity [33, 34] analysis of this real dataset, to understand at what level it is satisfied, based on the different classes and aggregation performed.

Table 5. Feature class aggregation vs. surprisal bits (assuming uniform distribution of classes), and vs. performance of ML modeling.

Feature	Number of classes per feature					
	user location	184	26	26	26	26
day of week	7	7	2	2	2	2
time of day	6	6	6	2	2	2
ad size	17	17	17	17	3	3
ad exchange	149	149	149	149	149	15
cookie syncing	2	2	2	2	2	2
IAB category	25	25	25	25	25	25
Surprisal Bits	29.9	27.0	25.2	23.7	21.2	17.8
ML performance vs. feature aggregation						
AUC of ROC	0.858	0.854	0.820	0.798	0.798	0.786

K-anonymity is used as privacy criterion in real applications such as the “Family Educational Rights and Privacy Act” (FERPA) of USA [35], and the “Guidelines for De-identification of Personal Data” of South Korea [36]. We consider a realistic scenario where the attacker has limited access per ad detected, without any knowledge of the user’s browsing behavior (i.e., does not have IAB categories reported). In this scenario, the attacker groups the entries found in the dataset based on the features of location, day of the week and time of the day (i.e., does not have access to other features). In this attack scenario and the different setups tested, a minimum k-anonymity of $k=6$ users is satisfied, and a median entry (i.e., ad) can be mapped to $k=35$ users. In general, these scores are within the range of $k=3-10$ reported in [37] and applied for electronic health records, lending support to the applicability in our scenarios as well. In future versions of the tool, we will also consider other techniques such as Differential Privacy [38], that can provide more tight privacy guarantees to the end user.

Finding: With feature aggregation, a median (30-45)-anonymity can be achieved.

7.4 User Anonymity vs. Price Modeling

We close this investigation by looking into the inherent trade-off between aggregation of feature classes, and modeling of prices at the server using machine learning (ML) methods. We use a decision tree model and 4 equidistant classes for the RTB prices detected and measure AUCROC, a standard ML performance metric. We use this model with features aggregated at different levels, as examined earlier. The results are shown at the bottom of Table 5. We find that the AUC is not greatly impacted by the aggregation performed (8.4% decrease) when comparing the no-aggregation scenario (1st column) vs. the fully aggregated scenario (last column). These results show that *YourAdvalue* can do intense feature aggregation at the client before reporting to the back-end, without significantly affecting the server’s ML model performance. Note here, that another formal way to achieve anonymization of user data with provable guarantees is by applying techniques such as Differential Privacy (DP) [38]. In such method, Laplacian or Gaussian noise can be added on each feature to be reported, so that the probability of existence of individual rows into the dataset is obfuscated up to a certain, and system-controlled degree. However, the problem with such technique, in relation to our setup, is that the full distribution of each feature must be known, before proper noise level is selected for each feature to be protected. For the reconstruction at the server of the full distribution in an anonymous fashion, complex and costly methods such as the one used in EyeWnder [39] can be employed. Alternatively, emerging, privacy-preserving ML methods such as federated learning coupled with differential privacy [40] can be used, to allow the server the building of a global ML

model across users, without the need to report raw or aggregated user data, but only pre-models with added DP-noise. We will investigate such techniques in the future.

Finding: The *YourAdvalue* client can perform high feature aggregation before reporting data with a minimal impact on the performance of the ML price model used in the system.

8 RELATED WORK

In [9], authors explore the cost advertisers pay to deliver an ad to the user in RTB auctions and how the personal data can affect this cost. The authors proposed a methodology to compute the total cost paid for the user even when advertisers hide the charged prices. Finally, they evaluated their methodology by using data from a large number of volunteering users. They find that advertisers paid a total of around 25 CPM to deliver ads to the average user across a year. In [25], authors use the same methodology to measure the costs of digital advertising on both the user's and the advertiser's side in an attempt to compare how fairly these costs are distributed between the two. In particular, they compare the cost advertisers pay in RTB with the costs imposed on the dataplan, the battery efficiency and the privacy of the specific user. In [8], authors detected RTB notification URLs and extracted the value of the auctioned ad. They made an extensive study on the RTB ecosystem and estimated the value of user's private data based on the cleartext price notification URLs. They found that the average price of an ad is 0.0001\$-0.004\$, depending on the user and ad campaign.

In [41], authors use a dataset of users' HTTP traces and provide rough estimates of the relative value of users by leveraging the suggested bid amounts for the visited websites, based on categories provided by the Google AdWords. FDTV [11] is a tool to inform users in real-time about the monetary value of the personal information associated to their Facebook activity. Contrary to our work, FDVT is obtaining prices from the Facebook AdPlanner and thus is relevant only for Facebook advertising. In [22], authors used a heuristic-based mechanism to detect information exchanged between advertisers through CS. They concluded that 97% of the users are exposed to CS at least once and ad-related entities participate in more than 75% of the overall synchronizations. In [42] authors demonstrated how this technique may leak user's cookie IDs and browsing history to a snooping ISP even when user uses TLS and secure VPN services.

Bashir et al. in [31], studied the diffusion of user tracking caused by RTB-based programmatic ad-auctions. Results of their study show that under specific assumptions, no less than 52 tracking companies can observe at least 91% of an average user's browsing history. In [24], the same group tried to enhance the transparency in ad ecosystem with regards to information sharing, by developing a content agnostic methodology to detect client- and server- side flows of information between ad exchanges and leveraging retargeted ads. By using crawled data, they collected 35.4k ad impressions and identified 4 types of information sharing between ADXs.

In [43], authors discuss the value of privacy after defining two concepts (i) *Willingness To Pay*: the monetary amount users are willing to pay to protect their privacy, and (ii) *Willingness To Accept*: the compensation that users are willing to accept for their privacy loss. Two user-studies [44, 45] have measured how much users value their own offline and online personal data, and consequently how much they would sell them to advertisers. In [46], they propose "transactional" privacy to allow users to decide what personal information can be released and receive compensation from selling them. In [4], authors measure the prevalence, pricing dynamics and characteristics of a different type of programmatic ad-buying namely header bidding. In [47] the authors conduct an analysis of the header bidding advertising ecosystem using a set of real users. Header Bidding is a relatively newly introduced ad-buying protocol which gains popularity among advertisers. In our present study, we focus on the widely used RTB protocol namely *waterfalling*, since it is more mature and commonly accepted among all publishers. In [48], authors measured how much value

online publishers derive from behavioral ad targeting. Their results show that publishers make just 4% more revenue when including trackers in their websites.

9 ETHICAL CONSIDERATIONS

The execution of this work has followed the principles and guidelines of how to perform ethical information research and the use of shared measurement data [49, 50]. In particular, this study paid attention to the following dimensions.

Respect for Persons: We received *informed consent* from users, who were provided with information about our study and tool by visiting our page, and before downloading and installing our plugin. These users were *volunteers* and expressed their *comprehension* to the requirements for participating in the study and allow us to collect their anonymized data. We treated all volunteers *equally and with justice*, and without prejudice to any gender, age or other demographic characteristic.

Respect for Law and Public Interest: In accordance to the GDPR and ePrivacy regulations, and by designing our tool to be privacy-by-design, we do not collect any kind of identifiers that can de-anonymize users of the tool. Also, we do not share with any other entity any data of users collected by our tool, and especially any sensitive personal data that may have inadvertently collected without our knowledge. Furthermore, user data have been stored in secured manner, according to GDPR policies.

Beneficence: During the design of the experimental tool and the decision on what data to be collected, we assessed the *benefits to society* from the findings of our study and made significant efforts to reduce, or at least *mitigate any harms* that could be experienced by any individual participant. In fact, and as explained in our system design section, we made significant efforts to design a tool that does not impact the users' experience while browsing the Web, nor has adverse effects to the ad-ecosystem while intercepting the notification URLs sent by ad-exchanges to DSPs for charging and information purposes.

10 SUMMARY & CONCLUSION

In this paper, we presented the design, implementation and deployment of *YourAdvalue*, a first-of-its kind, full fledged system that allows any user to learn in real time while browsing how much money the RTB ecosystem pays to show them ads. We operated the system for a year, and it was used by 200 users. During this time period, we collected RTB-price metadata, and using past reported numbers for RTB ad-prices in mobile and desktop, in this paper, we make the following findings regarding RTB prices and their evolution over time:

- Desktop RTB prices have increased 4.6× within 7 years.
- During weekends and Mondays, the ad-prices tend to be higher than weekdays.
- Ad-prices are higher in the morning hours than afternoon or evening hours.
- The most popular publisher categories (using IAB) are News, and Technology/Computing.
- The most expensive publisher categories (using IAB) are Sports, and Food/Drink.
- Advertisers pay 2.4× more to target women than men.
- Younger users (25-34 years old) cost 2.5× more to be reached than older groups (35-44 years old).
- USA, Switzerland, and Spain have higher median ad prices across the 11 countries in our dataset.
- Performing cookie syncing before an RTB auction leads to increased median prices by 19%.

In addition, we performed a privacy evaluation of the system, to identify possible threats against user's anonymity. We measured the limits of user anonymity with a uniqueness study via surprisal and k-anonymity analysis. Finally, we studied the trade-off of anonymity via feature aggregation vs. performance of price modeling with ML methods. In summary, the main takeaways from the privacy evaluation of this tool were:

- *YourAdvalue*'s privacy-preserving design protects users from typical and extreme de-anonymization attacks.
- With feature aggregation, the median surprisal bits under various distributions of classes (uniform or real) can be halved to 7.7bits, in comparison to no-aggregation scenarios.
- Location aggregation does not reduce user uniqueness as much as other features (e.g., time of day or day of week).
- Feature aggregation enables k -anonymity with median $k=30-45$.
- *YourAdvalue*'s client can do high feature aggregation before reporting with minimal impact on the ML price model.

We have demonstrated that the *YourAdvalue* tool can successfully capture the state of RTB across the market, by verifying our findings (historical or not) with industry and other technical reports. Interestingly, *YourAdvalue* does this without jeopardizing the end-user's privacy, but instead offers a looking glass on RTB via cheap and distributed experimentation. We envision that in the future, the tool will be further used by many end-users, privacy researchers and auditors, who can take advantage of its simple functionalities to increase transparency in the RTB ad-ecosystem and its obscure practices of user modeling and ad-costs.

ACKNOWLEDGEMENTS

This project received funding from the EU H2020 Research and Innovation programme under grant agreements No 830929 (CyberSec4Europe), No 830927 (Concordia), No 871793 (Accordion), No 871370 (Pimcity) and Marie Skłodowska-Curie grant agreement No 690972 (Protasis). These results reflect only the authors' view and the Commission is not responsible for any use that may be made of the information it contains.

REFERENCES

- [1] eMarketer. emarketer releases new global media ad spending estimates. <https://www.emarketer.com/content/emarketer-total-media-ad-spending-worldwide-will-rise-7-4-in-2018>, 2018.
- [2] Lauren Fisher. Us programmatic ad spending forecast 2019. <https://www.emarketer.com/content/us-programmatic-ad-spending-forecast-2019>, 2019.
- [3] Ola Rask. What is programmatic advertising? the ultimate guide (2020). https://www.match2one.com/blog/what-is-programmatic-advertising/#Who_uses_programmatic, 2020.
- [4] Michalis Pachilakis, Panagiotis Papadopoulos, Evangelos P. Markatos, and Nicolas Kourtellis. No more chasing waterfalls: A measurement study of the header bidding ad-ecosystem. In *Proceedings of the Internet Measurement Conference, IMC '19*, page 280–293, New York, NY, USA, 2019. Association for Computing Machinery.
- [5] John Cook, Rishab Nithyanand, and Zubair Shafiq. Inferring tracker-advertiser relationships in the online advertising ecosystem using header bidding. *Proceedings on Privacy Enhancing Technologies*, 2020(1):65–82, 2020.
- [6] Automatad Team. Programmatic direct - everything you need to know. <https://headerbidding.co/programmatic-direct/>, 2019.
- [7] Mat Bennett. Private marketplace deals (pmps): The essential guide. <https://oko.uk/blog/private-marketplace-deals-pmps>, 2020.
- [8] Lukasz Olejnik, Minh-Dung Tran, and Claude Castelluccia. Selling off user privacy at auction. In *21st Annual Network and Distributed System Security Symposium, NDSS, San Diego, California, USA, February 23-26, 2014*.
- [9] Panagiotis Papadopoulos, Nicolas Kourtellis, Pablo Rodriguez Rodriguez, and Nikolaos Laoutaris. If you are not paying for it, you are the product: How much do advertisers pay to reach you? In *Proceedings of the 2017 Internet Measurement Conference*, pages 142–156. ACM, 2017.
- [10] Nikolaos Laoutaris. Data transparency: Concerns and prospects [point of view]. *Proceedings of the IEEE*, 2018.
- [11] José González Cabañas, Angel Cuevas, and Rubén Cuevas. Fdvt: Data valuation tool for facebook users. In *Proceedings of the 2017 CHI Conference on Human Factors in Computing Systems*. ACM, 2017.
- [12] Pushkal Agarwal, Sagar Joglekar, Panagiotis Papadopoulos, Nishanth Sastry, and Nicolas Kourtellis. Stop tracking me bro! differential tracking of user demographics on hyper-partisan websites. In *Proceedings of The Web Conference 2020, WWW '20*, page 1479–1490, New York, NY, USA, 2020. Association for Computing Machinery.

- [13] InvestingAnswers. Cost Per Thousand (CPM). <http://www.investinganswers.com/financial-dictionary/businesses-corporations/cost-thousand-cpm-2917>, 2017.
- [14] William Vickrey. Counterspeculation, auctions, and competitive sealed tenders. *The Journal of finance*, 1961.
- [15] Interactive Advertising Bureau. IAB Tech - Lab Content Taxonomy. <https://www.iab.com/guidelines/iab-quality-assurance-guidelines-qag-taxonomy/>, 2017.
- [16] Statista: the Statistics Portal. Internet advertising spending worldwide from 2007 to 2022, by format. <https://www.statista.com/statistics/276671/global-internet-advertising-expenditure-by-type/>, 2019.
- [17] eMarketer. US Real-Time Bidding (RTB) Digital Display Ad Spending, by Segment, 2015-2021 (billions). <https://www.emarketer.com/chart/228035/us-real-time-bidding-rtb-digital-display-ad-spending-by-segment-2015-2021-billions>, 2019.
- [18] Christi Olson. The desktop's death has been greatly exaggerated: How it's holding its own in a mobile world. <https://searchengineland.com/desktop-isnt-dead-desktop-holding-mobile-world-268151>, 2017.
- [19] Google. Mobile path to purchase: Five key findings. https://ssl.gstatic.com/think/docs/mobile-path-to-purchase-5-key-findings_research-studies.pdf, 2013.
- [20] BusinessOfApps. Mobile app advertising rates (2018). <https://www.businessofapps.com/ads/research/mobile-app-advertising-cpm-rates/>, 2019.
- [21] Kenshoo Social. Men are cheap: 12-month study of 65 billion impressions reveals men cost less to reach with facebook ads and respond better. <https://www.kenshoo.co.uk/menarecheappr/>, 2012.
- [22] Panagiotis Papadopoulos, Nicolas Kourtellis, and Evangelos Markatos. Cookie synchronization: Everything you always wanted to know but were afraid to ask. In *The World Wide Web Conference*, pages 1432–1442, 2019.
- [23] Gunes Acar, Christian Eubank, Steven Englehardt, Marc Juarez, Arvind Narayanan, and Claudia Diaz. The web never forgets: Persistent tracking mechanisms in the wild. In *Proceedings of the 2014 ACM SIGSAC Conference on Computer and Communications Security*, 2014.
- [24] Muhammad Ahmad Bashir, Sajjad Arshad, William Robertson, and Christo Wilson. Tracing information flows between ad exchanges using retargeted ads. In *25th USENIX Security Symposium*, pages 481–496, August 2016.
- [25] Panagiotis Papadopoulos, Nicolas Kourtellis, and Evangelos P Markatos. The cost of digital advertisement: Comparing user and advertiser views. In *Proceedings of the 2018 World Wide Web Conference*, pages 1479–1489, 2018.
- [26] Fotios Papaodyssefs, Costas Iordanou, Jeremy Blackburn, Konstantina Papagiannaki, and Nikolaos Laoutaris. Web identity translator. In *ACM HotNets*, 2015.
- [27] Noel Agnew Kevin Flood. Iab tech lab announces final content taxonomy v2 ready for adoption. <https://iabtechlab.com/blog/iab-tech-lab-announces-final-content-taxonomy-v2-ready-for-adoption/>, 2017.
- [28] Mopub. Iab categories. <https://developers.mopub.com/publishers/ui/marketplace/iab-category-blocking/>, 2019.
- [29] Steven Englehardt and Arvind Narayanan. Online tracking: A 1-million-site measurement and analysis. In *Proceedings of the 2016 ACM SIGSAC Conference on Computer and Communications Security, CCS '16*, pages 1388–1401. ACM, 2016.
- [30] Emmanouil Papadogiannakis, Panagiotis Papadopoulos, Nicolas Kourtellis, and Evangelos P. Markatos. User tracking in the post-cookie era: How websites bypass gdpr consent to track users. In *Proceedings of the Web Conference 2021, WWW '21*, page 2130–2141, New York, NY, USA, 2021. Association for Computing Machinery.
- [31] Muhammad Ahmad Bashir and Christo Wilson. Diffusion of user tracking data in the online advertising ecosystem. *Proc. Priv. Enhancing Technol.*, 2018(4):85–103, 2018.
- [32] Peter Eckersley. How unique is your web browser? In *Proceedings of the 10th International Conference on Privacy Enhancing Technologies, PETS'10*, page 1–18, Berlin, Heidelberg, 2010. Springer-Verlag.
- [33] L. Sweeney. k-anonymity: A model for protecting privacy. *IEEE Security and Privacy*, 10:1–14, 01 2002.
- [34] Panagiotis Papadopoulos, Antonis Papadogiannakis, Michalis Polychronakis, Apostolis Zarras, Thorsten Holz, and Evangelos P Markatos. K-subscription: Privacy-preserving microblogging browsing through obfuscation. In *Proceedings of the 29th Annual Computer Security Applications Conference*, pages 49–58, 2013.
- [35] US Department of Education. Family educational rights and privacy act (ferpa). <https://www2.ed.gov/policy/gen/guid/fpco/ferpa/index.html>.
- [36] Ministry of the Interior and Safety. Personal data protection laws in korea. <https://www.privacy.go.kr/eng>.
- [37] Hyukki Lee, Soohyung Kim, Jong Wook Kim, and Yon Dohn Chung. Utility-preserving anonymization for health data publishing. *BMC Medical Informatics & Decision Making*, 17, 2017.
- [38] Cynthia Dwork. Differential privacy. In *33rd International Conference on Automata, Languages and Programming (ICALP)*, 2006.
- [39] Costas Iordanou, Nicolas Kourtellis, Juan Miguel Carrascosa, Claudio Soriente, Ruben Cuevas, and Nikolaos Laoutaris. Beyond content analysis: Detecting targeted ads via distributed counting. In *15th CoNEXT*, pages 110–122, 2019.
- [40] Robin C. Geyer, Tassilo Klein, and Moin Nabi. Differentially private federated learning: A client level perspective. In *31st Conference on Neural Information Processing Systems (NIPS)*, 2017.

- [41] Phillipa Gill, Vijay Erramilli, Augustin Chaintreau, Balachander Krishnamurthy, Konstantina Papagiannaki, and Pablo Rodriguez. Follow the money: Understanding economics of online aggregation and advertising. In *Proceedings of the 2013 Conference on Internet Measurement Conference*, 2013.
- [42] Panagiotis Papadopoulos, Nicolas Kourtellis, and Evangelos P. Markatos. Exclusive: How the (synced) cookie monster breached my encrypted vpn session. In *Proceedings of the 11th European Workshop on Systems Security*, EuroSec'18, pages 6:1–6:6. ACM, 2018.
- [43] Alessandro Acquisti, Leslie K John, and George Loewenstein. What is privacy worth? *The Journal of Legal Studies*, 2013.
- [44] Juan Pablo Carrascal, Christopher Riederer, Vijay Erramilli, Mauro Cherubini, and Rodrigo de Oliveira. Your browsing behavior for a big mac: Economics of personal information online. In *Proceedings of the 22nd international conference on World Wide Web*, 2013.
- [45] Jacopo Staiano, Nuria Oliver, Bruno Lepri, Rodrigo de Oliveira, Michele Caraviello, and Nicu Sebe. Money walks: A human-centric study on the economics of personal mobile data. In *Proceedings of the ACM International Joint Conference on Pervasive and Ubiquitous Computing*, 2014.
- [46] Christopher Riederer, Vijay Erramilli, Augustin Chaintreau, Balachander Krishnamurthy, and Pablo Rodriguez. For sale : Your data: By : You. In *Proceedings of the 10th ACM Workshop on Hot Topics in Networks*, 2011.
- [47] Waqar Aqeel, Debopam Bhattacharjee, Balakrishnan Chandrasekaran, P. Brighten Godfrey, Gregory Laughlin, Bruce Maggs, and Ankit Singla. Untangling header bidding lore: Some myths, some truths, and some hope. In *Proceedings of the 21st International Conference on Passive and Active Measurement*, PAM 2020, 2020.
- [48] Alessandro Acquisti Veronica Marotta, Vibhanshu Abhishek. Online tracking and publishers' revenues:an empirical analysis, 2019.
- [49] David Dittrich and Erin Kenneally. The menlo report: Ethical principles guiding information and communication technology research. Technical report, 2012.
- [50] Caitlin M. Rivers and Bryan L. Lewis. Ethical research standards in a world of big. *F1000Research*, 3, 2014.

Received August 2021; revised October 2021; accepted November 2021