

To bid or not to bid?

Measuring the value of privacy in RTB

Lukasz Olejnik¹ and Claude Castelluccia²

¹ INRIA, Grenoble, France, Email: lukasz.olejnik@inria.fr

² INRIA, Grenoble, France, Email: claude.castelluccia@inria.fr

Abstract. The use of Real-Time Bidding (RTB) is expanding. In this paper we analyze this technology from the economics of privacy angle. We introduce a privacy and transparency enhancing tool showing to users when RTB ads are displayed on a visited Web site and how much is paid by advertisers. Using this tool we obtain a significant dataset of RTB winning bids (17,289 bids with an average value of \$0.0012) from real, highly profiled users; we detect about 70 bidders.

We detect and dissect the numbers of exchanged advertisements as well as the monetary spendings of bidders, and earnings of Ad Exchanges. This analysis is complemented by the study of a direct measurement resulting in an unprofiled dataset. We analyze flows related to Real-Time Bidding and detect hostnames of Ad Exchanges and Bidders. We discuss the bidding strategy upon of which depends the value of users' private data. Finally, we unveil a close cooperation between one of the Ad Exchanges and a number of publishers, which results in security risks and privacy leaks.

1 Introduction

Real-Time Bidding (RTB) is already a vibrant and ubiquitous technology allowing the display of advertisements to users based on decisions made in real time. When a user visits a Web site which supports RTB, the RTB system holds an auction: it sends *bid requests* to its bidders. The auction's bidders submit their bids and the winner displays its advertisement, later issuing a payment for this benefit. The advertisement is displayed on the publisher's Web site.

Bidders make decisions based on data obtained from Ad Exchanges (AdXs, RTB systems) during the auction phase. This data usually contains information on the user, for example his *currently-visited site*, inferred gender or ethnicity. Thus users' private data are being exchanged.

Advertising in RTB is driven by algorithms and it happens in short time frames, typically in tens of milliseconds. It is perhaps closely related to High-Frequency Trading [18], with the exception that the exchanged goods are *the users' data*. The Web is now a large channel of advertisements dissemination and this will

continue to be so in the future. RTB will play a key role in this process.

RTB is actively used in the building of detailed profiles of users [21]. In a similar vein to all the other Big
According to a recent TRUSTe report, users' concerns over privacy is a growing trend, with 92% of US users expressing concerns. In particular, 58% of people recognize businesses sharing their data as a major problem to privacy. As a consequence, 83% of users are less likely to click on online ads and businesses will likely have to carefully study the users' perception of value of privacy [27].

Due to the users' data being a commodity and this process taking place in real time it is straightforward to notice the privacy risk [21]. Evaluation of this technology with privacy in mind is thusly of great importance and can bring significant insight into societal changes. This paper builds on the results of [21] in order to better understand the economics of privacy in RTB. Understanding the economics of privacy is important to the users, as well as businesses, law and policy makers. Businesses can align their services with users' expectations of privacy. Law-makers can shift the law accordingly. Privacy regulators can find the right balance and trade offs in their policies [2].

The main contributions of this work are summarized as follows:

- *RTB Tracker*. RTB is inherently nontransparent. Users do not know if and which of their data are being sent to the bidders. We introduce a transparency enhancing tool (TET) which shows to users when a RTB auction takes place (section 2.3). We also leverage RTB *price notification channel* to extract the prices paid by advertisers to display ads to a given user. The user can get access to these values, as well as to the average price. He can also compare *his value* (average price of his ads) against the values of other users, and see how "valuable" he is to advertisers, with respect to others.
- *Value of Privacy*. We leverage our RTB tracker tool to study pricing in Real-Time Bidding and the *value of privacy* for profiled users. We show that bidders value users differently according to their profile, and that the "average user value" (average of the advertisements' prices for a given user) varies from \$0.0001 to \$0.004 (section 3.2).
- *Advertisers' Bidding Strategies*. We study and compare advertisers' bidding strategies for profiled and unprofiled users. We detected that strategies differ significantly from one advertiser to another. Some companies (e.g. Media Math) serve ads very frequently, while others (e.g. AppNexus) tend to serve ads less frequently but bid higher prices for profiled users (section 3).
- *Understanding RTB Internals*. RTB is a technology which requires an interplay between Ad Exchanges auctioning advertising spaces, bidders who compete for them, and publishers who actually provide this space on their Web sites. We investigate the interplay between these three parties. We analyze interactions between Ad Exchanges and bidders (section 3). We also discuss cooperation between publishers and Ad Exchanges. To our surprise, we found evidence of strict *Ad Exchanges–Publishers* collaborations. In the encountered cases, this cooperation brings severe consequences to Web users, both in terms of privacy leaks and security risks. We analyze these practices and discuss the consequences.

1.1 Related Work

Real-Time Bidding is a new technology of serving advertisements [12, 7] which is quickly gaining market share [15]. To date, few research works directly address this problem. For example Yuan et al. [32] provide insight to pricing and a general analysis of the technologies. The authors use internal data of one Ad Exchange and they study the general behaviors of bids. However, they do not address privacy aspects.

Ad Exchanges facilitate the selling of advertising spaces and they serve as connectors between publishers and bidders. Muthukrishnan discusses them in [20], Yuan et al. describes Ad Exchanges, Bidders, Publishers their relationships in [31].

The work of Gill et al. studies money flows in the advertising business [11]. The authors use a dataset of users' HTTP connections. They estimate revenues by assuming the suggested bids prices of Google AdWords Contextual Advertising Tool assigned to certain categories of sites. The authors argue that if the users are willing to employ privacy protection techniques, the revenue of advertisers would significantly drop, by at least 30%". In our paper we use real bids which are based on the sites visited by users as well as the users' profiles.

Privacy issues of Real-Time Bidding are addressed in [21]. The authors describe how RTB works, describe the involved business model and establish lower bounds on the user's private data cost from the bidders point of view. They conclude it might be as low as \$0.0005. We build on their results and significantly and methodically complement their work by performing a more complete analysis of Web advertising ecosystem. We perform measurements and we study interactions and flows between Ad Exchanges and bidders. We discuss the fact that bidders bid for users visiting certain Web sites. Bidders obtain private data on these users by participating in RTB auctions. Moreover, we analyze a large corpus of user-supplied dataset on their profile evaluations from the bidders' perspectives. Our work provides insight into transparency of RTB in that we evaluate the value of privacy, we enlist the known Ad Exchanges that we encountered in our studies, as well as the bidders. Finally, we divulge and discuss a puzzling cooperation between an Ad Exchange and publishers and highlight the resulting risks.

Analysis of pricing in auctions is closely related to the value of privacy. Acquisiti et al. address the problem [2] using an economical framework of Willingness-to-pay (WTP) and Willingness-to-accept (WTA); both are based on user's perspectives. The authors perform a number of experiments, one resulting in a WTA/WTP ratio of 5.47. The authors point out that this ratio is significantly larger than in the case of same metric applied to ordinary goods. It is important to note that users might not be in a position to decide in privacy-sensitive cases due to lack of education [1].

Carrascal et al. evaluate the user's perception of the value of items in browsing history to be relatively high (7 EUR) [6]. Recently, Savage and Waldman estimate the user's perception of privacy value by studying how much users are willing to pay to conceal certain private data. Smartphone users are willing to pay \$2.28 to

conceal their browser history, list of contacts (for \$4.05), physical location (for \$1.19), etc. [26].

In our work, we provide a detailed view from the advertisers' perspectives, which unambiguously complements previous works on value of privacy.

1.2 RTB and Price Notification

When a user visits a *publisher's* Web site which supports RTB, the *RTB system* (*Advertising Exchange, AdX*) holds an *auction* for this *ad impression*. The auctions are sealed-bids according to Vickrey principles [29], where the second largest price (increased by a constant) is to be paid by the winner. The auction participants are composed from *bidders* who bid on behalf of the advertisers and/or perhaps different Ad Exchanges.

During the auction, *RTB systems* send *bid requests* to the *bidders*. These requests can contain information on the user, such as the *visited site*, the *IP address* (or its parts) in addition to others, even such as *ethnicity*³ and *income*; therefore the user's private data are being transferred to the participating parties. Participants of the auction appraise the received data and submit their bids to the auction holder. The whole process typically takes less than 100 milliseconds. The winner's advertisement is then displayed in the user's browser.

The user's browser requests the advertisement via a standard HTTP GET request of an URL. The request is executed by a script present in the winning bidder's ad snippet, supplied by the RTB. That ad snippet contains the elements such as HTML tags, responsible for performing these requests. This request's URL very often carry a price notification. The price notification is meant to inform the winner about the monetary value to be paid for the displaying of this advertisement, the *winning bid*. The URL is therefore a HTTP request of the form: *http://bidder.com?price=encrypted-price*, where *price* is the name of the price-containing parameter, and *encrypted-price* is the actual price notification.

The notification is very often in an encrypted form, conforming to the industry standard, the so-called *encrypted price* [13]. Some notifications are, however, sent in clear-text forms and the actual request might be of the form: *http://bidder.com?price=0.42*. In this example the paid price would be 0.42, expressed in Cost-per-Mille: the actual value in currency can be obtained by dividing this by 1000, \$0.00042 in this case. The price-containing parameters are identical between various ads requests and are directly related to a particular bidder. This means that for a specific bidder *A*, its price-containing parameter *P* is always the same whether the value is encrypted or not. This simple observation enables the detection of a list of bidders utilizing RTB systems, along with their *price-parameters* and in the end it is possible to detect the clear-text prices paid for the users in real time, the so called *clear-text price notification leak* [21].

Figure 1 shows a schematic diagram of interactions and involved parties.

³ OpenX can provide the following ethnicity information: *African American, Asian, Hispanic, White* and *Other* in its bid request




Fig. 1: Parties involved in RTB.

2 RTB Study: Methodology and Datasets

2.1 Goals of the Study

Our aim was to uncover the interplay between AdXs, bidders and publishers. We used *price notifications* described in [13] to detect that RTB auction took place. This, in turn, allowed us to make deductions about the interactions that happened between the auction holder (Ad Exchange) and the participants (bidders), *prior* to displaying the ad. It provided insight of the intensity and nature of communications between these parties.

The use of *clear-text price notification information leak* channel [21] allowed us to analyze the *value of privacy*. We leveraged two quantities: numbers of exchanged advertisements and the prices paid for them, to study the communication flows between Ad Exchanges and bidders. In the analysis we used two datasets composed from *unprofiled* and *profiled* bids.

2.2 Unprofiled Dataset: Direct Measurement (*Dataset₁*)

The first dataset was obtained using the following method of a *direct measurement*. We selected the top 5,000 sites from the Top 1M Alexa sites (www.alexa.com/topsites) that contained scripts enabling RTB ads of Doubleclick [9]. We then visited the sites with PhantomJS 65 times, as in [21], and extracted the 500 sites with the highest occurrences of clear-text price notifications. Our intention was to maximize the likelihood that an ad was served with a clear-text price notification, with a goal of studying flows between the involved parties. We then consecutively visited these sites 60 times. We saved all HTTP data transferred in this process, such as all the relevant headers belonging to the requests and responses, as well as the actual response bodies. We then extracted from this data, the information on the Ad Exchange, the winning bidder, and the winning bid of each visited site. All the experiments were performed in October and November, 2013. We denote this dataset as *Dataset₁*. We detected 23,344 price notifications in total, 32% of them were clear-text and 68% were encrypted.

We reseted any tracking means such as browser cookies after each visit of a site. So a visit was always treated as a *newly-seen user*, with no known tracking cookies; the employed technique is similar to the one in [21]. We did this to

ensure the prices were not affected by previous history. *Dataset₁* is composed from *unprofiled* bids.

2.3 Profiled Dataset: User-supplied Dataset (*Dataset₂*)

The [ANONYMIZED] plugins and project. We run a project aimed at studying the *value of users private data* from the bidders' perspectives. We created a transparency enhancing technology (TET) which shows to users the monetary value bidders in Ad Exchanges' auctions pay for a visit to a Web site. This TET leverages RTB *price notification channel* to extract the prices paid by advertisers to display ads to a given user.

We implemented the discussed transparency enhancing technology in form of Firefox and Chrome browser extensions. The extensions scan all the outgoing requests initiated by the user's browser, with a single aim. All the values of the requests' parameters are inspected against the industry-standard encrypted price notification, included in HTTP requests of advertisements served in Real-Time Bidding. Upon the plugin detecting a price notification, it saves the hostname of the bidder and the name of price-containing parameters. Subsequently, during browsing the plugin analyzes HTTP requests for the known bidders' hostnames. If a known hostname is detected, the plugin verifies whether a price-containing parameter is present and if it is a floating-point number. This value is then saved and displayed to the user as it is the winning bid the bidder is paying for displaying its ad to the user (see Figure 2 for an example).

The user can also get access to the average price of his advertisements. He can also compare *his value* (average price of his ads) against the values of other users, and see how "valuable" he is to advertisers, with respect to others, by visiting the following Web site: [ANONYMIZED] (see Figure 11 and 10 for examples).



Fig. 2: A notification about the price paid for an advertisement during the visit to www.accuweather.com. Price in CPM.

***Dataset₂* description.** We publicly advertised our project. As a result, the plugins were downloaded more than 300 times. In this paper, we analyzed data of 139 users, collected from September 2013 and January 2014. These users came from 20 countries, although the most significant numbers of them were from France (36.4%), the US (20.7%), Netherlands (15%) and Germany (7.9%). As a result of this experiment, we collected, for each user, the bidders' hostnames and their winning prices. We detected and analyzed 17,289 clear text price notifications in this dataset.

Privacy Considerations We strictly and rigorously considered and respected the privacy of users. Most importantly, our browser extension does not collect or save the user-visited sites. The only data collected by our extension are the bidders’ hostnames, paid prices and examples of encrypted prices and price-containing parameters. Moreover, the plugin is available from a Web site which clearly describes its functionality. Finally, users have to explicitly install it and therefore have to consent to it.

3 Bidders Analysis

This section analyses the differences in bidders’ strategies, i.e. how they issued the bids, using the unprofiled ($Dataset_1$) and profiled ($Dataset_2$) datasets. Although obviously we could not see the contents of encrypted bids, we do not believe they should differ from clear-text ones, in general. There is no reason for the bids to vary on the same site just on the basis of them being encrypted or not; encryption and clear-text forms of prices are just a matter of implementation in the Ad Exchanges’ and bidders’ systems. In fact, in the clear-text bids analyzes we detected that the bidders know how to attribute prices to specific sites. Specifically, different bidders tend to issue similar bids to display ads to users on the same sites.

3.1 Winning Price Analysis of Unprofiled Users ($Dataset_1$)

Table 2 refers to ads served in the RTB process that we encountered during our survey. In the case of $Dataset_1$, we detected 9 bidders which received clear-text price notifications and 13 bidders which received encrypted price notifications. Price notification is a mechanism on the Ad Exchange side, so bidders often receive them in different formats.

Winning Prices per advertisers. Second column of Table 2 shows the average price related to a specific bidder. It is striking to see that certain bidders are very active, for example Media Math (mathtag) served 33% of ads with clear-text price notifications, although the average bid price was low (0.25 CPM) and \$0.6 was paid for these ads in total. Others are more selective, for example AppNexus (adnxs) only served 9% of ads with clear-text prices and paid \$0.4 in total, during our measurement. AppNexus is clearly targeting more specific audiences (i.e. Web sites) in that less ads are served, but average winning bid price is higher (0.52 CPM). Just these two cases highlight how significantly different the strategies employed by the bidders are. It is interesting to note that Media Math also received the largest number of ads with encrypted price notifications (data not shown): we detected 6761 of these kinds of events; total numbers of exchanged ads are displayed on Figure 8b. If we assume the average price paid by Media Math in this case was identical to the one drawn from clear-text prices, it is possible to compute how much Media Math paid in total. Media Math served $2467 + 6761 = 9228$ ads and the known average CPM winning bid price per ad was 0.246 ($Dataset_1$ column in Table 2). This means Media Math

paid \$2.27 for serving these advertisements and acquiring data behind user’s Web browsing habits. We can conclude here that Media Math is tied to RTB providers by a large number of auctions it is involved in. On the other hand, AppNexus won a significant proportion of “*more valuable auctions*”: the average winning bid was higher, but the numbers of issued ads were lower.

Winning Prices per sites. We aimed to study the factors which determine bidders’ estimations of a given Web site. In other words, we aimed at understanding what is responsible for the differences in the average bids in Table 2 (*Dataset₁*).

The conclusion in [21] highlights the importance of the visited Web site’s content. But bidders still pay different figures on average. An underlying assumption behind the work from [19] is that bidders should know how to evaluate sites. In order to preserve a fair market in RTB in the presence of Cookie Matching [21] technologies, Web sites should be evaluated comparably by different bidders.

With the aim of studying this, we selected three bidders with substantial differences in average bids, for which we collected a significant number of bids: Doubleclick, Media Math, and Turn. For each site where ads of at least two of the analyzed bidders were seen, we computed the average price Avg_{th} issued by these bidders. We also computed the per-bidder, specific averages Avg_i , for $i = \{Doubleclick, Media Math, Turn\}$. During the analysis, we only considered bidders for which we detected more than 15 prices on a given site. So for some sites, it happened that we compared prices of only two bidders. The plotted data displaying a general average (blue line) for 68 of analyzed sites is shown on Figure 3. The dots on this figure relate to per-bidder averages on a specified site (X axis). In most of the cases, the absolute differences between the per-bidder average and the general average were very small. In fact, in case of 38 (56%) sites the absolute difference was smaller than or equal to 0.05, and in case of 52 (76%) sites, the absolute difference was smaller than or equal to 0.1. Since cookies were systematically cleared before visiting each site, bidders did not know the previously-visited sites. But different bidders still knew how to “evaluate the sites”. This balance likely emerges as a result of real-time auctions. If the bidders fail to win a satisfactory number of auctions, they might increase their bids in the future auctions. On the other hand, winning too many auctions might mean that they started to overpay. The value of advertisement spaces in RTB is volatile. The consequence for privacy is that *the value of private data also follows this trend* and is also subject to dynamic changes.

Winning Prices and Do Not Track. We tried to study whether Do Not Track (DNT) header affects pricing. For this, we treated the 68 sites from the previous paragraph as a *profile*. We then consecutively visited sites from this profile *without* removing cookies to allow the bidders to track this artificial profile. We performed this experiment 144 times (i.e. visited the full set of sites in this profile 144 times) considering two cases, when DNT header was on, and off.

We did not find any differences for profiles enabling DNT. The average winning bid in the case DNT header is off was 0.328 (std: 0.757, number of analyzed bids: 8804), and when we set DNT to on, the average was similar: 0.326 (std:

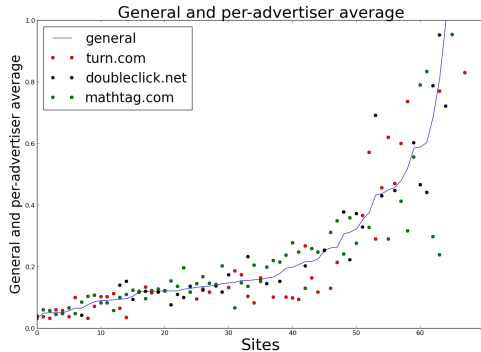


Fig. 3: Per-bidder averages are close to the general average on the studied sites ($Dataset_1$).

0.660, number of analyzed bids: 8706). Bidders do not evaluate the profiles basing on the DNT header, and this might be a result of the fact that in the bid request templates of the analyzed Ad Exchanges we did not find any indications of DNT header. So bidders do not know the value of DNT header during the auction.

3.2 Winning Price Analysis of Profiled Users ($Dataset_2$).

User Profiles and Winning Prices. We first computed, for each of the 139 users that participated in $Dataset_2$, their “average values”, i.e. the average of the prices paid for advertisements displayed to them during their Web browsing. The results are shown on Figure 4 (red line). Most of the users had an “average value” of less than \$0.001, but many were significantly more “expensive”.

It can also be seen from this figure that the number of prices varies between users (blue line). This is a consequence of the uncontrolled nature of the experiment: we did not encourage users to visit any pre-defined set of sites per day. We believe this led to a more natural behavior of users. In fact, some of the users kept submitting prices over several months. Users were profiled by advertisers. They did not use any ad blocking extensions, otherwise our plugin would have been unable to detect any price notification.

User Location and Winning Prices. Table 1a displays the per-bidder averages in three countries. It is seen that the largest prices for displaying ads, and acquiring users’ data, were paid in the US and this is especially acute for Doubleclick and Criteo; bidding strategies are clearly affected both by content of sites, countries and users. The averages in case of data from France and Netherlands were lower, although they still differed significantly. It is important to remember that the data are of profiled users so the bids were affected by the users’ past behaviors to large extents.

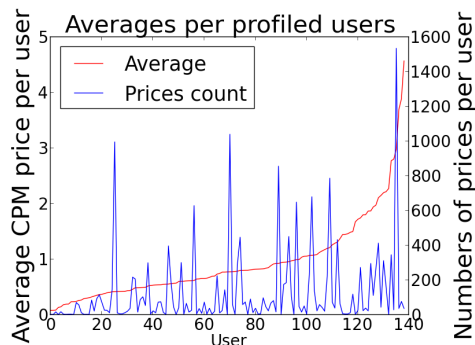


Fig. 4: Average per-user winning bid value and numbers of prices (*Dataset₂*).

Table 1b shows the average bids of users from three countries: France, The US and Netherlands. In general the prices in the US are very high. This likely reflects the fact that the RTB market in the US is more mature than else-where.

Bidder	FR		US		NL	
	Avg	Cnt	Avg	Cnt	Avg	Cnt
adnxs.com	0.68	581	1.98	406	1.5	1700
turn.com	0.95	847	1.35	269	1.3	499
mathtag.com	0.57	292	1.31	171	1.69	451
doubleclick.net	1.15	3204	2.07	271	1.23	253
criteo.com	0.61	1462	3.82	202	0.82	414

(a) Per-bidder.

Country	Avg	Std	Cnt
FR	0.92	1.23	7342
NL	1.38	1.47	3669
US	2.24	2.04	3074
all	1.2	1.53	17289

(b) General.

Table 1: Per-country analysis (*Dataset₂*).

Bidders and Winning Prices. We detected about 70 bidders as a result of the [ANONYMIZED] experiment, and 34 of those bidders received clear-text price notifications. A selection of them is shown in the third column of Table 2. Comparing data from the two datasets suggests an immediate conclusion that the influence of profiling, for example the past-visited sites, significantly affects the average bids. For example Media Math (mathtag.com) paid 0.246 CPM on average in case of unprofiled *Dataset₁*, but this increases about 5 times, to 1.29 in the case of profiled *Dataset₂*. The unprofiled experiments were performed from servers located in France, so it may be more appropriate to compare them with data in Table 1a, which shows results for French users of the experiment: in this case Media Math still bids more than 2-times higher than in the unprofiled case.

We highlight the fact that Doubleclick is a very active ads provider. From Table 2, it is responsible for over 15% of ads in *Dataset₁* and even more (27.1%)

in *Dataset₂*. This could be a pure coincidence as perhaps real users browsed sites of different nature. On the other hand, in case of Criteo (criteo.com), which specializes in retargeting, it is understandable that the bidding is always done for profiled users. Consequently, *Dataset₁* does not contain Criteo as it was not only an unprofiled measurement, but during the visiting of the sites we also visited only the main pages of the analyzed Web sites. So it was not possible to visit a site of any product which could then be advertised in retargeting. Similar comment applies to Sociomantic (sociomantic.com). Some of the bidders are only interested in profiled users and this confirms that any privacy study of Real-Time Bidding should be done from these two angles. The *Inc* column in Table 2 is a factor showing how many times the profiled (*Dataset₂*) average is larger than the unprofiled one (*Dataset₁*).

Bidder	<i>Dataset₁</i>			<i>Dataset₂</i>			Inc
	Avg	Std	Count	Avg	Std	Count	
rfihub.com	0.15	0.07	5 (0.1%)	1.54	1.95	300 (1.74%)	10.3
creativecdn.com	0.16	0.1	862 (11.5%)	0	0	0	n/a
mathtag.com	0.25	0.34	2467 (32.9%)	1.29	1.51	1267 (7.3%)	5.2
turn.com	0.37	0.51	2228 (29.7%)	0.89	1.43	2375 (13.7%)	2.4
adsrvr.org	0.42	0.33	41 (0.6%)	2.40	2.74	169 (1%)	5.7
doubleclick.net	0.45	0.69	1159 (15.5%)	1.09	1.37	4681 (27.1%)	2.4
w55c.net	0.51	0.38	26 (0.4%)	1.19	0.89	120 (0.7%)	2.2
zenoviaexchange	0.52	0.22	30 (0.4%)	0	0	0	n/a
adnxs.com	0.52	1.66	667 (8.9%)	1.27	1.45	3208 (18.6%)	2.4
adroll.com	0	0	0	0.49	0.54	387 (2.24%)	n/a
invitemedia.com	0	0	0	0.53	0.42	144 (0.8%)	n/a
pswec.com	0	0	0	0.84	0.53	167 (1%)	n/a
criteo.com	0	0	0	0.89	1.41	2332 (13.5%)	n/a
dotomi.com	0	0	0	2.62	1.82	876 (5.1%)	n/a
sociomantic.com	0	0	0	2.78	2.20	127 (0.7%)	n/a

Table 2: Winning bids of selected bidders. CPM averages (Avg), standard deviations (Std), numbers of winning bids notifications (Count, and percentage in a dataset), increase coefficient (Inc).

3.3 Bids Frequency Analysis

We also analyzed the variability of bids of each bidder, in order to study the differences between their bidding strategies. The probability mass functions on Figure 5a shows the differences between the bidders. We created this analysis for a selection of those bidders for which we were able to obtain substantial numbers of prices. The X axis on this figure shows bids intervals, while Y axis refers to the frequency of particular bids. The frequency peaks of low bids (near the point $X = 0.2$) lie very close in the case of Creative (creativecdn.com), Media Math

(mathtag.com), Turn (turn.com) and Doubleclick (doubleclick.net) and most of them were low. AppNexus and Doubleclick, however, also target more selective audiences (visitors of more high-profile sites) so they tend to bid higher. This is reflected in spikes close to the points $X = 0.4$ and $X = 0.45$ for AppNexus and Doubleclick respectively. Interestingly, these bidders either bid low ($0.05 - 0.1$) or high (0.4), largely excluding the interval between these points. It is then clear that different bidders target different users. For example AppNexus is often not bidding in the “less expensive” auctions.

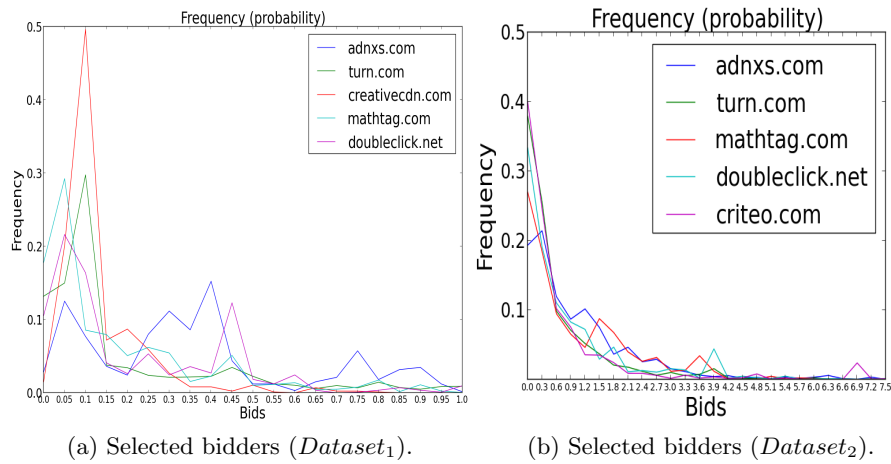


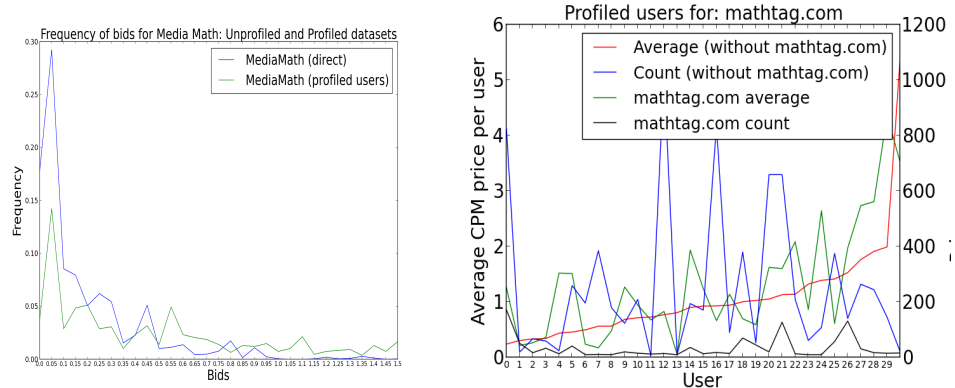
Fig. 5: Probability mass function for prices of selected bidders.

It is also interesting to discover the frequencies of bids for profiled users ($Dataset_2$), which are shown on Figure 5b. We observe a number of interesting differences to the unprofiled case. In overall, bids are higher. It is also clear that Media Math in this case bids higher as well (peak close to $bid = 1.5$ CPM): these results complement the ones shown on Figure 5a.

This case is interesting enough that we also present the Figure 6a which shows bids of Media Math for two cases: profiled and unprofiled. The average winning bid of Media Math was 5 times higher for profiled users, i.e. 1.3 CPM. Media Math’s strategy was clearly affected by profiling and it justified bidding in the “more expensive” auctions for the known users. Figure 6b shows the per-user average bids of Media Math as well as the general averages excluding this bidder (data shown only for users with at least 10 bids of Media Math). It is overall seen that bids of Media Math (green line) were higher than the average (red line).

Media Math mass bids on unprofiled, *previously unknown* users. Consequently, Media Math can monitor a significant fraction of Web traffic (via ob-

servation of large numbers of bid requests), and as a result, it can track large numbers of users in general. Media Math is also able to profile these users well enough to justify bidding in the auctions for known users.



(a) Bids frequency (FR, $Dataset_1$ and $Dataset_2$) (b) MediaMath prices of users ($Dataset_2$).

Fig. 6: Bids of Media Math.

Bid prices in RTB are determined via real-time algorithms that are affected by various factors related to users, for example, their location or previously visited sites [21]. This means that the value of users’ private data (*Web Browsing History*) is dynamic in nature. We browsed the sites of the bidders listed in Table 2 and saw that most of them explicitly specify that they support serving of both video ads and rich media content. So the format of ads contents may not necessarily be among the key factors and what is likely decisive and possible to observe by a network measurement is how they value the users. As argued in [21] previously visited sites matter and for example retargeted ads tend to be substantially more valued than the “generic” ones.

4 Ad Exchanges Analysis

In this section we focus on Ad Exchanges. We use $Dataset_1$ exclusively.

4.1 Ad Exchanges Activities and Incomes

We analyzed HTTP streams resulted while visiting the Web sites. If we detected an ad snippet (usually a HTML or JS code) containing a known bidder’s hostname (e.g. *turn.com*), its price-containing parameter (e.g. *acp*) or the actual price notification (e.g. 0.4584 for clear-text values), we concluded that the

AdX	Avg	Std	Count (%)
casalemedia.com	0.08	0.2	771 (10.3%)
lijit.com	0.22	0.34	1468 (19.6%)
360yield.com	0.22	0.16	153 (2%)
criteo.com	0.32	0.13	39 (0.5%)
zenoviaexchange.	0.34	0.33	185 (2.5%)
pubmatic.com	0.35	0.5	3604 (48.1%)
contextweb.com	0.42	0.3	92 (1.2%)
adnxs.com	0.46	0.52	727 (9.7%)
burstnet.com	0.49	0.3	276 (3.7%)

Table 3: AdXs serving clear-text prices (avg, std dev, numbers of bids and their overall share (%)), *Dataset₁*.

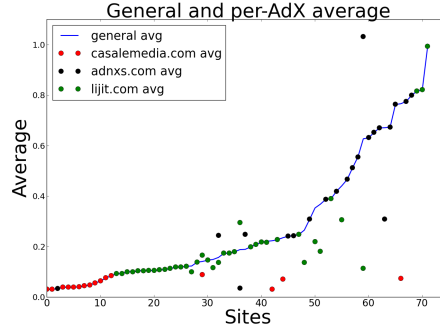


Fig. 7: Average winning bids for selected Ad Exchanges, *Dataset₁*.

hostname (e.g. *ad.doubleclick.com*) which provided this ad snippet was an Ad Exchange.

Table 3 displays the average price of ads served by a given AdX (data shown for AdXs which sent more than 10 winning price notifications). It can be seen, that both average AdXs income and numbers of transferred ads (advertising slots acquired by the bidders) differed significantly for various AdXs. For example PubMatic (*pubmatic.com*) sold a considerable number of advertisement spaces (3604) for a relatively medium price (0.354) and thus gained $\frac{3604 \times 0.354}{1000} = \1.28 for them, but has a lower average income, when compared to AppNexus (*adnxs.com*). AppNexus sold its ads slots for an average price of 0.46 CPM and the resulting income was \$0.33.

The incomes of Ad Exchanges are directly dependant on the winning bids of the bidders. As discussed in the previous section 3, bidders bid according to the sites and profiles. We performed a similar measurement for Ad Exchanges. We analyzed the winning bids paid to AdXs for ads on sites where we observed at least 10 bids attributed to an AdX. This analysis of dependency of prices on Web sites is displayed on Figure 7. It shows the general average (blue line) and per-AdX average for Casalemedia (red dots), Lijit (green dots) and AppNexus (adnxs.com; black dots) Ad Exchanges. It is seen that Casalemedia is holding auctions on *less valuable* Web sites, from the bidders' perspectives. On the other hand, Lijit and AppNexus are active on more valuable sites. Consequently, the income of AdXs depend on the bidders strategies and their ability of profiling the users.

We detected 7 Ad Exchanges that served encrypted price notifications: *bidsystem.com* (42 price notifications sent), *360yield.com* (196 resp.), *contextweb.com* (647 resp.), *openx.net* (696 resp.), *doubleclick.net* (3,916 resp.), *rubiconproject.com* (10,083 resp.). 15 Ad Exchanges served clear-text price notifications (Table 3). Doubleclick is also a bidder and consequently it is present in the Table 2 (section 3). As is seen on Figure 8b, Doubleclick also very often bids in its own auctions.

4.2 Ad Exchanges-Bidders Interaction Analysis.

During ads exchanging in RTB, bidders buy advertising spaces on the publishers' Web sites. In order to study the relationships between the ad slots suppliers (AdXs) and the bidders, we analyzed our data to expose these exact associations. Specifically, if we detected that for example PubMatic (pubmatic.com) supplied an ad snippet which then resulted in requests to Turn (turn.com, the bidder), we treated it as a relationship (flow): *pubmatic.com* \rightarrow *turn.com*. We performed two studies. First, we studied the average price paid by each bidder to Ad Exchanges. Second, we studied, for different bidders, the total number of ad snippets delivered by selected Ad Exchanges.

We have detected 40 flows between AdX's and bidders in the case of clear-text price notifications. The average CPM winning bid seen in these interactions could be as low as 0.051 and as high as 1.546. In the case of encrypted-price notifications we detected 41 flows. The numbers of exchanged advertisements could be anecdotal (9 for doubleclick.net \rightarrow rfihub.net) or as high as 5,634 (rubiconproject.com \rightarrow mathtag.com): MediaMath is a grand customer of Rubicon Project's service.

Money Flows. The graph on Figure 8a shows the communication flows. The rectangular shape of nodes means that this party could act as an AdX. For clarity reasons, only flows corresponding to more than 200 sent ads are shown. Direction of the edges shows the money flow: the average price paid by bidders to Ad Exchanges is shown on the edges labels. We can see that for example Turn paid 0.46 CPM to AppNexus (adnxs) on average, but AppNexus also paid 0.346 CPM to Lijit.

Ads Flows. The graph on Figure 8b displays the total numbers of exchanged ads, both encrypted and clear-text. Data is shown only if the total number of such events exceeded 400. Both Media Math and Turn were very active actors. They bought advertising spaces from many Ad Exchanges and as a consequence, these two bidders saw large numbers of bid requests.

Ad Exchanges have large numbers of bidders. For example we detected 11 unique winners of Doubleclick's auctions. We stress here that the numbers of detected bidders are not complete. We were only able to see the RTB auction's winning bidder. However, by analyzing the outcomes of multiple auctions we have seen a considerable number of bidders. Although the lists of bidders are certainly not complete, we treat them as a starting point in introducing transparency in RTB. Most RTB operators do not disclose the list of their bidders⁴. Our work highlights a large portion of the most active ones.

From Figure 8b we can also see the numbers of bidders involved in certain Ad Exchanges' auctions. For example Turn (turn.com) won an advertising space at 12 Ad Exchanges (we show only 7 for clarity reasons). Turn is definitely an active bidder in a large number of auctions. Similar comment applies to AppNexus (adnxs.com) and Media Math (mathtag.com). They serve as good examples of *major profilers*: MediaMath, Turn and AppNexus see large numbers of bid requests containing user-related data.

⁴ Pulse Point publishes a subset [24].

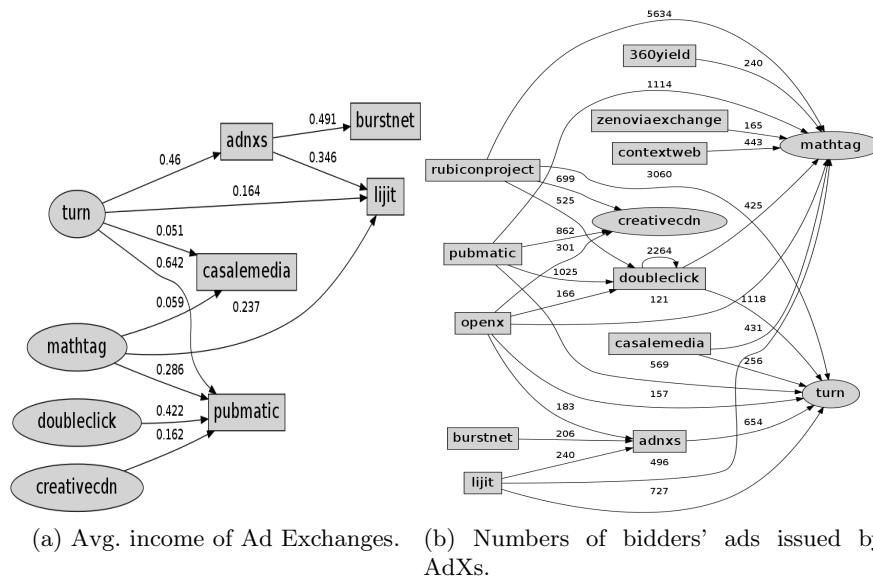


Fig. 8: Communication flows ($Dataset_1$).

Because of the fact that bidders (AdXs are also often bidders) bid in auctions held by several different Ad Exchanges, the advertising and user-data market is very complex. It is interesting to note that there are important differences between the data sent in bid requests of various Ad Exchanges [8, 25]. For example Doubleclick sends the Web user’s IP address in the “anonymized” form (without the last octet, i.e. $10.0.1.X$), while Pulse Point is sending the full IP address. These two information can then later be linked with data on the ethnicity and income of the user, which can be supplied in OpenX’s bid requests [22]. If one bidder, for example Media Math or Turn, takes part in several auctions, it is potentially possible to see the data relating to same users, several times. Consequently, it could be possible to link all the details obtained during the auctions of different AdXs. In the end it would allow the bidder to build a more comprehensive profile of a user.

4.3 Ad Exchange-Publisher Interaction Analysis

General Case In RTB, publishers’ sites such as *example.com* usually include third-party scripts to provide ad impressions and display ads. Technically this is often done by the use of a HTML `iframe` tag. An example domain name of such scripts can be *doubleclick.net*, in case of Doubleclick. Ad Exchanges usually set specific tracking cookies in the users’ browsers. For instance, Figure 9a shows that when the user visits *example.com*, a request to *doubleclick.net* is made; during this request, the user’s cookie controlled by Doubleclick is sent.

The winning bidder’s (Criteo in this case) ad snippet is then served and the procedure of serving ads can be initiated.




Fig. 9: Inclusion of ad snippets on publishers’ sites.

OpenX Case We noticed a possibly puzzling cooperation between one Ad Exchange and certain publishers. In this setting, the publisher `example.com` might include an `<iframe>` referring to the ads scripts from the domain `ox-d.example.com`. This subdomain, supposedly belonging to the visited site (`example.com`) in reality is a DNS alias. The actual server is controlled by a third-party, Ad Exchange: OpenX. This *hidden third-party* setting is an example of DNS aliasing [17, 16, 30]. The domain `ox-d.example.com` sets a cookie with a unique user ID, `OX_u` (example in Table 4). This means that when the user’s browser performs a request to this host, the cookie is included in this request’s headers. However, very often the first-party domain name, `example.com`, sets its own cookies as well. If these cookies have a *broad scope* [3], for example `Domain=.example.com`, they are consequently leaked to `ox-d.example.com`, which is a site operated by OpenX, an external entity. As we see in Figure 9b not only legitimate cookies, controlled by OpenX (belonging to `ox-d.example.com`), are transferred to the third-party server.

This means, that in this setting OpenX could have access to the cookies of `example.com`, due to their scope. We detected this to indeed be the case and cookies are leaking to OpenX systems on certain analyzed Web sites. Examples of the affected Web sites are `dailyherald.com` (leak to `ox-d.dailyherald.com`) and `popcrunch.com` (leak to `ox-d.popcrunch.com`). Whenever a user visits these sites, certain cookies are leaked to OpenX, depending on the cookie scoping.

Table 4 shows an example: `OX_u` is the per-user cookie of OpenX, while `_utmz` is a cookie related to Google Analytics. Its scope is `.popcrunch.com`, which means this cookie is sent during a HTTP request to any subdomain of `popcrunch.com`. Therefore, it is also sent to `ox-d.popcrunch.com`, a hostname controlled by OpenX.

Leaks Study. We crawled 1M Alexa⁵ sites and searched for requests of the form

⁵ <http://www.alexa.com>

ox-d.example.com, where *example.com* is the address of the visited site. We identified 127 sites. When performing a name resolution, all such hosts turned out to be operated by OpenX. Using PhantomJS browser, we subsequently visited each of these 127 sites and saved all the related cookies. We detected that OpenX’s cookie was set on about 20% of the sites (i.e. 26 sites).

We then analyzed cookies for each of these 26 Web sites *example.com*, in order to verify if cookie leaks take place. In case of 81% of these sites, we detected the leakage of cookies not belonging to OpenX. Those cookies commonly had a broad scope set. We found that *Google Analytics* cookies (e.g. *_utmz*) were leaked to OpenX servers in 70% of these sites. For example, the site *zam.com* leaked three cookies, two of them being related to Google Analytics.

Cookie Matching Potential. One of the potential consequences of this setting is that OpenX could create a custom Cookie Matching scheme [21], where cookies set by two different entities are mapped. OpenX could match their *user id* cookies with e.g., Google Analytics cookies belonging to these Web sites. As a result, OpenX could track the visitors of these Web sites based on the per-site cookies of Google Analytics. Example scenario could arise when the user removed OpenX cookies (or when they expired) but left the ones belonging to Google Analytics intact. OpenX could then regenerate the user’s profile using the unchanged Google Analytics cookie. In the case of example from Table 4, we manually verified that OpenX was not reproducing its tracking cookie out of Google Analytics cookies, after its removal. But linking the profile with the new cookie is still theoretically possible.

Name	Value	Scope
<i>_utmz</i>	236312704.1392366853.1.1.utmccn=(direct)— [...]	.popcrunch.com
<i>OX_u</i>	195e9f99-9b18-0991-06a2-98172b0d3651_m_1385039804	ox-d.popcrunch.com

Table 4: OpenX cookies (*OX_u*) and Google Analytics (*_utmz*) one. In this case *_utmz* leaks to OpenX due to cookie scoping.

Consequences of Such Setting. Cookies are often directly related to authentication in Web systems [23]. Consequently, the severity of this leak has both security and privacy implications [16, 30]. Krishnamurthy and Wills discussed such possibilities in [16]. Wills mentioned the risks of these settings where first-party site stores sensitive data such as e-mail address, in the cookies [30]. In the case analyzed by us, the problem might even be more complicated due to the fact that OpenX is an Ad Exchange and has real-time bidders who receive bids requests with information on the user.

By leveraging this very setting of close RTB-publisher collaboration, it is possible to evade host-based blacklists of advertising and tracker-blocking browser extensions, such as Ghostery. Moreover, this setting also serves as a *work-around over blocking of 3rd-party cookies*, a mechanism used by Safari browser and cur-

rently considered for inclusion to Firefox. We believe this might be a primary motivation in the encountered cases.

We performed a test with Firefox 26. We enabled *blocking of cookies from the sites the users did not visit* (3rd-party cookies). We detected that when visiting *popcrunch.com*, *OX.u* cookies were still set by *ox-d.popcrunch.com*, a consequence of how 3rd-party cookie blocking works in Firefox. We also installed Ghostery and enabled its blocking mode. Requests to *ox-d.popcrunch.com* were still executed and OpenX's cookie was set.

5 Conclusion

In this paper, we proposed a TET that allows users to understand how their data is being exchanged and sold in RTB systems. This tool is very useful to improve transparency, but also to inform people about the value of their data and the need to protect them.

We used this tool to study the internals of RTB, and understand the strong cooperations that exist between the different RTB components i.e. Ad Exchanges, Bidders, Publishers. We showed that these interactions are very complex and elaborate, and aim at optimizing their revenues, often at the expense of user privacy and security.

This paper also studied RTB from an economics point of view. We showed that advertisers have different bidding strategies and value users differently according to their profiles. Some advertisers tend to bid as often as possible, others are more selective and are only targeting some specific users.

Unfortunately, we believe that with the development of new companies that exploit personal data, commercial tracking, via RTB or other means, will become omnipresent. We hope that our tool and the results of this work will improve transparency and will contribute to increase users' awareness about the proliferation of these new commercial surveillance infrastructures.

References

1. A. Acquisti and J. Grossklags. Privacy and rationality in individual decision making. *Security & Privacy, IEEE*, 3(1):26–33, 2005.
2. A. Acquisti, L. John, and G. Loewenstein. What is privacy worth. In *Workshop on Information Systems and Economics (WISE)*, 2009.
3. A. Barth. HTTP State Management Mechanism. <http://tools.ietf.org/html/rfc6265>, 2011. RFC 6265.
4. Brandscreen. Brandscreen Q4 2013 real-time media insights report. <http://www.brandscreen.com/en/blog/brandscreen-q4-2013-real-time-media-insights-report>, 2013.
5. Business Insider. How real-time bidding will influence the 2016 election — and change Don Draper's job. <http://www.businessinsider.com/sc/real-time-bidding-helps-decide-elections-2014-1>, 14 January, 2014.

6. J. P. Carrascal, C. Riederer, V. Erramilli, M. Cherubini, and R. 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, pages 189–200. International World Wide Web Conferences Steering Committee, 2013.
7. DoubleClick. Doubleclick ad exchange real-time bidding protocol. <https://developers.google.com/ad-exchange/rtb/>.
8. DoubleClick. Processing the request – example bid request. <https://developers.google.com/ad-exchange/rtb/request-guide#example-bid-request>.
9. DoubleClick Help. Serve ads in a non-javascript environment. https://support.google.com/dfp_premium/answer/1638620?hl=en.
10. eMarketer. Advertisers continue rapid adoption of programmatic buying. <http://www.emarketer.com/Article/Advertisers-Continue-Rapid-Adoption-of-Programmatic-Buying/1010414>, 2013.
11. P. Gill, V. Erramilli, A. Chaintreau, B. Krishnamurthy, K. Papagiannaki, and P. Rodriguez. Follow the money: Understanding economics of online aggregation and advertising. In Proceedings of the 2013 Conference on Internet Measurement Conference, IMC '13, pages 141–148, New York, NY, USA, 2013. ACM.
12. Google. The arrival of RTB. http://static.googleusercontent.com/external_content/untrusted_dlcp/www.google.com/pl//doubleclick/pdfs/Google-White-Paper-The-Arrival-of-Real-Time-Bidding-July-2011.pdf.
13. Google. Decrypting price confirmations. <https://developers.google.com/ad-exchange/rtb/response-guide/decrypt-price>.
14. A. Halevy, P. Norvig, and F. Pereira. The unreasonable effectiveness of data. Intelligent Systems, IEEE, 24(2):8–12, 2009.
15. IDC. Real-time bidding in the United States and Western Europe, 2010–2015. http://info.pubmatic.com/rs/pubmatic/images/IDC_Real-Time%20Bidding_US_Western%20Europe_Oct2011.pdf, 2011.
16. B. Krishnamurthy and C. Wills. Privacy diffusion on the web: a longitudinal perspective. In Proceedings of the 18th international conference on World wide web, pages 541–550. ACM, 2009.
17. B. Krishnamurthy and C. E. Wills. Generating a privacy footprint on the internet. In Proceedings of the 6th ACM SIGCOMM Conference on Internet Measurement, IMC '06, pages 65–70, New York, NY, USA, 2006. ACM.
18. J. Loveless. Barbarians at the gateways. Commun. ACM, 56(10):42–49, Oct. 2013.
19. M. Mahdian, A. Ghosh, P. McAfee, and S. Vassilvitskii. To match or not to match: Economics of cookie matching in online advertising. In Proceedings of the 13th ACM Conference on Electronic Commerce, pages 741–753. ACM, 2012.
20. S. Muthukrishnan. Ad exchanges: Research issues. In Internet and network economics, pages 1–12. Springer, 2009.
21. L. Olejnik, M.-D. Tran, and C. Castelluccia. Selling off user privacy at auction. In In Proceedings of the 21st Annual Network and Distributed System Security Symposium (NDSS). The Internet Society, 2014.
22. OpenX. Openx market real-time bidding api reference. http://docs.openx.com/resources/rtb_api_ref_v13.pdf.
23. OWASP. Owasp periodic table of vulnerabilities - cookie theft/session hijacking. https://www.owasp.org/index.php/OWASP_Periodic_Table_of_Vulnerabilities_-_Cookie_Theft/Session_Hijacking, 2013.
24. Pulse Point. Private exchange. <http://docs.pulsepoint.com/display/RTB/Private+Exchange>.

25. PulsePoint. Real-time bidding protocol request examples. <http://docs.pulsepoint.com/display/RTB/Real-Time+Bidding+API#Real-TimeBiddingAPI-BidRequestParameters>.
26. S. J. Savage and D. M. Waldman. The value of online privacy. [Available at SSRN 2341311](#), 2013.
27. TRUSTe. Truste 2014 US consumer confidence privacy report. <http://www.truste.com/resources/?doc=447>, 2014.
28. Turn. Amplify. http://www.turn.com/downloads/brochure/TURN_brochure_print_LR.pdf [Access: 21.11.2013], 2013.
29. W. Vickrey. Counterspeculation, auctions, and competitive sealed tenders. [The Journal of finance](#), 16(1):8–37, 1961.
30. C. E. Wills. Identifying and preventing conditions for web privacy leakage. In [Proceedings of the W3C Workshop on Web Tracking and User Privacy](#), 2011.
31. S. Yuan, A. Z. Abidin, M. Sloan, and J. Wang. Internet advertising: An interplay among advertisers, online publishers, ad exchanges and web users. [arXiv preprint arXiv:1206.1754](#), 2012.
32. S. Yuan, J. Wang, and X. Zhao. Real-time bidding for online advertising: measurement and analysis. In [Proceedings of the Seventh International Workshop on Data Mining for Online Advertising](#), ADKDD, New York, NY, USA, 2013. ACM.

A *RTB Tracker*: reports for users

The *RTB Tracker* extension we introduce in this paper offers users a functionality of comparing the average bids for ads they receive, with average bids of other users.

An interested user could visit our site [**ANONYMIZED**] and discover several information behind his profile. The example possible scenario is shown on Figure 11. The Figure 10 shows how the site looks for a normal Web user, i.e. someone who did not use the extension.

How Much Are You Worth?

When you visit certain Web sites on the Internet, ad requests are sent to advertisers. They compete for a chance to serve ads to you. The bid prices they submitted to auctions are generally based on your information that advertisers possess, for example a profile inferred from your Web history, and your browsing context. The prices reflect how they evaluate your profile. We capture these prices to give you a quantification of your value from advertisers' perspective.

Results: We do not have any data associated with you. If you use Firefox or Chrome, please install the plugin and enable cookies. Note that our plugin does NOT work with Ad blocker extensions such as AdBlockPlus or any addons of these types. If you use Ad blockers, and still want to know how advertisers estimate your private data, you have three options:

- Deactivate Ad blockers, browse the Internet as usual, and reactivate your Ad blockers whenever enough prices are collected (i.e. our plugin starts showing you your average price and the number of prices is larger than 20-40). This could take a couple of days or just one day, depending on your browsing habits.
- Deactivate your Ad blockers, click on some (10 to 15, with e.g. several refreshes) of the links [example from this list](#), and reactivate your Ad blockers later once you are done. This option is faster, but might somehow affect the results.
- For Firefox create a [new Firefox profile](#) and browse the Web as usual with our plugin installed. Of course do not use Ad Blockers. For Chrome, create a new browser user profile: [see here](#) or alternatively check [this tutorial](#).

We show below some general highlights, if you do not use Firefox or Chrome, or do not wish to install the plugin.

The average price paid for user's private data (items in Web browsing history) for our users is:



\$0.001328

Summary of the results

Total number of users (with prices)	319
Average price of users	\$0.001328
Price of the cheapest user	\$0.000053
Price of the most expensive user	\$0.010000

Are you interested how much advertisers pay for your data? [Go and check!](#)

UPDATE: Using preliminary data we present the per-advertiser prices that are paid for your data: [see here](#). You can improve this data by continuing using our plugin. Thanks for your collaboration!

Fig. 10: The site for users without the extension.

How Much Are You Worth?

When you visit certain Web sites on the Internet, ad requests are sent to advertisers. They compete for a chance to serve ads to you. The bid prices they submitted to auctions are generally based on your information that advertisers possess, for example a profile inferred from your Web history, and your browsing context. The prices reflect how they evaluate your profile. We capture these prices to give you a quantification of your value from advertisers' perspective.

Items in your browsing history are worth (on average):

\$0.000455

Out of 169 users, you are ranked:

139th

Which means you have the 139th largest price in general. In your country this is 60th.

This value is averaged from prices advertisers had paid for you until now. Following is data from all examined users in our database.

Be sure to check the results of prices collection by our plugin regularly (e.g. in a few days) to see changes according to your browsing activities. You might be surprised! If the reported number of prices is small and you want to increase it to further see how you are valued, you might want to visit some sites from the [example list here](#).

Summary of the results

Your private data value	\$0.000455
Number of your detected prices	89
Total number of users (with prices)	169
Number of users worth more than you	138
Number of users worth less than you	31
Average price of users	\$0.001364
Price of the cheapest user	\$0.000053
Price of the most expensive user	\$0.010000

Fig. 11: The site for users with extension installed.

B Bidders

We provide a full list of detected RTB bidders: mythings.com, yieldmanager.com, dotomi.com, mxptint.net, sharethis.com, rtbidder.net, chango.com, afx.ms, pswe.com, netseer.com, struq.com, adcash.com, twenga.com, tubemogul.com, criteo.com, fastclick.net, twenga.fr, lucidmedia.com, liverail.com, doubleclick.net, twenga.nl, tidaltv.com, adsrvr.org, adscience.nl, quantcount.com, adap.tv, ctnsnet.com, pulsemgr.com, de17a.com, microad.jp, creativecdn.com, amazon-adsystem.com, everesttech.net, quantserve.com, sociomantic.com, adnxs.com, w55c.net, clickurl.net, adroll.com, gwallet.com, invitimedia.com, 254a.com, intelliad.com, ru4.com, creative-serving.com, zenoviaexchange.com, bidtheatre.com, kejet.net, turn.com, mathtag.com, rvty.net, roulartamail.be, esm1.net, owneriq.net, nxtck.com, wtp101.com, qservz.com, dyn-trk.com, rfihub.com, yieldoptimizer.com, revsci.net, capitaldata.be, rfihub.net, netmng.com, trgt.eu, m6r.eu, intelliad.de, semasio.net, impdesk.com.