The Value of Online Users: Empirical Evaluation of the Price of Personalized Ads

Miriam Marciel*, Jose Gonzalez†, Yonas Mitike Kassa‡†, Roberto Gonzalez*, Mohamed Ahmed*
*NEC Labs Europe
{miriam.marciel, roberto.gonzalez, mohamed.ahmed}@neclab.eu
†Universidad Carlos III de Madrid
{jgcabana@it.uc3m.es}@uc3m.com
‡IMDEA Networks Institute
{yonas.kassa}@imdea.org

***Abstract*—Ad networks use the behaviors of online users to associate them with preferences (features), and market these features to enable advertisers to target online users. Typical features associated with users include location, interests, gender, age, and etc. Furthemore, ad networks provide their clients with campaing creation tools to help to them to configure and run campains. In this paper, we study the pricing of ads using the ad campaing planning tools of ad networks. We develop tools to collect the suggested bid prices from two platforms: YouTube and Facebook. Analyzing these prices we find that United States is the most expensive country in both platforms. We also find that the most expensive preferences are different in YouTube and Facebook. In YouTube, the top preferences are related to Oil & Gas, while in Facebook are devices, ethnics or politics depending on the type of bidding. Finally, we do not find any price difference genders in Facebook.**

***Keywords*-online advertisement, bid auction, online users**

I. Introduction

Online advertising plays a vital role in today's Internet ecosystem, sustaining large service providers such as Facebook and Google to small user run websites. The Interactive Advertising Bureau (IAB) has recently reported that advertising generated $59.6 billion in 2015 in the US alone [1].

In contrast to traditional advertising such as newspapers or TV, online advertising gives advertisers more flexibility in selecting and targeting audience segments. Online advertising enables advertisers to define fine-grained profiles of target audiences. These profiles include features such as the location, age, gender, and interests of users. Furthermore unlike traditional media advertising, because the ad location is flexible, online advertising enables advertisers to re-target users as they go from website to website. For example, a user who has searched for flights, sees ads related to flights while reading the news across different sites. This practice is called Online Behavioral Advertising (OBA).

In order to infer user interests, ad networks track users while they browse the web and try to decipher the topics of the sites visited by users, as well as their interests. This information is then used to build user profiles that associate users with feature sets, such as their demographic and personal information, and their preferences. Finally, advertisers target users using these preferences and, users see ads that are judged to be 'relevant' to them (according to their profile). Since different user-segments define different interests/behaviors, they are priced based on the demand from advertisers. Prices are typically determined using auctions (see [2, 3] and references therein) and as such the price of any given user segment typically varies with the market (advertiser) demand for the segment [4].

Because user behaviors and the demand for audience segments change over time, the preferences assigned to users and the their value to the ad networks also changes. However, little is known about the pricing of the preferences assigned to users or how their prices change over time.

Understanding this would enable us to better understand how users are valuated by online systems. For example, YouTube recently released an ad-free subscription for 10$ per month in the US [5]. This leads us to believe that 10$ is approximately the revenue an average user would generate for YouTube, however there is no easy way to evaluate whether or not this is the case today. To the authors best knowledge, the only report available that ranks the values of the preferences assigned to users by ad networks is from the IAB advertising revenue report [1]. According to the report, "Retail", "Financial Services" and "Auto" are the preferences for which advertisers spent most money in 2015. But this report only gives aggregated numbers for the whole year.

The purpose of this paper is to shed light on the pricing of user preferences in ad networks. In order to collect prices, we take advantage of the "Creation Campaign" tools of different platforms. These webpages show an estimated bid price (using a range) of the price the advertiser would pay, depending on the targeting options selected. According to [4], in Facebook, these prices are sampled from the prices that advertisers pay.

We study the pricing of ads of two different platforms: YouTube and Facebook. We develop price collection tools that retrieve the estimated bid prices for different countries and preferences (for both platforms), and gender and type of bidding (for Facebook). Then, we analyze the prices collected by country, preference, gender and type of bidding. Our findings can be summarized as follows:

- Analyzing prices per country, on average the United States is the most expensive top country in both platforms. Among the most expensive 20 countries,

978-1-5090-0990-9/16 $31.00 © 2016 IEEE
DOI 10.1109/ARES.2016.89

we find that while most are developed countries, few are developing countries.

- In Facebook, prices vary aggressively depending on the type of bidding selected: Cost per Mile (CPM), Cost per Action (CPA) or Cost per Click (CPC).
- Ranking of top preferences is different for each platform. In YouTube, the top ranking preferences are related to Oil & Gas, while in Facebook, the ranking is different depending on the type of bidding.
- In Facebook, the top tanking preferences are related to mobile devices and browser for CPM, while for CPA and CPC, the top ranking preferences are related to ethics and politics.
- We do not find any price difference between genders in Facebook.

The remainder of this paper is organized as follows: Section II describes the related work, Section III describes the collection tools used and the datasets we collect, Section IV analyses the YouTube dataset, Section V analyses the prices in Facebook, and Section VI concludes the paper.

II. Related Work

The research community has paid particular attention to the online advertising industry and its tracking techniques over the last few years. Numerous works such as [6–9] haved performed crawls in order to understand the personalization of online ads. It has been shown that, from personalized ads, it is possible to guess users' preferences [10].

Ad networks use different technologies to identify and track users among webpages. The most common is HTTP cookies [11, 12] which usually includes unique user identifiers. In order to have more complete profiles, ad networks have created methods to allow them to share these identifiers. The most common is Cookie Matching [13, 14] in which different entities share cookie values between domains. HTTP Cookies can be deleted very easily from the browser, therefore, more complex technologies like evercookies (cookies using the HTML5 storage) [14] and browser fingerprinting [15] have appeared. These technologies are more persistent as they are more difficult to delete (evercookies) or more difficult to bypass (fingerprinting).

Fraud is also a problem of advertising [16–18]. It has been shown that fraudsters could take advantage about profiling users of ad networks, by adding high-valued user preferences [19] to profiles. Authors claimed that fraudsters could increase their benefit by 33%.

Related to value of users online, Saez-Trumper et al. [20] propose a model in which the value of users is defined by actions they make in an Online Social Network. In [21], authors make a survey to real users, asking them about the price they think they should receive to reveal information about their actions while browsing. More closely related to our work, authors in [13] study the prices of Real Time Bidding (RTB). RTB is a advertising system to sell ads in auctions. They find that prices change according to user's profile and time of the day. Finally, Liu et al. [4] study the pricing of ads in Facebook. They reverse engineer Facebook's pricing system, concluding that the estimated price shown is sampled from the last prices paid by advertisers.

	YouTube	Facebook
Measurement Period	10/Jun/15 - 18/Aug/15	13/Aug/15 - 25/Sept/15
# Categories	2043	584
# Countries	50	50
# Genders	1	3
# Entries	476K	18M

Table I: Description of our datasets for YouTube and Facebook

III. Price Collection Tools and Datasets

In this section we describe the different tools to collect the prices (suggested bids) from YouTube and Facebook. We also make a short description of the advertisement system of each of the platforms. Finally, we present the characteristics of our datasets.

A. YouTube Price Collector

In order to promote a product in YouTube, advertisers have to create a campaign in Google Adwords [22], the service to publish advertisements of Google. Google Adwords offers different options to target users such as geographical information (country, city, etc.), language, topics, preferences, gender, age range, parental status, etc.

When an advertiser creates a campaign, Google Adwords shows a suggested bid (price) that the advertiser will pay for one view of his ad in YouTube in Cost-per-View (CPV) [23]. The suggested bid is given as an interval (a maximum and a minimum value in the local currency). When the advertiser changes the targeting options, the suggested bid is updated in real time.

In order to collect these prices (suggested bids), we develop a webcrawler using Selenium Webdriver [24]. Selenium Webdriver is a library to automate testing on webpages, allowing actions like loading webpages, clicking on links or filling up forms. Our webcrawler is able to login into Google Adwords, change the targeting options in campaigns and collect the estimated bids for the configured options. All those bids are stored in a MySQL database. Our crawler collects prices using the country and the preference as targeting options. The preference is selected from the list of topics of Google [25].

B. Facebook Price Collector

Facebook has a similar tool for advertisers to promote a product. Advertisers have to create a campaign and, while they configure their targeting options, the suggested bids are updated. The targeting options are similar as YouTube: country, language, preferences, gender, age, etc.

Facebook has also different kinds of bidding: (i) Cost per click (CPC), (ii) Cost per Action (CPA), and (iii) Cost per Mile (CPM). In CPC, the advertiser only pays when the user has clicked on his ads. In CPA, the advertiser pays when the user has performed an action (like Page likes, installing an app, etc.). Finally, in CPM, the advertiser pays per 1K impressions of his ad.

Like YouTube, Facebook shows the suggested bids using a range for any of the kinds of bidding explained before. In this case, Facebook shows three values: maximum, minimum and median, and these values are given in the local currency.

In order to collect the prices of Facebook, we follow the methodology described in [4]. We send HTTP requests to the API of Facebook, setting the following targeting options: gender, maximum and minimum age, interests and country. Facebook API replies with the estimated prices (maximum, minimum, median) for those targeting options and for all the kinds of bidding described before. We store those prices in a MySQL database.

C. Datasets

Using the price collections tools explained in this Section, we collect two different datasets: one from YouTube and one from Facebook. The details of these datasets are described in Table I.

While we collected data (topics, interests and etc) for all the countries listed by both platforms, we restrict our analysis to 50 countries for which we were able to collect statistically significant quantity of data. These countries are distributed around the world: America (12), Europe (28), Asia (1), Africa (7) and Oceania (2). Moreover, we collect prices for different genders in Facebook (all, male and female). In the case of YouTube, there is only one gender (undetermined). In both platforms, we did not target any age, therefore we select all the age ranges.

In our datasets, an entry is defined as: date of the measurement, country, preference, gender, bidding type and value range. In total, the datasets have 476K and 18M entries for YouTube and Facebook, respectively. The difference in the number of entries is due to the collection method: YouTube is slower, since we use webcrawler, with Facebook we send hand-crafted HTTP requests. Furthemore for many categories, Google Adwords returns empty results, i.e. there is no interest in the category, and therefore we do not have any entries for those cases. Finally, in Facebook, we collect the prices for three genders while, in YouTube, their is only one.

As explained in Section III-A and III-B, YouTube gives the suggested bids as a range with a maximum and a minimum, while Facebook returns three values: maximum, median and minimum. For our analysis, we use the maximum value, since it is the common value for both platforms and, according to our measurements, shows more price variation.

IV. Study of YouTube Prices

In this section, we study the pricing of ads in YouTube and look at how it varies depending on users' profile: country and category.

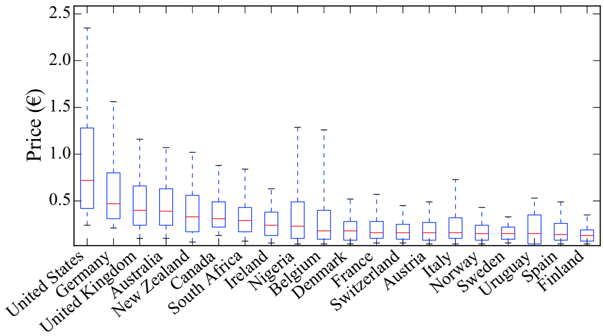

Figure 1: Comparison of YouTube prices among the top 20 countries.

A. Pricing of Countries

Geographical targeting is a common option used by advertisers, as they can create campaigns targeting a certain region, state or country. This feature is very important, for example, for small retailers, targeting users living in the same city; or big companies, creating a campaign to target a whole country (fixing also the language).

Google Adwords offers the option of targeting users depending on their location. The geographical options go from city to worldwide, listing options such as region/state, country or continent.

We analyze the data described in Table I, grouping all the prices by country. Figure 1 shows a boxplot with the ranking of top 20 countries analyzed. The whiskers of the boxplots are the 9th and 91st percentile.

The figure shows that 17 of the 20 top countries are countries with "Very High Human Development", according to the report of United Nations [26]. Only South Africa, Nigeria and Uruguay do not belong to this category, belonging to "Medium Human Development", "Low Human Development" and "High Human Development", respectively. This means that advertisers are also interested in paying a considerable amount of money for users of those countries, despite the fact that those countries are not listed as having good Internet connections or a large pool of net-citizens [26].

B. Pricing of Categories

Google Adwords enables advertisers to target users depending on their likes and preferences. This is done by tracking users while they are visiting websites, which are later mapped to the list of categories of Google [25].

The list of categories of Google Adwords [25] has a total 2043 categories and they are organized as a tree, containing 26 root categories. The categories include: Sports, Beauty, Computers, etc. Out of the 2013 categories, 847 are world locations, like countries or cities.

As in Section IV-A, we analyze the data grouping all the prices by category, not taking into account the country.

Table II shows the top 10 ranking categories, ordered by the median price. We see that there is a variety of categories and that belong to different branches of the

Rank	Category	Median Price (€)
1	Oil & Gas	2.57
2	Fuel Economy & Gas Prices	2.41
3	Asset & Portfolio Management	1.64
4	Customer Relationship Management	1.46
5	Web Hosting & Domain Registration	1.26
6	Electromechanical Devices	1.05
7	Debit & Checking Services	1.02
8	Murcia	0.92
9	Foreign Language Study	0.87
10	Telemarketing	0.86

Table II: Ranking of category prices order by highest median value.

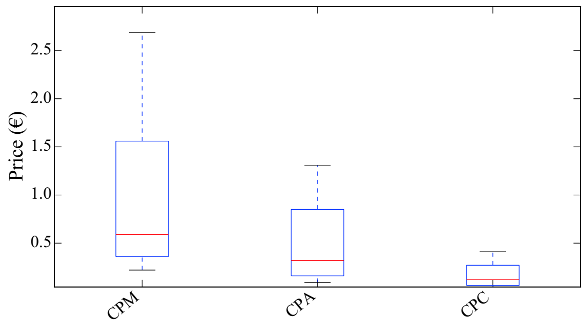

Figure 2: Comparison of Facebook prices for the different bidding types.

tree. It is surprising that the top 2 categories are related to Energy and Gas. These categories are almost 0.80€ more expensive than the third one, which suggests that advertisers pay at least that amount for a view of a user having those categories. Also, these similar categories belong to different branches of the tree, which could suggest that advertisers target both categories at the same time. We also see that the third and the fourth categories are related to management, however their price difference with respect to preferences ranked below them is smaller than with respect to the preferences ranked above.

V. Study of Facebook Prices

In this section, we analyze the prices of ads for Facebook. For this platform, we have more variables to analyze and, therefore, we can make a deeper analysis and look for differences between the types of bidding, countries, categories and genders.

A. *Pricing of type of bidding*

As discussed in Section III-B, Facebook has three different types of bidding (CPC, CPA and CPM), in contrast to YouTube that has only one (CPV). With these auctions, advertisers decide what action users have to perform in order to pay. For example, in CPC, advertisers pay when

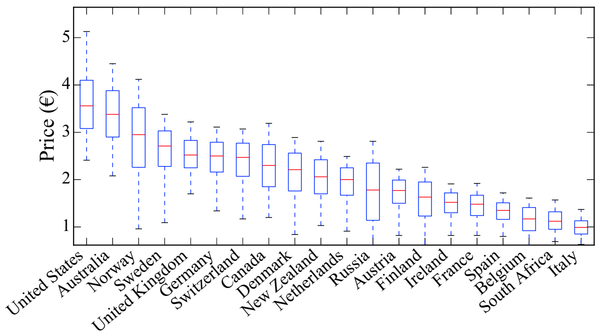

(a) CPM

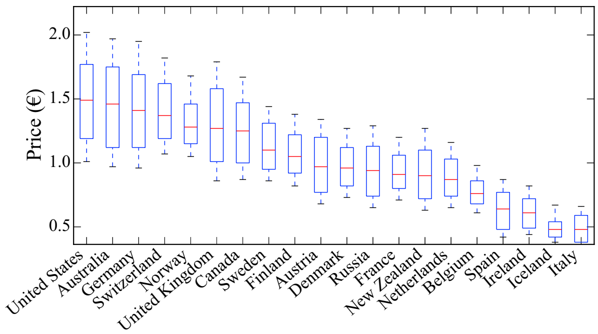

(b) CPA

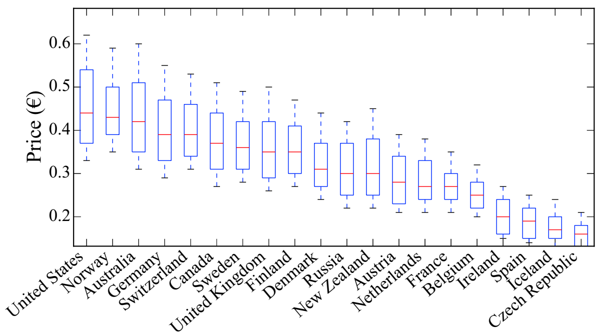

(c) CPC

Figure 3: Comparison of Facebook prices among the top 20 countries for different bid types.

a user clicks on an ad; in CPA, advertisers define the action the user has to do (i.e. like) and they pay according to it; and in CPM, advertisers pay per 1K ads.

Figure 2 shows the boxplot of the distribution of prices according to the type of bidding. Again the whiskers are the 9th and 91st percentiles. The results show that CPM is the most expensive type of bidding. This is reasonable since, using this type of bidding, advertisers reach 1000 users, while in CPA and CPC advertisers reach only 1. Therefore, the price per user in CPM is much lower than CPA and CPC.

We can also see that CPA is more expensive than CPC, which means that advertisers pay more for other actions than a user's click.

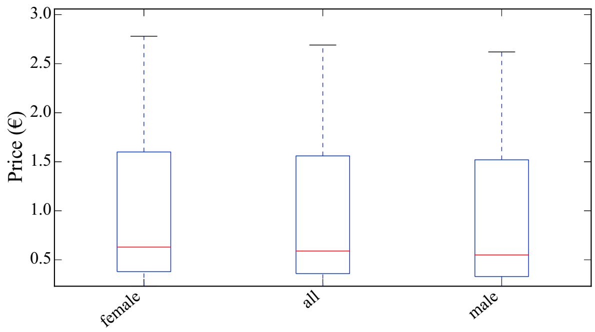

(a) CPM

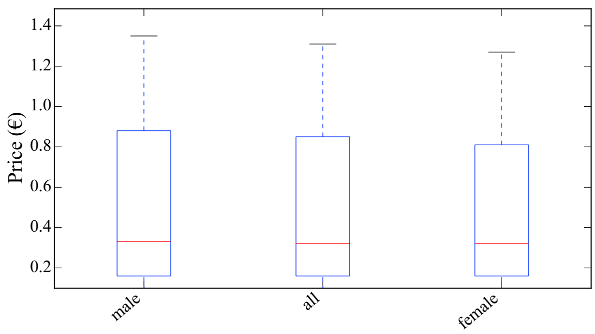

(b) CPA

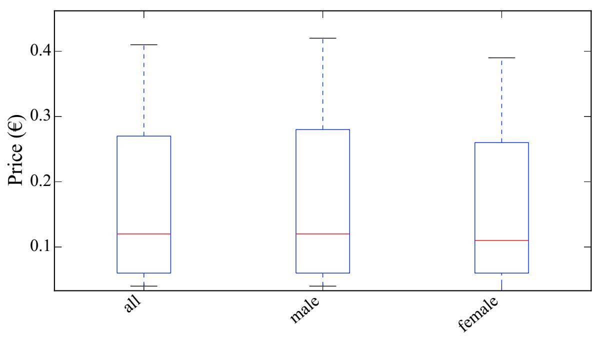

(c) CPC

Figure 4: Comparison of Facebook prices among genders for different biddings.

B. *Pricing of countries*

Like YouTube, Facebook offers advertisers the possibility to target users depending on their location, using the country, state/region or city. But also, it gives more fined-grained options such as postal code, address, designated market area or congressional district.

In order to keep the statistics simple and be able to compare to YouTube, we collect and analyze data per country. Also, for the analysis of this subsection, we only use the gender "All".

Figure 3 shows the boxplot of the top 20 countries according to the type of bidding. Again, the whiskers are the 9th and 91st percentiles.

As with YouTube, the top country independently of type of bidding is United States. Besides, most of the countries belong to "Very High Human Development" ranking of the United Nations [26]. The only exceptions are Russia, that is considered to be a country with "High Human Development", and South Africa (only in CPM), considered a "Medium Human Development" country.

Comparing the ranking of the three types of bidding, we see that the list of countries and the order of them does not vary much. Considering the set of 20 top countries in each of the biddings, we have a total of 22 countries. Furthermore, in the top 3 of all types of biddings, there are only four countries (United States, Australia, Norway and Germany). Therefore, we can conclude that the ranking of countries is independent of the type of bidding.

C. *Pricing of categories*

Facebook also offers the possibility to target users based on their preferences (a pre-defined set of categories).

Unlike YouTube, Facebook does not provide a public list of all the categories. In order to get the full list of categories, we analyze the traffic from a browser to Facebook, while creating a campaign. We discover that the browser downloaded 5 JSONs with targeting options. Two of these JSONs include information like the countries or languages available to target. The other three, include information to target users according to their preferences.

These three JSONs are organized in Demographics, including information about ethnics, relationship status or birthdays; Behaviors, including information like operating systems, devices, browsers or playing habits; and Preferences, that describes likes (i.e. Travel, Computers, etc.), more similar to the list of YouTube. In total, there are 595 categories, organized like a tree.

We analyze the data described in Section III-C for only the gender "All", grouping by preferences. Note that our dataset has 584 categories, when actually there are 595. This means that our collection tool was not able to retrieve information on 11 categories.

Table III shows the ranking of the top preferences ordered by the median for the different types of bidding. The first detail to notice is that the categories are totally different for CPM than for CPA and CPC. In CPM, the categories are mainly Behaviors (like browser and devices), while in CPA and CPC, the categories belong to Demographics (ethics and politics). The only common categories in the three biddings is "E-book readers", ranked 1st, 9th and 9th for CPM, CPA and CPC, respectively.

This fact could be explained because of the nature of the bidding: in CPM, advertisers are not able to collect any information about the user, as they pay per impression; while in CPC and CPA, they do. In CPC and CPA, users have to click or make an action in their webpage, giving the possibility to advertisers to collect information of the user. Therefore, advertisers are able to attach categories to users, that they would not be able to do in other case.

We also see that the top 9 categories are the same for CPA and CPC. This means that advertisers target the same category independently of the action the user has to make (although the user has to make an action).

Rank	CPM Category	CPM Median Price (€)	CPA Category	CPA Median Price (€)	CPC Category	CPC Median Price (€)
1	E-book readers	2.04	Asian American (US)	0.73	Asian American (US)	0.25
2	Primary Browser: Safari	0.88	FB Payments (Higher than average spend)	0.58	FB Payments (Higher than average spend)	0.20
3	iPhone 6 Plus	0.87	US Politics (Conservative)	0.57	US Politics (Conservative)	0.17
4	Parenting	0.85	US Politics (Very Conservative)	0.48	African American (US)	0.17
5	Children's clothing	0.83	African American (US)	0.47	FB Payments (Recent)	0.17
6	Currently traveling	0.83	Baby boomers (US)	0.47	Baby boomers (US)	0.16
7	Cruises	0.82	Hispanic (US - English dominant)	0.46	US Politics (Very Conservative)	0.16
8	Galaxy Note 3	0.82	FB Payments (Recent)	0.44	Hispanic (US - English dominant)	0.16
9	Apple Email Addresses	0.82	E-book readers	0.41	E-book readers	0.16
10	iPad Mini 3	0.82	Used travel app (1 month)	0.41	Played game yesterday	0.15

Table III: Ranking of categories order by highest median value for the different biddings CPM, CPA and CPC.

Analyzing deeper in each of the bidding, we see that in CPM the median price of the top category is more than double than the second category in the ranking. The difference between the 2nd and the 10th category of the rank is not very high (0.06€). In the case of CPA and CPC, the difference among the top preference and the last in the ranking is not so pronounced (0.32€ for CPA and 0.10€ in CPC).

D. Pricing of genders

Finally, we analyze the impact of the gender on the pricing of ads. Facebook offers three different genders to target: all, male and female.

In this case we analyze the whole dataset and, as we did before, we group the entries by the gender. Figure 4 shows the boxplots for CPM, CPA and CPC for the different genders. From the figure, we can see that the top gender is different is each of the biddings (female for CPM, male for CPA and all for CPC). But the difference is very small: the highest price difference between the first and the third gender in the ranking is 0.08 € for the CPM bidding. We can also see that the distribution is very similar among genders for any type of bidding.

Therefore, we can conclude that advertisers do not show a strong preference for a gender, meaning that there is no statistically significant price difference between genders.

VI. Conclusions

In this paper, we have studied the suggested bid prices for placing ads in two platforms: YouTube and Facebook. We developed two different tools to collect the prices shown on the web. We analyze the data finding that United States is the most expensive country in both platforms. In the top 20 countries we mainly found countries labeled as "Very High Human Development" according to the report of United Nations, but we also found countries not belonging to this category.

Studying the different prices among preferences, we found that YouTube and Facebook had totally different rankings. In YouTube, the top preferences are Oil & Gas. In Facebook, the ranking of preferences changes according to the type of bidding. For CPM, the top preferences are browsers and devices, while for CPA and CPC, are ethnics and politics related preferences.

Finally, we studied two other metrics in Facebook: type of bidding and gender. For the former, we found that CPM is at least two times for expensive in median that CPA and CPC. For the latter, we did not observe any price difference between genders.

VII. Acknowledgments

Many thanks to the anonymous reviewers who provided helpful comments. The research leading to these results has received funding from the European Union Horizon 2020 innovation action program under grant agreement No. 653449-TYPES.

References

[1] "IAB internet advertising revenue report." http://www.iab.com/wp-content/uploads/2016/04/IAB-Internet-Advertising-Revenue-Report-FY-2015.pdf.

[2] W. F. L. J. . Q. R. Yuan, Y., "A survey on real time bidding advertising," in *Service Operations and Logistics, and Informatics*, SOLI, IEEE, 2014.

[3] H. R. Varian, "Online ad auctions," *American Economic Review Papers*, vol. 99, no. 2, p. 430 âĂŞ 434, 2009.

[4] Y. Liu, C. Kliman-Silver, R. Bell, B. Krishnamurthy, and A. Mislove, "Measurement and analysis of osn ad auctions," in *Proceedings of the Second ACM Conference on Online Social Networks*, COSN '14, (New York, NY, USA), pp. 139–150, ACM, 2014.

[5] "Google to Launch Ad-Free YouTube Subscription Service." http://www.wsj.com/articles/google-to-launch-ad-free-youtube-subscription-service-144547224

[6] P. Barford, I. Canadi, D. Krushevskaja, Q. Ma, and S. Muthukrishnan, "Adscape: Harvesting and analyzing online display ads," in *Proceedings of the 23rd International Conference on World Wide Web*, WWW '14, (New York, NY, USA), pp. 597–608, ACM, 2014.

[7] S. Guha, B. Cheng, and P. Francis, "Challenges in measuring online advertising systems," in *Proceed-*

ings of the 10th ACM SIGCOMM Conference on Internet Measurement, IMC '10, (New York, NY, USA), pp. 81–87, ACM, 2010.

[8] C. E. Wills and C. Tatar, "Understanding what they do with what they know," in *Proceedings of the 2012 ACM Workshop on Privacy in the Electronic Society*, WPES '12, (New York, NY, USA), pp. 13–18, ACM, 2012.

[9] A. Hannak, P. Sapiezynski, A. Molavi Kakhki, B. Krishnamurthy, D. Lazer, A. Mislove, and C. Wilson, "Measuring personalization of web search," in *Proceedings of the 22Nd International Conference on World Wide Web*, WWW '13, (New York, NY, USA), pp. 527–538, ACM, 2013.

[10] C. Castelluccia, M.-A. Kaafar, and M.-D. Tran, *Privacy Enhancing Technologies: 12th International Symposium, PETS 2012, Vigo, Spain, July 11-13, 2012. Proceedings*, ch. Betrayed by Your Ads!, pp. 1–17. Berlin, Heidelberg: Springer Berlin Heidelberg, 2012.

[11] S. Englehardt, D. Reisman, C. Eubank, P. Zimmerman, J. Mayer, A. Narayanan, and E. W. Felten, "Cookies that give you away: The surveillance implications of web tracking," in *Proceedings of the 24th International Conference on World Wide Web*, WWW '15, (New York, NY, USA), pp. 289–299, ACM, 2015.

[12] T.-C. Li, H. Hang, M. Faloutsos, and P. Efstathopoulos, *Passive and Active Measurement: 16th International Conference, PAM 2015, New York, NY, USA, March 19-20, 2015, Proceedings*, ch. TrackAdvisor: Taking Back Browsing Privacy from Third-Party Trackers, pp. 277–289. Cham: Springer International Publishing, 2015.

[13] A. Ghosh and A. Roth, "Selling privacy at auction," *Games and Economic Behavior*, vol. 91, pp. 334 – 346, 2015.

[14] G. Acar, C. Eubank, S. Englehardt, M. Juarez, A. Narayanan, and C. Diaz, "The web never forgets: Persistent tracking mechanisms in the wild," in *Proceedings of the 2014 ACM SIGSAC Conference on Computer and Communications Security*, CCS '14, (New York, NY, USA), pp. 674–689, ACM, 2014.

[15] P. Eckersley, "How unique is your web browser?," in *Proceedings of the 10th International Conference on Privacy Enhancing Technologies*, PETS'10, (Berlin, Heidelberg), pp. 1–18, Springer-Verlag, 2010.

[16] B. Stone-Gross, R. Stevens, A. Zarras, R. Kemmerer, C. Kruegel, and G. Vigna, "Understanding fraudulent activities in online ad exchanges," in *Proceedings of the 2011 ACM SIGCOMM Conference on Internet Measurement Conference*, IMC '11, (New York, NY, USA), pp. 279–294, ACM, 2011.

[17] Z. Li, K. Zhang, Y. Xie, F. Yu, and X. Wang, "Knowing your enemy: Understanding and detecting malicious web advertising," in *Proceedings of the 2012 ACM Conference on Computer and Communications Security*, CCS '12, (New York, NY, USA), pp. 674–686, ACM, 2012.

[18] M. Marciel, R. Cuevas, A. Banchs, R. González, S. Traverso, M. Ahmed, and A. Azcorra, "Understanding the detection of view fraud in video content portals," in *Proceedings of the 25th International Conference on World Wide Web*, WWW '16, (Republic and Canton of Geneva, Switzerland), pp. 357–368, International World Wide Web Conferences Steering Committee, 2016.

[19] W. Meng, X. Xing, A. Sheth, U. Weinsberg, and W. Lee, "Your online interests: Pwned! a pollution attack against targeted advertising," in *Proceedings of the 2014 ACM SIGSAC Conference on Computer and Communications Security*, CCS '14, (New York, NY, USA), pp. 129–140, ACM, 2014.

[20] D. Saez-Trumper, Y. Liu, R. Baeza-Yates, B. Krishnamurthy, and A. Mislove, "Beyond cpm and cpc: Determining the value of users on osns," in *Proceedings of the Second ACM Conference on Online Social Networks*, COSN '14, (New York, NY, USA), pp. 161–168, ACM, 2014.

[21] 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*, WWW '13, (New York, NY, USA), pp. 189–200, ACM, 2013.

[22] "Google adwords." https://adwords.google.com.

[23] "Cost-per-view (CPV)." https://support.google.com/adwords/answer/2472735?hl=en.

[24] "Selenium webdriver." http://docs.seleniumhq.org/projects/webdriver/.

[25] "Google topic list." https://support.google.com/adwords/answer/156178?hl=en.

[26] "Human Development Report 2015." http://hdr.undp.org/en/data.