

The Accuracy of the Demographic Inferences Shown on Google's Ad Settings

Michael Carl Tschantz
ICSI
mct@icsi.berkeley.edu

Serge Egelman
ICSI and U.C. Berkeley
egelman@icsi.berkeley.edu

Jaeyoung Choi
ICSI and TU Delft
jaeyoung@icsi.berkeley.edu

Nicholas Weaver
ICSI and U.C. Berkeley
nweaver@icsi.berkeley.edu

Gerald Friedland
ICSI and U.C. Berkeley
fractor@icsi.berkeley.edu

ABSTRACT

Google's Ad Settings shows the gender and age that Google has inferred about a web user. We compare the inferred values to the self-reported values of 501 survey participants. We find that Google often does not show an inference, but when it does, it is typically correct. We explore which usage characteristics, such as using privacy enhancing technologies, are associated with Google's accuracy, but found no significant results.

ACM Reference Format:

Michael Carl Tschantz, Serge Egelman, Jaeyoung Choi, Nicholas Weaver, and Gerald Friedland. 2018. The Accuracy of the Demographic Inferences, Shown on Google's Ad Settings. In *2018 Workshop on Privacy in the Electronic Society (WPES'18)*, October 15, 2018, Toronto, ON, Canada. ACM, New York, NY, USA, 9 pages. <https://doi.org/10.1145/3267323.3268962>

1 INTRODUCTION

Google's Ad Settings offers users a window into the model that Google learns about them from online tracking and their account settings [8]. Users may see inferences Google made about them at <https://www.google.com/settings/ads>. Figure 1 shows a screen shot of the first author's settings from 2016, when we conducted this study (the page has since changed). The page provides two predictions each for the user's gender and age: one based upon the information Google uses for its web services and one based upon the information Google uses as part of its web-wide ad network. The page allows editing the inferences. Google provides some information about how it works [8], but questions remain about the accuracy of Google's profiles.

To study the accuracy of Google's predictions and how they are associated with user behaviors, attitudes, and usage of privacy enhancing technologies (PETs), we conducted a survey. We asked participants for their ages, genders, computer usage habits, attitudes, PETs usage, and for a copy-and-paste of the content of their Ad Settings page. We compared their supplied demographics to

the age range and gender provided by Google and examined how various factors are associated with accuracy.

We find that Google's predictions tend to be accurate when Google makes them, but that Google often makes no inference. In particular, Google rarely makes predictions for logged out users. While we document that Google's accuracy for some subgroups of users is far from its overall accuracy, we do not find statistical significance for any such association after adjusting for the large number of hypotheses examined in this exploratory work.

We believe we are the first look at the accuracy of Google's inferences on Ad Settings with a survey. We provide a new point of reference for understanding Google's ability to infer attributes of users. Additionally, we believe this paper is a reasonable starting point for larger-scale confirmation studies. We make additional information available in the appendix and data available at

<http://www.icsi.berkeley.edu/~mct/pubs/wpes18/>

2 RELATED WORK

Numerous studies have looked at how Google track users [e.g., 1–3, 6, 7, 9, 11–15]. Datta et al. experimented on Ad Settings to determine how they impact the ads shown and how browsing behaviors impact them [6]. Our work differs by looking at the accuracy of Google's stated inferences on real users.

Balebako et al. studied the effectiveness of PETs by examining how personalized the ads shown to browsers with PETs were compared to those shown to browsers without PETs installed [1]. In addition to differing by looking at real users, our work differs by looking at the Ad Settings interface instead of ads.

Small-scale anecdotal examinations of the accuracy of Ad Settings have appeared in the popular press [4, 10], as has a survey looking at the accuracy of Google's geo-location abilities [5].

3 METHODS

With IRB approval, we conducted a survey that consisted of three types of information collection. First, we provided participants standard questions to which they responded. We asked questions about their gender, age, browser usage, PETs usage, and opinions on the importance of privacy.

Second, we collected the inferred demographics displayed by Google's Ad Settings to compare to the participant's self-reported gender and age. We showed participants a screen shot of what this page looks like and included instructions on how to copy and paste the main content of the page. We asked them to paste it into a web form. We used scripts to extract various variables from

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

WPES'18, October 15, 2018, Toronto, ON, Canada

© 2018 Copyright held by the owner/author(s). Publication rights licensed to ACM.

ACM ISBN 978-1-4503-5989-4/18/10...\$15.00

<https://doi.org/10.1145/3267323.3268962>

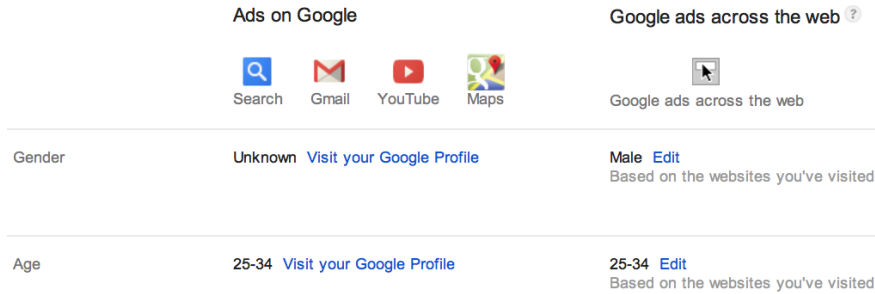


Figure 1: Screen shot of part of Google’s Ad Settings webpage while logged in with a Google account using Safari. ©Google

this page: inferred age from data from Google’s services, inferred gender from Google’s services, inferred age from data from across the web (Google’s ad network), inferred gender from across the web, whether the user was logged into Google, and whether the user opted out of Google’s interest-based ads.

Third, we conducted measurements of the participants’ web browsers. We used an invisible iframe to have their browsers visit our server, which ran a series of tests to determine whether or not first party or third party cookies were blocked, whether Google Analytics cookies were blocked, and whether their browser was transmitting the DNT header.

Our survey yielded four measures of accuracy along with numerous factors they could be associated with, making a myriad of comparisons possible. To compensate for the multiple testing problem, we split our survey responses into an exploratory set and a confirmation set. We used the exploratory set to identify associations that appear statistically significant under the χ^2 test without adjusting for multiple tests. We then tested just these associations on the confirmation data while adjusting for the total number of confirmation tests (nine) using the χ^2 test with a Bonferroni correction. For reasons of space, we report frequencies and accuracies over the whole data set despite comments on statistical significance referring to the two subsets of data.

4 RESULTS

We recruited for our survey using Mechanical Turk with an offer to pay 500 participants. On Oct. 29, 2014, 558 Turkers started our survey with 501 completing it. We eliminated 13 responses for not correctly providing us with a copy of their Ad Settings page and an additional 3 for not self-reporting gender and age.

Of the remaining responses, 4 of them were created using browsers with cookies disabled for which Google displayed a message saying as much and that it had no inferences for the person. To simplify the presentation, we eliminated these responses as well, although they could be considered additional cases of Google opting to not make an inference.

We took the first 289 of the remaining 481 responses to be our exploratory data set and the remaining 192 responses to be our confirmation data set. We use a temporal split of the data set to emphasize predictive ability.

Account Settings. Table 1 shows how many of the remaining respondents opted out of tracking on Google services or Google ads

	Google	Across the web
not opted out	455	467
opted out	26	14

Table 1: Number of respondents who opted out or in of various forms of tracking by Google

across the web. Additionally, we found that 13 of the respondents who opted out of Google ads on services also opted in for getting ads on YouTube, presumably overriding the more general opt out for that service.

We found 397 respondents to be logged into a Google account and 84 to not be. Since logged in users’ profiles are available, which makes inferences easier, we break down all further results along the lines of logged in and out users.

Overall Accuracy. Table 2 shows both the self-reported and inferred genders and ages. The values of “Unknown” and “N/A” are ones that Google lists, not comments from the authors on what we know or applicability. (“n/a” is ours.) Google gave one user an age range that overlapped with two other age ranges. We drop this range from further analysis.

Table 3 summarizes how often Google correctly stated the participants’ sexes and ages. We report the percentage of participants that Google got right, wrong, and skipped (by listing unknown or N/A). The results show that Google skipped 100% of participants who were logged out for Google services. In these cases, Google got 0% right, meaning Google is very inaccurate in one sense, but, in different sense, Google’s accuracy cannot even be evaluated since Google did not try to make inferences in these cases. Google also skipped over 70% of logged out participants for across the web. The results also show that Google rarely made a wrong prediction. From this, we conclude that Google is conservative in making predictions, but typically right when it does so.

Demographics. We checked whether the percentage Google got right was associated with actual gender or age of participants. Table 4 shows the results. Note that a column of all 0s reflects that Google did not attempt to make an inference for that combinations of factors (see Table 3), not Google guessing consistently incorrectly.

The results show that Google was right roughly as often for females as males, with the difference exceeding 10% for inferences about age for logged out users across the web, where the percentage

	Self-reported		Google inferred on			
			Services		Web	
	in	out	in	out	in	out
Female	159	35	102	0	130	11
Male	238	49	178	0	195	11
Unknown	n/a	n/a	78	0	48	40
N/A	n/a	n/a	39	84	24	22
18-24	69	8	49	0	45	2
21-35	n/a	n/a	1	0	1	0
25-34	186	41	144	0	148	8
35-44	92	17	70	0	80	2
45-54	21	11	21	0	28	2
55-64	20	5	9	0	17	2
65+	9	2	6	0	6	2
Unknown	n/a	n/a	58	0	48	44
N/A	n/a	n/a	39	84	24	22

Table 2: Number of respondents with each value broken by being logged in or out

	Right		Wrong		Skipped	
	in	out	in	out	in	out
Sex Google services	66	0	5	0	29	100
Sex Across the web	74	21	8	5	18	74
Age Google services	67	0	9	0	24	100
Age Across the web	65	11	16	11	18	79

Table 3: Google’s accuracy shown as the percentages Google got right, wrong, or skipped. We treat as skipped those Google called “Unknown” or “N/A”.

Google got right was 17% for females and just 6% for males, for an 11% difference. Given that Google only attempted to draw an age inference for 21% participants logged out for across the web, this difference might be just noise from the small number of attempts.

For age, the largest drop in the number right from the overall number is age for Google services for people who self-report an age of 65 or more, which could be just noise given the small number of participants in that age bracket. Focusing on the age brackets for which we have at least 50 participants (which cover ages from 18 to 44), we find the largest differences to be between the age brackets to be for logged out age across the web (25% vs. 6% right). The results include two statistically significant associations for age in the exploratory data set; neither of these differences proved significant in our confirmation data set.

Computer Usage. Table 5 shows the number of respondents with various usage conditions on the computer used to take our survey and the percentage of them about whom Google made correct inferences. Some of these activities intuitively makes it more difficult to correctly make inferences about any one user of the computer since they imply that the computer has multiple users, which could

	Sex services		Sex web		Age services		Age web	
	in	out	in	out	in	out	in	out
All	66	0	74	21	67	0	65	11
Female	61	0	75	26	63	0	69	17
Male	69	0	73	18	70	0	63	6
18-24	71	0	74	38	68	0	<u>57</u>	25
25-34	66	0	75	22	71	0	69	12
35-44	<u>62</u>	0	72	12	68	0	72	6
45-54	71	0	76	27	76	0	81	9
55-64	55	0	65	20	35	0	45	0
65+	78	0	78	0	11	0	11	0

Table 4: The percentage that Google got right for each reported gender and age range. Underlining shows associations with statistical significance in our exploratory data set.

pollute a model of any one of them. The decrease in accuracy is sizable in some cases and reaches statistical significance in our exploratory data set for three conditions involving the clearing cookies. However, none have a statistically significant association with Google’s error rate in our confirmation data set.

Attitudes. Table 22 in the appendix shows the associations between the respondents’ attitudes toward tracking and Google’s accuracy. Our exploratory analysis found no significant associations.

PETs. Table 6 shows the usage of various PETs and the number Google got right for users of each PET. Above the bar are the self-reported usage habits of PETs by respondents. Below the bar are our server’s measurements. For these measurements, “empty” means that our server did not detect a visit from the respondent (e.g., due to network loss). Unfortunately, the small number of users of some of the PETs limits our abilities to draw conclusions about them.

Looking at Adblock, in our exploratory data set, we found a statistically significant reduction in the accuracy of Google for data from across the web for sex both when logged in and out. A significant reduction in the accuracy for age across the web also exists, but only when logged in. The largest of these, for logged out sex across the web, was from 21% down to 7%, a drop of 14%.

The only other PET to get statistical significance is using webpage opt outs, and only in the case of age across the web when logged in. None of these differences proved statistically significant in our confirmation data set.

For two PETs, Ghostery and NoScript, the percentage that Google got correct is always 0% for logged out users. This is the best a PET can do in that logged in users may provide their demographics to Google directly, circumventing the PET. We cannot, from our observational data, conclude that these PETs caused the decrease.

5 DISCUSSION

We have no way of knowing whether the inferences shown on Ad Settings are the same as those actually used by Google for ad targeting, and prior work suggests that Ad Settings does not provide information about how Google remarkets to users based upon prior webpage visits [6] (a limitation made explicit on the Ad Settings page after the publication of [6]). Nevertheless, we

	Counts				Percent right							
	Yes		No		Sex services		Sex web		Age services		Age web	
	in	out	in	out	in	out	in	out	in	out	in	out
All (baseline)	n/a	n/a	n/a	n/a	66	0	74	21	67	0	65	11
Shared computer	62	19	335	65	63	0	77	21	71	0	69	11
Shared account	28	12	369	72	54	0	75	25	64	0	61	17
Other users in a week	107	24	290	60	62	0	74	21	69	0	69	17
Other users yesterday	61	17	336	67	59	0	67	18	66	0	64	18
Cleared cookies today/yesterday	58	36	339	48	64	0	<u>66</u>	8	64	0	59	3
Cleared cookies on close	33	17	364	67	61	0	64	18	<u>55</u>	0	<u>48</u>	6
Private mode	18	6	379	78	56	0	61	0	61	0	56	0

Table 5: The number of respondents with each computer usage characterization and the percentage that Google got right for just respondents with each computer usage characterization. For the percentages, underlining shows associations with statistical significance in our exploratory data set (i.e., a statistically significant difference from the baseline).

	Counts								Percent right							
	Yes		No		I don't know		empty		Sex services		Sex web		Age services		Age web	
	in	out	in	out	in	out	in	out	in	out	in	out	in	out	in	out
All (baseline)	n/a	n/a	n/a	n/a	n/a	n/a	n/a	n/a	66	0	74	21	67	0	65	11
AdBlock	216	45	138	28	40	11	3	0	66	0	<u>67</u>	<u>7</u>	65	0	<u>56</u>	4
Ghostery	18	8	333	69	45	7	1	0	56	0	33	0	61	0	28	0
NoScript	12	10	317	54	66	19	2	1	42	0	75	0	58	0	58	0
DoubleClick opt out	15	4	298	66	84	14	0	0	53	0	60	50	47	0	60	25
Webpage opt out	82	20	246	47	69	16	0	1	65	0	61	30	60	0	<u>57</u>	15
DNT set	100	34	217	33	80	17	0	0	57	0	68	18	65	0	65	6
DNT sent	53	22	339	55	n/a	n/a	5	7	57	0	68	23	62	0	57	9
1st cookies off	22	11	370	66	n/a	n/a	5	7	45	0	68	36	68	0	68	9
3rd cookies off	24	13	367	64	n/a	n/a	6	7	50	0	71	38	71	0	75	15
Google cookies off	39	17	353	60	n/a	n/a	5	7	54	0	64	24	59	0	59	6

Table 6: Number of users of each PET and the percentage that Google got right for users of each PET. Underlining shows associations with statistical significance in our exploratory data set. The percentages correct are broken by PET showing the percentage correct for just those participants who answered “yes” to having the PET or for whom we detected the PET.

find it noteworthy that Google rarely shows inferences for logged out users. We can only conjecture as to the reason, but perhaps one’s web browsing behavior is not as visible to or interpretable by Google as some fear. Unfortunately, since conducting our survey, Google has disabled Ad Settings for logged out users, precluding the possibility of further studying this phenomenon.

We relied upon self-reports of age and gender for ground truth, of PETs usage, and of behavior while looking for factors associated with Google’s accuracy. Self-reports of PETs usage, in particular, may be inaccurate due to the obscurity of PETs and the possibility that a shared browser may use one without the respondent’s knowledge. Furthermore, our server’s attempts to detect PETs usage by examining the behavior of respondents’ browsers could have measurement errors from factors such as network loss.

Mechanical Turkers might not be representative of standard web users. In particular, they may be more likely to use PETs or security measures due to the heavy use of their browsers for Turkling. Furthermore, they may visit an atypically large number of webpages unassociated with their demographics to fulfill their Turkling tasks.

Our exploratory results suggest that cookie clearing and AdBlock may be associated with decreasing Google’s accuracy. Using observational data, we cannot claim that they cause the decrease.

Future work includes running experiments to determine whether PET usage is the cause of such decreases in accuracy and conducting a larger-scale observational studies to bring larger number of PET users and cookie clearers into the sample. We hope this will allow us to find statistically significant associations, which in some cases appear unobtainable due to the small number of respondents with privacy-seeking behaviors (Tables 5 and 6).

ACKNOWLEDGEMENTS

This research was supported by the U.S. National Science Foundation (NSF) grants CNS 1065240 and CNS 1514509. The views and conclusions contained in this document are those of the authors and should not be interpreted as representing the official policies, either expressed or implied, of any sponsoring institution, the U.S. government or any other entity.

REFERENCES

- [1] BALEBAKO, R., LEON, P., SHAY, R., UR, B., WANG, Y., AND CRANOR, L. Measuring the effectiveness of privacy tools for limiting behavioral advertising. In *Web 2.0 Security and Privacy Workshop* (2012).
- [2] BARFORD, P., CANADI, I., KRUSHEVSKAJA, D., MA, Q., AND MUTHUKRISHNAN, S. Adscape: Harvesting and analyzing online display ads. In *Proceedings of the 23rd International Conference on World Wide Web* (Republic and Canton of Geneva, Switzerland, 2014), International World Wide Web Conferences Steering Committee, pp. 597–608.
- [3] BOOK, T., AND WALLACH, D. S. An empirical study of mobile ad targeting. *ArXiv 1502.06577* (2015).
- [4] BOSCH, T. Does Google accurately guess your age and gender? *Slate* (Jan. 2012). http://www.slate.com/blogs/future_tense/2012/01/25/google_ad_preferences_manager_does_it_accurately_guess_your_age_and_gender.html.
- [5] CAMERON, S. Does Google really know where you are? for nearly half of you, the answer is no. *Search Engine Land* (2015). searchengineland.com/google-really-know-230001.
- [6] DATTA, A., TSCHANTZ, M. C., AND DATTA, A. Automated experiments on ad privacy settings: A tale of opacity, choice, and discrimination. In *Privacy Enhancing Technologies (PoPETs)* (2015), pp. 92–112.
- [7] ENGLEHARDT, S., EUBANK, C., ZIMMERMAN, P., REISMAN, D., AND NARAYANAN, A. Web privacy measurement: Scientific principles, engineering platform, and new results. Manuscript posted at <http://randomwalker.info/publications/WebPrivacyMeasurement.pdf>, June 2014. Accessed Nov. 22, 2014.
- [8] GOOGLE. About Ads Settings. Ads Help webpage: <https://support.google.com/ads/answer/2662856?hl=en>, 2016.
- [9] GUHA, S., CHENG, B., AND FRANCIS, P. Challenges in measuring online advertising systems. In *Proceedings of the 10th ACM SIGCOMM Conference on Internet Measurement* (New York, NY, USA, 2010), pp. 81–87.
- [10] LARSON, S. I crawled my ad settings to see what Google really knows about me. *The Daily Dot* (2015). www.dailydot.com/debug/google-data-ad-settings/.
- [11] LÉCUYER, M., DUCOFFE, G., LAN, F., PAPANCEA, A., PETSIOS, T., SPAHN, R., CHAIN-TREAU, A., AND GEAMBASU, R. XRay: Increasing the web’s transparency with differential correlation. In *Proceedings of the USENIX Security Symposium* (2014).
- [12] LIU, B., SHETH, A., WEINSBERG, U., CHANDRASHEKAR, J., AND GOVINDAN, R. AdReveal: Improving transparency into online targeted advertising. In *Proceedings of the Twelfth ACM Workshop on Hot Topics in Networks* (New York, NY, USA, 2013), ACM, pp. 12:1–12:7.
- [13] NATH, S. Madscope: Characterizing mobile in-app targeted ads. In *Proceedings of the 13th Annual International Conference on Mobile Systems, Applications, and Services* (New York, NY, USA, 2015), ACM, pp. 59–73.
- [14] TSCHANTZ, M. C., DATTA, A., DATTA, A., AND WING, J. M. A methodology for information flow experiments. In *Computer Security Foundations Symposium* (2015), IEEE.
- [15] WILLS, C. E., AND TATAR, C. Understanding what they do with what they know. In *Proceedings of the 2012 ACM Workshop on Privacy in the Electronic Society* (New York, NY, USA, 2012), pp. 13–18.

A ADDITIONAL DETAILS

For the purposes of Table 5, we defined a “shared computer” to be one that respondent described as “Regularly used by multiple workers at a place of employment”, “Regularly used by multiple members of a family”, or “Regularly used by many people in a public place (library, Internet cafe, etc.)”, but not as “Regularly used only by me” nor as “None of the above”.

For the purposes of Table 22, we defined “Concerned about tracking” as answering with a 4 or 5 (very concerned) on a 5-point scale to the question “How concerned are you about online tracking of your behavior?” We defined “Confidence about avoiding it” as a 4 or 5 (very confident) to the question “If you have taken steps to prevent online tracking of your behavior, how confident are you that it prevents online tracking?”

B ADDITIONAL DATA

The following tables show the responses we received to various questions on our survey. Figure 2 provides a larger screen shot of Google’s Ad Settings.

	Right		Wrong		Skipped	
	in	out	in	out	in	out
Sex Google services	61	0	9	0	30	100
Sex Across the web	75	26	13	6	12	69
Age Google services	63	0	6	0	31	100
Age Across the web	69	17	18	11	13	71

Table 7: Google’s accuracy shown in percentages for females.

	Right		Wrong		Skipped	
	in	out	in	out	in	out
Sex Google services	69	0	2	0	29	100
Sex Across the web	73	18	5	4	22	78
Age Google services	70	0	10	0	20	100
Age Across the web	63	6	15	10	22	84

Table 8: Google’s accuracy shown in percentages for males.

Ads Settings

Settings for Google ads

Ads enable free web services and content. These settings help control the types of Google ads you see.

Ads on Google



Search



Gmail



YouTube



Maps

Google ads across the web [?]



Google ads across the web

Gender

Unknown [Visit your Google Profile](#)

Male [Edit](#)

Based on the websites you've visited

Age

25-34 [Visit your Google Profile](#)

25-34 [Edit](#)

Based on the websites you've visited

Languages

N/A

English [Edit](#)

Based on the websites you've visited

Interests

Unknown [Edit](#)
From your previous activity on Google

Air Travel, and 15 more [Edit](#)
Based on the websites you've visited

Advertisers' campaigns
you've blocked [?]

None
From your blocking activity

N/A

Opt-out settings

[Opt out](#) of *interest-based* ads on Google

[Opt out](#) of *interest-based* Google ads across the
web

Visit the [Consumer Ads Help Center](#) to learn more about how Google serves ads.

Google adheres to advertising industry privacy standards. To learn about these standards, including how you can opt out of interest-based advertising from Google and other participating companies, visit our [About Google Ads](#) page. If you want to permanently opt out of the DoubleClick cookie, you can install the [DoubleClick opt out extension](#).

[Send feedback](#)

Figure 2: Screen shot of Google's Ad Settings webpage while logged in with a Google account using Safari. ©Google

	in	out
Yes	389	80
I don't know	1	0
empty	2	0
No	5	4

Table 9: Did you use this computer yesterday?

	in	out
5 days	70	11
6 days	17	3
7 (every day)	289	63
I don't know	2	0
4 or fewer days	16	7
empty	3	0

Table 10: In the past week, on how many days did you use this computer?

	in	out
Regularly used only by me	331	65
Regularly used by multiple workers at a place of employment	6	1
Regularly used by multiple members of a family	53	17
Regularly used by many people in a public place (library, Internet cafe, etc.)	3	1
None of the above	4	0

Table 11: Which best describes this computer?

	in	out
3 or more days	57	14
0 days (no one else used it)	279	57
I don't know	11	3
2 days	18	5
1 day	32	5

Table 12: In the past week, on how many days did someone other than you use this computer?

	in	out
I don't know	12	3
No	124	22
Not applicable	218	45
Yes	28	12
Some	13	2
empty	2	0

Table 13: If anyone else used the computer you are currently using within the last week, did that person(s) use a different user account from the one you are currently using?

	in	out
Yes	61	17
I don't know	5	1
empty	0	1
No	331	65

Table 14: Did anyone else use this computer yesterday?

	in	out
1	62	9
0	123	21
3	86	25
2	71	15
5	4	1
4	51	12
empty	0	1

Table 15: If you have taken steps to prevent online tracking of your behavior, how confident are you that it prevents online tracking?

	in	out
inaccurate profile about you	53	7
accurate profile about you	134	25
empty	1	1
equally concerning	209	51

Table 16: Which is more concerning to you?

	in	out
1	14	3
3	95	15
2	44	7
5	51	21
4	190	38
empty	3	0

Table 17: How concerned are you about online tracking of your behavior?

	in	out
Week	83	11
Month+	96	10
I don't know	42	9
Never	56	11
Yesterday	41	19
Today	17	17
Month	62	7

Table 18: When was the last time you cleared the cookies of the browser you are currently using?

	in	out
Yes	18	6
I don't know	18	6
empty	1	1
No	360	71

Table 19: Are you currently using your web browser in "private browsing mode" (sometimes called "incognito")?

	in	out
Yes	33	17
I don't know	89	11
empty	5	2
No	270	54

Table 20: Does the browser you are currently using automatically clear cookies upon closing it?

	in	out
Some high school	4	2
Some college or Associate degree	135	32
Some graduate school	11	1
Master's degree	31	9
Bachelor's degree	147	23
Doctorate or professional graduate degree (Ph.D., J.D., M.D., etc.)	10	3
High school diploma or GED	59	14

Table 21: What is your highest completed level of education?

	Sex services		Sex web		Age services		Age web	
	in	out	in	out	in	out	in	out
All (baseline)	66	0	74	21	67	0	65	11
Concerned about tracking	62	0	71	25	66	0	65	10
Confidence about avoiding it	62	0	73	23	62	0	58	0
More concerned about accurate profiles	69	0	77	28	69	0	67	12
More concerned about inaccurate profiles	72	0	77	14	70	0	66	14

Table 22: The percentage that Google got right for respondents with each attitude