‘What Twitter Knows’ Extension – Dataset Exploratory Analysis

Computer Science Departmental Honors Thesis

Justin Goodman
Department of Computer Science
University of Maryland
College Park, MD, USA
jugoodma@umd.edu

Michelle Mazurek (advisor)
Department of Computer Science
University of Maryland
College Park, MD, USA
mmazurek@cs.umd.edu

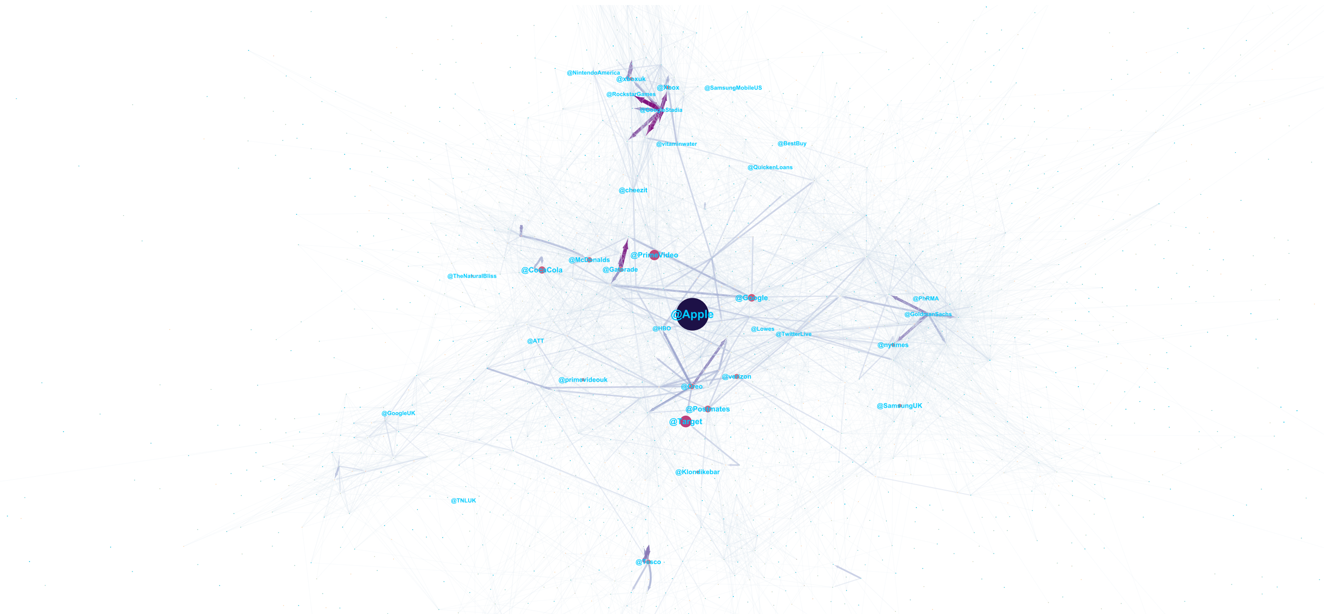

Figure 1: Follower-lookalike network structure

ABSTRACT

In a previous work, we conducted a study querying participants, using their personal downloaded Twitter data, on their opinions of advertisement targeting types and current advertisement explanations. In this work, we further explore the resulting 231-participant advertisement data and response set. Our work shows that merely shedding light and discussing one’s personal data makes the user more skeptical that the owners of said data sell it, and sheds new light on advertiser connectedness through FOLLOWER LOOKALIKE targeting.

1 INTRODUCTION

Social media is profitable by selling advertisements, and social media companies can accrue more revenue by helping advertisers target specific subsets of users [11, 13, 17]. Twitter, for example, collects and analyzes 12 terabytes of data every day (from a 2010 estimate [8]) in pursuit of helping advertisers more precisely target their ads. While collecting increasing amounts of data may be beneficial to social media companies, it puts at risk the privacy of their users. Storing user data carries privacy risk through potential data breaches [7] or identification of users in anonymized data [3]. Even if care is taken to protect against these threats, some information can be inferred through ML models that the user may not have intended to share with the company, for example, one’s pregnancy status [9]. While some governments around the world have passed legislation requiring increased transparency about the kinds of data companies collect about their users (e.g. the European Union’s General Data Protection Regulation [6]), it remains a research question the extent to which users are aware of or how they may feel about the data stored about them.

The present study seeks to better understand how consumers feel about the data stored about them by the popular ([5]) social media website Twitter. In a previous study of Twitter users, we collected each participant’s Twitter user data including targeted advertisement impressions and inferred demographics [19]. We then surveyed the participants on their opinions about relevance, fairness, and comfort with Twitter targeting ads using each targeting type. Notably, we asked the participants about their attitudes towards Twitter’s handling of user data both before and after showing them the data Twitter has stored about them. In this paper, we report how informing study participants about the data stored about them changes their beliefs about how Twitter uses their data. Although Twitter users can access their own user data at any time, we find that showing users samples of their data makes them more likely to believe Twitter sells their data. This paper presents evidence that social media companies should improve salience of the data they store about their users to ensure that individuals' preferences for data privacy are respected.

The remainder of this paper is organized as follows: Section 2 discusses related works and describes the original study, Section 3 presents our data analysis and findings, Section 4 discusses these findings, and Section 5 concludes. Appendix A contains targeting type definitions, Appendices B and C show sample Twitter data files, and Appendices D and E display the surveys from our original study.

2 BACKGROUND

Privacy is a central concern for internet users, although users frequently willingly sacrifice some privacy to use internet services (the *privacy paradox*) [14]. Many tools exist for maintaining internet privacy, including browsers (Firefox, Firefox forks, Chromium forks, Tor), extensions (ad-blockers and JavaScript blockers), virtual private networks (VPNs), internet search tools (DuckDuckGo), and messaging/mail services [25]. Further, some internet users report finding online behavioral advertising (OBA), "simultaneously useful and privacy invasive," [30]. Twitter implements OBA by allowing advertisers to run targeted advertisement campaigns. Advertisers can target users using numerous targeting types, outlined in Appendix A.

To better understand Twitter user sentiments towards data privacy, we conducted a study consisting of two surveys hosted on the Prolific survey-participant service. The first survey (Appendix D) presents the participant with a consent form, followed by instructions to download, verify, and upload to our server her Twitter advertisement data. Once the participant's data are uploaded, the program uses the upload to personalize survey questions. Only participants who manage to upload sufficient data are invited to continue to the second survey. This next survey (Appendix E) takes participants through questions involving targeting types, advertisement explanations, and demographics.

Our team concluded a few ideas from this study [19]. First, we described a full characterization of Twitter's advertisement ecosystem. Second, our respondents said that INTEREST (among others) targeting is fair, but FOLLOWER LOOKALIKE TARGETING, TAILORED AUDIENCE LISTS, EVENTS, and BEHAVIORS targeting is unfair. Third, our respondents generally agreed that current advertisement explanations lack detail.

In this paper, we extend our prior work by returning to our participants' uploaded data and survey responses. Our goal is to find patterns and provide a better understanding of what Twitter knows. We accomplish this through a series of observations, regressions, and network visualizations.

3 ANALYSIS

This section presents our methods and outcomes from our data analysis. We start by outlining the Twitter data we collected from our participants. We follow this by explaining how our survey affected participant perception on whether Twitter sells their data. Then we discuss advertisement sentiment and FOLLOWER LOOKALIKE network structure. We enlist the help of numerous `Python` packages in our analysis and visualization [4, 10, 12, 15, 18, 20, 22, 31].

3.1 Data Overview

In the original survey, we asked our participants to download their Twitter advertisement data, from which we collected three files:

- `ad-impressions.js`
- `personalization.js`
- `twitter-advertiser-list.pdf`

After pruning and cleaning, our dataset included the aforementioned three files for 231 participants. The first file shows every *ad impression* the participant saw while using Twitter in the 90 days leading to their download of the file. The second file displays demographic information that Twitter knows about the participant, as well as a list of inferred interests, partner interests, audiences, shows, and locations. The third file includes a list of *tailored audiences* and *similar audiences*, and is generated by Twitter using data from the 7 days leading to the participant's download of the file. Twitter defines both *tailored audiences* and *similar audiences* in this file. The *tailored audiences* list includes advertisers that have included the user in one or more tailored audiences, and the *similar audiences* list includes advertisers whose tailored audiences Twitter thinks the user are similar. Sample entries in the first two files are provided in Appendices B and C respectively.

In total, our participants saw 226303 advertisements from 9378 advertisers. In a given advertisement, there were an average of 8 matched targeting criteria with about 4 being unique targeting types. The FOLLOWER LOOKALIKES targeting type had the highest average number of occurrences in any particular advertisement, followed by LOCATIONS and AGE respectively. The BEHAVIORS targeting type occurred 158 times in a single advertisement, followed by KEYWORDS (147) and FOLLOWER LOOKALIKES (97). Table 1 contains the total number of occurrences of each targeting type observed in our set of advertisements. From this table we see that, for example, almost every advertisement targeted LOCATIONS but typically only used one or two values, whereas comparatively fewer advertisements targeted KEYWORDS but each advertisement that did used significantly more values.

We also have full survey (Appendix E) responses from our set of 231 participants. Our participants' demographics collected from the survey are presented in Table 2.

3.2 Survey Effect

Our original survey asked participants to rate their agreement with, "I believe that Twitter sells my data," both at the start and end. Participants' responses to this question generally moved towards more agreement, as reported in Table 3. 122 (of 231) total participants changed their answer – 95 switched to more agreement, 15 switched to more disagreement, 2 switched from unsure to neutral. The remaining 10 participants switched to unsure – 5 were previously neutral, 4 previously agreed, and 1 strongly disagreed.

These results imply that our survey likely had an effect on participant opinion that Twitter sells their data. Our agreement scale runs in $\{1, 2, 3, 4, 5\}$ with 1 being strongly agree and 5 being strongly disagree. For the following analysis, responses for *Don't Know* (6) either before or after the survey are dropped. We used linear regression to determine how much, on average, each participant changed their score following the survey. Our model fits

$$y_{it} = \beta_1 \alpha_t + \gamma_i + \epsilon_{it}$$

where i is the participant, t represents the time point before or after the survey, y_{it} is participant i's ‘Twitter sells my data’ response at time point t, α_t is 0 if t is before the survey and 1 if t is after, γ_i is the specific intercept for participant i, and ϵ_{it} is the participant-time-specific error not explained by the rest of the model. We find that, on average, users changed their reported opinions in favor of *agreement* by 0.4140 standard deviations (95% confidence interval (CI) = $[-0.538, -0.290]$ – negative because agreement is numerically lower than disagreement). Regression table with coefficients of interest are reported in Table 4. We run linear regression again to fit a second model

$$y_{it} = \beta_0 + \beta_1 \alpha_t + \mathbf{\Theta X_i} + \epsilon_{it}$$

where all variables are the same except $\mathbf{X}$ is a $k \times 1$ vector representing participant i's categorical and numeric demographic data, and $\mathbf{\Theta}$ is a $1 \times k$ vector of coefficients. We find that, on average, users changed their reported opinions in favor of *agreement* by 0.4140 standard deviations (95% confidence interval (CI) = $[-0.601, -0.227]$ Regression table with coefficients of interest are reported in Table 5.

Interestingly, our survey never asks the participant questions related to selling data. Our survey exclusively presents participants with, and asks questions about, their own data, and does not reference Twitter selling data other than the two start and end questions. On average, our participants shifted their score towards agreement by 17.50%, an effect that is statistically different from zero. We can conclude that the shifted rates in Table 3 are attributable to learning that takes place through the customized survey. Furthermore, all data shown to the study participants are readily available to them, yet several participants demonstrated their surprise by changing their scores. This finding suggests that merely providing users the option to examine data stored about them is insufficient for ensuring that a website's data collection practices are in line with their users' preferences.

3.3 Advertisement Sentiment

Tweet sentiment analysis is a well-studied area of natural language processing [16]. Twitter advertisements are shown in the form of tweets sent out by the advertiser. Given these, it seemed natural to perform sentiment analysis of our participants' advertisement tweet data.

We performed sentiment analysis using two methodologies – a `Python` `nltk` naive Bayes' classifier trained on a Twitter tweet sample dataset, and a `Python` natural-language processing library `TextBlob`. The first method returns “Positive” and “Negative” which we mapped to 1 and −1 respectively, whereas the second returns sentiments in the continuous interval $[-1, 1]$. A potential drawback of the first method is that training the NBC requires a random shuffle on the training and testing set, which may change classification accuracy. Our model achieved accuracy 99.67%. In repeated trials, however, advertisement sentiment scores did not shift significantly. Scores from the second method do not change.

Sentiment	−1	1
%	35.84%	64.16%

Table 6: Percentage of all advertisements classified as ‘Negative’ and ‘Positive’ by our Twitter sentiment naive Bayes' classifier

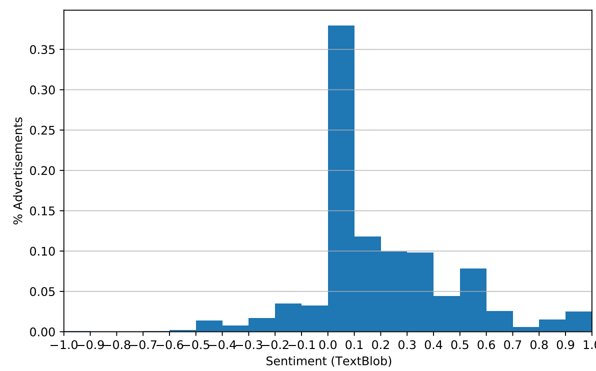

Figure 2: Histogram plotting percentage of all advertisements against binned `TextBlob` sentiment score intervals. There are 20 intervals, all of the form $[x_i, x_{i+1})$, with the last interval including 1.0

Sentiment	−1	1
%	34.30%	65.70%

Table 7: Percentage of *unique* advertisements classified as ‘Negative’ and ‘Positive’ by our Twitter sentiment naive Bayes' classifier

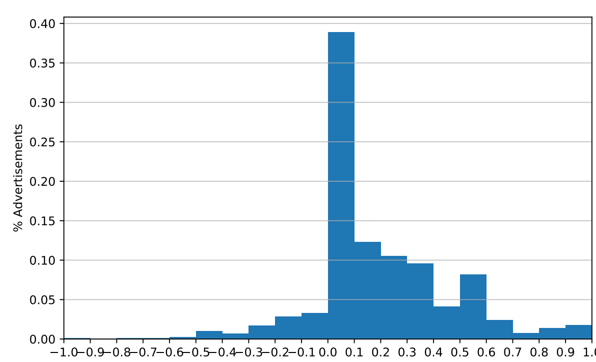

Figure 3: Histogram plotting percentage of *unique* advertisements against binned `TextBlob` sentiment score intervals. There are 20 intervals, all of the form $[x_i, x_{i+1})$, with the last interval including 1.0

We first notice that our Twitter classifier rates our advertisements as more positive (Table 6). Our TextBlob sentiment scores rate our advertisements as more neutral but leaning positive (Figure 2). These trends also hold when restricting to unique advertisements (Table 7, Figure 3).

The top advertisers, those who advertised the most to our participants, tended to use more neutral, leaning positive, advertisements. This held for both sentiment analysis methods (Figure 4). Table 8 lists the top advertisers' most frequent INTERESTS targeting value. Arguably, these topics give us a good descriptor for the advertiser. For example, @Apple produces the iPhone, iPad, Mac computers and desktops, and other assorted technologies [2]. We would expect, then, that @Apple would advertise to people interested in technology. We do, in fact, see this is true since they targeted our participants interested in *Tech news*. Similarly, @Target, a retail store, advertised to our participants interested in *Cooking* and *Skin care* – topics related to products one could buy at Target [24]. However, this does not give us a perfect summary of an advertiser. @Klondikebar solely targeted people interested in *Comedy*, yet Klondike sells ice cream bars and sandwiches [29]. They did, however, target people based on FOLLOWER LOOKALIKES (seemingly a list of comedians), KEYWORDS (mostly related to parenting), and CONVERSATION TOPICS (related to WWE wrestling). This gives us more context as to @Klondikebar's advertisement campaign (at the time of the original study), and less so context to @Klondikebar itself.

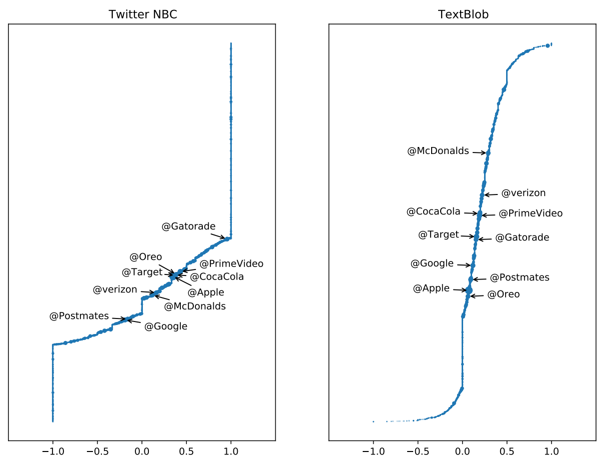

Figure 4: Advertiser plotted against average tweet sentiment score. Advertisers are listed along the y-axis in ascending sentiment score order. Point size reflects proportional number of advertisements. Top 10 advertisers, based on number of advertisements, are shown

In the first sentiment method (Twitter NBC), advertisements that targeted *Sunglasses; Women's bags; Health, mind, and body; Fine dining; Advertising; Biology; Liquor and spirits; Weddings; Dresses and skirts; Martial arts* INTERESTS were on-average tagged as negative, and *Canada; Italy; Africa; France; Language learning; Vintage cars; Data centers; Christian and gospel; Banking; Traveling with kids* INTERESTS were on-average tagged as positive. In the second sentiment method (TextBlob), advertisements that targeted *Liquor and spirits; Nonprofit; Politics and current events; Boxing; Romance; Fine dining; Face care; Linux; Canada; Vintage cars* INTERESTS all on-average leaned more negative, and *Martial arts; R&B and soul; Sunglasses; Women's bags; Language learning; Ice hockey; Weddings; Offroad vehicles; Babies and toddlers; Marketing* INTERESTS all on-average leaned more positive. There was no apparent correlation between tweet sentiment and the targeting types used to match said tweet.

3.4 Follower Lookalike Network

Advertisers can target Twitter users using the FOLLOWER LOOKALIKE targeting type, which allows advertisers to reach users who share interests with the given list of accounts. For example, the @Apple Twitter account's top FOLLOWER LOOKALIKE accounts are listed in Table 9. Some of @Apple's top FOLLOWER LOOKALIKE accounts are also advertisers. Twitter allows advertisers to list any Twitter account as a FOLLOWER LOOKALIKE – advertisers and regular users. This allows us to create a **directed** network $G = (V, E)$ where V is the union of the set of advertisers in our participant data and the set of observed FOLLOWER LOOKALIKES, and $(u, v) \in E$ with weight w if and only if v appears in $w > 0$ of u's advertisements. Figure 1 visualizes this graph. Darker and larger nodes represent more advertisements from that node, and darker and larger edges represent more occurrences of the destination node in the source node's advertisements as FOLLOWER LOOKALIKES. This graph gives us more visual intuition for how connected certain advertisers are with each other.

Finally, our *participants* might share matched FOLLOWER LOOKALIKE accounts in common. In fact, almost every pair of participants share at least one common matched FOLLOWER LOOKALIKE account. As such, an undirected graph with participants as vertices and number of shared FOLLOWER LOOKALIKE accounts as edge weights is nearly a complete graph. While analyzing this graph may not appear fruitful, we can instead compare participant demographics and Twitter advertisement information with shared FOLLOWER LOOKALIKES. We do this using a binned scatter plot with a log scale on both axes. In Figure 5, we calculate the percentage of shared FOLLOWER LOOKALIKE accounts of each pair of participants' data and plot that against the cosine similarity of the same pair. In the cosine similarity calculation, we consider both categorical data (end of survey, Appendix D) – gender, age range, education, technical background, income, and their responses to if they think Twitter sells their data – and numeric data – reported days on Twitter[1], reported average time per day on Twitter[1], number of interests[2], partner interests[2], audiences[2], advertisers[2], shows[2], and locations[2], number of advertisements seen[3], number of advertisements seen by their top advertiser[3], and their average advertisement tweet sentiment from both methods[3]. Categorical data is translated into a sparse vector (e.g.: a gender score ($\in \{1, 2, 3, 4, 5\}$) of 2 is translated to $[0, 1, 0, 0, 0]$), and numeric data is scaled out of the max from that column's observations. Percentage of shared advertisers

[1]Collected by us, Appendix E
[2]Collected from personalization.js files, Appendix C
[3]Calculated from ad-impressions.js files, Appendix B

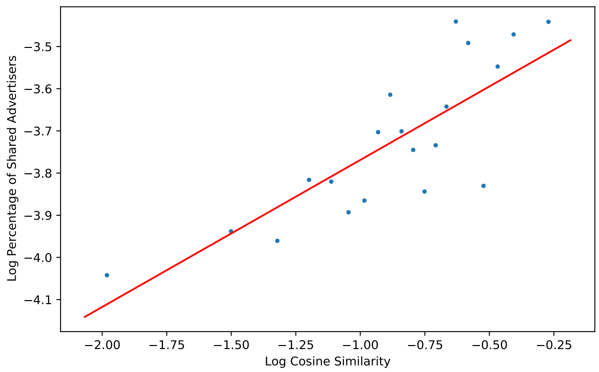

Figure 5: Pairwise participant data cosine similarity versus percentage of shared advertisers, as a binned scatter plot. Both axes plotted on a log scale to reduce bias. Cosine similarity originally $\in (0,1] \subset [-1,1]$ because nearly all data values were non-negative

is calculated as

$$\frac{A_i \cap A_j}{A_i \cup A_j}$$

where A_n is the set of advertisers for participant n. We see from this figure a linear relationship.

We also fit a linear model

$$y_{i,j} = \beta_0 + \beta_1 x_{i,j} + \epsilon_{i,j}$$

where $x_{i,j}$ is cosine similarity between users i and j scaled by the sample standard deviation, $y_{i,j}$ is percentage of shared advertisers between users i and j also scaled by the sample standard deviation, and $\epsilon_{i,j}$ is the user-pair-specific error not explained by the rest of the model. This yields $\beta_0 = 0.3902$ (95% CI = $[0.350, 0.430]$) and $\beta_1 = 0.1426$ (95% CI = $[0.128, 0.157]$). So a 1 standard deviation change in cosine similarity similarity is correlated with a 0.1426 standard deviation increase in percentage of shared advertisers. Our set of cosine similarities has $\mu = 0.4452$ and $\sigma = 0.1519$, and our set of percent shared advertisers has $\mu = 0.0535$ and $\sigma = 0.0662$.

4 DISCUSSION

The previous section inspires discussions about user privacy and advertiser networks. In this section, we delve into these discussions further, providing post-analysis commentary.

4.1 Does Twitter Sell Your Data?

A majority of our participants think that Twitter sells their data, and they are correct in their speculations. Twitter makes money by selling advertisement campaigns to advertisers (86.53% of revenue), and offering "data licensing" among other sources (13.47% of revenue) [26]. Twitter states in their privacy policy that,

> We use the information described in this Privacy Policy to help make our advertising more relevant to you, to measure its effectiveness, and to help recognize your devices to serve you ads on and off of Twitter. [28]

Their privacy policy never uses the term *sell*, instead opting to use the term *share*. However, Twitter sets pricing levels for their API access, which includes paid access to Twitter's tweet archive [27]. This gives payers access to all public tweet data – username, profile picture, tweet text, and tweet location, to name a few. We conclude, then, that Twitter's *data licensing and other revenue* includes selling user data.

This reminds us of the 2006 AOL search data dump, when supposedly anonymized search data were published but users could still be de-anonymized through their search queries (e.g.: user #4417749 was traced to Thelma Arnold by *The New York Times*) [3]. AOL's story shows us the importance of keeping private data secure, or at least anonymous, and reminds us of the importance of privacy. Although AOL's search data were private, Twitter explicitly only sells public data (and certain tracking data like user device IDs). While Twitter is not legally obligated to state in their privacy policy (that likely few users will read [23]) that they sell user data, the term *share* might be intentionally misleading.

Moreover, Twitter provides settings, outlined in their privacy policy, for controlling access to your data – setting your account to private, turning off advertisement personalization, and others [28]. Yet, users are unlikely to know about these settings and companies are unlikely to present users with usable transparency tools [21]. In fact, only 14.29% of our participants reported having visited their Twitter settings to examine or edit their advertising preferences. This further reiterates the need for data transparency, and making transparency tools well-known and easily accessible (without going too far [1]).

4.2 One-sided Advertisers

The transpose plot from Figure 5 is intuitive – Twitter allows advertisers to target advertisement campaigns based on demographics, so we expect similar users to be targeted by the same advertiser. On the other hand, the converse is not necessarily intuitive – we have no expectations for whether similar users will have more advertisers in common. We see, though, that this situation is correlational. This inspires the plausibility of *one-sided advertisers* – a situation, we coin, that occurs when a specific subset of users sees advertisements from the same advertiser(s).

The situation of one-sided advertisers is dangerous. For example, if a certain subset of users sees the same political advertisements over and over again, this is equivalent to political manipulation and causes polarization. Another example, a subset of users could see the same, potentially unwanted, advertisement about some socially relevant topic (e.g.: COVID-19, LGBTQ+, censorship, sports). Twitter explicitly states in their privacy policy that,

> In addition, our ads policies prohibit advertisers from targeting ads based on categories that we consider sensitive or are prohibited by law, such as race, religion, politics, sex life, or health. [28]

Twitter indeed has a policy explicitly banning certain advertisement subject matter – for example, advertisers cannot advertise illegal content, and only news publishers can advertise political content [28]. This does not explicitly ban socially relevant content (for example, health). If a subset of Twitter users indeed become subject to one-sided advertisers, and the advertisements they see are related to Twitter's banned-from-targeting sensitive topics (e.g.: socially charged advertisements), then in-essence these advertisements are being targeted on the topic of those advertisements. As a side-effect of targeted advertising, this breaks Twitter's privacy policy. This is not unheard of, however, as we noticed in our analysis during the initial study that some advertisers might already be breaking Twitter's privacy policy.

5 CONCLUSION

In this work, we provide an extended analysis from data collected in our previous study. We observe that our survey participants on average became more inclined to believe that Twitter sells their data, even though our survey only showed participants data readily available to them via Twitter. Further, we notice a mostly neutral, leaning positive, trend for advertisement sentiment. Finally, we see potential in analyzing FOLLOWER LOOKALIKE networks, and their implications for Twitter's rules on sensitive-topic targeting. We hope this work shows the fruitful nature in Twitter data analysis, and further justifies the need for clearer transparency tools and practices.

ACKNOWLEDGMENTS

Thank you to my brother for teaching me how to apply my knowledge of statistics. Thank you to Michelle, Leilani, and Dave for guiding me throughout my undergraduate career in the honors program. Thank you to everyone aiding society in getting us through COVID-19.

REFERENCES

[1] Acquisti, A., Adjerid, I., and Brandimarte, L. Gone in 15 seconds: The limits of privacy transparency and control. *IEEE Security Privacy 11*, 4 (2013), 72–74.

[2] Apple Inc. Apple. https://www.apple.com/, 2020.

[3] Barbaro, M., and Zeller Jr., T. A Face Is Exposed for AOL Searcher No. 4417749. *The New York Times* (2006).

[4] Bird, S., Loper, E., and Klein, E. *Natural Language Processing with Python: Analyzing Text with the Natural Language Toolkit.* O'Reilly Media, 2009.

[5] DataReportal. Global social media overview. https://datareportal.com/social-media-users, 2020.

[6] European Commission. *General Data Protection Regulation.* Luxembourg: Office for Official Publications of the European Communities, 2018.

[7] FTC. Equifax Data Breach Settlement. https://www.ftc.gov/enforcement/cases-proceedings/refunds/equifax-data-breach-settlement, 2020.

[8] Gura, D. All Those 140-Character Twitter Messages Amount To Petabytes Of Data Every Year. https://www.npr.org/sections/thetwo-way/2010/09/28/130199229/all-those-140-character-twitter-messages-yield-four-petrabytes-of-data-annually, 2010.

[9] Hill, K. How Target figured out a teen girl was pregnant before her father did. https://www.forbes.com/sites/kashmirhill/2012/02/16/how-target-figured-out-a-teen-girl-was-pregnant-before-her-father-did/, 2012.

[10] Hunter, J. D. Matplotlib: A 2d graphics environment. *Computing in science & engineering 9*, 3 (2007), 90–95.

[11] Kaplan, A. M., and Haenlein, M. The fairyland of second life: Virtual social worlds and how to use them. *Business Horizons 52*, 6 (2009), 563 – 572.

[12] Kluyver, T., Ragan-Kelley, B., Pérez, F., Granger, B., Bussonnier, M., Frederic, J., Kelley, K., Hamrick, J., Grout, J., Corlay, S., Ivanov, P., Avila, D., Abdalla, S., and Willing, C. Jupyter notebooks – a publishing format for reproducible computational workflows. In *Positioning and Power in Academic Publishing: Players, Agents and Agendas* (2016), F. Loizides and B. Schmidt, Eds., IOS Press, pp. 87 – 90.

[13] Knowlson, K. Targeted advertising, May 15 2003. US Patent App. 10/011,953.

[14] Kokolakis, S. Privacy attitudes and privacy behaviour: A review of current research on the privacy paradox phenomenon. *Computers & Security 64* (2017), 122 – 134.

[15] Loria, S. textblob documentation. *Release 0.15 2* (2018).

[16] Martínez-Cámara, E., Martín-Valdivia, M., López, L., and Montejo-Ráez, A. Sentiment analysis in twitter. *Natural Language Engineering 20* (01 2014), 1–28.

[17] McFarlane, G. How Facebook, Twitter, Social Media Make Money From You. https://www.investopedia.com/stock-analysis/032114/how-facebook-twitter-social-media-make-money-you-twtr-lnkd-fb-goog.aspx, 2020.

[18] McKinney, W., et al. Data structures for statistical computing in python. In *Proceedings of the 9th Python in Science Conference* (2010), vol. 445, Austin, TX, pp. 51–56.

[19] Miranda Wei, Madison Stamos, S. V. N. R. J. G. M. H. D. F. B. W. M. L. M. B. U. What twitter knows: Characterizing ad targeting practices, user perceptions, and ad explanations through users' own twitter data. (in press), 2020.

[20] Oliphant, T. E. *A guide to NumPy*, vol. 1. Trelgol Publishing USA, 2006.

[21] Schaub, F., Balebako, R., and Cranor, L. F. Designing effective privacy notices and controls. *IEEE Internet Computing* (2017), 1–1.

[22] Seabold, S., and Perktold, J. statsmodels: Econometric and statistical modeling with python. In *9th Python in Science Conference* (2010).

[23] Steinfeld, N. "i agree to the terms and conditions": (how) do users read privacy policies online? an eye-tracking experiment. *Computers in Human Behavior 55* (2016), 992 – 1000.

[24] Target Brands, Inc. Target : Expect More, Pay Less. https://www.target.com/, 2020.

[25] Taylor, S. Privacy Tools - How to Restore Your Online Privacy in 2020. https://restoreprivacy.com/privacy-tools/, 2020.

[26] Twitter, Inc. Form 8-k. https://investor.twitterinc.com/financial-information/sec-filings/sec-filings-details/default.aspx?FilingId=13891530.

[27] Twitter, Inc. Pricing – Twitter Developers. https://developer.twitter.com/en/pricing, 2020.

[28] Twitter, Inc. Twitter Privacy Policy. https://twitter.com/privacy, 2020.

[29] Unilever. Klondike® Ice Cream Bars and Sandwiches. https://www.klondikebar.com/, 2020.

[30] Ur, B., Leon, P. G., Cranor, L. F., Shay, R., and Wang, Y. Smart, useful, scary, creepy: Perceptions of online behavioral advertising. In *Proceedings of the Eighth Symposium on Usable Privacy and Security* (New York, NY, USA, 2012), SOUPS '12, Association for Computing Machinery.

[31] Van Rossum, G., and Drake, F. L. *Python 3 Reference Manual.* CreateSpace, Scotts Valley, CA, 2009.

Targeting type	Occurrences	Percentage of advertisements	Most frequent value
Follower look-alikes	553563	41.20%	@netflix
Keywords	333381	29.88%	parenting
Locations	232801	99.52%	United States
Age	167637	74.08%	18 and up
Conversation topics	124018	24.45%	Food
Tailored audiences (lists)	110616	12.36%	Lifetime suppression list (email)
Interests	81857	17.57%	Comedy
Behaviors	33836	3.11%	US - Household income: $30,000 - $39,999
Platforms	31498	13.92%	iOS
Languages	30773	13.60%	English
Gender	24521	10.84%	Female
Movies and TV shows	22095	6.38%	Love Island
Mobile audience targeting	21225	2.52%	Sign up Postmates - Local Restaurant Delivery & Takeout ANDROID All
Tailored audiences (web)	18061	4.10%	Quote Finish
Events	17353	5.04%	2019 Women's World Cup
Retargeting campaign engager	15218	3.95%	Retargeting campaign engager: 21155786
Retargeting engagement type	10166	4.49%	Retargeting engagement type: 2
OS version	7130	3.15%	iOS 10.0 and above
Device model	2746	1.21%	iPhone 8
Carriers	1371	0.61%	T-Mobile UK
Retargeting user engager	1370	0.61%	Retargeting user engager: 24742040
Followers of a user id	1294	0.57%	@nytimes
Tailored audience CRM lookalike targeting	1025	0.29%	amp_hva_xomcorp_yphisp_022519.csv_1_41230849
Flexible audience targeting	378	0.08%	iOS > Recently Installed (14days), No Checkout Initiated
New device targeting	212	0.09%	1 month
Retargeting custom audience lookalike targeting	189	0.07%	Tableau.com
Mobile audience lookalike targeting	140	0.05%	Install New York Times Crossword IOS All
WiFi only targeting	108	0.05%	WiFi-only
Flexible audience lookalike targeting	7	0.00%	All WBGs Android Purchase Events

Table 1: Observed targeting types in our set of participant seen advertisements, sorted by occurrences. An advertisement may target multiple values of the same targeting type. Each targeting type lists a percentage of advertisements that contained at least one value of that type, as well as the most frequent observed targeting value. Separation lines placed for readability

Female	Male	Non-binary
52.81%	45.45%	1.30%

(a) Gender

18-24	25-34	45-54	55-64	65+
31.60%	35.07%	19.91%	9.09%	3.90%

(b) Age

≤ High school	Some college	Trade/tech./voc.	Associate's	Bachelor's	Master's	Professional
12.12%	20.78%	3.90%	6.06%	42.42%	12.12%	2.60%

(c) Education

Tech.	Non-tech.
20.78%	76.19%

(d) Background

< $20,000	$20,000-$49,999	$50,000-$99,999	$100,000-$249,999	≥ $250,000
17.75%	34.63%	31.60%	11.26%	0.87%

(e) Income

Table 2: Survey participant demographic distribution. Abstained response percentages not displayed

	Strongly agree	Agree	Neither agree nor disagree	Disagree	Strongly disagree	Don't know
Before	12.99%	45.45%	19.48%	11.69%	2.16%	8.23%
After	32.90%	42.42%	8.23%	7.36%	2.60%	6.49%

Table 3: Percetage of participants' responses to, "I believe that Twitter sells my data," before and after the survey

Model $y_{it} = \beta_1 \alpha_t + \gamma_i + \epsilon_{it}$

	coef	**std err**	**z**	**P > |z|**	**[0.025**	**0.975]**
Participant ID	⋮	⋮	⋮	⋮	⋮	⋮
α_t	-0.4140	0.063	-6.562	0.000	-0.538	-0.290

No. Observations:	404	**Covariance Type:**	HC1
Df Residuals:	201	**Df Model:**	202
R-squared:	0.799	**Adj. R-squared:**	0.598
F-statistic:	nan	**Prob (F-statistic):**	nan
Omnibus:	41.727	**Durbin-Watson:**	2.957
Prob(Omnibus):	0.000	**Jarque-Bera (JB):**	243.739
Skew:	-0.000	**Prob(JB):**	1.18e-53
Kurtosis:	6.805	**Cond. No.**	14.3

Table 4: Linear regression predicting how much, normalized, our participants changed their agreement rating to *Twitter sells my data*

Model $y_{it} = \beta_0 + \beta_1\alpha_t + \Theta X_i + \epsilon_{it}$

	coef	std err	z	P > |z|	[0.025	0.975]
Intercept	-0.5551	0.376	-1.477	0.140	-1.292	0.181
Man	-0.2475	0.117	-2.114	0.035	-0.477	-0.018
Non-binary	-0.0447	0.241	-0.186	0.853	-0.517	0.427
(gender) Prefer not to say	0.4670	0.625	0.747	0.455	-0.758	1.692
(age) 25-34	-0.0038	0.137	-0.028	0.978	-0.272	0.264
35-44	-0.0003	0.179	-0.002	0.999	-0.352	0.351
45-54	0.2095	0.197	1.064	0.287	-0.176	0.595
55-64	-0.0400	0.240	-0.167	0.868	-0.511	0.431
65 or older	0.4186	0.326	1.286	0.199	-0.220	1.057
(education) High school	0.8033	0.351	2.290	0.022	0.116	1.491
Some college	0.3669	0.351	1.044	0.297	-0.322	1.056
Trade, technical, or vocational training	0.3113	0.394	0.791	0.429	-0.460	1.083
Associate's degree	0.4404	0.404	1.090	0.276	-0.351	1.232
Bachelor's degree	0.2208	0.331	0.667	0.505	-0.428	0.870
Master's degree	0.2838	0.363	0.783	0.434	-0.427	0.995
Professional degree	-0.2099	0.371	-0.565	0.572	-0.938	0.518
I do not have education/work in CS/related	-0.0531	0.155	-0.341	0.733	-0.358	0.252
(background/job field) Prefer not to say	-0.3613	0.507	-0.712	0.476	-1.355	0.633
$20,000 to $49,999	0.3542	0.187	1.889	0.059	-0.013	0.722
$50,000 to $99,999	0.3095	0.180	1.717	0.086	-0.044	0.663
$100,000 to $249,999	0.1674	0.213	0.786	0.432	-0.250	0.585
Over $250,000	-0.4829	0.406	-1.191	0.234	-1.278	0.312
(income) Prefer not to say	-0.1605	0.239	-0.672	0.501	-0.629	0.308
I have not used my Twitter ad settings	0.0068	0.168	0.041	0.968	-0.323	0.336
(Twitter ad settings) Don't know	0.2659	0.310	0.859	0.391	-0.341	0.873
α_t	-0.4140	0.095	-4.345	0.000	-0.601	-0.227
daysOnTwitter	4.798e-05	5.2e-05	0.922	0.356	-5.4e-05	0.000
avgTimePerDay	0.1315	0.072	1.822	0.068	-0.010	0.273
numInterests	-0.0004	0.001	-0.288	0.773	-0.003	0.002
numPartnerInterests	0.0006	0.001	0.707	0.479	-0.001	0.002
numAudiences	-0.0004	0.000	-1.849	0.065	-0.001	2.47e-05
numAdvertisers	0.0013	0.001	1.561	0.119	-0.000	0.003
numShows	-0.0004	0.002	-0.251	0.802	-0.004	0.003
numLocations	0.0648	0.039	1.644	0.100	-0.012	0.142
numAdsSeen	-6.812e-05	3.98e-05	-1.712	0.087	-0.000	9.88e-06
topAdvertiserNumAds	-0.0003	0.001	-0.455	0.649	-0.002	0.001
avgAdTweetSentimentTextBlob	1.1451	1.143	1.001	0.317	-1.096	3.386
avgAdTweetSentimentTwitter	-0.2138	0.316	-0.677	0.498	-0.833	0.405

No. Observations:	404	**Covariance Type:**	HC1
Df Residuals:	366	**Df Model:**	37
R-squared:	0.167	**Adj. R-squared:**	0.083
F-statistic:	3.448	**Prob (F-statistic):**	7.46e-10
Omnibus:	57.143	**Durbin-Watson:**	1.404
Prob(Omnibus):	0.000	**Jarque-Bera (JB):**	80.578
Skew:	0.956	**Prob(JB):**	3.18e-18
Kurtosis:	4.056	**Cond. No.**	1.28e+05

Table 5: Linear regression results – predicting *Twitter sells my data* agreement score, normalized, against categorical and numeric demographic data

@Apple (0.176, 0.071)	Tech news Music news and general info Music festivals and concerts Sports news Online gaming
@Target (0.117, 0.157)	Cooking Make-up and cosmetics Comedy Skin care Celebrity fan and gossip
@PrimeVideo (0.667, 0.192)	Sci-fi and fantasy Comedy Online gaming Movie news and general info Computer gaming
@Google (−0.276, 0.117)	Comedy Sports news Sporting events Technology Business news and general info
@CocaCola (0.393, 0.195)	Music festivals and concerts Comedy Sporting events Music news and general info Console gaming
@Postmates (0.700, 0.092)	Sporting goods Action sports Soccer Health news and general info Exercise and fitness
@Oreo (0.056, 0.062)	Comedy Entertainment awards Sports themed Foodie news and general info Cooking
@McDonalds (0.199, 0.230)	Comedy Technology Computer gaming Online gaming Console gaming
@Gatorade (0.745, 0.152)	Sports news NFL football Sporting goods NBA basketball Soccer
@verizon (0.023, 0.221)	Technology Music festivals and concerts Sporting events NFL football Music news and general info

Table 8: Top 10 advertisers shown with their top 5 INTERESTS targeting value. Scores (a, b) under each advertiser are average advertisement sentiment: a is from naive Bayes' tweet classifier, b is from `TextBlob`

Account	Count
@WIRED	154
@TEDTalks	141
@verge	130
@Spotify	112
@AppleMusic	111
@TechCrunch	108
@Twitter	106
@mashable	101
@Gizmodo	87
@kanyewest	83

Table 9: @Apple's top 10 FOLLOWER LOOKALIKE accounts. Bold-font accounts have also advertised to participants in our study

A DEFINITIONS OF TARGETING TYPES

This section features the precise terminology and definitions for the 16 targeting types investigated in the original user study, all taken (nearly) verbatim from Twitter Business's help pages (https://business.twitter.com/en/targeting.html).

- **Age targeting** allows advertisers to target people by age buckets, such as 18+ years old or 18-24 years old.
- **Behavior targeting** allows advertisers to target people based on inferred behavior, such as shopping and lifestyle habits or income.
- **Conversation topic targeting** allows advertisers to target people based on topics they have engaged with (e.g., Tweeted, clicked, Retweeted, replied, liked, viewed) on Twitter.
- **Event targeting** allows advertisers to target people based on events they are interested in or have engaged with (e.g., Tweeted, clicked, Retweeted, replied, liked, viewed) on Twitter.
- **Follower lookalike targeting** allows advertisers to target people who don't necessarily follow a given account, but have similar interests or demographics to the account's actual followers.
- **Gender targeting** allows advertisers to target people based on their self-reported or inferred gender.
- **Interest targeting** allows advertisers to target people based on inferred interests, as determined by who they follow on Twitter and their Tweets, Retweets, and clicks.
- **Keyword targeting** allows advertisers to target people based on words or phrases they have Tweeted about or searched for on Twitter.
- **Language targeting** allows advertisers to target people who use a certain language on Twitter.
- **Location targeting** allows advertisers to target people based on region, city, metro or zip code.
- **Mobile audience targeting** allows advertisers to target people who use their mobile app.
- **Movie and TV show targeting** allows advertisers to target people based on movies and TV shows they have watched or are likely to watch.
- **Platform targeting** allows advertisers to target people who use a certain platform, such as iOS or Desktop, to access Twitter.
- **Retargeting campaign engager targeting** allows advertisers to target people based on prior engagement with (e.g., Tweeting, clicking, Retweeting, replying, liking, or viewing) their company.
- **Tailored audience (list) targeting** allows advertisers to reach specific people on Twitter by uploading lists, which contain personal information (email addresses, phone numbers, or Twitter handles) that are matched to Twitter users' accounts.
- **Tailored audience (web) targeting** allows advertisers to target people who have visited their website.

B RAW AD IMPRESSIONS DATA EXAMPLE

This section features a sample of the author's personal `ad-impressions.js` file.


```
window.YTD.ad_impressions.part0 = [ {
  "ad" : {
    "adsUserData" : {
      "adImpressions" : {
        "impressions" : [ {
          "deviceInfo" : {
            "osType" : "Android",
            "deviceId" : <long string of characters>,
            "deviceType" : "LG V20"
          },
          "displayLocation" : "TimelineHome",
          "promotedTweetInfo" : {
            "tweetId" : "1095831177896886272",
            "tweetText" : "Brian and Michael wrote the playbook on making delicious cupcakes.
                           With Surface Pro 6, they can share them with the world.",
            "urls" : [ ],
            "mediaUrls" : [ ]
          },
          "advertiserInfo" : {
            "advertiserName" : "Microsoft Surface",
            "screenName" : "@surface"
          },
          "matchedTargetingCriteria" : [ {
            "targetingType" : "Interests",
            "targetingValue" : "Tech news"
          }, {
            "targetingType" : "Interests",
            "targetingValue" : "Technology"
          }, {
            "targetingType" : "Follower look-alikes",
            "targetingValue" : "@verge"
          }, {
            "targetingType" : "Follower look-alikes",
            "targetingValue" : "@mashable"
          }, {
            "targetingType" : "Age",
            "targetingValue" : "18 and up"
          }, {
            "targetingType" : "Locations",
            "targetingValue" : "United States"
          } ],
          "impressionTime" : "2019-02-25 13:12:22"
        }, { ... }
        ... ]
      }
    }
  }, { ... }
  ...
} ]
```

C RAW PERSONALIZATION DATA EXAMPLE

This section features a sample of the author's personal `personalization.js` file.

```
window.YTD.personalization.part0 = [ {
  "p13nData" : {
    "demographics" : {
      "languages" : [ {
        "language" : "English",
        "isDisabled" : false
      } ],
      "genderInfo" : {
        "gender" : "male"
      }
    },
    "interests" : {
      "interests" : [ {
        "name" : "Animals",
        "isDisabled" : false
      }, {
        "name" : "Animation",
        "isDisabled" : false
      }, {
        "name" : "Comedy",
        "isDisabled" : false
      }, { ... },
      ...
      ],
      "partnerInterests" : [ {
        "name" : "Auto > Auto service buyer"
      }, {
        "name" : "Auto > Body style: cross over vehicle"
      }, {
        "name" : "Auto > Body style: entry/economy/compact"
      }, { ... },
      ...
      ],
      "audienceAndAdvertisers" : {
        "numAudiences" : "414",
        "advertisers" : [ "@1981Marlen", "@5gum", "@ATT", ... ]
      },
      "shows" : [ "Black Mirror (Netflix)", "Christmas 2017", "Electric Daisy Carnival Las Vegas 2017", ... ]
    },
    "locationHistory" : [ "LOS ANGELES, USA", "MD, USA" ]
  }
} ]
```

D INSTRUCTIONS PROVIDED TO PARTICIPANTS FOR DATA REQUEST (PART 1 OF THE STUDY)

Below is the text we provided to explain to participants how to request their Twitter data. We included detailed and annotated screenshots highlighting each step of the process.

D.1 Consent Form

Study Title: Twitter Ad Transparency

DESCRIPTION: We are researchers at [redacted] doing research to better understand Twitter advertising transparency. In this survey, you will be asked about your experiences and opinions about Twitter. People who are age 18+ and live in the United States or United Kingdom are eligible to participate. Additionally, you must have an active Twitter account. Participation consists of two parts: first, a short 5-minute preliminary survey, and then the main survey, which should take about 25 minutes.

RISKS and BENEFITS: The risks to your participation in this online study are those associated with basic computer tasks, including boredom, fatigue, mild stress, or breach of confidentiality. The only benefit to you is the learning experience from participating in a research study. The benefit to society is the contribution to scientific knowledge.

COMPENSATION: Participants who complete all tasks will be compensated $7.86: $0.86 for Part 1 and $7.00 for Part 2.

CONFIDENTIALITY: No personally-identifiable information will be collected from you. Any reports and presentations about the findings from this study will not include your name or any other information that could identify you. In some cases, you might provide personal stories or beliefs that we might quote or paraphrase as part of our research findings – any personally identifying information will be removed to protect your privacy. We may share the data we collect in this study with other researchers doing future studies – if we share your data, we will not include information that could identify you.

SUBJECT'S RIGHTS: Your participation is voluntary. You may stop participating at any time by closing the browser window or the program to withdraw from the study.

[Additional content removed for anonymity]

Please indicate below, that you are at least 18 years old, have read and understand this consent form, and agree to participate in this online research study.

I am at least 18 years old.

◯ Yes ◯ No

I have read and understood this consent form.

◯ Yes ◯ No

I agree to participate in this research study.

◯ Yes ◯ No

D.2 Introduction and Instructions

Thank you for your participation in our study.

In Part 1 of this study (today), you will log into your Twitter account and request two data downloads. On the next page, you will be guided through the process of requesting your Twitter data.

NOTE: The information we collect in this study will not include your personal information. We will NOT ask for your Twitter username, messages, tweets, etc.

We are only interested in data about ads you have seen on Twitter. You will request your entire Twitter data archive today, but in Part 2 of this study, we will provide instructions for uploading only the data we need for our research.

There are two downloads that you need to request in this part of the study.

How to make the first request:

1) Log into Twitter: `https://twitter.com` (opens in a new tab).
2) Click on "More" at the bottom left. On narrower screens, it may only display the icon with three dots, without the word "More".
3) Click "Settings and privacy".
4) Click "Your Twitter data" at the bottom of the menu on the right side.
5) If prompted, enter your Twitter password.
6) Scroll to the bottom of the page. In the "Download your data" section, click "Request data" in the Twitter row.

If you do not see a button that says "Request data" where the red box appears above, this means you have already requested your data. Continue to the instructions below.

Twitter will email you when your download is ready. There is no need to do anything with this data until Part 2.

To verify that you have successfully requested your data, please copy the text immediately to the left of the "Retrieving data" button, and where the gray box appears in the screenshot below. Paste the text in the text box below.

No text where the gray box appears in the screenshot? You may have previously requested your Twitter data. Instead, please write out the two words on the button that appears instead of the "Retrieving data" button.

How to make the second request:

To make the second request, begin with the same first 5 steps.

1) Log into Twitter: `https://twitter.com` (opens in a new tab).
2) Click on "More" at the bottom left.
3) Click "Settings and privacy".
4) Click "Your Twitter data" at the bottom of the menu on the right side.
5) If prompted, enter your Twitter password.
6) Scroll to the bottom of the page. Now, click "Interests and ads data".
7) Click "Tailored Audiences".
8) Click "Request advertiser list".
9) On the pop-up, click "Request".

Twitter will email you when your download is ready. There is no need to do anything with this data until Part 2. That's it for the second request!

To verify that you have successfully requested your data, please copy the text shown where the gray box appears in the screenshot below. Paste the text in the text box below.

Thank you for making the data requests. For today's last task, please find the summary statistics shown in your Twitter settings on the "Interests and ads data" page.

Please enter the summary statistics into the fields below as numbered in the screenshot. (Fields 1, 2, 3, 4)

D.3 Conclusion

Thank you for completing Part 1 of our study.

It may take a few hours or days until your Twitter data is ready to download.

You will be invited back for Part 2 via Prolific in a few days. In Part 2, you will be given instructions on how to download your Twitter data and upload it to the study.

Part 2 will be a survey that takes 25 minutes to complete.

(Optional) Do you have any final thoughts or comments?

E SURVEY INSTRUMENT (PART 2 OF THE STUDY)

This section provides the survey instrument for the main part of the original user study.

E.1 Introduction and General Questions

Thank you for your participation in our study. This survey will take about 35 minutes.

This survey has 4 sections. The first section will ask a few general questions about your data and Twitter.

If a company writes in their privacy policy that "we do not sell your data," what does that mean to you? In your explanation, please include at least one example of a specific thing you think they would not be allowed to do. *text-area*

Please rate your agreement with the following statement: I believe that Twitter sells my data.

◯ Strongly agree ◯ Agree ◯ Neither agree nor disagree ◯ Disagree ◯ Strongly disagree ◯ Don't know

E.2 Companies

This is the 2nd section (of 4).

In this section, we use the data you uploaded from your own Twitter account. You will be asked about advertisers and up to 4 different advertising methods on Twitter.

The list below shows some companies that showed you an ad on Twitter in the last 3 months.

Please select all of the companies, if any, you remember seeing ads from.

- None of the below
- *[Company 1]*
- *[Company 2]*
- *[Company 3]*
- *[Company 4]*
- *[Company 5]*
- *[Company 6]*
- *[Company 7]*
- *[Company 8]*
- *[Company 9]*
- *[Company 10]*

E.3 Targeting Types

[We repeated this section 4 times for a random selection of 4 *[targeting types]* (e.g., "keywords") and associated specific *[instances]* of that type (e.g., "my cat is my best friend") from the participant's Twitter data. We first asked about the targeting type in the abstract, then about a specific instance from the participant's Twitter data, and then about the targeting type more generally with a selection of frequent and infrequent instances of that type from the participant's Twitter data.]

E.3.1 Abstract.

What does the term **[targeting type]** in the context of online advertising mean to you? If you have never heard this term before, please write your best guess. *text-area*

This next section is about *[targeting type]*.

[targeting type] [definition of targeting type].

Prior to this survey, I would have expected that advertisers currently target ads on Twitter using *[targeting type]*.

◯ Strongly agree ◯ Agree ◯ Neither agree nor disagree ◯ Disagree ◯ Strongly disagree ◯ Don't know

E.3.2 Specific.

On this page, we will give you a specific example of *[targeting type]* from your Twitter data.

According to your Twitter data, you are interested in *[instance]*.

Please rate your agreement with the following statements:

I can think of a reason why I Twitter would conclude that I am [interested in, located in or around, would be added to a list of mobile app users by, etc.] *[instance]*.

◯ Strongly agree ◯ Agree ◯ Neither agree nor disagree ◯ Disagree ◯ Strongly disagree ◯ Don't know

Being [interested in, a speaker of, in the age group, etc.] *[instance]* describes me accurately.

◯ Strongly agree ◯ Agree ◯ Neither agree nor disagree ◯ Disagree ◯ Strongly disagree ◯ Don't know

Assume the number of ads you see doesn't change.

I want some of the ads I see to be chosen for me based on being interested in *[instance]*.

◯ Strongly agree ◯ Agree ◯ Neither agree nor disagree ◯ Disagree ◯ Strongly disagree ◯ Don't know

I am comfortable with Twitter allowing advertisers to target me based on being interested in *[instance]*.

◯ Strongly agree ◯ Agree ◯ Neither agree nor disagree ◯ Disagree ◯ Strongly disagree ◯ Don't know

E.3.3 General.

Overall, in the last three months, advertisers have targeted up to *[#]* ads using *[targeting type]*. In two sentences, please describe your initial reaction to the data above. *text-area*

This section will ask you to consider how you feel about advertisers using *[targeting type]* in general. Please rate your agreement with the following statements:

Assume the number of ads you see doesn't change.

I want some of the ads I see on Twitter to be chosen for me using *[targeting type]*.

◯ Strongly agree ◯ Agree ◯ Neither agree nor disagree ◯ Disagree ◯ Strongly disagree ◯ Don't know

I am comfortable with *[targeting type]* being used to choose ads for me.

◯ Strongly agree ◯ Agree ◯ Neither agree nor disagree ◯ Disagree ◯ Strongly disagree ◯ Don't know

I believe it is fair that Twitter allows advertisers to choose ads for me using *[targeting type]*.

◯ Strongly agree ◯ Agree ◯ Neither agree nor disagree ◯ Disagree ◯ Strongly disagree ◯ Don't know

Please explain your answer to the previous question. If you believe it is fair, why? If you do not believe it is fair, why not?

text-area

E.4 Ad Explanations

This is the 3rd section (of 4).

E.4.1 Introduction.

This section will ask for your opinions about potential explanations for why you received a particular ad on Twitter.

To your knowledge, does Twitter have a feature that explains why you received a particular ad?

◯ Yes ◯ No ◯ Don't know

If an ad explanation on Twitter did not include all reasons an ad was shown to you, which reasons would be most important for you to see?

text-area

E.4.2 Explanations.

[We repeated this section 6 times for the six ad explanations in randomized order. Each was associated with an ad that the participant had been shown according to their Twitter data, alongside the matched targeting criteria for that ad.]

Imagine a Twitter feature that explains how a particular ad was chosen for you.

In this next section, you will see up to 6 different ads that Twitter has shown you before on its platform, each followed by a different explanation. Then, you will answer questions about what you thought of each ad explanation.

[the Twitter ad]

[the tested ad explanation]

What was the most memorable part of this ad explanation? *text-area*

What information, if any, did you feel was missing from this ad explanation? *text-area*

I think this ad explanation shows me **all** of the information used to target the ad to me.

◯ Yes ◯ No ◯ Don't know

I feel that this ad explanation was useful.

◯ Yes ◯ No ◯ Don't know

I feel that this ad explanation gave me enough information to understand how the ad was chosen for me.

◯ Strongly agree ◯ Agree ◯ Neither agree nor disagree ◯ Disagree ◯ Strongly disagree

I would want an ad explanation similar to this one for all ads I see on Twitter.

◯ Strongly agree ◯ Agree ◯ Neither agree nor disagree ◯ Disagree ◯ Strongly disagree

Seeing this ad explanation made me more concerned about my online privacy.

◯ Strongly agree ◯ Agree ◯ Neither agree nor disagree ◯ Disagree ◯ Strongly disagree

Seeing this ad explanation increased my trust in the advertiser who displayed this ad.

◯ Strongly agree ◯ Agree ◯ Neither agree nor disagree ◯ Disagree ◯ Strongly disagree

Do you have any additional comments about this ad explanation? *text-area*

E.4.3 General Opinions.

For your reference, the previous ad explanations that you've seen in this study will appear below:

[all 6 ad explanations]

Please describe your ideal explanation for ads on Twitter. You are not limited to the things you have seen in this study. Feel free to think big!

text-area

E.5 Demographics

This is the 4th section (of 4). Almost done!

In this section, you will be asked about your Twitter usage and demographics.

Please rate your agreement with the following statement: I believe that Twitter sells my data.

◯ Strongly agree ◯ Agree ◯ Neither agree nor disagree ◯ Disagree ◯ Strongly disagree ◯ Don't know

Please explain your answer to the previous question. If you believe Twitter sells your data, why? If you believe Twitter does not sell your data, why not? *text-area*

What month and year did you join Twitter?

On average, about how many hours do you spend on Twitter each day?

◯ Less than 1 hour ◯ 1-2 hours ◯ 2-4 hours ◯ 4-6 hours ◯ More than 6 hours

Have you ever gone to your Twitter account's settings to look at or make changes to your advertising preferences?

◯ Yes ◯ No ◯ Don't know

Did you look at any of the Twitter files you requested in Part 1 of this study before beginning Part 2?

◯ Yes ◯ No ◯ Don't know

What is your gender?

◯ Woman ◯ Man ◯ Non-binary ◯ Prefer to self-describe ◯ Prefer not to say

What is your age?

◯ 18-24 ◯ 25-34 ◯ 35-44 ◯ 45-54 ◯ 55-64 ◯ 65 or older ◯ Prefer not to say

What is the highest degree or level of school you have completed?

◯ Some high school ◯ High school ◯ Some college ◯ Trade, technical, or vocational training ◯ Associate's degree
◯ Bachelor's degree ◯ Master's degree ◯ Professional degree ◯ Doctorate ◯ Prefer not to say

Which of the following best describes your educational background or job field?

◯ I have an education in, or work in, the field of computer science, engineering, or IT.
◯ I do not have an education in, or work in, the field of computer science, engineering, or IT.
◯ Prefer not to say

What is your annual household income?

◯ Less than $20,000 ◯ $20,000 to $49,999 ◯ $50,000 to $99,999 ◯ $100,000 to $249,999
◯ Over $250,000 ◯ Prefer not to say

When people work on tasks, they are sometimes in situations that can be distracting. How distracted were you while completing this survey?

◯ Not distracted at all ◯ Somewhat distracted ◯ Very distracted

(Optional) Do you have any final thoughts or comments? *text-area*