

RESEARCH NOTE

## Quota sampling using Facebook advertisements

Baobao Zhang<sup>1\*</sup>, Matto Mildenerger<sup>2</sup>, Peter D. Howe<sup>3</sup>, Jennifer Marlon<sup>4</sup>, Seth A. Rosenthal<sup>4</sup> and Anthony Leiserowitz<sup>4</sup>

<sup>1</sup>Department of Political Science, Yale University, New Haven, CT, USA, <sup>2</sup>Department of Political Science, University of California, Santa Barbara, Santa Barbara, CA, USA, <sup>3</sup>Department of Environment and Society, Utah State University, Logan, UT, USA and <sup>4</sup>School of Forestry and Environmental Studies, Yale University, New Haven, CT, USA

\*Corresponding author. Email: baobao.zhang@yale.edu

### Abstract

Researchers in different social science disciplines have successfully used Facebook to recruit subjects for their studies. However, such convenience samples are not generally representative of the population. We developed and validated a new quota sampling method to recruit respondents using Facebook advertisements. Additionally, we published an R package to semi-automate this quota sampling process using the Facebook Marketing API. To test the method, we used Facebook advertisements to quota sample 2432 US respondents for a survey on climate change public opinion. We conducted a contemporaneous nationally representative survey asking identical questions using a high-quality online survey panel whose respondents were recruited using probability sampling. Many results from the Facebook-sampled survey are similar to those from the online panel survey; furthermore, results from the Facebook-sampled survey approximate results from the American Community Survey (ACS) for a set of validation questions. These findings suggest that using Facebook to recruit respondents is a viable option for survey researchers wishing to approximate population-level public opinion.

Survey researchers face a fundamental cost-bias tradeoff in respondent recruitment. Fielding opinion polls on high-quality probability samples entail substantial costs that limit the scale and breadth of research activity. Recognizing the limits of using university-based convenience samples, many researchers have examined whether web-based crowd-sourcing tools like Amazon's Mechanical Turk and Google Consumer Surveys can be cost-effective methods to survey diverse target populations (Berinsky, Huber and Lenz 2012; Huff and Tingley 2015; Santoso, Stein and Stevenson 2016). The empirical properties of these Internet-based samples remain a subject of active study.

In this article, we evaluate Facebook's potential as a platform for survey recruitment. We propose a quota sampling method using Facebook advertisements to generate public opinion estimates that approximate national averages efficiently. Using a proof-of-concept study on US climate opinions, we demonstrate that researchers can cost-effectively recruit respondents through quota sampling using Facebook advertisements, at a fraction of the cost of hiring an online survey firm (we sampled over 2000 respondents at about \$4 per response). Using Facebook for quota sampling also offers a comparative advantage because it gives researchers control over how they recruit subjects and allows researchers to target specific subpopulations.

Our method contributes to an emerging literature that uses Facebook to recruit respondents into research studies. For example, psychologists enrolled over 4 million Facebook users as subjects by creating a Facebook application that allowed users to take psychometric tests (Kosinski et al. 2015). The researchers subsequently linked these tests to users' Facebook profiles to predict personal attributes and personality traits from users' social media behavior (Kosinski,

Stillwell and Graepel 2013; Youyou, Kosinski and Stillwell 2015). Medical researchers have similarly used Facebook to recruit subjects from specific subpopulations, such as young adults who smoke cigarettes (Ramo and Prochaska 2012; Ramo et al. 2014) or middle-aged American women (Kapp, Peters and Oliver 2013).

Political science researchers have used Facebook to recruit respondents in Brazil for survey experiments (Samuels and Zucco 2013, 2014), survey likely voters in state primaries (Hirano et al. 2015), deliver political advertisement treatments to diverse respondent pools (Broockman and Green 2014; Ryan 2012), and survey political activists in Germany and Thailand (Jäger 2017). In Samuels and Zucco (2013, 2014), researchers raffled off an iPad to one randomly selected survey-taker. For a cost of \$1.86 per valid survey completion, the researchers obtained 3286 responses. Although the Facebook sample was less representative than a national sample from a survey firm, the results of the Facebook-sampled survey experiment were similar, especially after weighting the responses through post-stratification. Likewise, Jäger (2017) used Facebook samples of local party activists in Germany to accurately predict the outcome of a party leadership race and approximate results from a representative local leader survey.

We contribute to this literature by developing a platform-specific quota sampling technique. Our method uses the Facebook Marketing API to semi-automate a quota sampling process. This method enabled us to recruit 2432 respondents at the cost of \$4.05 per completed response. This approach allows quota sampling on conditional strata themselves rather than simply ensuring the sample marginals approximate the population marginals. Results from the Facebook-sampled survey approximate results from a survey conducted by a reputable survey research firm, GfK,<sup>1</sup> as well as results from the American Community Survey (ACS). To accompany this research note, we are releasing an R package, *fbsample*, that automates the quota sampling process using the Facebook Marketing API.<sup>2</sup>

## Methodology

To obtain the most representative sample given a constrained advertisement budget, we employ quota sampling to target a series of strata, or demographic subgroups. Convenience sampling methods that do not set demographic quotas might potentially produce very non-representative samples. Facebook states that “if you optimize for link clicks, your ads are targeted to people in your audience who are most likely to click the ads’ links.”<sup>3</sup> When a user first launches an advertisement optimized for link clicks, Facebook will send the advertisement to different types of people to learn who is most likely to click on the link. After this “learning phase” (after about 50 link clicks), Facebook targets the advertisements to those predicted to click on the link.<sup>4</sup> For public opinion researchers, Facebook’s advertisement delivery optimization might prove problematic because it could lead to the recruited subjects being homogeneous. For instance, if many young white males clicked on the advertised link during the “learning phase,” Facebook may target mostly young white males. To avoid this “homogenization” of their sample, public opinion researchers must target a diversity of demographic strata.

The Facebook Marketing API makes it relatively easy to target specific demographic subgroups. Before 2015, advertisers had to set up advertisements that targeted specific groups manually. However, after Facebook made its Marketing API publicly available in 2015, advertisers gained the ability to write programs to target many pre-specified groups at once. Our

<sup>1</sup>GfK recruits respondents for its online panel through address-based probability-sampling. Organizations that have used GfK’s panel include the American National Election Study, the U.S. Federal Reserve, and Pew Research Center.

<sup>2</sup>*fbsample* is available online at <https://github.com/13bzhang/fbsample>.

<sup>3</sup>Facebook, “About the Delivery system: Optimization.” [https://www.facebook.com/business/help/355670007911605?helpref=faq\\_content#](https://www.facebook.com/business/help/355670007911605?helpref=faq_content#)

<sup>4</sup>Facebook, “About the delivery system: Learning phase.” [https://www.facebook.com/business/help/112167992830700?helpref=faq\\_content](https://www.facebook.com/business/help/112167992830700?helpref=faq_content)

current method uses this capability to quota sample many demographically diverse strata simultaneously.

However, efforts to estimate national public opinion from Facebook quota samples depend on a series of assumptions, some of which may not hold in reality. Define  $R_i$  as an indicator variable for whether member  $i$  of the population took the survey ( $R_i = 1$  means  $i$  took the survey and  $R_i = 0$  otherwise);  $\mathbf{D}_i$  as respondent  $i$ 's characteristics that researchers used to construct strata for quota sampling;  $\mathbf{X}_i$  as respondent  $i$ 's characteristics not used to construct strata but are measured in the survey; and  $Y_i$  as person  $i$ 's response to a survey question.

**Assumption 1** :  $Y_i \perp\!\!\!\perp R_i \mid \mathbf{D}_i = \mathbf{d}_i, \mathbf{X}_i = \mathbf{x}_i, \forall \mathbf{d}_i \in \text{Supp}(\mathbf{D}_i) \text{ and } \forall \mathbf{x}_i \in \text{Supp}(\mathbf{X}_i)$

**Assumption 2** :  $Pr(R_i = 1 \mid \mathbf{D}_i = \mathbf{d}_i, \mathbf{X}_i = \mathbf{x}_i) > 0, \forall \mathbf{d}_i \in \text{Supp}(\mathbf{D}_i) \text{ and } \forall \mathbf{x}_i \in \text{Supp}(\mathbf{X}_i)$

Assumption 1 implies that conditional on strata and observed respondent characteristics, responses of those who took the survey would be the same as those who did not take the survey. While this assumption might not hold in reality, the degree to which violations of Assumption 1 will bias efforts to use Facebook advertisements to estimate national public opinion is an empirical question that this research note is designed to address.

Assumption 2 asserts that conditional on strata and observable characteristics, each person in the population has a non-zero probability of taking the survey. Of course, not everyone in a given population has a Facebook account or uses Facebook regularly. According to a 2018 Pew Research Center study, 68 percent of American adults reported using Facebook (Smith and Anderson, 2018). Even so, Facebook allows one to reach a far greater proportion of the US adult population than other ways researchers recruit convenience samples. For instance, researchers recruiting respondents using Mechanical Turk could reach only about 7300 workers in any quarter year (Stewart et al. 2015).

## Study design

To demonstrate the validity of the quota sampling method described above, we conducted a study to compare the results of a quota-sampled Facebook survey with a high-quality probability-sampled GfK online panel survey. The two surveys, which shared 25 identical questions, were conducted a few months apart in 2016. Apart from the questions asking about respondent demographics, both surveys focused on climate change attitudes and policy preferences. Also, to validate the Facebook sample against the ACS, we asked three questions from the ACS about veteran status, home ownership, and country of birth. (These results are reported in the Online Supplementary Information).

The Facebook Marketing API allows researchers to write code that target several thousand highly specific demographic groups at once. For our quota sampling, we generated 544 strata according to demographic characteristics (e.g., gender, age group, race, level of education, and nine US Census regions). We chose these demographic characteristics because crosstabs containing frequencies conditional on these demographics are readily available through the US Census.<sup>5</sup>

Next, we assigned an advertising budget to each stratum according to proportional allocation for most strata; we allocated greater budget to strata that contain very small sub-populations nationally for two reasons. First, Facebook requires advertisers to set a daily budget of \$5.00 per

<sup>5</sup>Facebook also allows advertisers to target users by users' political ideology and party identification. Users' ideologies are imputed from their social media behavior (Bond and Messing 2015). While we did not do so in this project, it is feasible for researchers to construct strata based on Facebook users' politics. Nevertheless, researchers should not rely on Facebook to accurately provide respondents' demographic information. While the site's internal algorithms can predict age group, political ideology, and general geographic location with high accuracy, the algorithms are less accurate when predicting one's party identification or whether one identifies as black (Sances 2018).

advertisement, although it is not required to spend the entire daily budget. Second, since it is more difficult to recruit from small sub-populations, we had to allocate a larger budget to generate the minimum number of respondents.

Our advertisements informed respondents that by taking the survey they could see how their views on climate change compared to other Americans; no rewards were promised. To maximize viewership, we promoted the advertisements on all Facebook platforms (e.g., news feed, right column advertisement space, etc.), except on Instagram.<sup>6</sup> Over a period of two weeks in July 2016, we recruited 2432 respondents who completed the survey from the 7642 Facebook users who clicked on our advertisement. At the end of the recruitment period, we filled or overfilled 218 strata, partially filled 61 strata, and failed to recruit anyone for 157 strata. On average, the Facebook survey cost \$4.05 per complete.<sup>7</sup>

Our probability-sampled online panel study was conducted in March 2016. This survey sampled 2459 respondents with 1346 completions for a completion rate of 54.7 percent.<sup>8</sup> This study used GfK's Knowledge Panel, with a household recruitment rate for our survey of 12.3 percent (American Association for Public Opinion Research Response Rate 3). Overall, our online panel responses cost nearly six times as much as our Facebook user responses.

## Results

Results from the Facebook survey are similar to results from the online panel survey and the ACS even without weighting. For questions where the Facebook and online panel survey results differed, the Facebook respondents produced answers indicating greater concern about climate change.

Table 1 displays the summary statistics for the demographic groups in the GfK, Facebook, and ACS surveys. Contrasted with the GfK and ACS samples, the Facebook sample is younger and less white. To make both surveys more nationally representative, we used inverse probability weighting to weight each sample to the March 2016 Current Population Survey (CPS).<sup>9</sup>

We combined each sample and the CPS, then we used logistic regression to estimate the probability of being included in the sample. Covariates used in our propensity score model include gender, age group, level of education, race, geographic region, whether the respondents lived in a metropolitan area, and the interaction between region and the metropolitan indicator variable. The final weights are the inverse of the estimated probabilities normalized such that the sum of each sample's weights equals the sample's size.<sup>10</sup>

As Table 1 demonstrates, the weighted demographic summary statistics of the two surveys are very similar to each other and to the ACS estimates. Comparing the political variables, the Facebook sample contains a greater proportion of Independents and a smaller proportion of Republicans than the GfK sample.

The results for the eight questions about climate change are presented in Figure 1. For the unweighted results, the mean difference is 5.1 percentage points with a standard deviation of 4.2 percentage points. Regarding these unweighted results, the Facebook-sampled respondents

<sup>6</sup>During the respondent recruitment period, we activated the advertisements in the morning and deactivated them after 8 p.m. EST. After each day of recruitment, we removed advertisements for strata whose quotas had been filled. This allowed some strata to be overfilled, wasting some advertisement budget. We discuss potential solutions to this problem in the Online Supplementary Appendix.

<sup>7</sup>We did not offer monetary incentives. Instead, we told respondents that by completing our survey they would discover how their views on climate change compare with views of other Americans. The average cost per link click was \$2.16 (SD \$1.73); the distribution of costs per link click is reported in the Online Supplementary Information.

<sup>8</sup>GfK dropped cases where more than half the survey questions were left blank and/or where respondents completed the survey in under 7 minutes, generating a final 1317 observations.

<sup>9</sup>We take an agnostic stance on the ideal weighting strategy. Our assessment is thus driven by the empirical properties of our estimates during validation, rather than evaluating the merits of different weighting strategies.

<sup>10</sup>The correlation between the inverse probability weights for the GfK sample and the weights supplied by GfK is 0.64.

**Table 1.** Summary Statistics of Survey Respondents

	GfK	Facebook	ACS	GfK	Facebook	ACS
Weighted	No	No	No	Yes	Yes	Yes
Prop. male	0.50	0.48	0.48	0.48	0.47	0.49
Prop. age 18–24	0.08	0.19	0.11	0.12	0.12	0.12
Prop. age 25–34	0.14	0.16	0.15	0.17	0.20	0.18
Prop. age 35–44	0.14	0.15	0.15	0.17	0.17	0.16
Prop. age 45–64	0.39	0.29	0.35	0.35	0.33	0.34
Prop. age 65 and above	0.26	0.21	0.24	0.19	0.18	0.20
Prop. with college degree	0.21	0.23	0.18	0.20	0.19	0.18
Prop. with graduate degree	0.16	0.16	0.11	0.11	0.12	0.11
Prop. Hispanic	0.11	0.13	0.13	0.15	0.13	0.16
Prop. black	0.09	0.06	0.10	0.11	0.11	0.12
Prop. white	0.74	0.61	0.78	0.66	0.68	0.74
Prop. Democrat	0.43	0.42		0.44	0.41	
Prop. Republican	0.36	0.27		0.33	0.27	
Prop. Independent	0.10	0.18		0.10	0.17	
Mean political ideology	3.05	2.77		3.04	2.76	
Prop. registered to vote	0.83	0.81		0.79	0.81	

The GfK survey included 1317 respondents; the Facebook survey included 2432 respondents; the 2016 American Community Survey (ACS) Public Use Microdata Sample include 2,503,750 respondents. In the GfK and Facebook surveys, respondents who indicated they “lean Democrat” were counted as Democrats; likewise, those who indicated they “lean Republicans” were counted as Republicans. Political ideology was measured using a 5-point scale such that 1 is very liberal and 5 is very conservative.

expressed slightly less skepticism about climate change, greater concern about climate change, and greater support for policies that mitigate climate change. After weighting both the Facebook survey and the online panel survey to the CPS, the mean difference in responses is 4.3 percentage points with a standard deviation of 3.4 percentage points.


One reason that the Facebook-sampled respondents may be more concerned about climate change is that our advertisement explicitly contained language about climate change and included a photo of the Earth. Respondents who are pro-climate action might have been more motivated to click on the advertisement than those who are unconcerned about climate change. A breakdown of the survey results by partisanship reveals that Democrats, Republicans, and Independents within the Facebook survey provided similar responses to Democrats, Republicans, and Independents in the GfK survey, respectively. However, respondents in the Facebook survey who said they are uninterested in politics were significantly more pro-climate action than their counterparts in the GfK survey. This difference between the two samples suggests that within each stratum, Facebook users may have not as-if randomly sorted into the recruitment sample. One possible strategy to avoid self-sorting that affects survey results is to create advertisements that do not discuss the content of the survey being advertised, although this strategy will require careful attention to advertisement wording and design.<sup>11</sup>

As a further robustness check, we also compared three results from the Facebook survey with the 2016 ACS One-Year Estimates. The Facebook sample is also able to somewhat approximate the ACS benchmarks, as detailed in our Online Supplementary Information.

## Conclusion

This study produced valuable lessons for best practices in quota sampling using Facebook advertisements; this method can generate results that approximate high-quality probability-based national opinion surveys. Further improvements to our proposed method would result from systematic attention to the factors that drive Facebook users to click on advertisements. We highlight some considerations for future researchers. First, researchers might inadvertently

<sup>11</sup>During the testing phase of this study, we launched advertisements for a generic public opinion survey. These advertisements received very few clicks. We decided to launch instead the advertisements that contained information about climate change to generate a sufficient number of clicks.



**Figure 1.** Comparing Facebook survey with GfK survey: climate change public opinion. The plots above report the point estimate and the 95 percent confidence interval for each outcome measure. The confidence intervals are calculated from heteroscedasticity-consistent standard errors.

recruit particular types of respondents by advertising the content of the survey. This form of self-selection bias may not be eliminated by strata targeting or by conditioning on observable demographic characteristics. To avoid this problem, researchers can advertise their survey using vague language without revealing its core content. Additional design effort may be required, however, to persuade users to click on such advertisements.

Future researchers could also further reduce bias in their estimates by measuring other characteristics of the respondents and using those characteristics to weight their Facebook-sampled survey. In particular, we suggest that researchers obtain respondent characteristics that are not available through Facebook or are inaccurately predicted by Facebook to improve re-weighting efforts.

More broadly, our findings suggest that using Facebook to recruit respondents is a viable option for survey researchers seeking to approximate public opinion estimates for some populations at significantly lower cost. While our method costs a fraction of recruiting respondents through survey panels that use probability sampling, it is not dramatically cheaper than using other “opt-in” online survey panels. Nevertheless, our method provides greater flexibility to researchers in deciding who to recruit. Furthermore, quota sampling

using Facebook advertisements may also be particularly useful for generating targeted samples of geographic or demographic subpopulations, for which national panels may prove to be inadequate or prohibitively expensive. For instance, Sances (2017) has used Facebook advertisements without quota sampling to recruit voters from US municipalities. For future research, we plan to investigate ways to further reduce the cost of this quota sampling method.

**Supplementary Material.** To view supplementary material for this article, please visit <https://doi.org/10.1017/psrm.2018.49>

**Acknowledgments.** We would like to thank Erin Hartman, Devin Caughey, Peter Aronow, Solomon Messing, and M. Kent Jennings for their feedback on earlier drafts of this research. We also appreciate feedback from participants of the following workshops: the New Faces in Political Methodology IX Workshop (Pennsylvania State University), the Oxford Internet Institute Summer Doctoral Programme, and the Southern California Methods Workshop. We also thank the Skoll Global Threats Fund and the MacArthur Foundation for their support of the project. BZ's work on this project was supported by the National Science Foundation Graduate Research Fellowship.

## References

- Berinsky AJ, Huber GA and Lenz GS** (2012) Evaluating Online Labor Markets for Experimental Research: Amazon.com's Mechanical Turk. *Political Analysis* **20**(3), 351–368.
- Bond R and Messing S** (2015) Quantifying Social Media's Political Space: Estimating Ideology from Publicly Revealed Preferences on Facebook. *American Political Science Review* **109**(01), 62–78.
- Broockman DE and Green DP** (2014) Do Online Advertisements Increase Political Candidates' Name Recognition or Favorability? Evidence from Randomized Field Experiments. *Political Behavior* **36**(2), 263–289.
- Hirano S, Lenz GS, Pinkovskiy M and Snyder JM** (2015) Voter Learning in State Primary Elections. *American Journal of Political Science* **59**(1), 91–108.
- Huff C and Tingley D** (2015) “Who Are These People?” Evaluating the Demographic Characteristics and Political Preferences of Mturk Survey Respondents. *Research & Politics* **2**(3), 2053168015604648.
- Jäger K** (2017) The potential of online sampling for studying political activists around the world and across time. *Political Analysis* **25**(3), 329–343.
- Kapp JM, Peters C and Oliver DP** (2013) Research Recruitment Using Facebook Advertising: Big Potential, Big Challenges. *Journal of Cancer Education* **28**(1), 134–137.
- Kosinski M, Matz SC, Gosling SD, Popov V and Stillwell D** (2015) Facebook as a Research Tool for the Social Sciences: Opportunities, Challenges, Ethical Considerations, and Practical Guidelines. *American Psychologist* **70**(6), 543.
- Kosinski M, Stillwell D and Graepel T** (2013) Private Traits and Attributes Are Predictable from Digital Records of Human Behavior. *Proceedings of the National Academy of Sciences* **110**(15), 5802–5805.
- Ramo DE and Prochaska JJ** (2012) Broad Reach and Targeted Recruitment Using Facebook for an Online Survey of Young Adult Substance Use. *Journal of Medical Internet Research* **14**(1), e28.
- Ramo DE, Rodriguez TM, Chavez K, Sommer MJ and Prochaska JJ** (2014) Facebook Recruitment of Young Adult Smokers for a Cessation Trial: Methods, Metrics, and Lessons Learned. *Internet Interventions* **1**(2), 58–64.
- Ryan TJ** (2012) What Makes Us Click? Demonstrating Incentives for Angry Discourse with Digital-Age Field Experiments. *The Journal of Politics* **74**(4), 1138–1152.
- Samuels DJ and Zucco C** (2013) Using Facebook as a Subject Recruitment Tool for Survey-Experimental Research. Working Paper, available at SSRN 2101458.
- Samuels DJ and Zucco C** (2014) The Power of Partisanship in Brazil: Evidence from Survey Experiments. *American Journal of Political Science* **58**(1), 212–225.
- Sances MW** (2017) Ideology and Vote Choice in US Mayoral Elections: Evidence from Facebook Surveys. *Political Behavior* **40**(3), 737–762.
- Sances MW** (2018) Missing the Target? Using Surveys to Validate Social Media Ad Targeting. Working Paper, Unpublished paper.
- Santoso LP, Stein R and Stevenson R** (2016) Survey Experiments with Google Consumer Surveys: Promise and Pitfalls for Academic Research in Social Science. *Political Analysis* **24**(3), 356–373.
- Smith A and Anderson M** (2018) Social Media Use in 2018. Survey Report, Washington, DC.
- Stewart N, Ungemach C, Harris AJ, Bartels DM, Newell BR, Paolacci G and Chandler J** (2015) The Average Laboratory Samples a Population of 7,300 Amazon Mechanical Turk Workers. *Judgment and Decision Making* **10**(5), 479.
- Youyou W, Kosinski M and Stillwell D** (2015) Computer-Based Personality Judgments Are More Accurate Than Those Made by Humans. *Proceedings of the National Academy of Sciences* **112**(4), 1036–1040.