

Online Surveys and Digital Demography in the Developing World: Facebook Users in Kenya

Katherine Hoffmann Pham
Stern School of Business
New York University
New York, USA

Francesco Rampazzo
University of Southampton
and Max Planck Institute for
Demographic Research
Southampton, United Kingdom

Leah R. Rosenzweig
Institute for Advanced Study
and MIT GOV/LAB
Toulouse, France

KEYWORDS

social media, Facebook, digital demography, computational social science, online survey, audience estimates

1 MOTIVATION

Digital platforms have transformed the study of human behavior. Large online platforms such as Facebook, Twitter, and Wikipedia enable researchers to test behavioral theories on user communities “in the wild”, whereas crowdsourcing platforms such as Amazon Mechanical Turk (MTurk) offer new opportunities to create synthetic online laboratory environments. These platforms have opened up new subject pools for academic research, allowing investigators to study more diverse populations than the conventional university undergraduate samples.

While the use of these online platforms to recruit subjects has been well documented in the US and Europe, there is less evidence on whether these platforms can also be used to reach respondents in developing countries. As more of the world’s citizens “get online”, we expect that they will be easier to count, interview, and represent with digital platforms and survey research. Nevertheless, limited internet availability and computer literacy may restrict access to these platforms, raising the question of who, exactly, is represented. To investigate this question, we focus on the Facebook advertising platform and begin with a case study of Kenya.

Facebook’s advertising platform has gained traction among social science researchers as a result of two core capabilities. First, Facebook allows advertisers to produce “audience estimates” of the number of platform users with given demographic or behavioral characteristics. The potential of these estimates has been recognized by demographers, who have used this tool to study topics such as the digital gender divide and international migrant populations [7, 19]. Second, Facebook allows advertisers to reach these tailored audiences with targeted ads. Researchers have used this capability to recruit respondent samples for opinion and attitude surveys in the US, India, and Brazil [5, 14, 15]. To our knowledge, however, no research has systematically investigated the potential of Facebook advertising for representative research in an African country.

We assess the value of these tools in Kenya using two primary strategies. First, we collect subnational Facebook audience estimates and compare these with official population data from the 2009 census. Second, we use Facebook’s targeting tools to recruit a pilot sample of 957 respondents, whose demographics and survey responses we compare to the results of the 2016 Afrobarometer (Round 7) and the results of a canonical choice experiment (Tversky

and Kahneman’s 1981 Asian disease problem). This pilot data was collected in preparation for a full-scale survey that was launched in late September, immediately following the Round 8 Afrobarometer survey deployment and after the completion of Kenya’s most recent census round in August 2019.

2 RELATED WORK

2.1 New data sources for the developing world

While the use of Facebook to recruit respondents in the developing world is relatively new, researchers have explored how other big data sources might supplement costly or infrequent survey data collection in resource-poor settings. Early work focused on mobile phone data, using Call Data Records (CDR) to study phenomena such as the distribution of poverty, the spread of disease, or population mobility in response to natural disasters [4, 13, 18]. More recently, research has focused on remote sensing, including satellite and nighttime lights imagery, to study poverty and even political favoritism [9, 12]. Along these lines, we compare passively collected information from a novel data source (Facebook) to official statistics that are typically considered to be “ground truth”. Rather than simply observing behavior remotely, we also have the ability to reach out to study subjects and survey them directly.

2.2 The representativeness of online platforms

This research also draws on work on the demographics of online platforms and how online samples compare to more established data sources. These studies have drawn attention to a number of important considerations in the use of online platforms to recruit survey respondents. For example, Zhang et al. [20] explore the potential of using targeted ads to generate representative surveys on attitudes towards climate change. They conduct stratified sampling of Facebook audiences in the US using age, education, gender, race, and region. While their results approximate their “ground truth” data sources, they note two key limitations of the Facebook data: first, they were unable to recruit any respondents in 157 of their 544 strata; and second, they observed within-strata imbalances on other characteristics (such as political orientation), which they attribute to the fact that Facebook strategically serves ads to those who are most likely to respond, exacerbating selection bias.

In the context of crowdsourcing, Difallah et al. [6] note a similar challenge with selection bias on the MTurk platform, since respondents may elect tasks that appeal to them. They also observe that the population on the platform is dynamic over time, with the entry




Figure 1: Annual growth in internet users and mobile phone subscribers in Africa; Kenya is shown in black. Source: International Telecommunications Union [10].

and exit of workers; similar patterns may appear on Facebook, as users join or leave the platform.

A final interesting challenge has been raised by Goel et al. [8] when conducting a comparison between MTurk and established survey sources. The authors observe that on a series of opinion questions, there are notable discrepancies *between* survey sources, making it difficult to establish a “ground truth” to use as a benchmark. In such settings, online platforms may provide valuable additional insight into conflicting survey results.

Despite these concerns, Zhang et al. [20] and Goel et al. [8] find promising results: given the appropriate stratification and/or post-stratification weighting, data collected from online platforms is often able to approximate established, nationally representative data sources. A key question inspiring our research is whether the limitations described above are prohibitive in the context of developing countries like Kenya.

3 CONTEXT

We selected Kenya as a case study for two main reasons. First, Kenya represents an average example of mobile and internet use on the African continent (see Fig. 1). Kenya’s telecom networks have seen a strong demand for mobile services: Safaricom’s popular M-Pesa platform was a leading pioneer of mobile money, and Airtel participates in Facebook Free Basics to offer costless browsing on the social network. However, Kenya also has large populations without regular access to the internet or even electricity (e.g., 52% of Afrobarometer respondents do not have a connection from the main electrical grid to their home), raising the question: how representative is the Facebook audience?

Second, Kenya presents an excellent opportunity of timing. Specifically, two major survey efforts coincided: the decennial census from 24-31 August 2019, and the Round 8 Afrobarometer in September 2019. We planned our Facebook survey for September 2019, ensuring that any differences with the census and Afrobarometer would not be due to changes in the population or attitudes over time. While Facebook studies should not replace such rigorous, nationally representative surveys, these surveys are infrequently

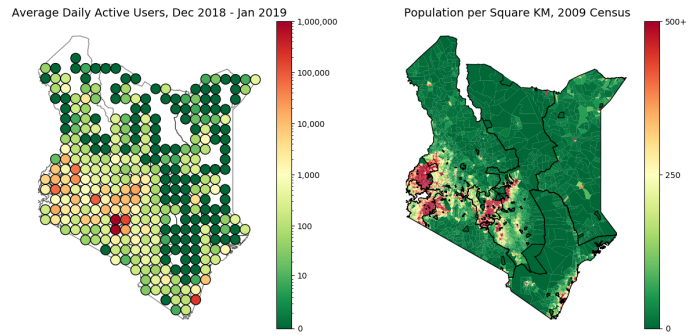


Figure 2: A comparison between Facebook’s daily active user estimates for late 2018/early 2019, and population density estimates from the 2009 census. Missing circles indicate that Facebook estimated no users.

collected¹ and often delayed in their release. In contrast, Facebook data can be collected in near-real time. Thus, we examine the extent to which Facebook data collection can be used to complement existing sources with more frequent collection.

4 DATA COLLECTION

4.1 Facebook audience estimates

Audience estimate queries were submitted to the Facebook Marketing API using the PySocialWatcher Python package [1, 2]. We collected estimates for Facebook users using two key geographic targeting strategies. First, we divided the country into a *grid of non-overlapping circles* with 20km radii, and collected audience estimates of Facebook users within each radius. This allowed us to examine geographic variation in the density of Facebook users across the country. Second, we generated estimates of the number of users within a 20km radius of each of *Afrobarometer’s survey clusters*. This enabled us to study the number of Facebook users in areas targeted by the Afrobarometer, in order to understand whether Facebook samples could feasibly be recruited there.

4.2 Facebook surveys

To supplement these audience estimates and understand how they translate into researchers’ ability to engage users, we also conducted a pilot survey on the Facebook advertising platform. Between 6 and 9 June, we advertised the opportunity to participate in a 15-minute survey in exchange for a 50 Ksh (about 50 cents) mobile airtime credit. During this period, we were able to recruit 1,190 Kenyans on Facebook to take the survey, with a total cost of \$500 in advertising fees and \$630 in respondent incentives and administrative costs. After removing duplicate respondents, we have a sample of 957 unique respondents who spent a reasonable amount of time on the survey (≥ 5 minutes).

We explored two different ad targeting strategies. First, we targeted advertisements to users residing in each of Kenya’s eight *provinces*. This approach was designed to force geographic variation in respondent locations, while minimizing the number of

¹Kenya’s last census was in 2009, and Kenya’s most recent Afrobarometer rounds were in 2011, 2014, and 2016.




Figure 3: A comparison of demographics between Afrobarometer and our Facebook sample

separate campaigns. Second, we targeted advertisements to users residing within a 40km radius of select *Afrobarometer clusters*. For each province, one Afrobarometer cluster was selected as a median-case example cluster based on the audience estimates provided by the Facebook API. Furthermore, we selected two worst-case clusters for which Facebook estimated less than 1,000 monthly active users (and zero daily active users). This approach was designed to explore the cost and difficulty of recruiting respondents near these clusters.

5 PRELIMINARY RESULTS

5.1 Facebook audience estimates

As shown in Fig. 2, Facebook has an active user base across most of the country. Facebook audience sizes generally reflect underlying population density: there are many users in the Nairobi and Mombasa areas, as well as in the population-dense Western parts of the country. Unsurprisingly, Facebook is underrepresented along Kenya’s northern borders, which consist of relatively remote and dry regions shaped in part by the military conflict with Somalia.

5.2 Facebook surveys

5.2.1 Logistical aspects of recruiting Facebook respondents. Our advertisements reached 159,743 Facebook users in a four-day period. The performance of our advertisements varied by province. Unsurprisingly, users in Nairobi were responsive, with a click-through rate of 5%, a completion rate of 40%, and an advertising cost of \$0.05 per survey. Our worst performance was in North Eastern province, with a 1% click-through rate, a 7% survey completion rate, and an advertising cost of \$2.65 per survey. Targeting at the cluster level increased advertising costs. Across the test clusters with median audience sizes, these costs ranged from \$0.36 per survey in Rift Valley to \$9.71 in the North East. In one cluster – the worst-case Rift Valley cluster – we were unable to recruit any respondents.

5.2.2 Geography of Facebook respondents. Although we used a limited targeting strategy for the pilot, respondents reported coming from all 47 counties of Kenya, and had 13 different tribal/ethnic




Figure 4: Replication of the Tversky & Kahneman 1981 Asian Disease Experiment

affiliations. Generally speaking, respondents’ self-reported locations did *not* closely match a given ad’s target location. It is possible that respondents have ties to multiple locations; we have added additional survey questions to the full-scale survey to address this possibility. However, others have documented similar issues with Facebook’s geographic ad targeting in the US context [16].

5.2.3 Demographics of the Facebook respondents. Relative to the Afrobarometer sample, our respondents were more likely to be male, university-educated, and connected to the electrical grid. They were less likely to be employed, or to leave their compound to access water or sanitation. We identified comparable rates of phone and vehicle ownership across the samples (see Fig. 3).

5.2.4 Experimental results. We replicated a canonical behavioral experiment – Tversky and Kahneman’s 1981 Asian disease problem – that has been conducted in a wide range of contexts over time [17]. Interestingly, our results are similar to the original sample and other Western convenience samples [3]. When the problem is framed in terms of the number of people who will die from the disease, a majority (67%) of respondents prefer the risky option. When the problem is framed in terms of the number of lives saved, a similar majority (68%) prefer the certain policy (see Fig. 4).

6 CONCLUSION AND NEXT STEPS

We conduct one of the first systematic studies of Facebook as a research tool in the developing world. We assess the potential of Facebook for producing population estimates and recruiting survey respondents. We present preliminary results from a pilot survey of 957 respondents, which suggest that Facebook can be used to recruit respondents from diverse areas of Kenya at reasonable costs. Although the sample exhibits predictable biases in terms of gender, access to infrastructure, and education levels, our sample is less privileged than expected: over one fifth of respondents said that their home did not have a connection to the electrical grid, and over one quarter did not have piped water inside their home.

Our full-scale survey tests whether a more sophisticated sampling strategy, involving stratification on gender and 25 separate

regional clusters, will address some of the biases in our sample. We also plan to explore whether post-stratification can further reduce remaining imbalances in demographics, before using the weighted sample to compare attitudes and behaviors with those from nationally representative surveys. Finally, we are also considering different strategies for extending our sample. These could include standard approaches such as respondent-driven sampling, or developing a custom objective function for the ad targeting algorithm, similar to the approach of Ipeirotis and Gabrilovich [11].

Digital inclusion and social media offer exciting new opportunities for experimental and survey research in the developing world. We build a foundation for such research by documenting the type of respondent sample that can be recruited via an online platform. We expect that this will form part of a larger line of research describing the strengths, limitations, and biases of these new tools. In general, we are optimistic about the use of these platforms for social science research. We anticipate that our findings will be of broad interest to researchers working in developing countries where an increasing number of their populations of interest are accessing social media.

7 ACKNOWLEDGEMENTS

We thank the Summer Institute for Computational Social Science (SICSS), the Russell Sage Foundation, and the Alfred P. Sloan Foundation for financial support. We thank Chris Bail, Matt Salganik, Anne Helby Petersen, Julien Migozzi, and Tina Law for their assistance in shaping this project. We thank Ingmar Weber, Alessandro Sorichetta, Dennis Feehan, and researchers at the Busara Center for Behavioral Economics and Afrobarometer for helpful comments. We thank Nelson Ngige, Eunice Williams, and Kibuchi Eliud, Warsama Abdifitah, and Ahmed Hared for translation assistance.

REFERENCES

- [1] Matheus Araujo. 2019. Social Watcher on Facebook Marketing API. (Aug. 2019). <https://github.com/maraujo/pySocialWatcher>
- [2] Matheus Araujo, Yelena Mejova, Ingmar Weber, and Fabrizio Benevenuto. 2017. Using Facebook Ads Audiences for Global Lifestyle Disease Surveillance: Promises and Limitations. *arXiv:1705.04045 [cs]* (May 2017). <http://arxiv.org/abs/1705.04045>
- [3] Adam J. Berinsky, Gregory A. Huber, and Gabriel S. Lenz. 2012. Evaluating online labor markets for experimental research: Amazon.com’s Mechanical Turk. *Political analysis* 20, 3 (2012), 351–368. 02669.
- [4] Joshua Blumenstock, Gabriel Cadamuro, and Robert On. 2015. Predicting poverty and wealth from mobile phone metadata. *Science* 350, 6264 (2015), 1073–1076.
- [5] Taylor C. Boas, Dino P. Christenson, and David M. Glick. 2018. Recruiting large online samples in the United States and India: Facebook, mechanical turk, and qualtrics. *Political Science Research and Methods* (2018), 1–19.
- [6] Djellel Difallah, Elena Filatova, and Panos Ipeirotis. 2018. Demographics and Dynamics of Mechanical Turk Workers. In *Proceedings of the Eleventh ACM International Conference on Web Search and Data Mining (WSDM ’18)*. ACM, New York, NY, USA, 135–143. <https://doi.org/10.1145/3159652.3159661>
- [7] Masoomali Fatehikia, Ridhi Kashyap, and Ingmar Weber. 2018. Using Facebook ad data to track the global digital gender gap. *World Development* 107 (July 2018), 189–209. <https://doi.org/10.1016/j.worlddev.2018.03.007>
- [8] Shirad Goel, Adam Obeng, and David Rothschild. 2017. *Online, Opt-in Surveys: Fast and Cheap, but Are They Accurate?* Technical Report. Working Paper, Stanford University, Stanford, CA.
- [9] Roland Hodler and Paul A. Raschky. 2014. Regional favoritism. *The Quarterly Journal of Economics* 129, 2 (2014), 995–1033.
- [10] International Telecommunications Union. 2019. Statistics. (2019). <https://www.itu.int/en/ITU-D/Statistics/Pages/stat/default.aspx>
- [11] Panagiotis G. Ipeirotis and Evgeniy Gabrilovich. 2014. Quizz: targeted crowd-sourcing with a billion (potential) users. In *Proceedings of the 23rd International World Wide Web Conference (WWW)*. ACM, 143–154.
- [12] Neal Jean, Marshall Burke, Michael Xie, W. Matthew Davis, David B. Lobell, and Stefano Ermon. 2016. Combining satellite imagery and machine learning to predict poverty. *Science* 353, 6301 (2016), 790–794.
- [13] Xin Lu, Linus Bengtsson, and Petter Holme. 2012. Predictability of population displacement after the 2010 Haiti earthquake. *Proceedings of the National Academy of Sciences* 109, 29 (July 2012), 11576–11581. <https://doi.org/10.1073/pnas.1203882109>
- [14] David J. Samuels and Cesar Zucco. 2013. *Using Facebook as a Subject Recruitment Tool for Survey-Experimental Research*. SSRN Scholarly Paper ID 2101458. Social Science Research Network, Rochester, NY. <https://papers.ssrn.com/abstract=2101458>
- [15] Michael W. Sances. 2018. Ideology and Vote Choice in U.S. Mayoral Elections: Evidence from Facebook Surveys. *Political Behavior* 40, 3 (Sept. 2018), 737–762. <https://doi.org/10.1007/s11109-017-9420-x> 00004.
- [16] Michael W. Sances. 2019. Missing the Target? Using Surveys to Validate Social Media Ad Targeting. *Political Science Research and Methods* (March 2019), 1–8. <https://doi.org/10.1017/psrm.2018.68>
- [17] Amos Tversky and Daniel Kahneman. 1981. The framing of decisions and the psychology of choice. *Science* 211, 4481 (1981), 453–458.
- [18] Amy Wesolowski, Nathan Eagle, Andrew J. Tatem, David L. Smith, Abdisalan M. Noor, Robert W. Snow, and Caroline O. Buckee. 2012. Quantifying the Impact of Human Mobility on Malaria. *Science* 338, 6104 (Oct. 2012), 267–70. <http://science.sciencemag.org/content/338/6104/267>
- [19] Emilio Zagheni, Ingmar Weber, and Krishna Gummadi. 2017. Leveraging Facebook’s Advertising Platform to Monitor Stocks of Migrants. *Population and Development Review* 43, 4 (Dec. 2017), 721–734. <https://doi.org/10.1111/padr.12102>
- [20] Baobao Zhang, Matto Mildenerger, Peter D. Howe, Jennifer Marlon, Seth A. Rosenthal, and Anthony Leiserowitz. 2018. Quota sampling using Facebook advertisements. *Political Science Research and Methods* (2018), 1–7. <https://doi.org/10.1017/psrm.2018.49>