Extended Abstract - Targeting Development Aid with Machine Learning and Mobile Phone Data: Evidence from an Anti-Poverty Intervention in Afghanistan

Emily L. Aiken
UC Berkeley
emilyaiken@berkeley.edu

Guadalupe Bedoya
World Bank
gbedoya@worldbank.org

Aidan Coville
World Bank
acoville@worldbank.org

Joshua E. Blumenstock
UC Berkeley
jblumenstock@berkeley.edu

ABSTRACT

Recent papers demonstrate that non-traditional data, from mobile phones and other digital sensors, can be used to roughly estimate the wealth of individual subscribers. This paper asks a question more directly relevant to development policy: Can non-traditional data be used to more efficiently target development aid? By combining rich survey data from a "big push" anti-poverty program in Afghanistan with detailed mobile phone logs from program beneficiaries, we study the extent to which machine learning methods can accurately differentiate ultra-poor households eligible for program benefits from other households deemed ineligible. We show that supervised learning methods leveraging mobile phone data can identify ultra-poor households as accurately as standard survey-based measures of poverty, including consumption and wealth; and that combining survey-based measures with mobile phone data produces classifications more accurate than those based on a single data source. We discuss the implications and limitations of these methods for targeting extreme poverty in marginalized populations.

KEYWORDS

program targeting, mobile phone metadata, machine learning, poverty

ACM Reference Format:
Emily L. Aiken, Guadalupe Bedoya, Aidan Coville, and Joshua E. Blumenstock. 2020. Extended Abstract - Targeting Development Aid with Machine Learning and Mobile Phone Data: Evidence from an Anti-Poverty Intervention in Afghanistan. In *ACM SIGCAS Conference on Computing and Sustainable Societies (COMPASS) (COMPASS '20), June 15–17, 2020, , Ecuador*. ACM, New York, NY, USA, 2 pages. https://doi.org/10.1145/3378393.3402274

Link to Full Paper:
http://www.jblumenstock.com/files/papers/jblumenstock_ultra-poor.pdf

Permission to make digital or hard copies of part or all of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for third-party components of this work must be honored. For all other uses, contact the owner/author(s).
COMPASS '20, June 15–17, 2020, , Ecuador
© 2020 Copyright held by the owner/author(s).
ACM ISBN 978-1-4503-7129-2/20/06.
https://doi.org/10.1145/3378393.3402274

1 INTRODUCTION

Program targeting—the task of determining who is eligible and who is ineligible for humanitarian aid—is a major source of inefficiency in anti-poverty program administration [7]. Typically, aid programs in developing countries make targeting decisions based on survey-based measures of assets or consumption [9]. In many developing countries, however, reliable data for targeting do not exist and would be prohibitively expensive to collect [10]. Over the past several years, a handful of studies have shown that non-traditional "digital trace" data—behavioral indicators recorded in everyday interactions with technology—are predictive of wealth in developing contexts [e.g. 5, 11]. There is optimism in the machine learning and development communities that these data could provide a quick and low-cost alternative to standard field-based targeting methods [e.g. 3, 8].

This paper evaluates the extent to which digital trace data can be used for program targeting. Specifically, we match mobile phone transaction logs (call detail records, or CDR) to household survey data from a World Bank-led impact evaluation of the Afghanistan government's Targeting the Ultra-Poor (TUP) program [1]. In the TUP program, ultra-poor households are targeted for aid based on the combination of a community wealth ranking and a measure of multiple deprivation. We evaluate the accuracy of machine learning methods leveraging CDR data in comparison to two standard targeting methods—asset-based wealth and consumption expenditure—for differentiating the ultra-poor households deemed eligible for the TUP intervention from ineligible non-ultra-poor households.

2 THE TUP PROGRAM

Our ground-truth survey data comes from the Targeting the Ultra-Poor (TUP) program implemented by the government of Afghanistan with support from the World Bank in 80 of the poorest villages of Balkh Province, Afghanistan. The TUP program included a randomized controlled trial evaluating the impact of a "big push" program for lifting the ultra-poor out of poverty with multi-faceted interventions [1]. TUP households were targeted for aid based on being ultra-poor, defined by a community wealth ranking and a follow-up verification to check that ultra-poor households met several qualifying criteria relating to living conditions, education, marital status, and other measures of vulnerability and deprivation. We use survey data for both ultra-poor and non-ultra-poor households from the baseline TUP survey conducted from November 2015 to

April 2016, covering 2,899 households. We focus on three measures of poverty recorded for each household: asset-based wealth (the first principal component in variation in ownership for a number of different assets), consumption expenditure, and the designation of households as ultra-poor or non-ultra-poor.

We obtain informed consent from TUP households to match their survey responses to their Call Detail Records (CDR). The CDR are provided by one of Afghanistan's largest cell providers; 537 households match between the TUP survey and our CDR for November 2015 to April 2016. Unmatched households either do not own a mobile phone (80% of households surveyed in the TUP baseline report owning at least one phone), subscribe to a mobile network operator other than the one for which we have CDR (our operator's market share is estimated at around 30%), or do not use their phone in this 6-month time period.

The households in our sample participated in hundreds of thousands of CDR transactions in these months, including phone calls, SMS messages, and airtime top-ups. From these CDR, we compute hundreds of behavioral indicators capturing aggregate aspects of each individual's mobile phone use, including features relating to an individual's overall behavior (for example, average call duration and percent initiated conversations), their network of contacts, their spatial patterns based on cell tower locations, and their top-up patterns. We then train machine learning methods to predict whether or not a household is ultra-poor based on their patterns of mobile phone use, and evaluate our methods out-of-sample via cross validation.

3 OVERVIEW OF RESULTS

Our core analysis compares the ability of three different methods to identify ultra-poor households: an asset-based wealth index; consumption expenditure; and a new machine learning approach based on features extracted from phone data. We evaluate targeting on ultra-poverty headcount using standard targeting metrics including accuracy, coverage, and leakage. To evaluate the trade-off between inclusion errors and exclusion errors resulting from selecting alternative targeting thresholds, we also evaluate methods based on receiver operating characteristic (ROC) curve.

Our key finding is that in our sample of 537 phone-owning households in a set of poor villages in one province of Afghanistan, machine learning methods leveraging behavioral indicators computed from CDR are approximately as accurate as standard asset- and consumption-based methods for identifying ultra-poor households. We also evaluate a simple ensemble learning method to combine information on assets and consumption with CDR, and find that this combined method performs better than targeting based on any single data source.

To evaluate a CDR-based method for identifying the ultra-poor under conditions of incomplete phone ownership, we generate a sample with phone ownership reflective of the demographics of phone ownership in the overall TUP sample by adding synthetic households in proportion to the size of the non-phone-owning ultra-poor and non-ultra-poor. We consider methods that use CDR-based targeting for those with phones and classify the remaining households as either all ultra-poor or all non-ultra-poor. We find that classifying all non-phone-owning households as ultra-poor maintains a similar standard of accuracy to our original benchmark, while classifying all non-phone-owning households as non-ultra-poor yields significantly worse targeting outcomes.

4 POLICY IMPLICATIONS

These results extend past work on wealth estimation from mobile phone data to suggest that CDR and other digital trace data could be used in practice to target anti-poverty programs or other development interventions. Moreover, recent reviews of standard field-based poverty targeting schemes find that targeting is limited by low-quality ground truth data on poverty across programs and regions [6, 7]. CDR-based methods like the one presented here could provide a lower-cost complement to standard targeting methods without sacrificing accuracy, and could be particularly useful in times of conflict or humanitarian crisis when field-based targeting is infeasible [4].

Of the many limitations of this approach, we highlight three. First, CDR-based targeting applies only to households that own a mobile phone, so other options are required to reach households without phones. Second, important ethical and privacy considerations arise when personal data is used for any purpose, including humanitarian ones [12]. Finally, we worry that basing benefits on phone use may lead to strategic gaming, particularly if the algorithm is made transparent for legal or ethical reasons [2]. Given these limitations and our promising results on combining CDR data with standard survey measures for increased classification accuracy, CDR-based methods may be best deployed in conjunction with standard targeting methods so that survey-based data on poverty is complemented by digital trace data.

REFERENCES

[1] G. Bedoya, A. Coville, J. Haushofer, M. Isaqzadeh, and J. Shapiro. 2019. No household left behind: Afghanistan Targeting the Ultra Poor impact evaluation. *World Bank Policy Research Working Paper* 8877 (2019).
[2] Daniel Björkegren, Joshua E. Blumenstock, and Samsun Knight. 2020. Manipulation-Proof Machine Learning. *arXiv:2004.03865 [cs, econ]* (April 2020).
[3] Joshua Blumenstock. 2016. Fighting poverty with data. *Science* 353 (2016), 753–754.
[4] Joshua Blumenstock. 2020. Machine learning can help get COVID-19 aid to those who need it most. *Nature* (May 2020). https://doi.org/10.1038/d41586-020-01393-7
[5] J. Blumenstock, G. Cadamuro, and R. On. 2015. Predicting poverty and wealth from mobile phone data. *Science* 350 (2015), 1073–1076.
[6] C. Brown, M. Ravallion, and D. van de Walle. 2018. A poor means test? Econometric targeting in Africa. *Journal of Development Economics* 134 (2018), 109–124.
[7] D. Coady, M. Grosh, and J. Hoddinott. 2004. Targeting outcomes redux. *The World Bank Research Observer* 19, 1 (2004).
[8] M. De-Arteaga, W. Herlands, D. Neill, and A. Dubrawski. 2018. Machine learning for the developing world. *ACM Transactions on Management Information Systems (TMIS)* 9, 2 (2018).
[9] R. Hanna and B. Olken. 2018. Universal basic incomes versus targeted transfers: Anti-poverty programs in developing countries. *Journal of Economic Perspectives* 32 (2018), 201–226.
[10] Peter Jerven. 2013. *Poor Numbers.* Cornell University Press.
[11] E. Sheehan, C. Meng, M. Tan, B. Uzkent, N. Jean, D. Lobell, M. Burke, and S. Ermon. 2019. Predicting economic development using geolocated Wikipedia articles. *KDD* (2019).
[12] Linnet Taylor. 2016. No place to hide? The ethics and analytics of tracking mobility using mobile phone data. *Environment and Planning D: Society and Space* 34 (2016). Issue 2.