

Urban Networks and Targeting: Evidence from Liberia

By LORI BEAMAN, NIAL KELEHER, JEREMY MAGRUDER, AND CARLY TRACHTMAN*

In many contexts, informal social networks are critical for smoothing away frictions (e.g. labor market search) or filling in for formal institutions which do not exist (e.g. insurance). To do so, information must flow within networks and be accessible when needed. Policymakers have a similar goal when targeting cash transfers. In the absence of high quality administrative data on residents, policymakers often leverage pre-existing social and political institutions to identify beneficiaries and enhance the legitimacy of the program. In this paper, we ask residents of Monrovia, Liberia to identify poor neighbors and target an unconditional cash transfer to reveal the quality of information that can be extracted from this urban network.

We build on results from three important studies. First, Alatas et al. (2012) show that community members are relatively efficient in identifying poor households when compared to a proxy means test in Indonesia. Second, Alatas et al. (2016) show that the quality of knowledge deteriorates with social distance, but community members are a useful repository of information for targeting.¹ Finally, Banerjee et al. (2019) demonstrate that people in rural India identified highly central people in their community when asked to nominate others to help spread information. If the highly central have useful information about the network

as in Alatas et al. (2016), this suggests a targeting strategy: simply ask who would be a useful person to assist in targeting. We contrast the information and targeting performance of several key people in social networks: local leaders, random neighbors, and nominated neighbors.

I. Experimental Design and Setting

There are four key steps in our experimental design:

- 1) Data on household welfare and social connections was collected in 2018 by means of a household census (92% survey rate with 2,433 surveys completed) in 13 urban neighborhoods of Monrovia;
- 2) During the household census, we asked all households to nominate a member of their community to assist with targeting a social protection program;
- 3) We invited a subset of community members to first provide information on the poverty status of fellow community members (community knowledge interview), and then nominate up to 2 community members to receive a cash transfer ;
- 4) We then selected 280 households to receive a one-time transfer of \$80.

We capture household welfare with two main measures: per capita daily expenditure (PCE) and a proxy-means test based on asset ownership (PMT).² These measurements are correlated but different: in our sample the correlation between the PCE and PMT measures is approximately 0.247.

²See Online Appendix for a detailed description of the methodology used to generate the proxy-means test and details on implementation of the cash transfer.

* Beaman: Department of Economics, Northwestern University (email: l-beaman@northwestern.edu); Keleher: School of Information, University of California Berkeley (email: nkeleher@berkeley.edu); Magruder: Agricultural and Resource Economics, University of California Berkeley (email: jmagruder@berkeley.edu); Trachtman: Agricultural and Resource Economics, University of California Berkeley (email: ctrachtman@berkeley.edu). We thank Ellen Bates-Jefferys and the IPA-Liberia team. The RCT was registered as AEARCTR-0001164.

¹Hussam, Rigol and Roth (2020) find that community members have even more subtle information such as marginal returns to capital.

In Appendix Table 2, we present summary statistics on our sample. The households we survey are poor: the average per capita daily consumption is about \$2.51 (PPP), and the median person in our sample consumes about \$1.63 per day. They also experience frequent, serious shocks. Half of the sample has experienced an income shock in the past year (either a lost source of income or lost property or major asset), and 32% has experienced a serious health shock which led to a work stoppage or death. Extremely few households report receiving any assistance from government or NGOs. Residence in communities is rather stable; the average person has lived in their community for over 10 years. These facts together suggest that we may well expect community networks to exist and play a pivotal role in boosting welfare during periods of hardship.

A. Social Networks in Monrovia

In our household census, we asked households to list up to five people in their social network who they spent the most time with (by choice) in the last 14 days.³ This elicitation allows us to observe who is central in these networks. Appendix Figure 1 shows the correlates of eigenvector centrality in these community networks. Larger households, households with a community leader,⁴ and households who have been in the community longer tend to be more central in these networks. In addition, the most central people tend to be relatively poor: per capita consumption (and even total household consumption) are negatively correlated with eigenvector centrality. This suggests that the poor are well integrated into these networks, which again supports the validity of a community-targeting approach.

³We elicited only connections within administratively-defined neighborhoods. People may have important social connections in other communities, but we are interested in whether we can elicit good information from local community members.

⁴We define leaders as anyone that is a self-reported member of community government or holds a leadership position in a social, youth, health, or religious organization.

B. Identifying Nominated Neighbors

Building on Banerjee et al. (2019), we asked randomly-selected community members to help us identify central people in the network. Specifically, we asked community members “If we want to spread information about a social assistance program, to whom do you suggest we speak?”⁵

From our census and these referrals, we then identified three lists of community members. First, we selected 175 community residents at random (random neighbors) from our census to provide a benchmark. Second, we identified 345 community leaders, and selected 48 at random. Third, we identified 820 people nominated by another community resident, and selected 254 of these nominated neighbors at random.⁶ On average, random neighbors received 0.99 (s.d. = 2.07) nominations, community leaders received 3.1 (s.d. = 5.6), and nominated neighbors received 3.8 (s.d. = 4.1) nominations.

Leaders and nominated neighbors have similar characteristics, but are different from the population as a whole (Appendix Table 3). Compared to random neighbors, leaders and nominated neighbors are wealthier, with larger households, more assets, and greater expenditures. They are also more central in the network, in both in-degree and eigenvector centrality. This may reflect that leaders were quite likely to be nominated.⁷

We conclude that (i) the process of decentralized nomination of local community members and (ii) self-identification of local leaders both selected individuals who are central and elite in these networks.

⁵We randomly varied some aspects of the elicitation script; we do not observe large differences between neighbors nominated under different elicitations. See Online appendix table 4.

⁶Nominated neighbors were selected proportional to the number of nominations they received.

⁷This is consistent with evidence from Indonesia: when asked who might be a good representative to make targeting decisions, about 67% of respondents chose a government leader of some type and around another 22% of respondents choose a religious, tribal, or other community leader (Alatas et al., 2012 and author calculations).

TABLE 1—NEIGHBOR TARGETING QUALITY

Panel A	Random	Leader	Nominated	Random	Leader	Nominated
% Households Known	32.0% (1.4)	41.3% (1.8)	35.8% (1.0)			
Panel B	<i>Per Capita Expenditures</i>			<i>Proxy Means Score (Assets)</i>		
% Poor Households Known	32.0% (1.7)	42.4% (2.1)	35.9% (1.3)	32.3% (1.7)	42.4% (2.1)	34.5% (1.2)
% Poverty Status Rated Correctly	60.1% (2.1)	62.5% (2.4)	62.1% (1.4)	61.0% (1.5)	63.9% (1.6)	64.6% (1)
% Cash Nominations to Poor	27.2% (2.7)	25.4% (3.0)	23.0% (2.1)	30.6% (2.8)	32.1% (3.2)	26.8% (2.2)

Note: Rows 1-3 show targeting neighbor’s knowledge of randomly selected households within the same neighborhood. Row 4 shows the percentage of households nominated by a neighbor that are among the poorest 20% of households within a neighborhood, as defined by the poverty metric. The poverty metric in columns 1-3 is self-reported per capita daily expenditures and is proxy means score in columns 4-6. Columns correspond to the neighbor type: “Random” indicates a randomly selected neighbor; “Leaders” indicate community leaders; and “Nominated” indicates a neighbor nominated by at least one resident of the neighborhood. Standard errors are in parenthesis.

II. Do Neighbors Identify the Poorest?

In each community knowledge interview, we completed two tasks. First, we presented a neighbor with a list of about 30 community residents.⁸ We asked the neighbor whether they knew each household. If they did know a household, we asked whether the household was poor, which we defined for them as being in the bottom quintile of the community. Second, after completing the community knowledge interview, we informed the neighbor that we would be giving a grant of \$80 USD to a number of poor households in the community, and then asked for their help in identifying recipients. This was a meaningful decision: the grants were equal to about 2 weeks of household expenditures for the median household, and poor households had a 15% chance of receiving one of these grants.

In Table 1, Panel A, we find that leaders know the most households, followed by nominated neighbors and then random neighbors (41% versus 36% and 32%, respectively, $p < .05$ for all comparisons). Panel B shows that poor community members - captured by either the PMT or PCE -

are known by all types of neighbors at similar rates to the non-poor. This is consistent with the descriptives in section I.A that poor households are not socially isolated in these communities, as well as with previous evidence that leaders are often quite capable of accurate targeting, despite the potential distortion of elite capture (Basurto, Dupas and Robinson, 2020).

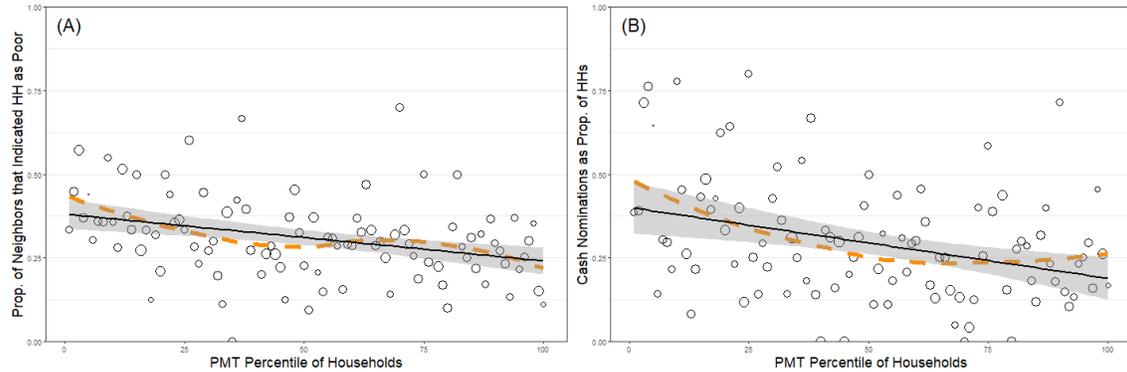
At the same time, all types of neighbors are inaccurate in assessing the poverty of the neighbors they know. Overall, neighbors accurately assessed whether a household was poor about 63% of the time (based on PCE) or 64% of the time (PMT). PCE accuracy is not statistically differentiable between leaders, nominated, or random neighbors, while nominated neighbors slightly outperform random neighbors in PMT rankings.⁹ While poverty was defined as the bottom quintile, neighbors reported that about 30% of their neighbors were poor; if they randomly assigned 30% of households as poor regardless of their underlying poverty status, we would expect to see a very similar 62% correct assignment.¹⁰ Could different neighbors do bet-

⁸In practice, the number of community residents ranged between 17 and 48 depending on the size of the community.

⁹65% versus 61%, $p = 0.044$. Given the number of comparisons (6) this may be attributable to chance.

¹⁰Note that the actual accuracy was in fact worse

FIGURE 1. IDENTIFYING POOR NEIGHBORS & ALLOCATING TRANSFERS BASED ON ASSET POVERTY



Note: The figure shows the relationship between targeting information provided by neighbors and poverty of households in the neighborhood. Panel A shows the relationship between the proportion of neighbors who voted households as poor and the percentile of households in the within-neighborhood proxy means score distribution. Panel B shows the relationship between the proportion of households nominated for a cash grant and the percentile of households in the within-neighborhood proxy means score distribution. The black line shows the weighted-least squares regression line. The shaded area shows the 95% confidence interval. And the orange dashed line depicts the local polynomial regression fit with smoothing parameter, α , equal to 0.75.

ter? In Appendix Table 5, we verify that the same qualitative results hold for three different types of neighbors: those who were nominated by many of their neighbors, those who are themselves poor, and those who were frequently nominated by the poor. None of these groups are particularly effective at characterizing their neighbors' poverty.

Figure 1 (Panel A) shows the non-parametric relationship between a household's actual poverty status - as measured by the PMT - and the proportion of neighbors who indicated that household was poor. The x-axis is the within-block percentile of a household in the distribution of consumption or assets, and the y-axis indicates the probability that that household is marked as poor. There is a statistically significant relationship between assessed poverty and PMT-defined poverty. The magnitude is small, however; moving from the 10th to the 90th percentile of assets is associated with only a 17.6% decrease in the probability of being identified as a poor household (on a base of 35.7% for the 10th percentile).¹¹

than picking at random if neighbors correctly assigned 20% of households as poor.

¹¹Appendix Figure 2 shows a still weaker relationship between knowledge and poverty measured by PCE.

III. Are Poor Households Targeted for Transfers?

Whether or not neighbors have the information to identify poor households, do neighbors select poor households to receive transfers? The last row of Panel B of Table 1 explores this. We find that all three groups allocate the transfer to a poor household somewhat more frequently than could be explained by chance: if we use consumption (asset) metrics to define poverty, neighbors chose a poor beneficiary about 25% (29%) of the time, with small and statistically insignificant differences across types of neighbors (random, leaders, and nominated neighbors). Both in terms of consumption and assets, neighbors' nominations are statistically different from the 20% accuracy ($p < .01$) which would be achieved by directing transfers at random, but no one group of neighbors outperforms another in a statistical sense. Panel B of Figure 1 shows the non-parametric relationship between asset-based poverty of a household and the likelihood the household was nominated for a transfer. We see a steeper relationship than in panel A. In Appendix Table 2, we see similar rates of poor households selected for transfers among nominated neighbors, poor neighbors, and neighbors nominated by the

poor. We are again unable to identify a group of neighbors who direct transfers to the poor at higher rates.

These results show that despite the relatively poor knowledge that neighbors have about most households in the community, they target poor households to receive transfers somewhat more often than would be attributable to chance. In other words, while neighbors may be unable to identify the landscape of poverty in their neighborhood, many can identify one or two poor households, and will refer those households for social support. Of course, 25% success (or 29%) is far from an ideal 100% of transfers going to poor households. It is also worse than PMT targeting: 35% of households identified as poor by the PMT are poor according to PCE as well. Thus, either information frictions or incentive problems may leave a large gap between realized and ideal community targeting in Monrovia.

IV. Lessons for Targeting in Urban Neighborhoods

To target transfers at poor households in Monrovia, we solicited the help of the community network. We did so in three different ways: asking people at random who should receive the transfer; asking leaders who should receive the transfer; and asking random people to nominate neighbors to identify who should receive the transfer. All three of these approaches led to better targeting than would be achievable by random selection. The gains, however, were modest in each of these approaches, and none excelled over the others. In a broad sense, we also found that all three groups have limited knowledge of community poverty, as demonstrated by their low accuracy at ranking community members as “poor”. This happened despite two key characteristics of the networks: (i) stability: most people, including these neighbors, have resided in the same community for about a decade, and (ii) the fact that the poor are well integrated in the networks.

In this study, we highlight that not all networks have information that can be easily leveraged for community targeting. This

poor information may have broader implications for how well networks function, including risk sharing. Our setting is an impoverished urban environment, and our results are consistent with other work which suggests networks in urban areas may have poorer information than those in rural areas.¹² Identifying whether information frictions limit the effectiveness of urban networks remains a key priority for future research.

REFERENCES

- Alatas, Vivi, Abhijit Banerjee, Arun G. Chandrasekhar, Rema Hanna, and Benjamin A. Olken.** 2016. “Network Structure and the Aggregation of Information: Theory and Evidence from Indonesia.” *American Economic Review*, 106(7): 1663–1704.
- Alatas, Vivi, Abhijit Banerjee, Rema Hanna, Benjamin A. Olken, and Julia Tobias.** 2012. “Targeting the Poor: Evidence from a Field Experiment in Indonesia.” *American Economic Review*, 102(4): 1206–1240.
- Banerjee, Abhijit, Arun G Chandrasekhar, Esther Duflo, and Matthew O Jackson.** 2019. “Using gossips to spread information: Theory and evidence from two randomized controlled trials.” *The Review of Economic Studies*, 86(6): 2453–2490.
- Basurto, Maria Pia, Pascaline Dupas, and Jonathan Robinson.** 2020. “Decentralization and efficiency of subsidy targeting: Evidence from chiefs in rural Malawi.” *Journal of Public Economics*, 185: 104047.
- Hussam, Reshmaan, Natalia Rigol, and Benjamin Roth.** 2020. “Targeting High Ability Entrepreneurs Using Community Information: Mechanism Design In The Field.” *HBS Working Paper*, 20-082.

¹²For example, in the Indonesian sample from Alatas et al. (2012), neighbors were more accurate in identifying PCE poor households in rural hamlets than urban hamlets. See the Online Appendix for details.