

Chapter 4

Identifying Program Beneficiaries

Vivi Alatas

Asa Kreativita

Benjamin A. Olken

Massachusetts Institute of Technology

Abhijit V. Banerjee

Massachusetts Institute of Technology; Abdul Latif Jameel Poverty Action Lab (J-PAL); Paris School of Economics

Rema Hanna

Harvard Kennedy School

I. INTRODUCTION

Should governments target assistance? Governments face this question in addressing both temporary and permanent poverty. In some cases, there can be compelling political, economic and/or practical arguments for dividing a given social assistance budget equally and providing transfers to everyone. For example, if everyone is eligible for a program or policy, it may raise political support for a program or policy and increase tax collection, and thus the budget available for the program. Indeed, some government programs are universal: everyone, or everyone within a certain category (e.g., children, new parents, the elderly), receives the program.

On the other hand, with a limited budget, universal programs may end up channeling substantial resources to those who need it less, mechanically leaving fewer resources available for those most in need of help, i.e., those who are food insecure, those with disabilities, pregnant women, etc. (Hanna and Olken 2018). This may be particularly true in low- and middle- income countries where tax collection rates are low, so government budgets are lower as a share of GDP, and at the same time, there are many competing budgetary needs (e.g., health, infrastructure, education). Thus, many government programs worldwide often seek to channel their limited resources to those who need them the most, or to provide

additional assistance to those who are poorer. Doing so requires a system for identifying who should receive transfers in this case that is effective, feels fair and politically legitimate.¹

This chapter explores the second case: what are the challenges in identifying who should be eligible for a program in the cases when the government has decided to target eligibility to a particular sub-population. The particular sub-population may, of course, vary from program to program depending on the program's goals. In some cases, the goal is to identify the very poor, many of whom work but earn incomes that keep them below the poverty line – e.g., the 719 million people worldwide who live on less than \$2.15 a day. In other cases, the goal is to identify households where assistance can help households escape a 'poverty trap,' whether in a single generation (e.g., Banerjee, Duflo, and Sharma 2021) or through investments in children that can break inter-generational cycles of poverty (Aizer et al. 2016; Cahyadi et al. 2020). In yet other cases, the government may aim to provide basic assistance to those who cannot work, for reasons such as disability, mental health or old-age, and who may need long-run assistance to provide a basic standard of living. Making it even more complex, the goal may be about improving other forms of well-being, such as nutrition or child development, which are not always fully aligned with income levels.

¹ More discussion about this debate is include in Chapter 6.

There are different sets of methods to target, and there are always tradeoffs when choosing among this set. In this chapter, we focus on how the unique circumstances in low- and middle-income countries – particularly, high levels of informality and consequent limited income information available to the government – shape targeting choices. We begin with a short discussion of these information challenges. We, then, review the current state of the literature on targeting methods and propose open research questions for the future.

II. INCOME TARGETING

A. The challenge

What would a government try to target? In the classic formulation by Mirrlees (1971), the goal for redistribution is to target ‘earning ability,’ with the idea that this is a characteristic of an individual and does not depend on choices, such as labor supply. Since earning ability is not directly observed, targeting in high-income countries is most frequently based on income, which is strongly related to earning ability (albeit affected by labor supply decisions) and is more observable. Income-based redistribution is done through the progressive tax system and through a wide variety of means-tested programs.

However, income-based targeting requires three important steps, of which gathering data on incomes is only just one. First, governments typically create and

assign a unique identifier to each person. In the United States, for example, this system of identifiers – Social Security Numbers – was developed as part of the introduction of a social protection system, the Social Security social pensions, so the government could track each individual's earnings over time and tie their benefits at retirement to the social security taxes.

Second, governments need to match information on people's income to that identifier. In high income countries, most employment occurs through formal firms. Double-reporting by those firms– firms tell both the government and the individuals how much each person was paid – allows governments to reliably measure most people's incomes (Kleven et al. 2011; Kleven, Kreiner, and Saez 2016). For example, in the United States, pandemic aid to individuals was targeted based on the previous years' tax returns, and the Earned-Income Tax Credit, which provides assistance to working families, also relies on income tax data. This system can also be extended to cover income shocks – for example, if a person loses their job, there is a paper trail through payroll tax payments that can be used for unemployment insurance.

Third, governments need an identification system and distribution channel to ensure that the people claiming benefits are the same people who were deemed eligible. There is a need to authenticate (ensuring that the person is who he/she

claimed to be) and to deduplicate (identify duplicate records), and then to have a financial system in place that could safely make transfers.

On net, targeting based on incomes is not a panacea – and indeed, there is a literature in high income countries about potential labor supply responses when benefits phase out very rapidly as income grows (see, e.g., Saez 2002; Keane 2011) – but the ability to observe incomes and match them to specific individuals gives high-income governments a large advantage over low and middle income countries that lack this information.

In short, low- and middle-income countries face challenges at each of these three steps – establishing a unique identifier, linking incomes to that identifier, and verifying identity at program receipt. We discuss each of these in turn.

B. Establishing a unique identity

Designing an identification system that is both comprehensive – i.e., that includes everyone – and that is unique – i.e. one ID per person – is challenging. In many countries, many people lack formal birth records (or death certificates), which makes it hard to ensure that the identification systems are inclusive and up-to-date. Low- and middle- income countries may also lack the technological capacity to conduct proper authentication and deduplication, and to ensure that someone who presents an ID is in fact the person in question.

To address these challenges, many countries are investing heavily and creating biometrically authenticated identification systems. The idea behind biometrics is that they can solve both challenges at once – they can be used at the first step to de-duplicate identity systems, creating ‘uniqueness’ in the unique identity system – and that they can also be used in the third step to verify in person that someone claiming benefits is in fact the eligible person, thereby increasing the effectiveness of aid (Muralidharan, Niehaus, and Sukhtankar 2016).² Examples include India’s Aadhaar card, Pakistan’s Computerized National Identity Card (CNIC), and the Nigerian National Electric Identity Card. Togo, for example, was only able to launch emergency cash aid during the pandemic (the NOVISSI program) because they had just created a biometric voting identification card for voting purposes.

A key challenge with relying on biometric identity card systems for the distribution of benefits is that they need to be complete, since anyone *without* such a card could wrongly end up excluded from the system. This is unfair to those who could not get registered yet, those who lack the documentation that the government requires to establish citizenship, and many others, including those with missing or partially eroded fingerprints that make the biometrics not work well (Muralidharan, Niehaus, and Sukhtankar forthcoming). Understanding how to make sure these

² Some of the biometric systems countries are experimenting with include fingerprints, retina scans, and facial recognition.

systems are inclusive, or, alternatively, designing systems that allow exceptions and maintain access for those who need it without undermining the identification requirement, is important.

C. Figuring out who needs assistance

The other piece of the puzzle – linking in information about people’s incomes – is potentially even more challenging in low- and middle-income countries. This is largely because informal work and self-employment are common, and so there are no verifiable records of employment or income. In fact, about 70 to 90 percent of households are not even included in the tax net for this reason (Jensen 2022), making it impossible to use the tax system to identify who needs assistance. The fundamental lack of verifiable income information means that most of the strategies used in high-income contexts do not work in low and middle-income contexts.

One option is to use unverified, self-reported income data. Perhaps most notably, Brazil bases eligibility for its cash transfer program on self-reported, unverified income data, though it supplements this by doing some cross-checks with administrative data (for those in the formal sector) and flags people with predicted consumption that substantially exceeds reported income (Lindert et al. 2005). However, many other countries are skeptical of using this strategy due to fear that there could be substantial under-reporting of incomes; the degree to which

this happens in practice in different contexts (say where administrative data is not as good as Brazil) is as an open question for future work.³

These challenges in income verification have led to the development of alternative methods for channeling benefits to those with low incomes (Hanna and Karlan 2017)—this includes geographic targeting, proxy-means testing, categorical targeting, self-targeting, and community targeting. We next discuss each of these methods. However, note that while we discuss each individually for ease of exposition, many governments often choose to use some combination of the various methods to try to take advantage of each method’s unique advantages.

Geographic targeting—in which one targets programs to poor regions within a country—is one of the simplest targeting methods as it does not rely on large-scale, individual level data. Instead, one predicts which areas have the highest incidence of poverty from a simple, representative dataset, and then make the programs available to everyone who lives within these chosen areas (Baker and Grosh 1994; Elbers et al. 2007). In recent years, as newer remote-sensed or administrative datasets have become available – from mobile phone metadata to satellite data to Google Street Views – as well as innovations in statistical prediction techniques, it is now possible to improve prediction and to consider even smaller

³ Theoretically, if everyone mis-reported by the same amount, the government could correct for this – the problem arises when there is unobserved heterogeneity in the propensity to misreport; see Banerjee et al. (2022).

geographic areas (e.g., Jean et al. 2016; Blumenstock, Cadamuro, and On 2015; Naik et al. 2017). The challenge in doing so is showing their accuracy for fine geographic units (Tarozzi and Deaton 2009). Understanding what levels of geographic disaggregation these newer methods can provide with what type of statistical accuracy, as well as understanding whether these new administrative datasets provide a good picture of need, is important for future research.

Geographic targeting has clear benefits: it is easy to administer and can be simple for governments to explain in terms of the transparency of programs (although, maybe less so if even finer geographic locations are used). It is also a form of providing benefits to all in a particular location, which may help with politics and fairness concerns (at least within those areas). They may also work well if there are areas with high poverty density. However, these approaches tend to have higher targeting error than approaches that use individual data (i.e. since rich people in poor regions may get access, and poor people who live in richer areas will not). It is also an open research question whether these methods would allow for real-time changes in vulnerability over time, as people or places experiences shocks. Moreover, there is concern that these types of place-based policies may discourage migration to areas where more jobs or opportunities may exist, as households would be worried to migrate and lose their benefits. For example, Imbert and Papp (2020) show that this has been the case in India's public work program, Mahatma Gandhi National Rural Employment Guarantee Act, which was

targeted to rural areas.⁴

An alternative method is *Proxy-Means Test (PMT)*. This method uses a “proxy” to determine a person’s true economic status (e.g., per-capita household consumption, income, etc.).⁵ It is more data intensive than geographic targeting in that it requires large-scale data on individuals or households. Specifically, it relies on two datasets. The first data set is typically a nationally representative sample that is collected for research or general statistical purposes and has information on both what the government is trying to target on (let’s say, per capita household consumption) and a set of potential proxies for it that may be more observable (e.g. assets such as car or TV ownership, as well as immutable demographics such as the age and education level of the household head). Using these data, one predicts per capita consumption with the set of proxies (e.g. through OLS or new machine learning methods) and generate a formula that can be used to predict consumption from the proxies. Next, the government needs to collect a second dataset on the proxy variables for the entire relevant population through a household census of the entire country, or a ‘social register’ that contains information from anyone who may

⁴ Given this concern, governments are experimenting with finding ways (e.g. through better data systems, through the use of digital transfers) to allow people to keep their benefits even if they temporarily migrate to different areas within the country (Baseler et al. 2023) – but this means that the program is no longer geographically targeted to the same extent.

⁵ To the best of our knowledge, modern PMTs were first used in the early 1980s in Chile for the targeting of its Ficha CAS program (Coady, Grosh, and Hoddinott 2004), though at some level using observable proxies for income is a much older idea – see, for example, the discussion of Britain’s ‘window tax,’ in place from 1696 to 1851 (Oates and Schwab 2015).

plausibly be eligible. The proxy variables that are collected can be used to generate a predicted per capita consumption score for each household. Anyone household that scores below a pre-specified cutoff (e.g. based on a poverty line, based on program budget, etc.) then receives the program.

PMTs have several features and challenges. First, given that PMTs aim to use observable assets and demographics in the prediction formula, they tend to capture the more permanent component of household income. Moreover, given the scope of the data collection effort involved, the data collection—and, hence, the eligibility determination—typically only occurs every few years. Given both of these features, PMTs are often most appropriate for programs that aim to address chronic poverty, e.g. long-run cash or conditional cash transfer program, long-run food subsidies, health insurance subsidies, rather than programs that are designed to provide short-term help against economic shocks. Second, there are a number of prediction challenges: the variables used are often limited to both what is in the sample survey and census, there may be small sample issues if one wants to allow the relationship between assets and predicted consumption to differ in smaller geographic areas, and the prediction is often better if you are trying to predict a larger group (i.e. trying to predict the bottom 40% rather than the bottom 5%).

It is also important to note that PMTs, if the models are trained to target *consumption*, may also have a slightly different objective than income-based

targeting methods, which are based on *incomes*. Going back to the conceptual measurement discussed above – that earning ability may be ultimately what policy makers are trying to target - both consumption and incomes are related to earning ability, but each involve different choices in between – how much labor to supply, how much to consume today vs. to save for tomorrow.

Recently, there have been increasing attempts to use administrative data to conduct PMTs rather than collecting data periodically through large-scale surveys. For example, Blumenstock, Cadamuro, and On (2015) use anonymized data from mobile phone networks, while Aiken et al. (2023) use mobile phone metadata. The fact that these kinds of administrative datasets are automatically updated suggest that there is potential to use them to predict when a household has a shock and is in need of temporary assistance. Gentilini et al. (2020) document how governments, such as Pakistan and Togo, used these kinds of approaches to roll-out social assistance programs quickly during the 2020 COVID-19 pandemic. Moreover, there is often discussion in many countries about using land ownership records or tax records to improve the targeting by “knocking off” the very rich. However, there are many legal and ethical issues regarding the use of administrative data of these sorts to consider, and there are concerns for distortion of these data if it gets repeatedly used for targeting (i.e. people not paying taxes, people having more than one electricity meter, people having multiple SIM cards). Nonetheless having a better understanding whether they could be more flexible in helping to address

shocks is important.

As with every targeting approach, there are pros and cons of using PMTs. In terms of pros, these methods have a clear set of rules, and everyone can be considered without location restrictions (unlike in geographic targeting methods). However, like most predictive algorithms, the PMT formulas have prediction errors (Hanna and Olken, 2018; Brown, Ravallion, and van de Walle 2018).⁶ Moreover, people may have different preferences in terms of the types of assets that they want to own, and the PMT formula cannot fully account for this kind of heterogeneity across individuals; for example, someone may buy a motorbike since they need it to travel to work, but their income level could be below others who do not. Importantly, as we discussed above, the exclusion error (those who should be included but are erroneously not included) tends to increase the more narrowly these programs are targeted. This is an issue for many countries, as many current cash transfer programs are funded to only cover only a very small portion of those in need.

Given some of the PMT challenges, another form of targeting is *categorical targeting*, which provides transfers for people who exhibit a certain characteristic (i.e. transfers for those of old age, transfers for those with kids, support for pregnant

⁶ Researchers have tried to use more sophisticated machine learning methods to reduce the errors, but most papers typically find little improvement relative to common prediction methods (e.g. McBride and Nichols 2018; Baez, Kshirsagar, and Skoufias 2020; Areias and Wai-Poi 2022).

women). Within these categories, programs can be provided universally—anyone who qualifies based on the characteristic is eligible—or some means-based targeting could be done within the category. Categorical targeting often has a benefit when aiming to reach people for a particular goal (i.e. providing subsidized nutrients to pregnant women). Like geographic targeting, it is also easy to understand from a transparency standpoint. There are interesting questions relating to stigma—some have argued, for example, providing school meals to all kids (rather than just low-income children) can reduce the stigma of school meals programs. On the other hand, others have argued that categorical targeting—support for AIDS orphans—generated other forms of stigma beyond income status. We discuss categorically targeted programs in several chapters within this book, including chapter 7, which discusses programs geared at child health and development, and chapter 14, which discusses pension programs.

Two other targeting methods are also more commonly used to provide the government with “soft” information, and to potentially introduce more flexibility into the timing of when one can apply for and receive social assistance. The first is a *self-selection mechanism*, where a program applicant must perform an action that is potentially more costly or less desirable for those who are richer, to screen out those who do not need the assistance. For example, an individual may need to physically apply and stand in line for many hours, which gainfully employed people may not have the time to do. Or, the government may provide lower quality food

(e.g. lower grade rice or smaller portions of milk) that richer households may not necessarily want (Nichols and Zeckhauser 1982).

A key disadvantage of this method is that it can impose a cost on applicants. For example, one common form of self-targeting is public work programs, which require individuals to work to receive the payments.⁷ Public works programs can help target those in need if the work component screens out those that do not need assistance (since those who are already employed would not have the time or the desire to do these jobs). These programs may also offer more flexibility than other targeting methods since individuals can potentially sign up for a number of days of paid labor exactly when they need it—e.g. they have lost their jobs, their business has gone under, it is the lean season and there is no employment, or there has been a crop failure. However, if the work is unpleasant and not productive (i.e. digging ditches in the hot sun that have no productive use), this, in the end, may be a very costly way to target. Moreover, one challenge in many low- and middle- income countries is that the program wage is often set at the government's official minimum wage, but with much of labor being informal, actual wages are often much lower. Thus, with the higher minimum wage, public works programs may crowd out productive, private-sector work (Murgai, Ravallion, and van de Walle

⁷ It dates back to at least the 19th century in England, where transfers were granted through residence in a workhouse (Besley and Coate 1992), and to the United States' Civilian Conservation Corps and Works Progress Administration in the 1930s (Aizer et al. 2020).

2016; Bertrand et al. 2021).

Given the potentially high costs of public works programs for beneficiaries, an important question is whether you can generate selection, but with smaller costs. A common approach is to let people apply for programs, rather than enroll people automatically, under the argument that only those who need it will apply. Sometimes a screening mechanism is added on, when people apply—such as a PMT or categorical screen—that adds an additional layer designed to screen out those who may not need the program as much (or to act as a deterrent to apply if you know you do not need the program). Alatas et al. (2016) find that the self-targeting, combined with a PMT, reduced inclusion error relative to an automatic enrollment PMT. The reason is that many richer households did not apply, perhaps because they expected that they would not pass the screening mechanism. Though most of them would have been screened out had they applied, some small fraction of them would have (erroneously) passed the screening mechanism, the fact that they self-selected out improved targeting on net. Perhaps more surprisingly, having people apply also reduced exclusion error, as poor households that were missed under the government census to automatically enroll people came out and got included.⁸

However, it is worth noting that one needs to be careful in designing these

⁸ It is an important open research questions to understand how the advertising content and alternative screening mechanisms work to screen in those more in need.

kinds of application mechanisms to make sure that they do not screen out poor households. For example, Gupta (2017) shows that having complex application procedures that involve filling in long forms may be challenging for those who are illiterate or semi-literate and that needing local government approval for your application could screen out those who may feel intimidated by having to interact with the state. This suggests a need to carefully design these kinds of mechanisms to elicit the income screening features without screening out those who are vulnerable on other dimensions.

A second alternative method to elicit “soft” information is *community-based targeting*. Community-based targeting aims to utilize local information (Alderman 2002; Galasso and Ravallion 2005). For example, a PMT may look at two households and assume that they are of similar income status if they both have a 2-room house and a motorbike. However, one of the households may have lost their business recently or have a family member who is sick. The PMT may miss this kind of soft information, but it is possible that the community may know and bring in this kind of information. And even beyond that, communities may have a somewhat different, locally grounded notion of what poverty means. Alatas et al. (2012) find that, experimentally comparing a PMT with community-based targeting in Indonesia, the PMT test was somewhat better at identifying households based on per-capita consumption, although the households who were included by one and excluded by the other were mostly sufficiently close to the poverty line that both

approaches would perform similarly in terms of most social welfare functions. However, community targeting did substantially better in terms of identifying households who self-identify as poor, suggesting that indeed they were matching a local poverty metric that was highly correlated with, but not exactly the same as, per-capita consumption.⁹ Note, however, that Trachtman, Permana, and Sahadewo (2022) find (in a lab in the field study in Indonesia) that community targeting results tended to match long-run poverty rather than short-run income shocks, suggesting that a path for future research to see whether one could use this method to target programs that provide short-run assistance.

Another aspect of community targeted programs that is worth discussion is the political economy considerations. Targeted programs could have negative political implications if some people are receiving a program and others are not, but community targeting could negate these negative implications by giving people more voice in the process. On the other hand, involving communities could be negative if community members argue over rankings and this causes discord, and/or if elites capture the process (or there is a perception of elite capture even if it is not true). Exploring elite capture, both Alatas et al. (2019) and Basurto, Dupas, and

⁹ Premand and Schnitzer (2021) find similar results that while the PMT did better at matching PMT than a community based approach in Niger, community-based targeting does a better job matching self-assessed welfare status. Several other studies also find that community-based targeting does worse than proxy-means tests on matching per-capita consumption, though these typically studies do not examine self-assessed welfare as an outcome (Stoeffler, Mills, and del Ninno 2016; Beaman et al. 2021; Dupas, Fafchamps, and Houeix 2022).

Robinson (2020) find, in Indonesia and Malawi, that while elite capture exists, it is very small and thus the welfare consequences are likely small. This may explain, in part, why the programs are viewed favorably by community members: Alatas et al. (2012) find that community targeting led to higher levels of satisfaction than when the PMT was conducted. In other contexts, however, there is a risk that community-based targeting becomes politicized and clientalistic, and so only those who are connected to powerful local politicians receive benefits (Haseeb and Vyborny 2022). Given the heterogeneity in results, there are a lot of open questions in this space. For example, it is important to understand which factors predict where community targeting will work well and to understand how community targeting works in repeated settings, where community members may update on the methods over time. Moreover, it is important to better understand whether and how community targeting can augment other targeting approaches (e.g. administrative data, PMTs).

In short, there are numerous targeting methods that have been used in practice, and a growing empirical literature that is helping us understand how to make tradeoffs across them. Even beyond the tradeoffs discussed above, we believe that there are several broader issues worth considering. First, all of these methods – particularly when trying to target the extreme poor – have substantial exclusion error. How can these methods be modified to substantially eliminate this error? Can grievance mechanisms be implemented in fair ways to fix some of the challenges

in these systems?

Second, and related: income-based targeting methods have the advantage that the rules are clear, transparent, and establish a ‘right’ to a program: if your income is below \$X, you are eligible for the program. None of the targeting methods here (other than geographic targeting) are so quite so clean and transparent. Clear established rights matter: for example, *conditional* on being deemed eligible, governments can notify eligible citizens of their right to receive a program, which Banerjee et al (2018) show in a randomized trial, can lead to those eligible households receiving much more of what they were entitled to. But with complex PMT formulas, or soft information from community targeting, how can governments convey ex-ante who has a right to a program? Are there ways of simplifying eligibility to increase transparency, increasing citizens’ rights and potentially increasing legitimacy, without substantially sacrificing predictive power?

Third, many of the methods here are fundamentally static: targeting is done infrequently and based on of long-running predictors of poverty. Yet, incomes and needs can change rapidly. Figuring out how to adapt these various methods to capture an individual’s changing circumstances is an important direction for future work.

Finally, how can these methods work in concert with each other? For

example, one may want to do universal transfers in high poverty density neighborhoods, but do PMT in low poverty density ones—would this reduce error? How would people think about the fairness and politics of using different methods within the same country, even if the contexts within the country are quite different?

III. TARGETING ON MARGINAL EFFECTS

The discussion thus far mainly focused on how to identify households that have persistently low-income levels for long-run assistance, or to identify those who have an income shock. However, the discussion of community targeting hinted that there could be other objectives. Communities in Indonesia, for example, have their own distinct notion of who most needs aid. Part of this could be that communities can observe earning ability, rather than earnings, in the spirit of Mirrlees (1971), and indeed, there is some evidence in the Indonesian case that communities do adjust for earning ability. Part of it could also be that communities have different notions of marginal utility – if there is heterogeneity in utility functions, just because someone's consumption is low does not necessarily mean that their marginal utility is higher than someone whose consumption is a bit higher but has great additional uses for more resources. And, part of it could also be that, in some contexts, certain households may benefit more from particular types of programs than from others.

Recent work has explored the idea that one may want to explicitly target on

program treatment effects. For example, if you wanted to run a graduation program, who are the people who would benefit the most from the program? If you have a subsidy program for preventive health, who are the people who are most likely to take-up and follow-through with the program? If you have a subsidized training program, who will most benefit from it?

A key challenge is that, unlike poverty, which is (in principle) observable, here for any given household what their treatment effect *would be* if they received the program is an unobservable counterfactual. However, many of the same methods that we discussed above—data approaches/machine learning techniques, self-selection, and community-targeting—are now being used to try to predict outcomes. We highlight some of the existing research below, but we also note that this remains a very open area for future research.

In terms of data-driven approaches, several studies try to determine the observable characteristics that would best predict the outcome one cares about, increasingly using machine learning to estimate these effects. For example, subsidized bednets for malaria prevention often go unused. In the context of limited supplies, this could leave people who actually want to use them without any bednets, with many bednets wasted. Bhattacharya and Dupas (2012) consider who should receive subsidies for bednets, accounting for the probability of the use of bednets conditional on receipt. Haushofer et al. (2022) also consider a related

approach in thinking about who can best make use of large, one-time cash transfers (that can be used for starting a business, etc.).

Self-targeting is another approach that can be used to elicit who would benefit from the program. For example, Dupas et al. (2016) experimentally compare free chlorine distribution with a voucher where recipients have to actually go to a nearby shop to pick up the chlorine. They find the same level of chlorine use across both groups, but that the voucher program was much cheaper because only those who intended to use the chlorine went to pick it up. Understanding how different ordeal design mechanisms select on treatment outcomes, particularly when there is heterogeneity in the cost of participating, different types of need, and different valuations of benefits, is an important area for future research.

Finally, community targeting methods may also have the potential to be used to target on treatment effects. For example, Hussam, Rigol, and Roth (2022) explore whether community targeting could help identify high-growth micro-entrepreneurs for a cash grant in India. They find that it does identify the entrepreneurs with the highest returns even when there is already a rich list of observable characteristics available about these entrepreneurs. However, they caution that it was important to design the information elicitation process in way that facilitated truth-telling.

Learning how to best target on these outcomes, for different types of programs or contexts, remains a fruitful area for future research. However, it is

worth noting that there is a tradeoff: if potential outcomes are higher for those with higher incomes, targeting on potential outcomes could mean that poorer households, with lower treatment effects but higher need, could be left out. In principle, one could do a more systematic approach, in which the government maximizes a combination of need and treatment effects across a set of different people and programs.

IV. TARGETING FOR EMERGENCIES

The discussion thus far has focused on ‘normal times’ – how governments can design social protection systems to channel assistance to those who need it given everyday life. But, what about when disasters strike?

It is worth noting that several of the approaches that we discussed above can be used even in times of crisis. For example, when COVID-19 hit, the government of Togo needed to set up a system to distribute cash assistance to poor households when they had not yet created a systematic PMT-based social registry. Instead, they used mobile-phone metadata to target, which Aiken et al. (2022) report substantially reduced exclusion errors relative to geography based targeting. For refugees, Altındağ et al. (2021) show using data from Syrian refugees in Lebanon that PMTs can be applied using the basic administrative data collected by refugee agencies such as UNHCR.

However, there are other questions that come up that are specific to times of emergencies. For example, as climate-related shocks, such as floods and extreme storms, become increasingly common, an important question is whether assistance can be targeted *in advance* of the shock occurring, rather than waiting for the shock to occur. If so, they can use the assistance for preventive behaviors (e.g. moving away from the area the shock will hit, protecting assets) which can improve long-run outcomes. Pople et al. (2021) examine one such program in Bangladesh comparing households that received anticipatory transfers before a flood and those that (given funding distribution delays) did not receive any assistance during the study period. They find that those who received cash in advance were substantially less likely to go without eating during the flood, experienced lower asset loss because of the flood, and were better off three-months later. On the other hand, if climate prediction is imperfect, this approach may be more costly than an ex-post approach. This may be compounded if there are limited resources, the resources are sent to the “wrong areas,” and there is not enough funding for the areas that are actually hit by the emergency. More research is needed to understand the relative tradeoffs of anticipatory cash transfers versus more standard humanitarian aid.

Finally, it is important to note that building social protection systems – such as establishing identification cards, identifying beneficiaries, and so on – can also make assistance more effective in times of crisis. In the COVID-19 crisis, many countries, for example, built on whatever systems they had in place before hand,

and the World Bank notes that the countries that did the best in addressing poverty during the pandemic were those that had pre-existing systems in place for the targeting (i.e. id cards, unified databases, etc.) and distribution of social protection aid (World Bank 2021). And indeed, the same study of Togo discussed above notes that the mobile-phone machine-learning approach, while better than geographic targeting, did not perform as well as a regular PMT-based measure would have had it been put in place earlier (Aiken et al. 2022). Investing in the infrastructure for identifying beneficiaries in normal times can substantially improve shock-responsiveness, particularly if the system incorporates administrative data (e.g., mobile phone usage, electricity consumption, and so on) and/or on-demand approaches that can respond rapidly and pick up changes due to changing circumstances.

V. CONCLUSION

In many contexts, governments seek to target assistance to a sub-population that needs it the most. This chapter has outlined the key approaches and tradeoffs in doing so in the low information environments present in many low- and middle-income countries. While the most common approach relies on categorical targeting or on data on assets to predict incomes for targeting, increasingly countries are also

using administrative data, community-based approaches, and self-targeting approaches as well.

We note a few important areas for future research. First, as we discuss, a key goal is to increase the net benefits from a program – which is a combination of the treatment effects of a program per se with the marginal utility for a household. Approaches to better target this objective – perhaps by using self-selection in combination with prediction techniques – seem important. Second, while we note that many existing machine learning approaches have not delivered radical improvements in accuracy over and above logistic or other simple regression approaches, the continual advances in AI-type approaches suggest that there may yet be important opportunities in this area in coming years. Third, there is increased awareness that social protection systems may play an important role in response to climate-induced shocks. Understanding how to effectively adapt these systems to that context remains an important area for future research.

Author Acknowledgements: We thank our numerous coauthors and policy partners who contributed to much of the research that we discuss in this chapter. AEA-Randomized author order. The views here are those of the authors alone.

A full set of references for all Handbook chapters, including this one, can be found at this link: <https://www.dropbox.com/scl/fi/9lqs2mdrawkjdrv4m648e/References-Social-Protection-Handbook.pdf?rlkey=jt0f8kute31mhdke77aoiw99d&st=kd7l8ff1&dl=0>