CHAPTER 9

Heterogeneity in Target Populations and Locations: Reflections on the Challenges for Poverty Targeting

Rachel Sabates-Wheeler

Despite long-held aspirations of increased productivity for small-scale agriculture in Africa south of the Sahara (Lipton 1977, 2009), productivity gains in many countries of the region have failed to materialize, especially for the rural poor (Jayne et al. 2003, Ellis 2010, 2012). Positive changes have been unevenly felt. These outcomes, together with significant population growth in the poorest regions, changes in climate, and limited opportunities for off-farm employment, mean that rural populations in Africa will continue to be substantial and will still need to sustain their living primarily from agriculture for the foreseeable future (Losch, Fréguin-Gresh, and White 2012). This situation is exacerbated by a lack of opportunities for economic diversification, limited options for employment outside agriculture and the informal sector, and widespread poverty (whereby median incomes are estimated to range between US$0.50 and US$2.00 in purchasing power parity per person per day) (Losch, Fréguin-Gresh, and White 2012).

Over the past 15 years a response to this precarious situation has been to increase the coverage of social protection, not only as a way of protecting lives but also as a means of risk insurance and livelihood promotion. The extent of need as well as the limited budgets available for funding these programs means that, depending on the objective, program implementers need to decide on how to target the available resources, and to whom. If there is a reason to believe that poverty correlates with the age profile or dependency profile of a household, then program implementers might choose to target elderly people or households with a greater number of children. However, in the context of widespread poverty, many programs prefer to provide general transfers to the poorest households or those most in need.

Nevertheless, identifying the poorest and most vulnerable for selection into social programs is a perennial challenge facing program implementers and continues to be a source of lively debate in social protection design and delivery. In addition to tight budgets, other reasons for rationing (and therefore targeting) social programs include the desire to ensure that the most vulnerable are reached, to maximize the poverty-reducing impact of the program, and to attract or retain the support of key constituencies.[1] Poverty-targeted social transfers, mainly in the form of predictable and regular cash payments (but also in the form of food, assets, and vouchers) are the instrument of choice for many development partners and governments of lower-income countries for addressing predictable food insecurity and hunger (Grosh et al. 2008; Adato and Hoddinott 2010; Hulme, Hanlon, and Barrientos 2012). As reported by Honorati, Gentilini, and Yemtsov (2015), 130 low- and middle-income countries have at least one noncontributory unconditional cash transfer (UCT) program (including poverty-targeted transfers and old-age social pensions), with growth in program adoption especially high in Africa, where 40 countries out of 48 in the region now have a UCT, the number having doubled since 2010.

Despite the growing popularity of cash transfers, several studies have shown that the targeting mechanisms frequently used within these programs lead to substantial inefficiencies and can often be ineffective at enabling a program to deliver on its intended outcomes. Work by Ellis (2012) argues that in the context of deep and widespread poverty in rural Africa, poverty-targeted transfers can create significant social tensions between the "included" and the "excluded." Using a rather arbitrary targeting eligibility cutoff in the face of tight budgets and in a context in which everyone is poor calls into question the social acceptability, as well as the political attractiveness, of targeting.

There are additional reasons why poverty targeting might not deliver on the intended outcomes, even when errors of inclusion and exclusion are small, that relate to assumptions of homogeneity in heterogeneous target

1 See Devereux et al. (2017) for a full discussion of targeting rationale.

groups and misguided assumptions about the individualized use of cash transfers. This chapter aims to illustrate these challenges by drawing on in-depth knowledge of a limited number of recent and active social protection programs implemented in eastern Africa. Specifically, it focuses on three challenges inherent to poverty targeting that constrain the achievement of program objectives: (1) the difficulty of identifying the poorest from among the poor, particularly in a context in which a large proportion of households holds the view that "we are all poor here"; (2) heterogeneity in household characteristics within a target population that is assumed to be relatively homogeneous; and (3) provision of "individual/household" transfers in diverse social and cultural contexts.

The chapter is framed around these three challenges, illustrated by drawing upon four cases with which the author has substantial familiarity: the Productive Safety Net Program (PSNP) in Ethiopia; the Hunger Safety Net Program (HSNP) in Kenya; and the Vision 2020 Umurenge Program (VUP) and the Concern Worldwide Graduation Program, both in Rwanda.[2] These examples were chosen because they share commonalities in social protection provision. First, Ethiopia, Kenya, and Rwanda all have relatively well developed, nationally embedded social protection systems in place. Second, the programs reviewed are all targeted at poor households. Third, all include cash transfers as the primary intervention but are complemented by other support interventions such as training, public works, and financial services. Fourth, monitoring or evaluation studies have been undertaken for all these programs. Finally, the programs chosen are similar in design to other large-scale programs being implemented throughout eastern and southern Africa, and therefore lessons will be largely applicable and transferable. Drawing on these cases, we discuss the implications of the three challenges and conclude by offering suggestions for policy.

Targeting the Poor

The targeting challenge is how to accurately and cost-effectively identify and register households or individuals who are eligible to receive resource transfers, thereby screening out those who are defined as ineligible. Of course, the corollary of this challenge is to successfully deliver the social resource to the eligible households. Targeting is frequently an expensive and time-consuming activity, and typically trade-offs need to be made between targeting accuracy and targeting costs. Suboptimal targeting can result in large inclusion errors (whereby noneligible people are included) and exclusion errors (whereby eligible people are excluded), which represent a significant waste of scarce resources (often public) and may undermine the program's effectiveness and longer-term political support.[3]

The term *poverty* in *poverty targeting* defines the intended eligible population. The obvious challenge when poverty defines the target group is how to measure poverty and where the threshold for eligibility will be drawn (that is, what separates the poor from the nonpoor, or the ultra poor from the poor). The poverty-targeting approach requires selection criteria that successfully identify those most in need of the limited resources that governments may make available for social cash transfers. This requirement usually means that poverty-targeted transfers rely on proxy indicators of need to accomplish beneficiary selection.

In a context of high administrative capacity and data availability, verified means testing using accurate personal income data (usually from administrative and occupation-based records) would be the obvious way to target the households and people who are most in need, and this is precisely the method used in many higher-income countries. This type of testing is not often used in low- and middle-income countries due to its high demand

2 The latter is the only one of these programs implemented by a nongovernmental organization. The others are nationally owned and supported programs.

3 The need to target social transfers (as opposed to providing universal coverage) is typically justified in the context of tight budget parameters and limited political appetite for large-scale spending on long-term support for the poorest.

for data and for the administrative capacity to verify the data (Coady, Grosh, and Hoddinott 2004). Furthermore, accurate means testing of semi-subsistence rural households or households making a living in large, unregulated informal sectors is an impossibility. Therefore, other targeting mechanisms are frequently employed, such as the following:[4]

- *Categorical targeting* identifies specific demographic groups who display a higher level of poverty or risk of vulnerability. These include specific groups of children, elderly people, or people living with disabilities. Categorical targeting is a popular method, particularly in combination with other methods, such as community validation.
- *Geographic targeting* targets locations with high levels of poverty and vulnerability for support. It is a relatively straightforward method with low administrative costs but often goes hand in hand with high targeting errors.
- *Proxy means tests* (PMTs) use multiple indicators to identify the poor and vulnerable. Typically, a PMT collects data on observable characteristics of an individual or household to obtain a score that proxies the resources available at the household level. Commonly included indicators are housing materials, ownership of durable goods, and the educational attainment of household members. This method is considered to be more accurate than geographic or categorical targeting, but it is more demanding in terms of data and administrative capacity.
- *Self-selection* (for instance, for the labor requirement in public works programs) relies on the program design to ensure that only the most vulnerable and those in need benefit from the program. It does so by offering low transfer levels or making it difficult to obtain transfers. Although it is considered effective, there may be considerable social and psychosocial costs associated with self-targeting (White 2017).
- *Community-based targeting* (CBT) delegates beneficiary identification and selection to community members, often a group of community representatives or elders who use their local knowledge to inform decisions about who is to benefit from the cash transfers. Because most of these community mechanisms work on a voluntary basis, administrative costs are low. Results are mixed in terms of accuracy, and the potential for incurring social costs is relatively high.

Overall, the empirical evidence confirms that targeting does benefit the poor. In a review of benefit incidence of more than 100 programs, Coady, Grosh, and Hoddinott found that targeted programs provide "approximately 25 percent more resources to the poor than would random allocations" (2004, 2). However, they also found that the way targeting is conducted matters for both the effectiveness and the efficiency of outcomes. Coady, Grosh, and Hoddinott (2004) found that when ranked in terms of their ability to reach the poor, targeting mechanisms performed differently. Self-targeting on public works was the most accurate in terms of identifying the poor, followed by geographic targeting and means testing. After these, proxy means testing, CBT, and categorical targeting achieved small pro-poor benefits. Categorical targeting of older people performed worst. In the years since their study, many more poverty-targeted programs have been implemented and some evaluated. The next section reviews some more recent evidence that corroborates the work of Coady, Grosh, and Hoddinott (2004).[5]

4 The description of targeting mechanisms and methods draws on Sabates-Wheeler, Hurrell, and Devereux (2015), Coady, Grosh, and Hoddinott (2004), and Devereux et al. (2017).

5 It also includes a short review on the findings on mechanisms in Devereux et al. (2017).

The Challenges of Identifying the Poorest from among the Poor

Given the need to target and the range of mechanisms available, why is it that targeting the poorest of the poor is such a difficult, and sometimes elusive, goal? As convincingly demonstrated by Ellis (2012) in his review of the implications of poverty-targeted cash transfers for social, economic, and political difference in rural Malawi, Zambia, and Ethiopia, the process of differentiating the target group of the poorest from other, almost as poor, households is an incredibly difficult, at times impossible, task. In many rural contexts, particularly in Africa south of the Sahara, the income and consumption distribution of poor households is so flat that the task of deciding the cutoff threshold for program eligibility means that targeting can be both socially and politically divisive. Consider that a cash transfer targeted to the poorest 5–10 percent of a community might well have poverty-reducing impacts for them and change their poverty status (that is, their place in the local income distribution) relative to the less poor households comprising the 10th to 20th percentiles of the income distribution. Ellis (2012) discusses the likely resentment from poor but noneligible households in these situations.

Empirical work by Sabates-Wheeler, Yates, Wylde and Gatsinzi (2015) on poverty targeting in Rwanda shows that, even putting social and political tensions aside, the measurement and identification of proxy indicators that correlate with extreme poverty is technically no small feat. The VUP began in 2008 and builds on the same basic design as Ethiopia's PSNP. The VUP aims to reduce extreme poverty in Rwanda through providing regular and predictable cash transfers, either as payment for short-term public work or through unconditional transfers to labor-constrained households. It also offers financial assistance, via low-interest loans to program beneficiaries, for investing in productive enterprises. Households eligible for the VUP must be extremely poor.

In the first few years of the VUP, cash transfers were targeted using a traditional community poverty mapping process, known as *Ubudehe*. As a concept, *Ubudehe* has its roots in the traditional practice of working together to solve problems. As a social categorization process, it first took place in 2001. Communities categorize themselves into one of six well-being groups and use this analysis to discuss local development priorities. Although the categorization got off to a promising start in the early days, with high levels of participation, over time the process has become increasingly administrative. Villages no longer produce social maps; rather, data collected by the *Ubudehe* committee is entered manually into a table and processed by computer. Moreover, research findings from a study supported by the Rwandan Civil Society Platform (Sentama 2013) found that of 250 households interviewed, 83.6 percent stated that they were very dissatisfied with the category in which they had been placed. Only 6 percent indicated that they were satisfied or very satisfied with their categorization.

A quantitative analysis verified the unsatisfactory targeting results by correlating consumption poverty and *Ubudehe* status (Sabates-Wheeler, Yates, Wylde and Gatsinzi 2015). It showed that the richest quintile was equally as likely to be allocated the lowest *Ubudehe* status as the poorest (indeed, the likelihood was fairly consistent across all consumption quintiles). The analysis suggested that 62 percent of the extreme poor would be excluded from the national poverty-targeted cash transfer program—the VUP—because they were not in the lowest *Ubudehe* categories.

Due to the increasingly nontransparent process whereby households were being placed into *Ubudehe* categories, along with the heightened incentive for households to belong to a lower category (because of the number of social benefits attached to categories 1 and 2), the government of Rwanda decided, in 2014/2015, to introduce "objective" indicators as a complement to the community approach at the village level. A first step in this direction was to correlate a range of objective poverty indicators, such as housing quality, electricity access, type of latrine, level of education, livestock ownership,

TABLE 9.1—LIGHTING SOURCE AND TOILET TYPE BY CONSUMPTION QUINTILE, PERCENTAGE OF RURAL HOUSEHOLDS, RWANDA

Quintile	Asset				
	Electricity	Oil lamp	Battery	Pit latrine (improved)	Flush toilet
Lowest	2	10	21	23	1
Second	1	13	23	20	6
Third	6	18	22	23	6
Fourth	15	27	22	23	10
Highest	76	32	12	11	77
Total	100	100	100	100	100

Source: Sabates-Wheeler, Yates, Wylde, and Gatsinzi (2015).

and so on, with consumption poverty. The findings demonstrated that the variables, on their own, are not strong predictors of poverty. For instance, Table 9.1 shows that although electricity use for lighting or flush toilet ownership is helpful for identifying the richest, these variables do not allow for differentiation among the poor because there is insufficient variation across quintiles to be useful.

The indicators shown in the table either only separate the richest quintile from the rest or are too evenly distributed across quintiles to be relevant. Unfortunately, there are simply no easily identifiable variables (or combinations of variables) that can clearly and accurately distinguish the poorest from those who are better off. Even the construction of simple asset indexes was unable to help with poverty targeting because so many poor households were scored identically. A frequently overlooked reason for this difficulty is the fact that poverty is so widespread among the population, with very little variation in poverty levels at the lower end of the distribution, that any "simple" form of targeting, amid high levels of rationing, is bound to lead to large errors of inclusion and exclusion.

A similar poverty profile threw up comparable challenges for targeting the poorest in northern Kenya's HSNP, a UCT program (Merttens et al. 2013; Merttens et al. 2017b). Launched in 2009, the HSNP aims to reduce extreme poverty by delivering regular cash transfers to some 300,000 poor and vulnerable individuals in four districts—Mandera, Marsabit, Turkana, and Wajir. Targeting in this context presents considerable challenges, not just logistical but also in terms of defining an appropriate and identifiable target population: appropriate in terms of being consistent with the program's objective to reduce extreme poverty and identifiable in terms of exhibiting specific observable and verifiable characteristics.

Impact evaluations and targeting effectiveness evaluations by Sabates-Wheeler, Hurrell, and Devereux (2015), Silva-Leander and Merttens (2016), and Merttens et al. (2017a, 2017b) showed that the extent and uniformity of poverty in areas targeted by the HSNP (Figure 9.1) makes it very

FIGURE 9.1—DISTRIBUTION OF HSNP DISTRICTS' POPULATION BY CONSUMPTION, KENYA

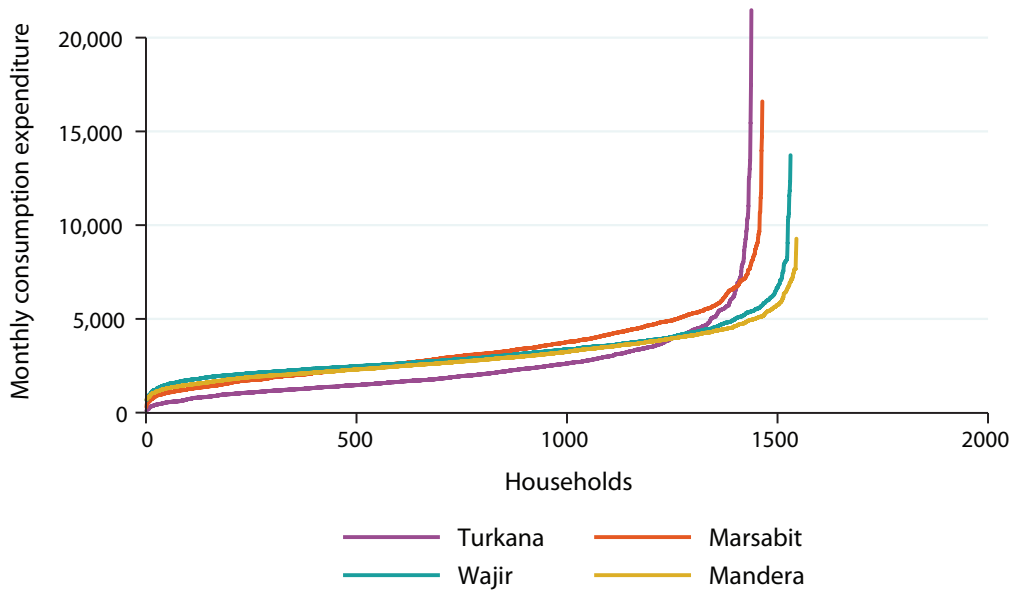

Source: HSNP2 Impact evaluation survey 2016.

difficult for the program to accurately identify the poorest households. The overall finding was that HSNP Phase 1 targeting was mildly pro-poor on aggregate, with HSNP beneficiaries 13 percentage points more likely to fall into the bottom 51 percent of the consumption distribution (the program coverage rate in evaluation areas). Exclusion and inclusion errors in the HSNP were found to be very high in both Phase 1 and Phase 2—roughly similar to what would have been achieved under a random targeting rule—and targeted beneficiaries were not considerably worse off than nonbeneficiaries in terms of monetary poverty (Silva-Leander and Merttens 2016).

During the first phase of the HSNP program (2010–2014), three different poverty-targeting methods were trialed to determine which would be most effective in identifying the poorest households. Two of the mechanisms were variants on categorical targeting (the dependency ratio and the presence of older persons), and the third was CBT. An analysis of the relative accuracy and effectiveness of these three mechanisms by Sabates-Wheeler, Hurrell, and Devereux (2015) showed that in terms of the comparative accuracy of the indicators as proxies for poverty, CBT came out as performing the best. Furthermore, CBT was more likely to be perceived as a fair process by households and communities. Old-age targeting was the least effective of the mechanisms trialed in northern Kenya simply because old age was not strongly associated with poverty. So, although the program might have performed well in identifying older people for an old-age targeting criterion (96 percent of beneficiaries fulfilled this eligibility criterion), the fact that older people in northern Kenya are not generally poor means that this was not a good poverty proxy.

Of the three targeting mechanisms being piloted, Sabates-Wheeler, Hurrell, and Devereux (2015) judged CBT to be the most pro-poor (with beneficiaries 17 percentage points more likely to fall into the bottom 51 percent of consumption), but its performance was dependent on context; in some places, CBT was regressive. Due to CBT's reliance on relative rankings, rather than objective or absolute poverty criteria, it is not sensitive to variations in poverty levels across localities and regions. In northern Kenya, poverty and food insecurity vary substantially across districts. Similarly, in Rwanda, poverty varies markedly between rural and urban areas. Yet localized relative rankings mean that CBT will frequently identify the same proportion of poor households in all districts, regardless of the district's overall relative wealth or poverty.

Sabates-Wheeler, Hurrell, and Devereux (2015) identified other challenges to using CBT, including the possibility of non-deliberate exclusion of certain groups and households in the targeting process—perhaps due to disabilities associated with old age, which can lead to elderly people being absent from critical targeting meetings. A second challenge was the deliberate exclusion of newcomer groups or poor individuals, either by the community process or due to elite capture of the program benefits. Such forms of deliberate exclusion can be seen in the targeting study of the HSNP (Sabates-Wheeler, Hurrell, and Devereux 2015) and in recent evaluations of the PSNP in Ethiopia (Berhane et al. 2016). Without external, independent verification, elite capture of benefits can plague social protection programs, particularly those using CBT. Our findings resonate with those of Coady, Grosh, and Hoddinott (2004), who found that CBT achieved the most variable results of all mechanisms.

In an attempt to improve the accuracy of targeting for Kenya's HSNP program at the end of Phase 1, the analysis by Sabates-Wheeler, Hurrell, and Devereux (2015) simulated targeting outcomes using a simple PMT methodology and compared them with the simulated results of three other methods. The simulation exercise showed that a PMT would have performed better than single categorical indicators such as the old age and dependency ratio targeting methodologies. Based on Phase 1 recommendations, a new targeting system was developed that used a PMT and community-based validation. A Phase 2 targeting analysis by Silva-Leander and Merttens (2016) showed no major improvement in the accuracy of

targeting the poorest, again due to the difficulty of targeting in northern Kenya—in particular, due to the flat consumption distribution and the wide extent of poverty. Silva-Leander and Merttens offered the following assessment:

> The extent and uniformity of poverty in areas targeted by HSNP2 made it difficult to accurately identify the poorest households using either the PMT or CBT targeting mechanisms. Exclusion and inclusion errors are very high, and targeted beneficiaries are not considerably worse off than nonbeneficiaries in terms of monetary poverty. Importantly, the targeting performance appears to be very close to what would have been achieved if a random targeting rule had been used. This raises serious questions about the cost-effectiveness of the current targeting mechanism (2016, 6).

The study concluded that "geographic targeting is the most efficient way to target the poor en masse, but pure geographic targeting (i.e., deriving beneficiary quotas based on geographic poverty rates alone) has proved not to be politically feasible" (Silva-Leander and Merttens 2016, 4). Within-district geographic poverty targeting, seen as more palatable, is now being piloted.

In summary, as shown in the cases of the VUP and the HSNP, accurate poverty targeting in contexts of very high poverty levels proves extremely difficult to implement. In Rwanda, asset indicators are unable to effectively distinguish the poorest quintiles from the less poor. The same phenomenon is illustrated by the consumption expenditure indicator in Kenya—the flat distribution means that it is very difficult to identify a cutoff threshold for targeting the poorest.

Although there appears to be a general consensus in the global reviews that a combination of targeting mechanisms that includes objectively verifiable indicators (such as simple PMT scorecards) as well as community-based validation would work well, the foregoing review of recent targeting efforts clearly shows the importance of understanding the contextual nature of poverty before choosing an appropriate targeting mechanism or proxy(ies) for poverty.

Targeting Assumptions about the Homogeneity of Poor Households Undermine Program Objectives

Poverty targeting at scale and where means testing is not possible typically requires that assumptions be made about the correlation of location or identifiable characteristics, such as age, gender, or disability, with poverty. For instance, dependency ratio targeting is usually based on evidence or assumptions that households with high dependency ratios are likely to be poorer. The same can be said for targeting older persons, although the evidence presented above from the HSNP in Kenya showed that old age was not a good predicator of poverty in that context. Even in cases in which a proxy, such as a disability or asset ownership, does correlate with income or consumption poverty, the blueprint implementation of a social protection program for the identified population still does not account for heterogeneity within that population, and this shortcoming can strongly undermine the achievement of program objectives. Put simply, assumptions of homogeneity across program-eligible households—that is, that they are all equally disadvantaged or vulnerable—lead to unnecessary inefficiencies and ineffectiveness in social protection programs. This section considers how these assumptions in targeting lead to suboptimal and unintended impacts.

Using panel survey data from a model graduation program implemented by a nongovernmental organization (NGO) in Rwanda, Sabates-Wheeler, Sabates Aysa, and Devereux (2018) tested some hypotheses around enablers and constrainers that predict different livelihood pathways.

TABLE 9.2—SUMMARY OF DETERMINANTS OF LIVELIHOOD TRAJECTORIES IN DIFFERENT OUTCOME AREAS FOR PARTICIPANTS OF A GRADUATION PROGRAM IN RWANDA

Determinant	Female-headed	# of adults	Initial assets	Dependency ratio	Initial livestock	Initial land	Co-op member	Location
INDICATOR: Trajectory								
BASIC NEEDS: Improvement versus crashing out	NEGATIVE	POSITIVE	ns	NEGATIVE	ns	ns	ns	ns
LIVESTOCK: Improvement versus crashing out	ns	POSITIVE	ns	ns	POSITIVE	POSITIVE	POSITIVE	POSITIVE
ASSETS: Improvement versus crashing out	ns	POSITIVE	POSITIVE	NEGATIVE	ns	POSITIVE	ns	POSITIVE

Source: Summary of data in Sabates-Wheeler, Sabates Aysa, and Devereux (2018).
Note: "ns" indicates that this variable was not statistically significant in explaining the trajectory of change.

They found that only 28 percent of beneficiaries managed to sustain program benefits two years postprogram across all three outcome indicators. The program provision was not tailored appropriately to households' needs, the authors indicated, because heterogeneity in program participants' households and circumstances was typically not taken into account during targeting, implementation, or beneficiary removal from the social protection program. In terms of intended program outcomes, overlooking this heterogeneity in household and context characteristics led to muted effects.

Sabates-Wheeler, Sabates Aysa, and Devereux (2018) identified four livelihood trajectories for program participants—improvement, decline, crashing out,[6] and late improvement—and econometrically evaluated the comparative performance of households on these different trajectories in relation to three outcome areas: food security and basic needs, livestock holdings, and household assets. The analysis (summarized in Table 9.2) showed clearly that some household characteristics, such as gender of household head, labor availability, and initial livestock/asset ownership, affect the household's trajectory of change. Table 9.2 shows that the odds of being on an improvement trajectory (measured in terms of food security and basic needs) are lower for female-headed households than for male-headed households. The raw results indicated that the probability of a female-headed household's being on an improvement trajectory was significant and approximately half that of a male-headed household. The amount of effective labor (measured by the number of adults in the household) was positive and significant for the improvers in comparison with those crashing out across all outcome indicators. Specifically, the results showed that for an increase in labor availability of one adult, the odds of being on an improvement trajectory relative to crashing out were 1.70 times for the basic needs outcome, 1.97 times for the asset outcome, and 1.94 times for the livestock outcome. Furthermore, the findings showed location to be a critical factor in enabling households to effectively use any social protection to their advantage. The location dummy variable proxied for a range of factors that were not specifically identified in the survey data, such as limited market linkages, worse service delivery, and physical remoteness. Sabates-Wheeler, Sabates Aysa, and Devereux (2018) concluded that "graduation" from social protection programs requires a full understanding of the heterogeneity of

6 Crashing out refers to those who either never rose above the targeting threshold (during the transfer) and have declined in welfare indicator since leaving, and those who improved but have, since the end of the transfer, declined below baseline threshold levels.

beneficiaries being targeted in different programs as well as the context in which people live and work.

Similar conclusions can be drawn from evaluations of Kenya's HSNP, which illustrate how program impacts are felt differentially by different population groups. In addition to estimating the overall average program impact, which was shown to be mildly pro-poor, the Phase 1 impact evaluation (Merttens et al. 2013) assessed the degree to which program impact varied across different types of households.[7, 8] This impact heterogeneity analysis assessed the variation in program impact across a number of dimensions:[9]

1. Consumption expenditure: Is program impact stronger for relatively poorer households?
2. Household size: Because the transfer value is not indexed to household size, the effective per capita value of the transfer is larger for smaller households; is the program impact therefore stronger for smaller HSNP households?[10]
3. Mobility status: Is the program having a differential impact on HSNP households depending on whether they are partially or fully mobile or fully settled? The trend observed at Follow-Up 1 was that mobile households tended to show increased impact.

The trends observed in Phase 1 showed that although the impacts after 1 year on consumption and poverty were not significant overall, HSNP households that were initially poorer, mobile, or smaller did experience an impact.[11] Heterogeneity analysis revealed that the impact on dietary diversity was most marked for households that were poorer, smaller, or mobile, and, again, there was a positive impact on dietary diversity for relatively poorer households at Follow-Up 2. As with consumption expenditure, the results showed an increased impact on food expenditure for poorer and smaller households. After two years, the program was found to have a significant impact on consumption expenditure and poverty, with HSNP households some 10 percentage points less likely to fall into the bottom national income decile. In addition, a larger and significant impact on poorer and smaller households was found. In other words, the impact on poverty was being driven by HSNP households that were relatively poorer or smaller.[12]

Basing programs on the implications of the above findings is not easy, inasmuch as it would be prohibitively costly and administratively impossible to tailor support to the needs of each and every household. Yet, where there are obvious patterns, such as in regard to labor availability, initial asset profiles, and gender of household head, different groups can be targeted to receive adapted packages of support for variable time frames. Only with some thought given to appropriate support for specific groups will a program stand some chance of achieving resilience-building program objectives.

---

7 The Phase 1 impact evaluation data were collected over the course of three rounds comprising a baseline round (August 2009 – November 2010), a follow-up 1 round (November 2010 – November 2011), and follow-up 2 round (February 2012 – November 2012).

8 For the Phase 1 impact analysis, there were 20 treatment and 20 comparison (control) sublocations, with 1,224 HSNP households (the treatment group) and 1,212 control-group households across four districts of northern Kenya from 2009 to 2012 (Merttens et al. 2013).

9 Variations in impact between targeting mechanisms were analyzed at Follow-Up 1—one year later—but did not reveal any systematic differences across the targeting mechanisms. This finding is not surprising, because the targeting report shows a large degree of overlap in the characteristics of beneficiaries targeted by age, dependency ratio, and CBT, so it makes sense that the HSNP impact did not vary by targeting mechanism.

10 Households were classed as small if they had fewer members than the median household size found at baseline.

11 Due to sample attrition, which particularly affected mobile households, and the reduction in overall sample size, this analysis could not be performed at Follow-Up 2.

12 A heterogeneity analysis was also performed for households that had received a larger cumulative per capita value of transfer. The impact was significant and positive for these households also.

The Difficulties of Targeting in Diverse Social and Cultural Contexts

Despite the variety of mechanisms available for targeting, in almost all social protection programs, the eligible participant is identified as an individual (or carer for an individual if the eligible person is a child or has a severe disability) or a household (although even when a household is the eligible unit, an individual[s] must be named as the recipient of the transfer). The "individualization" of entitlement delivery and receipt can create challenges in contexts of social norms and living arrangements predicated on sharing and reciprocal support; when the named recipient is highly mobile, such as in the case of pastoralists, internally displaced populations, or migrants; or when the person receiving the transfer is not the best placed to use the transfer for the purposes for which it is intended. This section looks at each of these pitfalls in turn.

Cultures of Sharing and Resultant Transfer Dilution

The notion of setting a transfer value to cover a consumption deficit or food gap, or to provide small asset packages at the individual or household level, can be problematic in some contexts. Take, for instance, the 120 million pastoralists and agro-pastoralists worldwide, of whom 50 million live in Africa south of the Sahara. These groups dwell and move in contexts in which social networks and relationship-based access to social provision are the predominant forms of distribution. "Network-based affiliation and distribution provide important functions under high levels of uncertainty, such as subsistence assurance and risk mitigation" (Sabates-Wheeler, Lind, and Hoddinott 2013, 2). These social norms can undermine intended program effects related to building resilient livelihoods if the transfer amount intended for building the resilience of one person or household is shared across multiple individuals or households.

The extension of the PSNP, in Ethiopia, from the highlands to the lowlands illustrates this very well. When the PSNP was first introduced in 2005, the program was confined to the highland regions, made up primarily of poor rural households that were sedentary in terms of residence but working across the rural-urban space. The PSNP, a poverty-targeted program for food-insecure households in chronically food-insecure areas, using geographic indicators and community validation for targeting, was designed for this population. Districts can choose food, cash, or a mixture of both as a transfer modality, although an increasing preference for cash over the years has meant that the majority of PSNP beneficiaries in the highlands are receiving cash. The objective of the program is to enable households to escape food insecurity, accumulate assets, and over time, build sustainable livelihoods that will allow them to move off the program ("graduate").

In 2008, the PSNP was scaled out to cover the lowland areas, which are characterized by pastoral and agro-pastoral livelihoods. Only food, not cash, was provided in these regions, mainly due to the difficulty of setting up financial delivery services as well as the lack of food markets in some areas. The geography and culture in the lowlands is very different from that of the highlands. Drawing on extensive mixed-methods survey data, Sabates-Wheeler, Lind, and Hoddinott (2013) described how practices of sharing within mutual support networks and the role of informal authority structures in the targeting process moderated and muted intended PSNP impacts. Intense pressure to give support within horizontal networks of the very poor resulted in the dilution of the food transfers provided to any one household. The analysis showed that payment levels were low in the majority of the nine *woredas* (districts) studied. In six of the *woredas*, transfers were 10 kg or less per capita (whereas the per capita transfer entitlement should have been 15 kg). As coverage rates rose, per capita grain payments fell (in one woreda to as low as 2 kg per person), reflecting the tension of ensuring that individuals and households got the "right" amount as set by the program and the pressure to provide food support to as many community dwellers as possible. In fact, in four *woredas*, coverage exceeded

70 percent, and in one, coverage was almost universal. Sabates-Wheeler, Lind, and Hoddinott (2013) argued that this situation was reflective not only of sharing cultures but also of the pressure on local-level targeting staff to include as many of their community members as possible in the PSNP.

Although the program could be argued to provide an important safety net against hunger and severe deprivation, this transfer dilution undermined the program objectives for resilience building. In other words, the sharing of transfers is a critical cultural practice for ensuring survival for many (a fundamental objective of a safety net), but it is a problem for the program vision of asset and resilience building.

In communities where sharing and reciprocity are endemic, individualized targeting of transfers, however "perfect," is an inadequate tool for promoting significant economic transformation of livelihoods. The intended per capita transfer value is split across multiple households and individuals, thus muting any proposed impacts (calculated on the assumption that an individual receives a specific amount). In these contexts, where communities display strong horizontal bonding, local universal support tailored to community needs is likely to be more effective in building economic resilience than individualized, targeted support. More recent evidence from successive rounds of PSNP evaluation data over a six-year period (2010–2016) shows that the targeting of the PSNP in the lowlands has failed to improve, with the better off as likely to be PSNP beneficiaries as the poorest. The analysis shows that in addition to pressure to dilute transfers, suboptimal targeting results reflect local notions of "fairness" in these areas—that everyone should benefit regardless of wealth (Lind et al. 2018).

High levels of vulnerability and widespread poverty in pastoral areas imply that there is a great need to provide social protection in these areas. However, it appears that program design and delivery suited to the agrarian-based livelihoods of the highlands are not appropriate for the lowlands. Other forms of social protection that find their rationale in community- or clan-based provision may be more appropriate.

Targeting People on the Move

Another challenge is the difficulty of targeting mobile populations—such as pastoralists, migrants, and internally displaced persons. In the lowlands of the Horn of Africa, livelihoods have historically been based on pastoralism and transhumance. In this livelihood context, the culture of sharing is intricately related to the mobility of the adult male household members with the livestock over large tracts of land. These long periods of male absence from the larger family, combined with smaller, polygamous household units across multiple locations, means that resource distribution and sharing takes on necessarily complex forms. Although livelihoods in the lowlands are much more mixed these days—including agro-pastoralism, pastoralism, farming—the PSNP experience of targeting food transfers to mobile and semi-mobile groups has illustrated two seemingly contradictory challenges: first, the problem of locating mobile groups for targeting processes and provision of food or cash, and second, the interference with herding patterns and transhumance caused by the requirement to be present for the targeting process and at collection points.

Setting up the PSNP in the lowland areas required that all vulnerable households be registered and that the designated household representative (typically the male head of household) be available to come to a "local" collection point on a bimonthly basis to pick up the food transfer. Difficulties in registering households emerged due the absence of household heads (given that heads are usually out with the herds), the movement of households, and the high levels of insecurity and conflict in a number of the districts. Furthermore, if households prioritize receipt of the food transfer, the requirement that household heads be present at a specific collection point severely disrupts the usual mobility and herd pasturing patterns. On the other hand, some mobile groups might deliberately choose to opt out of the program because their livelihoods are not conducive to sedentarization (Scott 2009). Moreover, any survey evaluation that requires interviews with

households is very complicated and costly, due in large part to the process of locating the migrants.

This relatively straightforward case illustrates the large challenges that human mobility/migration poses for the design of social protection. The case of pastoralists is an example of a voluntary internally mobile group. The problems this group faces are exacerbated for the internationally mobile and forcibly displaced, a specific challenge being the lack of coordination and harmonization between different delivering agencies. Cash transfers frequently use some form of software or digital transfer system to distribute cash. They often use a combination of software systems, phones, and vouchers or cash. However, these cash transfer systems are mainly bespoke, disconnected, and not interoperable. This situation is exacerbated when people cross national administrative boundaries (as in the case of displaced populations or migrants) or international borders, particularly in the event of humanitarian and protracted crises. Such movement can mean that some families receive several entitlements, each through a different modality (cash or voucher) and transfer system (mobile, card, software, or other), and from a different agency. Duplication of this sort exists in several protracted crises where harmonization and common platforms have been slow to emerge. Innovations in cash transfer systems for populations on the move are being piloted in a number of countries (Hagen-Zanker, Ulrichs, and Holmes 2018 describes the experience of cash transfers for refugees in Jordan). More work needs to be conducted in this area.

Conclusions

Over the last 10 to 15 years social protection has been heralded as an answer for both the protection of lives and the promotion of livelihoods—that is, as an answer to food security and as a development paradigm for supporting economic growth through building livelihood resilience. Due to substantial need and limited funding, the majority of social protection for the poorest relies on targeting. Therefore, accurate targeting as a form of rationing becomes a critical element of both food security and livelihood support for the poorest.

Drawing on recent work, this chapter has highlighted some real targeting challenges facing social protection programs. These challenges, which are embedded in the respective economic, social, and political contexts, are as important for the predicted success of a program as the technical design and implementation infrastructure. Specifically, this chapter has illustrated the difficulties of identifying the poorest from the poor, the problems for envisaged program impact when heterogeneity of the target population is not taken into account, and the problems for program objectives when endemic sharing is the norm or when populations are mobile. These are some challenges that present themselves in the context of four programs in three different East African countries. However, they are also general challenges facing similar programs (many of which have emerged throughout Africa) in similar contexts throughout the continent (such as the lowland, pastoralist regions of the Horn of Africa, or where populations are on the move or displaced).

This chapter has shown that targeting the poorest works badly where income or asset distributions are flat, making it difficult to distinguish the poorest from the poor. In these situations, a combination of targeting methods appears to work best, such as an objective PMT combined with CBT. Although many development partners and national governments opt for a combination of proxy means testing and community-based validation or targeting due to excessive concern over free riding and inclusion problems, a strong case can often be made for progressive, blanket geographic coverage of entire communities. This approach would substantially reduce the costs of deciding which combination of targeting mechanisms will work best, if at all; minimize exclusion errors; reduce the social tensions created when the poorest of the poor are, overnight, catapulted to income levels higher than those of the moderately poor; and be a more ethical solution in the context of local development. Under a progressive political

agenda, budgetary commitments could follow the geographic expansion of the program across the country.

In fact, this review points to certain factors characterizing intended target populations and their contexts that call into question the justification for targeting at all. As discussed by Devereux (2016), critics of targeting point out that targeting the poor is frequently based on an arbitrary poverty line, whereas (as shown above) the difference in income between those who are identified as eligible and those who are ineligible is negligible—if there is any difference at all. Furthermore, arguments that poverty targeting can create social tension in poor communities lend weight to the proposition that in contexts of widespread poverty, poverty targeting is likely to be very difficult, as well as socially and politically challenging. A more effective way to support the poor is likely to be universal targeting at a localized level, where poverty and location are highly correlated. However, the obvious constraint to the uptake of universal targeting as a policy suggestion is the political milieu of a specific country and culture. For instance, currently in Ethiopia, approximately 8 million people are served by the PSNP, yet there are a further 8 million in need of yearly support through humanitarian emergency appeals. A large proportion of this latter group would be eligible for the PSNP, yet in the context of economic growth, prior political commitments to vulnerability reduction, and budget forecasts, the government will not consider increasing the PSNP caseload.

When universal targeting is not up for discussion (as illustrated in the point above), some recommendations for social protection programming fall out of this brief review. First, for *targeting* to be effective—in the sense that it supports and facilitates program objectives—attention to context, culture, and population characteristics will be critical. Second, *support* delivered through the program must be appropriate and sensitive to the different contexts and livelihoods. Third, *delivery* should be fitting to context. So, for cultures in which sharing is the norm, delivery may need to be provided to clans or communities rather than individual households. Or, where people are on the move, delivery (registration of target populations and designation of payment points) will need to adapt to mobility patterns and locations. E-payments provide obvious delivery advantages for people on the move.

Blueprint roll-out of features in social protection designs frequently includes specific time limits on how long an eligible household can stay on the program—often one to three years for NGO-implemented projects. Another feature is identical packages of support at preidentified times during program participation—for instance, financial literacy training, coaching, savings, and asset transfers. A third feature is identical provision regardless of the location and characteristics of participating households. However, this review shows that heterogeneity in household type, in location, and in population group means that a one-size-fits-all intervention is unlikely to work, especially in regard to targeting eligible households as well as identifying households ready to exit or graduate from a program. Assumptions about similarities within a target group can be misplaced, leading to inappropriate provision for some households and premature graduation for others. Different households need diverse types of support for different lengths of time. A "leave-no-one-behind" agenda requires that we seek to coordinate and deliver the appropriate combination of interventions to different population groups in different contexts.