Best Practices for Targeting Social Benefits to Ensure Food Security

By: Carly Trachtman

FSP Best Practices 07
FEBRUARY 2025

A best practice is a method or technique that has been generally accepted as superior to any alternatives because it produces results that are superior to those achieved by other means or because it has become a standard way of doing things. This document is one of a series of reports from the Food Security Portal on best practices for emerging topics in agriculture and food security policy.

Introduction

Social protection programs, which provide beneficiaries with cash, food, or other in-kind transfers, comprise a key component of the national food security strategies of many countries, including Ethiopia, Malawi, Uganda, and Yemen (FAOLEX 2025). Such programs are used to help households smooth consumption in the face of adverse shocks and to help chronically poor populations maintain a base standard of well-being. A critical question for policymakers designing these social protection programs is which populations should be targeted by the program: in other words, which households should be chosen to receive benefits. Answering this question involves carefully balancing program objectives, political considerations, resource availability, and logistical concerns. In this brief, we discuss best practices for the targeting of social programs, highlighting the key policy decisions involved and identifying factors to consider when making these decisions.

Targeted vs. Universal Transfers

The first question a program implementer needs to consider is whether to target benefits at all. In lieu of targeting a specific subset of the population, benefits could instead be made universal. However, if the program has a fixed budget, there is a trade-off between the number of beneficiaries that can be reached and the transfer amount that can be provided to each beneficiary. Transfers that are too small may not be particularly effective; for example, they may not provide enough additional resources for food-insecure households to purchase the staples they need to survive. Even if a program does not have a fixed budget and could maintain a larger transfer size while adding more beneficiaries, there is always an opportunity cost of using additional funds for social protection that could instead go toward another program or initiative. Moreover, providing additional resources to households that do not need them will likely not contribute to the goal of increasing food security.

However, in some cases, there may be a compelling argument for providing universal benefits rather than targeted ones. For instance, programs may be more politically palatable, and hence more likely to be voted for and instituted, if they serve more people. Additionally, the targeting process itself can consume both time and resources. If the resources necessary to target beneficiaries are more than the additional resources needed to provide universal benefits, targeting may not make financial sense. Further, for programs in which benefits are meant to help households cope with a significant shock, such as a widespread harvest failure or natural disaster, the priority may be to get benefits to affected households as soon as possible. In such cases, universal or near-universal transfers may make the most logistical sense.

In practice, most social assistance programs have at least some targeting, even in the case of emergency relief programs. For instance, during the COVID-19 pandemic, out of the 509 new cash transfer programs introduced globally to help households cope with the crisis, only 12 offered untargeted, universal benefits (Gentilini et al. 2021). Unless there are circumstances that make targeting prohibitively expensive, which could be the case in some fragile, conflict-affected settings or if a program will not be able to be implemented unless benefits are universal, it is generally a best practice to have at least some degree of benefit targeting.

Defining Targeting Objectives

Once a program implementer decides that targeting beneficiaries is appropriate, a natural next question is who should be targeted. Answering this question involves specifying a targeting objective, which defines the population the program aims to reach. Note that targeting objectives are not the same as program eligibility criteria, which describe who is allowed to receive benefits in practice. Targeting objectives are generally influenced by program objectives; for instance, a program seeking to improve food security might aim to target the poorest or most food-insecure households.

Defining a Welfare Metric

Targeting objectives must be well-defined, as notions like being the "worst off" or "most in need" can mean different things to different people and in different contexts. Specifying a measurable welfare metric in the statement of the objective can both resolve confusion about who the program seeks to reach and allow a program to easily assess its performance against its objective. Programs interested in reaching the poorest or neediest households may in theory want to target households based on per capita income, given that income is the means to secure food. For example, programs like *Progresa* in Mexico and *Bolsa Familia* in Brazil target households with low per capita income (Dávila Lárraga 2016, Soares et al. 2010). In practice, however, this can be challenging in many settings in which household income is not easily documented because households earn their income from small-scale agriculture or other informal activities. In such settings, per capita consumption measured through a standard household survey is a preferable substitute for income (Deaton 2003). Many programs, like Livelihoods Empowerment Against Poverty (LEAP) in Ghana and *Program Sembako* in Indonesia, target beneficiaries based on per capita consumption levels (de Groot 2016, Banerjee et al. 2021).

Neither per capita income nor per capita consumption may fully capture all notions of welfare or food security that could be useful in targeting a program, however. Thus, policymakers could consider using either more holistic or more specific metrics, based on program objectives. For instance, Ethiopia's Productive Safety Net (PSNP) program specifically targets households that have experienced continuous food shortages over the preceding three years, though this metric can be challenging to observe and verify in practice (Sharp et al. 2006).

Defining a Poverty Line

Once a welfare metric is identified, program implementers will also want to think about program coverage: that is, how many people will be included. Programs may consider having eligibility criteria based on either an absolute or a relative poverty line. An absolute poverty line specifies a value of the welfare metric, such as a per capita income level of $2/ day, and deems everyone who falls below the threshold as eligible for the benefit. A relative poverty line instead looks to benefit a set percentage of the population with the lowest income levels. The choice of a relative versus an absolute poverty line may depend on the philosophy of the program. If the ethos behind the program is that everyone living below a certain standard needs the resources to raise their welfare to the minimum standard of living, then targeting based on an absolute poverty line is most sensible. Notably, as individuals' welfare shifts over time, the number of individuals eligible for the program may also change. Having a good understanding of such potential shifts is critical for budgetary planning. If the program has a fixed pool of resources, on the other hand, and hopes to reach those who need benefits most, it makes more sense to have a relative poverty line targeting a fixed number of beneficiaries or percentage of the population.

If the fraction of individuals living under the poverty line stays relatively fixed over time, then the practical distinction between an absolute poverty line and a relative poverty line is limited. However, it may still affect the way targeting accuracy is evaluated. Targeting accuracy is often measured in terms of inclusion errors (when individuals who are not intended to receive benefits get them anyway) and exclusion errors (when individuals who are intended to receive benefits do not get them). With a relative poverty line, inclusion and exclusion errors are symmetrical mechanically; assuming all benefits are distributed, any individual not in the poorest x% who ends up getting included necessarily implies that an individual who is in the poorest x% is being excluded. If inclusion and exclusion benefits are equally unfavorable to a policymaker, this may not be troublesome. However, some policymakers may be more concerned about minimizing one of these error types. Exclusion errors may be seen as more harmful than inclusion errors; exclusion can mean that program objectives of ensuring food security among all poor households are not being met, while inclusion just means that some resources are accidentally going to families that may not need it. (In practice, accidentally included families may be close to the eligibility threshold anyway.) When an absolute poverty line is used, there is the possibility of designing targeting processes that prioritize avoiding exclusion errors over inclusion errors or vice versa because errors are not mechanically related.

Regardless of the type of poverty line chosen, it is also important to consider that poverty line's threshold value, which likely will depend on program objectives. If the program aims to only help people in the most acute distress, threshold values will be lower than if it aims to help any household experiencing occasional food insecurity. Note that the choice of threshold values can also affect targeting accuracy; it may be easier to identify households that are very poor or very rich than those that are somewhat poor or somewhat rich, meaning that more extreme (either very high or very low) poverty lines may be associated with higher targeting accuracy. This does not imply that a program should necessarily change its poverty threshold to improve targeting accuracy, but it is useful to keep in mind when comparing accuracy between programs.

Specifying an Objective

Taken together, a welfare metric and poverty line (e.g., the program aims to target the 10% of households with the lowest income) constitute a well-defined targeting objective. Programs can also have multi-dimensional objectives, which identify either multiple criteria that the targeted population should meet or target multiple populations. This may be the case when the program has additional objectives beyond simply helping the most food insecure households. For instance, social programs that aim to improve both food security and children's health outcomes may target households in the poorest 10% of the income distribution that also have children under the age of 5. Alternatively, programs could target households that are in the poorest 10% of the income distribution and/or have a child under the age of 5.

Targeting Methods

The next decision for a program implementer is to determine how to identify beneficiaries in a way that satisfies the targeting objective. This process can be resource-intensive, and there is often a trade-off between cost and accuracy. The optimal targeting method for a given program depends on the targeting objective, program features, and contextual factors. Here we describe some commonly used targeting methods and discuss conditions in which each might be the appropriate method.

Categorical Targeting

Categorical targeting involves choosing beneficiaries based on one or more demographic characteristics, such as gender, age, or disability status. In this case, anyone who fits in the specified category is eligible for the program, while those not in the category are ineligible. Generally, categorical targeting is most sensible when the eligibility criteria are easily verifiable. For instance, it is straightforward for a program to target adults 65 and older if most individuals have some sort of government identification to verify their age. Additionally, categorical targeting can be efficient in scenarios in which the welfare metric specified in the targeting objective is highly correlated with belonging to the targeted category. If, for example, the objective is to target low-income households, and many households containing a senior citizen are low income, then categorical targeting based on age may be appropriate. A main upside of categorical targeting is that it is generally relatively cheap to identify and verify the eligibility of potential beneficiaries. However, the drawback is that targeting errors can be significant if targeting objectives are not highly correlated with the eligibility categories. In the example above, targeting households only with senior citizens misses any low-income households which do not contain senior citizens and includes wealthy households containing senior citizens.

Geographical Targeting

Geographical targeting is a special case of categorical targeting in which eligibility is based on a household's location. Geographical targeting can be advantageous when beneficiaries who satisfy the targeting objective are spatially correlated. Household location may also be the cheapest category to verify in settings where individuals do not tend to have formal identification documents. However, like any other categorical targeting method, simple criteria can translate to a high rate of targeting errors. Moreover, with geographical targeting, a program implementer should also consider whether the transfer size is large enough to cause price disturbances in the local economy. If many households in the same market experience a simultaneous increase in their demand for food, the price of goods may increase, diminishing the benefits' impacts on households' purchasing power. Yet empirical evidence that geographically concentrated transfers affect local prices is mixed and may depend on whether benefits are given as cash or in-kind (Cunha et al. 2019, Egger et al. 2022).

Means Testing

Means testing involves verifying whether an individual meets the welfare criteria specified in the targeting objective by measuring or observing the welfare metric for the household. For instance, in a program which targets households based on their income, this would involve calculating each household's income and identifying those which satisfy the targeting objectives. Means testing is most appropriate in settings where there is at least some level of formal documentation of households' income, such as through employment or tax records, which can be used to identify eligible households. In settings where such records don't exist, verifying the eligibility of an entire population is likely prohibitively expensive, as it may involve administering complicated surveys and income verification exercises to everyone. The big upside of means testing (in places where it is possible) is that by design, targeting errors are minimal. However, even relatively comprehensive social registries may systematically exclude potential beneficiaries, such as informal workers or refugees. Hence there can still be significant opportunities for inclusion errors to arise.

Proxy Means Testing

Proxy means testing (PMT) exploits the fact that proxies that are relatively cheap to verify (such as household demographic characteristics and ownership of visible assets) may be good predictors of welfare metrics like income or consumption, which are more expensive to measure. To conduct proxy means testing, a program implementer must first locate data from a comprehensive survey of either a subset of the population of interest or a similar population which contains information on both the targeted welfare metric and the proxies of interest. This data is then used to estimate a formula that describes the relationship between the proxy variables and targeted welfare metric. Program implementers can then predict the value of the targeted welfare metric (often called a proxy means score) in the population of interest using more easily obtainable information on the proxies.

Proxy means formulas can vary in terms of their complexity and the numbers/types of proxies used as inputs. They can be as simple as the Poverty Probability Index (PPI), which uses only 10 proxies and combines them with a simple, additive formula, or as complex as 60+ variables and cutting-edge machine learning methods (IPA 2022, Wobcke and Mariyah 2023). There is often a trade-off between simplicity and accuracy. Adding more proxies to the formula increases the costs of collecting proxies from each potential beneficiary but also likely increases the accuracy of the predictions. PMT methods are most appropriate in settings where program implementers can calculate a highly predictive formula with proxies that are easy to measure. Notably, when judging a proxy means formula's performance, program implementers should consider not just in-sample performance (how well the proxies predict the targeted metric in the data sample which is used to estimate the formula) but also out-of-sample performance (how well the proxies and formula predict the targeted metric in a distinct sample). Only considering in-sample performance can lead to "overfitting," in which variables that happen to be correlated by chance in the sample used to estimate the formula influence the formula, causing it to perform poorly in other samples (Clark 2004).

We are also starting to see the use of non-survey-based "big data" proxies to predict income, such as satellite imagery (Huang et al. 2021) and call data records (Aiken et al. 2022). Such methods have the potential to provide predictive welfare proxies without having to undertake expensive surveys (which can also be prone to reporting errors), as this data becomes easier and cheaper to obtain and methodological advances make working with big data more feasible. These methods have been piloted in a few cases but have yet to be widely implemented. Like any other type of targeting database, it is critical to consider who might be excluded in such systems. Satellite imagery-based targeting may omit individuals who are not homeowners or landowners, while call data records omit individuals who do not own cell phones.

Community-Based Targeting

The previous methods all involve program implementers directly gathering information (of varying levels of complexity) from individuals or households to assess eligibility. An alternative approach is to decentralize the process, outsourcing targeting to informants who can more cheaply observe the welfare status of potential beneficiaries. As these informants tend to be geographically proximate local leaders or neighbors, this method is referred to as "community-based targeting." In practice, there are many ways to conduct community-based targeting, ranging from asking a local government official to identify poor households from a village roster to convening the entire community and having them rank all households from richest to poorest. Involving the entire community can increase collective agency, though targeting exercises also involve time and possibly psychological costs (if it is unpleasant to discuss one's welfare publicly).

A key benefit of community-based targeting is that in situations with poor income documentation, it can be significantly cheaper than even conducting a proxy means test (Alatas et al. 2012). Hence this method may be appropriate for programs with limited resources to spend on targeting and relatively strong local government capacity. However, as with other methods, this decrease in cost has the possibility of translating into an increase in targeting errors. Indeed, existing evidence suggests that regardless of the process particulars and context, community-based targeting has at best a similar error rate and, at worst, a much higher error rate than proxy means testing (Alatas et al. 2012, Karlan and Thuysbaert 2019, Stoeffler et al. 2016, Premand and Schnitzer 2021).

Perhaps the more troubling aspect with community-based targeting is that the process through which targeting errors may arise is less clear. With a proxy means test, errors can only arise because the proxy means formula does not perfectly predict poverty status or because incorrect values of the proxies are used. In community-based targeting, errors could be due to any potential bias, conscious or unconscious, that affects informants' knowledge or beliefs about others' well-being. Moreover, it can be difficult to ensure that local informants' targeting objectives perfectly align with the program's targeting objectives; informants may have their own interpretation of what the targeting objectives mean or may even have an ulterior political motive driving how they allocate benefits. Hence, program implementers should be mindful of any of these types of errors that they might find particularly objectionable; for instance, bias against a specific social group may be mitigated in part by making sure that all relevant social groups in a community are represented among the set of targeting informants. It is also useful to consider cultural context and norms more broadly. Community-based targeting is most appropriate in societies where people would know when others are experiencing hardships, perhaps because communication about such topics is less stigmatized.

Self-Targeting

Self-targeting includes even further decentralization, outsourcing the task of targeting beneficiaries to the potential beneficiaries themselves. This involves the program introducing some "hurdle" or cost necessary to obtain benefits that is only worth incurring if an individual truly needs the benefits. In practice, hurdles could consist of beneficiaries having to wait in a long line to receive benefits or to provide low-wage labor on a public works project. Self-targeting is one of the least resource-intensive methods and may be a good option when there are few resources to be devoted to targeting or when there is a dual program objective, such as increasing productive employment (for a public works program). However, self-targeted programs can also suffer from high rates of both inclusion and exclusion errors, depending on their design. Notably, some individuals may face disproportionate costs to overcome the hurdle and thus be more likely to be excluded; for instance, standing in line for hours may be particularly challenging for senior citizens, individuals with disabilities, or mothers with young children. Hence self-targeted programs may need to be complemented with other programs that can reach excluded groups. On a moral level, programs may struggle with the idea of placing additional burdens on the very populations they seek to help.

Hybrid Methods

Many programs use a combination of the methods discussed to leverage the different benefits each method offers. For instance, programs like Kenya's Older Persons Cash Transfer Programme and Rwanda's Vision 2020 *Umurenge* Programme use community-based targeting to identify individuals who likely qualify for benefits and then use a proxy means test to verify eligibility (Chepngeno-Langat et al. 2021, Williams et al. 2020). This combination leverages the low cost of community-based targeting to identify a subset of potential beneficiaries and the higher accuracy of proxy means testing to minimize inclusion errors. Similarly, programs often filter the subset of potential beneficiaries using geographical or categorical targeting and then verify eligibility using a proxy means test or further community verification.

Additional Concerns

Compliance

Targeting programs in which cash or other benefits are being distributed are vulnerable to corruption, elite capture, and other forms of manipulation. This is especially true when there are multiple levels of bureaucracy between where targeting objectives are being set and where benefits are being targeted and distributed. Regardless of the targeting method(s) implemented, requiring the entire process to be well-documented and targeting decisions to be justifiable based on verifiable information can mitigate non-compliance. This allows targeting decisions to be monitored and audited. Additionally, re-verification of eligibility at regular intervals can help identify households whose beneficiary eligibility status has changed and households who were wrongly included or excluded at the last eligibility verification check.

Public Perceptions and Involvement

When benefits are not given to everyone, there is the possibility for perceptions of "unfairness" to arise. Process transparency can help clarify both targeting objectives and eligibility decisions to participants, thus alleviating some of the perceptions of unfairness. At the same time, too much transparency can make the system "gameable." For instance, a household that knows ownership of specific assets affects their proxy means test score may attempt to hide those assets to obtain benefits they are not actually eligible for.

A related question is whether perceptions of fairness are improved when potential beneficiaries are involved in targeting decisions, such as through community-based targeting. Some evidence from Indonesia suggests that including the community in the targeting process increases potential beneficiaries' overall satisfaction with the targeting process (Alatas et al. 2012). Yet other evidence from Niger suggests that individuals who ultimately do not end up receiving benefits are more satisfied with formula-based methods (like a proxy means test) than with community-based targeting (Premand and Schnitzer 2021), as the latter can be seen as more susceptible to manipulation.

Program Design Features and Targeting

Design features of a program that are not directly related to the explicit targeting process still have the potential to induce additional unintended effects on the selection of beneficiaries. Details as simple as the language(s) and delivery modalities of the program's promotional materials, the locations/times of day when benefits are distributed, and whether supplemental technology such as a smart phone or mobile money account is required to receive benefits can all induce additional self-targeting. Hence it is prudent for program implementers to carefully consider the targeting implications of all program elements, not just those directly related to the selection of beneficiaries.

Conclusion

There are many possible ways to target benefits for social assistance programs that aim to ensure food security, and multiple methods may be appropriate in any given context. Specifying clear targeting objectives and carefully choosing a targeting method based on the program objectives, targeting objectives, and context will increase the probability of success.

References

Aiken, E., Bellue, S., Karlan, D., Udry, C. and Blumenstock, J.E. 2022. Machine learning and phone data can improve targeting of humanitarian aid. Nature, 603(7903), pp.864-870.

Alatas, V., Banerjee, A., Hanna, R., Olken, B.A. and Tobias, J. 2012. Targeting the poor: Evidence from a field experiment in Indonesia. American Economic Review, 102(4), pp.1206-1240.

Banerjee, A., Hanna, R., Olken, B.A., Satriawan, E. and Sumarto, S. 2021. Food vs. food stamps: Evidence from an at-scale experiment in Indonesia. NBER Working Paper (No. w28641). National Bureau of Economic Research.

Chepngeno-Langat, G., Van Der Wielen, N., Falkingham, J. and Evandrou, M. 2023. Targeting cash transfers on the "poorest of the poor" in the slums: How well did the Kenya's older persons cash transfer programme perform? Journal of Aging & Social Policy, 35(1), pp.107-124.

Clark, T.E. 2004. Can out‐of‐sample forecast comparisons help prevent overfitting? Journal of forecasting, 23(2), pp.115-139.

Cunha, J.M., De Giorgi, G. and Jayachandran, S. 2019. The price effects of cash versus in-kind transfers. The Review of Economic Studies, 86(1), pp.240-281.

Dávila Lárraga, L.G. 2016. How does Prospera Work?: Best practices in the implementation of conditional cash transfer programs in Latin America and the Caribbean.

De Groot, R. 2016. Ghana LEAP 1000 impact evaluation: Targeting effectiveness. The Transfer Project Research Brief 10.

Deaton, A. 2003. Household surveys, consumption, and the measurement of poverty. Economic Systems Research, 15(2), pp.135-159.

Egger, D., Haushofer, J., Miguel, E., Niehaus, P. and Walker, M. 2022. General equilibrium effects of cash transfers: Experimental evidence from Kenya. Econometrica, 90(6), pp.2603-2643.

FAOLEX Database. 2025. Food and Agriculture Organization of the United Nations.

Gentilini, U. 2022. Social protection and jobs responses to COVID-19: A real-time review of country measures.

Huang, L.Y. Hsiang, S.M. and Gonzalez-Navarro, M., 2021. Using satellite imagery and deep learning to evaluate the impact of anti-poverty programs. NBER Working Paper (No. w29105). National Bureau of Economic Research.

Innovations for Poverty Action. 2022. About the PPI: A poverty measurement tool.

Karlan, D. and Thuysbaert, B. 2019. Targeting ultra-poor households in Honduras and Peru. The World Bank Economic Review, 33(1), pp.63-94.

Premand, P. and Schnitzer, P. 2021. Efficiency, legitimacy, and impacts of targeting methods: Evidence from an experiment in Niger. The World Bank Economic Review, 35(4), pp.892-920.

Schnitzer, P. and Stoeffler, Q. 2024. Targeting social safety nets: Evidence from nine programs in the Sahel. The Journal of Development Studies, 60(4), pp.574-595.

Sharp, K., Brown, T., and Teshome. A. 2006. Targeting Ethiopia's Productive Safety Net Programme (PSNP). Overseas Development Institute.

Soares, S., Perez Ribas, R. and Veras Soares, F. 2010. Targeting and coverage of the Bolsa Família Programme: Why knowing what you measure is important in choosing the numbers. International Policy Centre for Inclusive Growth Working Paper 71.

Stoeffler, Q., Mills, B. and Del Ninno, C. 2016. Reaching the poor: Cash transfer program targeting in Cameroon. World Development, 83, pp.244-263.

Timothy, P., Nzahabwanayo, S., Lavers, T. and Ndushabandi, E. 2020. Distributing social transfers in Rwanda: The case of the Vision 2020 Umurenge Programme (VUP).

Wobcke, W. and Mariyah, S. 2023. Machine learning and data augmentation in the proxy means test for poverty targeting. Statistical Journal of the IAOS, 39(4), pp.961-977.

Carly Trachtman, c.trachtman@cgiar.org
International Food Policy Research Institute, Markets, Trade, and Institutions Unit