10. Aid targeting

Ryan C. Briggs[1]

BASICS OF TARGETING

Targeting is the attempt to control who does (and does not) receive a benefit. This is done to concentrate benefits on people that are considered poor, are thought to benefit most from a program, or are deserving of special attention for some other reason.[2] There are a number of rationales for targeting social programs. First, if one has a fixed budget and assumes that some people will benefit more from a transfer than others, then one can likely do more good by targeting the transfer to those who will benefit most. Equivalently, targeting lets one help more of some fixed subset of the population at lower cost than providing a universal benefit.

Second, putting aside direct cost-benefit calculations, targeting benefits can allow one to achieve broader political goals. For example, one could target specific groups of people or places based on equity concerns. More crassly, when a benefit is targeted and some groups are excluded then this creates avenues for aiming the transfers according to a political logic. For example, politicians may be tempted to aim transfers in ways that are vote maximizing. While not socially useful, we will see that politicians sometimes find this kind of behavior privately useful.

Targeting requires information, with more precise targeting requiring more precise information.[3] At one extreme is individual targeting using individual assessments. If the data are available, one can means test a benefit by checking to see if an individual or household's adjusted income falls under some threshold. A less data-demanding approach is a proxy means test, where applicants to the program are given a score based on observable attributes like housing quality or demographic characteristics. The score should correlate with income and can be used to identify poor people or households.

The other extreme is self-selection, where a benefit is set up such that only people intended to be targeted should find the benefit worth the hassle or opportunity cost of attaining it. This can be done by designing a program that is, for example, demanding on one's time so that only people with a low opportunity cost of time enroll. It could also be done by offering an inferior good, such as low quality food that is less desirable as income increases. There are obvious issues here with stigma and wasting people's time that will be discussed later.

In between these two extremes is, first, targeting that focuses on whole categories of people based on easily observable attributes. This could include targeting people above a certain age or people living in a certain place. A second option is deferring choices of who to target to communities that then nominate households.

While the advantages of targeting are obvious, it faces a number of quite serious challenges in practice. First, targeting is always imperfect. This means both that benefits will sometimes end up going to people not in the targeted group and that some people in the targeted group will not receive the benefit. Thus, even if one only wants to reach a subset of the population, a universal benefit will generally do that more completely than a targeted benefit (though obviously at higher cost).

Second, targeting itself costs money, and resources directed to collecting information for targeting and making targeted transfers are resources that could otherwise fund benefits. Third, the public may conceptualize targeted and universal benefits differently, with universal programs being thought of as more legitimate.

So far, this discussion has been theoretical but given that there are both benefits and costs to targeting, the question of whether or not – or when – targeting is useful is ultimately an empirical one. Fully evaluating the large literature on targeting effectiveness and targeting costs is beyond the scope of this chapter. Instead, in the following two sections I will give an overview of empirical research on the costs and benefits of targeting.

Benefits of Targeting

One of the first questions one has when evaluating targeting is how well it is able to concentrate benefits on the targeted population. In a particularly useful study, Coady et al. (2004) examined 122 targeted antipoverty interventions in 48 low- and middle-income countries. Their key measure of effectiveness is a ratio of how well the program concentrated benefits relative to a uniform (or random) allocation of the same amount. They find that the median program transferred 25 percent more to the bottom 40 percent of the population than would random allocation. The range in performance, however, was very wide. The best programs transferred 300–400 percent more than would random allocation, but a quarter of all programs were regressive and transferred less to the bottom 40 percent than would random allocation.

There was also wide variation within each method of targeting, with the specific type of targeting only explaining about one-third of the variation in effectiveness. Nevertheless, means testing, geographic targeting, and self-selection based on a work requirement performed better than average.[4] Other kinds of self-selection, such as giving inferior goods, were among the least effective forms of targeting. Targeting was more effective in richer countries, perhaps because they have better data and implementation capacity or perhaps because the dollar gaps between the poor and the middle are larger in these countries and so the bottom 40 percent is genuinely easier to identify. It also was more effective in countries where governments were more accountable and in countries with higher levels of inequality.

Another way to approach the question of targeting effectiveness is to ask the extent to which a targeted program is able to reach everyone it is targeting. This is a question of coverage.[5] An analysis of 42 social protection schemes examining coverage showed that it was often quite low (Kidd and Athias 2019, 2020).[6] The best performing program excluded 40 percent of their target population and the worst performing excluded 97 percent. Some of this is due to obvious trade-offs between exclusion errors and inclusion errors. If you loosen selection criteria then you reduce the share of people wrongly excluded, but you increase the number of people wrongly included.

Two other factors are important in explaining this poor performance. First, short-run poverty churn can be quite large while targeting criteria are basically fixed in the medium-term. This means that even with an ideal means test, as the time between the test and the transfer grows possibly large shares of people will move in or out of poverty and this degrades the ability of the test to target transfers to the poor.[7] Second, in lower income countries the poorest 40 percent often has a level of consumption that is not much lower than the next 40 percent. While this means that it is difficult to have a high enough fidelity of information to target only the poorest, it also means leakage to richer deciles is less of a problem. If targeting is deemed necessary due to budget constraints, the authors suggest trying to exclude the affluent (e.g. top 20 percent), who are noticeably better off and thus easier to identify, rather than trying to target resources to the absolute poorest (Kidd and Athias 2019, 2020).

One note of caution in nearly all studies of targeting effectiveness is that they assume that poor households contain only (and all of) a country's poor people. It is difficult to test if this assumption is reasonable because most targeting and evaluation data is collected at the household level, but recent work has shown that this assumption may be unreasonable. For example, under-nourished people are spread quite widely across households in many countries in Africa, where "roughly three-quarters of underweight women and undernourished children are not found in the poorest 20 percent of households" (Brown et al. 2019, p. 631). A similar result also holds for consumption-based poverty in Bangladesh, where within-household inequality means that around 40 percent of poor people are not found in poor households (Brown et al. 2021). These are early but important results because household targeting becomes less attractive the more that poor people are not contained within poor households.[8]

Costs of Targeting

We must consider the costs of targeting alongside its possible benefits. This terrain is well summarized by two sources on which this section is based. The first is a systematic review of 85 papers analyzing the costs of targeting social programs in Devereux et al. (2017) and the second is a recent edited book by Grosh et al. (2022b).

The first cost, which is borne by the organization doing the targeting, is the administrative cost of doing the targeting instead of making a universal transfer. While there is variability here, typical ranges of administrative costs for cash or near cash programs are around 4–9 percent (Devereux et al. 2017, p. 183). Means testing, including proxy means testing, is more costly than targeting whole categories of people. Thus, while targeting costs can easily make up most of the administrative costs for a project, targeting is rarely prohibitively costly.[9]

The other costs are borne more diffusely and are harder to quantify, but this does not make them less important. *Private costs* are costs to individuals that want to participate, such as time and money used to demonstrate eligibility or time to visit a registration center. *Psycho-social costs* are also private and represent decreases in the mental well-being of recipients of a transfer, typically by increasing stigma or shame. *Incentive-based costs* occur when individuals act in otherwise suboptimal ways in order to meet targeting criteria. They have a private and a social component.

Social costs are costs that are borne by a community, including a loss of cohesion or trust, or in extreme cases the outbreak of conflict, due to the division of people into those that qualify for a transfer and those that do not. This can be exacerbated when community-based targeting is used, as it gives elites opportunities for capture and can increase the perception that the targeting was unfair. *Political costs* are again social and relate to the sustainability and quality of a program. On the former, the basic idea is that targeted programs are less likely to enjoy support from a broad cross-section of society and so may be more easily rolled back than universal programs. On the question of quality, Sen (1995, p. 14) put the issue best: "benefits meant exclusively for the poor often end up being poor benefits."

Given that these costs are difficult to quantify and are borne by different actors, Devereux et al. (2017) do not attempt an adding up of costs and benefits of various targeting approaches.[10] Grosh et al. (2022a, p. 120) again acknowledge the difficulty of quantifying these sort of costs, but note that private transaction costs create barriers that "can be significant" for "at least some members of the target population."[11]

In sum, while we know a lot about the possible successes of targeting and can enumerate many potential costs, we have little evidence about how well targeting works at reducing poverty in practice. One study that addresses this question is Ravallion (2009), which examines a means-tested Chinese cash transfer program where targeting effectiveness varied across municipalities. The main, striking, finding is that "cities that did a better job excluding the non-poor tended to do less well in reducing poverty" (p. 225). This was because they also excluded more poor people. This shouldn't be read as saying that targeting is naturally ineffective, but rather that common ways of measuring targeting performance do not naturally lead to better poverty reduction effects, even on a per-dollar basis. Targeting can still be useful, but one should track end goals such as poverty reduction directly and not attempt to infer them based on targeting performance.

HOW AID TARGETING IS UNIQUE

This chapter has so far discussed issues related to the targeting of social programs *in general*. However, *aid* targeting foregrounds some otherwise hidden issue and raises some entirely new issues. This section covers both.

The main hidden issue that is foregrounded when considering aid is one of spatial targeting. In brief, nearly all social programs require people to go a place at some point in time and this creates avenues for targeting even ostensibly universal benefits. This is easy to see with aid because aid often provides local public goods like schools or health clinics, but this issue applies to nearly all social programs. The new issues raised by aid targeting mostly have to do with the fact that donors and recipients, who often have different interests and incentives, must cooperate in order to target aid.

Access Points

Most discussions of targeting social programs assume a range of targeting granularity that runs from universal access on one side to individually-targeted transfers on the other. Policymakers are imagined to select a degree of granularity from this range. However, many universal programs require people to go to specific places in order to access the program, and the presence of these *access points* means that universal programs can in essence be geographically targeted.

Consider the extreme case of voting rights. While voting is a right for adult citizens living in a democracy, one must visit an access point (polling place) to exercise this right. This dependence creates opportunities to spatially target access to this right. By changing the density of polling places or the amount of resources (workers, voting booths) per polling place, one can limit access. This can be done strategically to shape the de facto electorate. This occurs in the United States, for example, where black people wait much longer to vote than do white people (Pettigrew 2017; Chen et al. 2022). This has the expected effect on turnout, as voters who waited longer in line in past elections are less likely to vote in future elections (Pettigrew 2021). This is an extreme example, but it shows how a dependence on physical access points enables targeting of what is ostensibly a universal right. This applies to any "universal" program that requires one to visit a place in order to access the program, and it means that universal programs like free primary school or health care remain open to targeting.

The issue of how access points are targeted is even more important when considering foreign aid, because typically aid is used to fund the development rather than recurrent part of a recipient's budget. For example since 2015, project aid, which has a geographic scope by definition, has consistently made up over 60 percent of all aid reported to the OECD. Budget support never exceeded 10 percent. So while some aid can and does support individualized transfers or national-level programs, a large portion of it funds the creation or improvement of roads, health clinics, hospitals, electricity lines, or schools. These are always placed *somewhere*, and so are always geographically targetable. Further, if poverty reduction is the goal and resources are limited, then these programs can generally benefit from careful targeting because in many parts of the world poverty is geographically concentrated (Cohen et al. 2019).

Donors and Recipients Target Jointly

Aside from geographic targeting being more important with aid, aid targeting is also unique in that donors and recipients must cooperate in order to target aid. This differs from the normal theoretical case of targeting, where a unitary actor makes all targeting decisions.

The key starting point for this analysis should be familiar to those who think about foreign aid: donors and recipients must cooperate for aid to be disbursed and implemented but they typically have different interests and face different incentives. This creates a back and forth "development dance" (Swedlund 2017b) from which flows all manner of aid modalities, from attempts at control like conditionality and carefully tranched aid to more hands off modalities like budget support. These dynamics replicate themselves when it comes to aid targeting, with donors and recipients trying to exercise control over outcomes and processes. Below I describe three different ways that this development dance can influence targeting.

The first way that donor and recipient dynamics can be at play in aid targeting is over who should be targeted. Here both donors and recipients would be expected to make choices that are in part political. For donors, that could mean wanting to target populations that produce good optics with their home populations (e.g. girls in low-income families that are seeking an education). For recipients, that could mean targeting populations that are important for staying in power, such as key groups of voters or populations that are otherwise politically important. Donors obviously have the money in this situation, but recipients typically have much more knowledge about their population. Thus, one can find situations where recipients have taken donor interests but operationalized them in ways that achieve their political goals. For example, in the early 1990s in Kenya the major donors were very much at odds with the Moi regime. Despite this, the Kenyan government was still able to direct a great deal of aid to politically important parts of the country that shared Moi's ethnicity. This was done by taking the high-level donor goal of pro-poor aid and operationalizing the spatial targeting in a way that led to aid concentrating in the parts of Kenya that formed Moi's ethnic coalition (Cohen 1995; Briggs 2014). Donors went along with this subnational distribution of aid despite their frosty relationship with Moi's government. This shows that even when donors do not want to support a recipient government, they can still end up doing so if recipients shape the information that donors receive.

Another classic aid concern that applies to targeting is that of fungibility. Recipients share control over aid with donors but fully control their own budget. Thus, if donors only want to target a population that the recipient already serves, then recipients can acquiesce to the donor's request and then cut their budgetary allocation to that population, thus offsetting the increase in aid. For a host of reasons having largely to do with imperfect substitutions and transaction costs, concerns around fungibility are likely to be smaller in practice than in theory. Nevertheless, this is a place where differing donor and recipient incentives can make targeting less effective. For examine, Seim et al. (2020) ran an experiment on 460 politicians in Malawi where they randomly provided some with information on aid projects in local schools. They then tracked how these politicians made decisions over future resource allocations to the schools, and found that politicians who learned which schools received aid reduced future transfers to these schools by around 25 percent. Politicians did not re-allocate these freed up resources to needier schools or schools in politically important areas. Rather, they seem to have done the re-allocation in a way that produced a more even spread of educational resources across schools. Regardless of one's normative feelings towards fungibility, its presence undermines the case for targeting.

A final issue has to do with the ways that recipients may have to balance the priorities of donors and the preferences of their public. As noted in the first section, one concern around targeting is that targeted programs may receive less public support than broad-based programs. Thus, given a free hand in program design a recipient government may want to set up a broad-based program. However, in many cases donors would prefer to fund programs targeted towards charismatic or marginalized groups. This puts recipients in the position of having to balance donor interests for targeting against their public's interest for broad-based transfers. One could see these sorts of dynamics play out in, for example, discussions around foreign aid programs for South Sudanese refugees in Uganda. In this case donors wanted to fund programs for refugees but the Ugandan public wanted broader programs that also provided services for Ugandans. Donors should be aware that if they have a preference for narrow targeting then this may put recipients in a difficult position.

This section has covered four ways that aid targeting creates new issues that layer on top of the regular issues with targeting social programs. The next section summarizes empirical work on where aid has been targeted, and the final section summarizes some future avenues for targeting research and practice.

WHERE HAS AID BEEN TARGETED?

At the highest level, one can consider aid targeting across countries. Here one simple result is that aid flows to poorer and more populous countries. In 2019, a 1 percent increase in recipient GDP per capita correlated with a 0.6 percent aid decrease while a 1 percent increase in recipient country population correlated with a 0.4 percent aid increase. Together, these two variables explained 60 percent of the variation in aid flows across countries.[12]

Still at the cross-national level, donors tend to target aid to recipient countries that are more closely linked to them through trade, migration, or historical ties (Bermeo 2018). Focusing on interconnections between donor and recipient helps to explain some otherwise puzzling variation in cross-national poverty targeting across donors. For example, Japanese aid has been quite insensitive to poverty since the 1990s and this seems to be due to the fact Japanese aid clusters on its region, which has become much richer over time (Nunnenkamp and Thiele 2006).

Aid from multilateral donors, Scandinavian donors, and the UK tends to be more sensitive to poverty (Nunnenkamp and Thiele 2006). "New" donors like Brazil, Saudi Arabia, UAE, or South Korea tend to be less sensitive to recipient poverty than the "traditional" OECD DAC donors (Dreher et al. 2011). There are instances of donors withholding aid because of governance issues (e.g. Girod et al. 2009), but the cross-national pattern between aid and governance is ambiguous (Neumayer 2003).[13] Turing to politics, Chinese aid in particular is not more politically motivated than that of OECD DAC donors (Dreher and Fuchs 2015).

While descriptive cross-national analyses of aid and poverty are useful, they increasingly mask important aspects of poverty targeting because many of the world's poorest people live in middle-income countries (Kanbur and Sumner 2012) and because poverty is often geographically concentrated within countries (Cohen et al. 2019). Unfortunately, if we want to examine where aid goes at a level that is more fine-grained than countries, data limitations mean that one can generally only examine project aid. While this is unfortunate, since 2015 project aid has consistently made up over 60 percent of all aid reported to the OECD and is about four times larger than the second-largest category.[14]

While research on subnational aid targeting is fairly new, a number of consistent results have emerged. First, geographic poverty targeting of aid is weakest in Africa, and in fact in Africa aid often flows to the places where *richer* rather than poorer people live (Briggs 2017, 2018a, 2018b; Öhler et al. 2019; Odokonyero et al. 2018; Borghi et al. 2018) or to places with *lower* levels of infant mortality (Kotsadam et al. 2018).[15] Globally, Öhler and Nunnenkamp (2014) found that World Bank and African Development Bank aid was insensitive to regional levels of infant mortality, maternal health, and malnutrition. Briggs (2021b) has shown that after conditioning on population, World Bank aid flows to the parts of countries that are brighter (and so likely richer) and that this result holds across all world regions and does not appear to be changing over time. Öhler et al. (2019) evaluated a sample of 58 countries and shown that in 41 of them World Bank aid correlates positively with the share of the population that is in the poorest 40 percent of the income distribution. This result was weakest in Africa where for many countries the correlation was negative. BenYishay et al. (2022) find that World Bank aid targets better off (more urban) subnational districts, and that this targeting does not change as the aid budget shrinks. They also examine differences in implementation costs across districts and find cost differences do not explain pro-rich aid targeting.

The above results are only partially in harmony, with the primary disagreement being how pro-poor subnational aid is outside of Africa. To examine this further, I re-analyzed the Öhler et al. (2019) data to see which countries had pro-poor regional aid and which did not. The results are in Figure 10.1, where each dot is one of 58 countries. The x-axis shows the degree to which people in the bottom decile live in the same regions as people in the top 60 percent and the y-axis shows the degree to which aid flows to regions with more people in the bottom decile.[16]

Figure 10.1 shows that aid flows more to the places where the poor live so long as the poor and those who are better off live in the same places. When the poor are spatially segregated from the better off, aid is no longer pro-poor. The graph also reveals why some people think that pro-rich regional aid flows are an African phenomena. If we look at the bottom half of the graph, 17 of the 21 countries with a negative correlation between the location of the bottom decile and aid are African. If we look at the left side of the graph, we also see that African countries also have more spatial segregation of poverty at the ADM1 level as 17 of the 27 countries with negative correlations between the bottom 10 percent and top 60 percent are

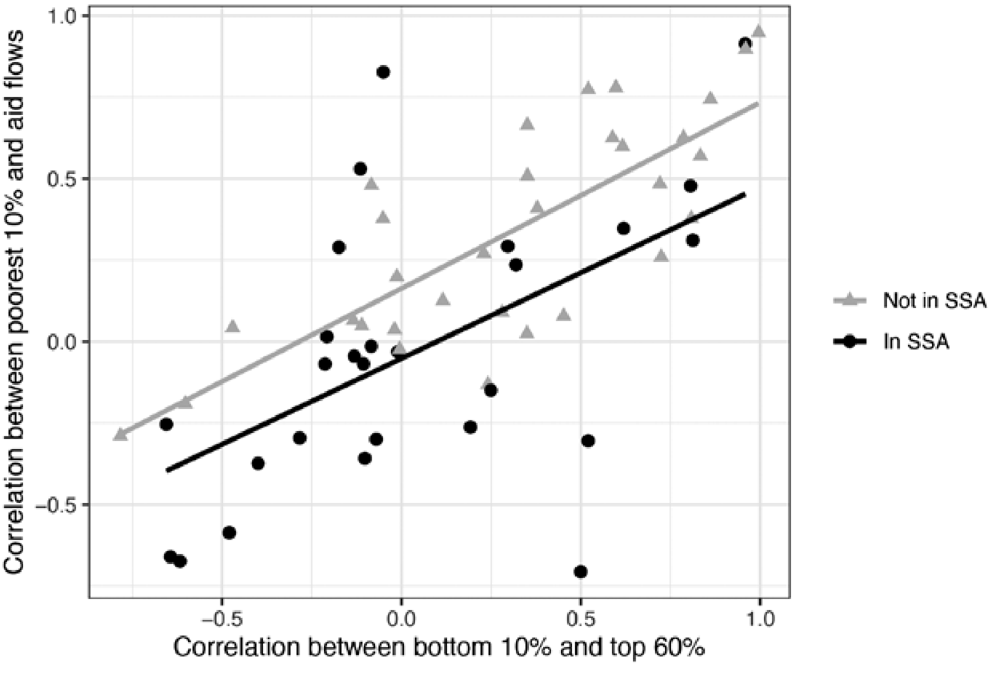

Figure 10.1 Pro-poor aid flows and the spatial segregation of poverty

African. However, the trend lines show that the relationship between the spatial segregation of poverty and pro-poor aid targeting is the same inside and outside Africa.[17]

What is likely going on is that donors are targeting populous places like cities or the areas around cities.[18] This could be because aid might work better in cities, or it could be that ease of access and ease of implementation are important factors in shaping where social investment goes (Francken et al. 2012; Brass 2012; Briggs 2021b; Harris and Posner 2020), or it could be due to some other reason. Regardless, in many countries a few populous regions hold large shares of people from across the entire income distribution, and if aid flows to these places then at the regional level it will appear to reach everyone. This is less likely to happen in Africa, probably because Africa has many more rural poor people than other parts of the world. When the poorest people in a country are spatially segregated into remote or rural regions then when aid concentrates on populous urban regions it will miss the poorest. While this happens more often in Africa, other regions show the same relationship between the spatial segregation of the poor and pro-rich aid targeting. For example, both Timor-Leste and Belarus exhibit spatial segregation of the poor and pro-rich regional aid targeting.

Putting aside the above re-analysis, past research has provided some evidence *against* some explanations for weak geographic poverty targeting of aid. First, in a survey of World Bank operational staff, Briggs (2021b) has shown that donor staff think: that client governments are not biased against targeting aid to the poor, that aid to poorer areas is relatively easy to get approved through the World Bank's bureaucracy, that aid to poorer rural areas gets better evaluation ratings, and that such pro-poor aid is better for development. They did, however, think that aid was harder to implement in rural, remote, and poorer parts of countries.[19]

Moving away from survey evidence, Noemie Zurlinden examined World Bank health aid and shows that it not only reduces infant mortality but that it works better in poorer and more rural regions (Zurlinden 2021, p. 157). This all suggests that pro-rich aid targeting is probably not due to aid working better in richer areas, and that it might be due to difficulties in reaching the poorest with aid.

Shifting to other factors that influence aid targeting, a growing literature suggests that foreign aid can be targeted politically by recipient governments. Jablonski (2014) and Briggs (2014) show that aid favored areas of core government support in Kenya. Briggs (2019) looks across most donors in Nigeria, Senegal, and Uganda and finds aid favors core areas in the first two but found no evidence of favoritism in Uganda. Dreher et al. (2019) show that birth regions of African leaders receive more Chinese foreign aid, but they do not final a similar bias in World Bank aid. This is in contrast to Öhler and Nunnenkamp (2014), who find that both WB and AfDB aid favors birthplaces of leaders.

While the dominant finding is that aid favors areas of core support, there is also evidence of other forms of political targeting. Briggs (2021a) shows that one of Ghana's political parties targeted (largely aid-funded) village electrification projects to core areas while the other party targeted swing areas. Swing targeting has been seen in governmental transfers in Senegal (Caldeira 2012) and Ghana (Banful 2011). Masaki (2014) find evidence of aid favoring opposition areas in Zambia. Finally, political targeting is not omnipresent. For example, there is no evidence that WB projects in India (Nunnenkamp et al. 2017) or aid projects implemented by NGOs in Uganda (Springman 2022) were targeted politically. In general, we lack compelling explanations for when governments choose to target aid (or other resources) according to a core voter or swing voter (or opposition voter) logic.

One can also ask not where recipients or donors try to direct aid, but why donors choose to target aid at all. If one believes aid is primarily given as a bribe or side payment from donors to recipients, then the presence of anything that reduces the value of the payment is puzzling. There is no such puzzle if donors care about producing development outcomes, so the presence of many targeted aid projects can in the first instance be read as showing donors do care in part about outcomes. Some donors also behave in ways that accord with this logic. The World Bank, for example, gives project aid with a wider geographic scope when recipients are better governed (Winters 2010). This may be because more precise geographic targeting leads to less aid capture (Winters 2014). Donors also direct more aid through NGOs (and so around recipient governments) when recipient governance is worse (Dietrich 2013).

THE FUTURE OF AID TARGETING

This chapter concludes by describing two developments that are likely to influence research and practice on aid targeting in the future. The first is the application of machine learning to targeting and the second is research on the organizations and people that target aid.

It seems plausible that targeting over the next decade will be strongly shaped by machine learning.[20] This could happen in at least three ways. First, machine learning models can produce better data for traditional targeting algorithms. For example, machine learning can help to produce maps of global population at a high level of spatial detail (Tiecke et al. 2017) or more geographically disaggregated measures of wealth or poverty (Chi et al. 2021; Lee and Braithwaite 2022).

Second, machine learning models can be used to classify who is poor. This could allow policymakers to make use of high-frequency data sources to identify people who are presently poor, avoiding a common problem where slow moving data sources make it hard to identify the poor when people regularly churn in and out of poverty. For example, metadata from mobile phones can be used to predict wealth (Blumenstock et al. 2015) and when combined with survey data it may identify poor households better than traditional approaches (Aiken et al. 2021a). In the wake of COVID-19 the government of Togo and the organization GiveDirectly, along with a group of researchers, applied machine learning models to a mixture of more traditional survey data, satellite imagery, mobile phone metadata, and high resolution population maps (themselves produced by machine learning) in order to decide where best to target cash transfers (Aiken et al. 2021b). Cash transfers were then made directly to the mobile phones of the targeted people.

It seems likely that the future will see many more applications of machine learning to targeting along these lines. This has the potential to enable faster and more accurate targeting. However, the application of machine learning to targeting social programs also raises many questions around fairness, rights, and the legitimacy of targeting choices. All of these issues stem from the fact that current best-in-class machine learning techniques produce uninterpretable models. If one cannot interpret or audit a model then one cannot fully explain it to the public (or anyone else) and it is challenging or impossible to guarantee that it will have certain properties, such as not discriminating against protected classes or guaranteeing transfers to certain groups in accordance with legislation.

Machine learning interpretability is an active area of research and we may see advances on this problem (e.g. Olah et al. 2018; Hohman et al. 2019), but in the near term we are likely to see applications of machine learning run ahead of our ability to understand these models and so we will have to grapple with the wisdom (or lack thereof) of handing targeting choices over to a decision-maker that outperforms humans or simple models but cannot fully explain its choices. One can imagine targeting algorithms as lying along a legitimacy-accuracy frontier, with more traditional approaches being more transparent, explainable, and thus legitimate while machine learning approaches lack these qualities but are more accurate. Given this setup, one can easily envision policy-makers prioritizing legitimacy over accuracy.

The third way that machine learning research can inform targeting is related to concerns around interpretability and is known as algorithmic fairness. This topic is common in machine learning research but is less familiar in social science. While definitions of algorithmic fairness are contentious, one intuitive way to think about fairness is that the errors of a classification algorithm – and all targeting rules are classification algorithms – should not be related to a list of features of people such as their race or gender. Machine learning researchers have done more work on this than social scientists because, as noted above, their models are often quite opaque and one cannot learn about how they make predictions by simply reading a list of coefficients. However, having a simple and interpretable targeting algorithm does not guarantee algorithmic fairness and so social scientists can learn from machine learning research on this subject.[21] For example, Noriega et al. (2018) look at conditional cash transfer programs in Mexico, Ecuador, and Costa Rica. They show that normal targeting rules can be quite unfair. Poor urban households were more than twice as likely to be wrongly classified as not-poor than rural households, and poor elderly households were again more than twice as likely to be wrongly classified as not-poor than traditional nuclear families. It seems possible that equity concerns will drive more attention to this issue going forward.

Turning from machine learning to human behavior, we have a new stream of research on how information affects decision-makers who exercise discretion over targeting. This is important as many local public goods are targeted not according to set rules but rather by people making constrained decisions with limited information. BenYishay and Parks (2019) review this new literature and show that providing geotagged information about aid, population needs, and project performance to public decision-makers can improve allocation decisions and performance. Mehmood et al. (2021) ran a randomized control trial on Deputy Ministers in Pakistan and show that a short training course in econometrics creates large and persistence (present at 6 months post-treatment) effects on their perceived importance of quantitative analysis. Treated Deputy Ministers were also twice as likely to choose policies for which there was experimental evidence. The aforementioned Seim et al. (2020) study on fungibility is also in this vein.

Into this stream we can also place survey experimental work on donor staff (Swedlund 2017a; Briggs 2021b; Dietrich et al. 2021). In taking the knowledge, preferences, and constraints of decision-makers seriously, this research may allow for tweaks to trainings or incentive structures that produce better targeting outcomes.

This chapter has summarized the costs and benefits of targeting social programs, the ways that aid targeting is different from other kinds of targeting, recent research on how aid is targeted to poverty and how it is targeted politically, and new avenues for research on aid targeting. While the theoretical benefits of targeting are clear, in practice there are numerous costs to targeting. The upshot of this review is that even if one takes a narrow view that cares only about sending resources to a poor segment of the population, precise targeting is not obviously the right choice. If one expands the scope of considerations to include concerns of dignity or public support for programs, then precise targeting looks worse still. This can push people to prefer universal programs, but many universal programs require visiting access points to use or enroll in the program and so in this sense targeting is hard to avoid. This is especially true when one considers the typical kinds of investments that are funded with foreign aid. Future targeting is likely to benefit from advances in machine learning and from advances in our understanding of how policymakers actually make decisions about where to target access points to universal goods or local public goods.

NOTES

1. I'd like to thank Desh Girod, participants at APSA 2022, and the editors for feedback.
2. Note that the people who benefit most from a transfer may not be the poorest people, and so there can be a tension between targeting poverty and targeting expected impact (Haushofer et al. 2022).
3. The following summary paragraphs draw heavily on Coady et al. (2004).
4. For skepticism on the value of proxy means tests specifically, see Brown et al. (2018). For an analysis of geographic targeting plus either proxy means testing or community based targeting applied in six LIC countries in the Sahel, see Schnitzer and Stoeffler (2021). The authors show that no approach is clearly better than random allocation across all outcomes and countries, though in some cases proxy means testing outperforms community based targeting.
5. For a summary of the variety of ways of measuring targeting efficacy, see Ravallion (2015, p. 272).
6. This point is made for proxy means tests in particular in Brown et al. (2018).

7. For example Merfeld and Morduch (2022) show that monthly poverty churn is large in India and they show that taking seasonality into account when targeting can improve targeting performance.
8. Readers interested in measuring the performance of targeting should also consult Leite and Kanth (2022).
9. See also Grosh et al. (2022a, pp. 113–119).
10. "This review demonstrates insufficient evidence in the literature on information required for optimal decisions about whether or not to target and the choice of targeting methodology" (Devereux et al. 2017, p. 194). For another good discussion of the costs of targeting, see Alik-Lagrange et al. (2021, pp. 161–162).
11. Grosh et al. (2022b) also devote substantial attention to ways of reducing these costs in practice and interested readers should consult their book.
12. This comes from a regression of the log of net ODA on the log of population and the log of GDP per capita (adjusted for purchasing power). The data are for 2019 and was limited to non-high income countries. Indonesia, China, and Thailand have negative net ODA and were dropped, leaving 120 countries for the analysis. All data were drawn from the World Development Indicators.
13. For further discussion, see the literature referenced in Swedlund (2017a).
14. The second largest category, "Core contributions and pooled programs and funds" was never more than 15 percent of all aid and the third largest category, budget support, never exceeded 10 percent. So while a focus on project aid is an unfortunate data constraint, it is still a plurality of all aid and is many times larger than the second-largest category.
15. These results are all descriptive. I am ignoring research where people ran "garbage can" regressions and then incorrectly interpreted the values on the control variables as showing evidence that aid (unconditionally) goes to richer or poor places as this is an example of the "table 2 fallacy" (Westreich and Greenland 2013). Notably, this result of a lack of pro-poor aid subnational targeting in Africa exists across many donors and is not limited to the World Bank, which is commonly studied because it produces more accessible data.
16. The underlying WB data are private and I can only work with the released aggregates, which is why I cannot use a finer category than "top 60 percent." The results are not sensitive to using the bottom decile and are similar if I use the bottom 40 percent instead of bottom decile.
17. The pattern is in fact similar within every world region.
18. This pattern is not driven by capital cities and still exists, though becomes somewhat weaker, when all regions holding capital cities are dropped.
19. This finding is consistent with Harris and Posner (2022), who show that difficulty in reaching poor people plausibly explains why many Constituency Development Fund projects are not targeted to the places where poorer people live in Kenya.
20. This topic is also covered in Ohlenburg (2020) and Areias and Wai-Poi (2022), both of which I recommend to interested readers.
21. For example, variables in the model may correlate with protected but unobserved characteristics in the population, and so a model that seems "fair" from a regression table may produce targeting choices that seem perverse.

REFERENCES

Aiken, E. L., Bedoya, G., Coville, A., and Blumenstock, J. E. (2021a). Targeting humanitarian aid with machine learning and mobile phone data: Evidence from an anti-poverty intervention in Afghanistan. Working Paper.

Aiken, E., Bellue, S., Karlan, D., Udry, C. R., and Blumenstock, J. (2021b). Machine learning and mobile phone data can improve the targeting of humanitarian assistance. Technical report, National Bureau of Economic Research.

Alik-Lagrange, A., Dreier, S. K., Lake, M., and Porisky, A. (2021). Social protection and state–society relations in environments of low and uneven state capacity. *Annual Review of Political Science*, 24, 151–174.

Areias, A. and Wai-Poi, M. (2022). Machine learning and prediction of beneficiary eligibility for social protection programs. In M. Grosh, P. Leite, M. Wai-Poi, and E. Tesliuc (eds.), *Revisiting Targeting in Social Assistance* (pp. 507–539). Washington, DC: World Bank.

Banful, A. B. (2011). Do formula-based intergovernmental transfer mechanisms eliminate politically motivated targeting? Evidence from Ghana. *Journal of Development Economics*, 96(2), 380–390.

BenYishay, A., DiLorenzo, M., and Dolan, C. (2022). The economic efficiency of aid targeting. *World Development*, 160, 106062.

BenYishay, A. and Parks, B. (2019). *Can Providing Local Data on Aid and Population Needs Improve Development Decision-Making? A Review of Recent Experimental Evidence*. AidData Publication.

Bermeo, S. (2018). *Targeted Development: Industrialized Country Strategy in a Globalizing World*. Oxford: Oxford University Press.

Blumenstock, J., Cadamuro, G., and On, R. (2015). Predicting poverty and wealth from mobile phone metadata. *Science*, 350(6264), 1073–1076.

Borghi, J., Munthali, S., Million, L. B., and Martinez-Alvarez, M. (2018). Health financing at district level in Malawi: An analysis of the distribution of funds at two points in time. *Health Policy and Planning*, 33(1), 59–69.

Brass, J. N. (2012). Why do NGOs go where they go? Evidence from Kenya. *World Development*, 40(2), 387–401.

Briggs, R. C. (2014). Aiding and abetting: Project aid and ethnic politics in Kenya. *World Development*, 64, 194–205.

Briggs, R. C. (2017). Does foreign aid target the poorest? *International Organization*, 71(1), 187–206.

Briggs, R. C. (2018a). Leaving no one behind? A new test of subnational aid targeting. *Journal of International Development*, 30(5), 904–910.

Briggs, R. C. (2018b). Poor targeting: A gridded spatial analysis of the degree to which aid reaches the poor in Africa. *World Development*, 103, 133–148.

Briggs, R. C. (2019). Receiving foreign aid can reduce support for incumbent presidents. *Political Research Quarterly*, 72(3), 610–622.

Briggs, R. C. (2021a). Power to which people? Explaining how electrification targets voters across party rotations in Ghana. *World Development*, 141, 105391.

Briggs, R. C. (2021b). Why does aid not target the poorest? *International Studies Quarterly*, 65(3), 739–752.

Brown, C., Calvi, R., and Penglase, J. (2021). Sharing the pie: Undernutrition, intra-household allocation, and poverty. *Journal of Public Economics*, 200, 104460.

Brown, C., Ravallion, M., and Van de Walle, D. (2018). A poor means test? Econometric targeting in Africa. *Journal of Development Economics*, 134, 109–124.

Brown, C., Ravallion, M., and Van de Walle, D. (2019). Most of Africa's nutritionally deprived women and children are not found in poor households. *Review of Economics and Statistics*, 101(4), 631–644.

Caldeira, E. (2012). Does the system of allocation of intergovernmental transfers in Senegal eliminate politically motivated targeting? *Journal of African Economies*, 21(2), 167–191.

Chen, M. K., Haggag, K., Pope, D. G., and Rohla, R. (2022). Racial disparities in voting wait times: Evidence from smartphone data. *The Review of Economics and Statistics*, 104(6), 1341–1350.

Chi, G., Fang, H., Chatterjee, S., and Blumenstock, J. E. (2021). Microestimates of wealth for all low- and middle-income countries. *arXiv preprint arXiv:2104.07761*.

Coady, D., Grosh, M., and Hoddinott, J. (2004). Targeting outcomes redux. *The World Bank Research Observer*, 19(1), 61–85.

Cohen, J. L., Desai, R. M., and Kharas, H. (2019). Spatial targeting of poverty hotspots. In H. Kharas, J. McArthur, and I. Ohno (eds.), *Leave No One Behind: Time for Specifics on the Sustainable Development Goals*. Washington, DC: Brookings Institution.

Cohen, J. M. (1995). Ethnicity, foreign aid, and economic growth in sub-Saharan Africa: The case of Kenya. Technical report, Harvard Institute for International Development, Development Paper No. 520.

Devereux, S., Masset, E., Sabates-Wheeler, R., Samson, M., Rivas, A.-M., and te Lintelo, D. (2017). The targeting effectiveness of social transfers. *Journal of Development Effectiveness*, 9(2), 162–211.

Dietrich, S. (2013). Bypass or engage? Explaining donor delivery tactics in foreign aid allocation. *International Studies Quarterly*, 57(4), 698–712.

Dietrich, S., Hardt, H., and Swedlund, H. J. (2021). How to make elite experiments work in international relations. *European Journal of International Relations*, 27(2), 596–621.

Dreher, A. and Fuchs, A. (2015). Rogue aid? An empirical analysis of China's aid allocation. *Canadian Journal of Economics/Revue canadienne d'économique*, 48(3), 988–1023.

Dreher, A., Fuchs, A., Hodler, R., Parks, B. C., Raschky, P. A., and Tierney, M. J. (2019). African leaders and the geography of China's foreign assistance. *Journal of Development Economics*, 140, 44–71.

Dreher, A., Nunnenkamp, P., and Thiele, R. (2011). Are new donors different? Comparing the allocation of bilateral aid between nonDAC and DAC donor countries. *World Development*, 39(11), 1950–1968.

Francken, N., Minten, B., and Swinnen, J. F. (2012). The political economy of relief aid allocation: Evidence from Madagascar. *World Development*, 40(3), 486–500.

Girod, D. M., Krasner, S. D., and Stoner-Weiss, K. (2009). Governance and foreign assistance: The imperfect translation of ideas into outcomes. In A. Magen, T. Risse, and M. A. McFaul (eds.), *Promoting Democracy and the Rule of Law* (pp. 61–92). Basingstoke: Palgrave Macmillan.

Grosh, M., Alas, C. P. R., Zafar, U., Wai-Poi, M., Tesliuc, E., and Leite, P. (2022a). Unpacking the empirics of targeting in low- and middle-income countries. In M. Grosh, P. Leite, M. Wai-Poi, and E. Tesliuc (eds.), *Revisiting Targeting in Social Assistance* (pp. 81–148). Washington, DC: World Bank.

Grosh, M., Leite, P., Wai-Poi, M., and Tesliuc, E. (eds.) (2022b). *Revisiting Targeting in Social Assistance*. Washington, DC: World Bank.

Harris, J. A. and Posner, D. N. (2020). Does decentralization promote poverty alleviation? Evidence from Kenya's constituencies development fund. AidData Working Paper 106.

Harris, J. A. and Posner, D. N. (2022). Does decentralization encourage pro-poor targeting? Evidence from Kenya's constituencies development fund. *World Development*, 155, 105845.

Haushofer, J., Niehaus, P., Paramo, C., Miguel, E., and Walker, M. (2022). Targeting impact versus deprivation. CEGA Working Paper Series WPS199.

Hohman, F., Park, H., Robinson, C., and Chau, D. H. P. (2019). Summit: Scaling deep learning interpretability by visualizing activation and attribution summarizations. *IEEE Transactions on Visualization and Computer Graphics*, 26(1), 1096–1106.

Jablonski, R. (2014). How aid targets votes: The impact of electoral incentives on foreign aid distribution. *World Politics*, 66(2), 293–330.

Kanbur, R. and Sumner, A. (2012). Poor countries or poor people? Development assistance and the new geography of global poverty. *Journal of International Development*, 24(6), 686–695.

Kidd, S. and Athias, D. (2019). Hit and miss: An assessment of targeting effectiveness in social protection. Working Paper.

Kidd, S. and Athias, D. (2020). Hit and miss: An assessment of targeting effectiveness in social protection. Summary version with updated analysis. Working Paper.

Kotsadam, A., Østby, G., Rustad, S. A., Tollefsen, A. F., and Urdal, H. (2018). Development aid and infant mortality. Micro-level evidence from Nigeria. *World Development*, 105, 59–69.

Lee, K. and Braithwaite, J. (2022). High-resolution poverty maps in sub-Saharan Africa. *World Development*, 159, 106028.

Leite, P. and Kanth, P. (2022). Measuring the performance of targeting methods. In M. Grosh, P. Leite, M. Wai-Poi, and E. Tesliuc (eds.), *Revisiting Targeting in Social Assistance* (pp. 467–506). Washington, DC: World Bank.

Masaki, T. (2014). The political economy of aid allocation in Africa: Evidence from Zambia. Working Paper.

Mehmood, S., Naseer, S., and Chen, D. L. (2021). Training policymakers in econometrics. NBER Working Paper.

Merfeld, J. D. and Morduch, J. (2022). Poverty at higher frequency. Working Paper.

Neumayer, E. (2003). *The Pattern of Aid Giving: The Impact of Good Governance on Development Assistance*. London: Routledge.

Noriega, A., Garcia-Bulle, B., Tejerina, L., and Pentland, A. (2018). Algorithmic fairness and efficiency in targeting social welfare programs at scale. Paper presented at Bloomberg Data for Good Exchange Conference.

Nunnenkamp, P., Öhler, H., and Sosa Andrés, M. (2017). Need, merit, and politics in multilateral aid allocation: A district-level analysis of World Bank projects in India. *Review of Development Economics*, 21(1), 126–156.

Nunnenkamp, P. and Thiele, R. (2006). Targeting aid to the needy and deserving: Nothing but promises? *The World Economy*, 29(9), 1177–1201.

Odokonyero, T., Marty, R., Muhumuza, T., Ijjo, A. T., and Owot Moses, G. (2018). The impact of aid on health outcomes in Uganda. *Health Economics*, 27(4), 733–745.

Ohlenburg, T. (2020). *AI in Social Protection: Exploring Opportunities and Mitigating Risks*. Deutsche Gesellschaft für Internationale Zusammenarbeit (GIZ) GmbH.

Ohler, H., Negre Rossignoli, M., Smets, L., Massari, R., and Bogetic, Z. (2019). Putting your money where your mouth is: Geographic targeting of World Bank projects to the bottom 40 percent. *PloS ONE*, 14(6), e0218671.

Ohler, H. and Nunnenkamp, P. (2014). Needs-based targeting or favoritism? The regional allocation of multilateral aid within recipient countries. *Kyklos*, 67(3), 420–446.

Olah, C., Satyanarayan, A., Johnson, I., Carter, S., Schubert, L., Ye, K., and Mordvintsev, A. (2018). The building blocks of interpretability. *Distill*, 3(3), e10.

Pettigrew, S. (2017). The racial gap in wait times: Why minority precincts are underserved by local election officials. *Political Science Quarterly*, 132(3), 527–547.

Pettigrew, S. (2021). The downstream consequences of long waits: How lines at the precinct depress future turnout. *Electoral Studies*, 71, 102188.

Ravallion, M. (2009). How relevant is targeting to the success of an antipoverty program? *The World Bank Research Observer*, 24(2), 205–231.

Ravallion, M. (2015). *The Economics of Poverty: History, Measurement, and Policy*. Oxford: Oxford University Press.

Schnitzer, P. and Stoeffler, Q. (2021). Targeting for social safety nets. World Bank Policy Research Working Paper.

Seim, B., Jablonski, R., and Ahlbäck, J. (2020). How information about foreign aid affects public spending decisions: Evidence from a field experiment in Malawi. *Journal of Development Economics*, 146, 102522.

Sen, A. (1995). The political economy of targeting. In D van de Walle and K. Nead (eds.), *Public Spending and the Poor: Theory and Evidence* (pp. 11–24). Baltimore: Johns Hopkins University Press.

Springman, J. (2022). Aid flows and incumbency advantage: Evidence from NGO projects in Uganda. *World Politics*, 74(4), 523–563.

Swedlund, H. J. (2017a). Can foreign aid donors credibly threaten to suspend aid? Evidence from a cross-national survey of donor officials. *Review of International Political Economy*, 24(3), 454–496.

Swedlund, H. J. (2017b). *The Development Dance: How Donors and Recipients Negotiate the Delivery of Foreign Aid*. Ithaca, NY: Cornell University Press.

Tiecke, T. G., Liu, X., Zhang, A., Gros, A., Li, N., Yetman, G., Kilic, T., Murray, S., Blankespoor, B., Prydz, E. B., et al. (2017). Mapping the world population one building at a time. *arXiv preprint arXiv:1712.05839*.

Westreich, D. and Greenland, S. (2013). The table 2 fallacy: Presenting and interpreting confounder and modifier coefficients. *American Journal of Epidemiology*, 177(4), 292–298.

Winters, M. S. (2010). Choosing to target: What types of countries get different types of World Bank projects. *World Politics*, 62(3), 422–458.

Winters, M. S. (2014). Targeting, accountability and capture in development projects. *International Studies Quarterly*, 58(2), 393–404.

Zurlinden, N. (2021). Essays in development economics. PhD thesis, University of St. Gallen.