

Targeting social transfer instruments

The objective of this chapter is to explain how social transfer targets should be determined. The first part of the chapter outlines the potential benefits and costs of targeting. The rest of the chapter explains different types of targeting and examines their effectiveness, drawing upon an international evidence base.

Benefits of targeting

Cash-based social transfers can either be universal, where everyone within a category – such as children or the elderly – is eligible or they can be explicitly targeted to people who are identified as poor or vulnerable. The main benefit of targeting the poor is that it potentially saves money by reducing the “inclusion error” of universal programmes – the distribution of transfers to people who are not poor. Effective targeting makes sure scarce resources go to those who need them most.¹ Box 8.1 provides an example of how targeting can reduce the cost of transfers.

Economists often argue that the real cost of targeting is what the same resources could achieve in a universal programme. The standard against which targeting should be evaluated, therefore, is the impact of a comparably funded universal programme.² What option will reduce poverty more, social transfers targeted to the poor or transfers provided universally? The answer depends on the cost of targeting, which in turn is determined by the political, social, administrative and economic factors discussed above. For example, the universal approach may be particularly relevant for low income countries. A recent study of fifteen African countries found little difference between universal provision and perfect targeting.³

In addition to the potential fiscal savings offered by targeting, there are

Box 8.1. Effective targeting can potentially reduce the cost of social transfers

Targeting potentially reduces the cost of social transfers. Consider a programme that aims to deliver the equivalent of \$30 per month to each poor individual. In a population of 20 million people, of which 60% are poor, the cost of a transfer delivered to everyone (universalism) would be \$600 million dollars per month. If it were possible to exactly identify the poor and deliver the transfers only to them (perfect targeting), the cost of the transfers would fall to \$360 million dollars – a savings of \$240 million.

The approximate gross savings from targeting are proportional to the percentage of the population that is not poor. In this example, with a 60% poverty rate, the savings equal 40% of the cost of a universal grant.

This simple calculation, however, ignores the costs of targeting. Targeting saves on the cost of transfers but imposes other types of costs – administrative, economic, political, social – which are discussed below. The decision to target is informed by weighing the benefits against the costs.

two other possible indirect benefits. First, the perceptions by policy-makers and the public's perceptions of the targeting mechanism may improve political acceptance of the programme. Second, when the conditions are used to target they may be socially productive. For example, requiring that children attend school and that caregivers complete health visits may reduce the number of beneficiaries while potentially improving the human capital of the targeted group. Politically, a requirements that poor households ensure their children's school attendance can satisfy the mindset some policy-makers hold of the "deserving poor". Economically, the improvements in education and health help break poverty traps. Conditional cash transfer programmes recognise these benefits and they are critical contributors to the success of these types of initiatives. (However, it is important to note that even unconditional transfer programmes have been found to improve human capital accumulation.⁴)

On the other hand, targeting involves direct and indirect costs, which vary from country to country and depend on the targeting method chosen. The direct cost is the administrative expense incurred by implementing and complying with the targeting mechanisms, both by the government, the beneficiaries and third parties. Indirect costs include political, economic, and social losses. The following sections discuss the important costs of targeting.

Costs of targeting

Exclusion error

No targeting process is perfect – any attempt to direct social transfers to the poor will likely entail two types of error.⁵ Inclusion error, as mentioned above, is the mistake of providing the social transfer to someone in a household that is not poor. Exclusion error is the failure to provide a transfer to a targeted household that is poor. The reduction of inclusion error is the potential benefit of targeting, exclusion error is part of the cost. Inclusion and exclusion errors are not easily comparable. An unwarranted social transfer (inclusion error) is at best an inadvertent tax rebate (with the associated costs⁶), and at worst a waste of money. On the other hand, depriving poor households of a source

of social investment (exclusion error) can trap generations in poverty, with a social cost many times the unutilised fiscal expenditure. Some social policy analysts have suggested weighting exclusion errors several times that of inclusion errors.⁷

International evidence indicates that poverty targeting in developing countries always has high errors, and it is common for over half of eligible beneficiaries to be excluded from programmes. Exclusion error can nearly completely negate the potential social protection benefits of transfers. Only about six out of a hundred of the poorest (bottom fifth) eligible households in Bangladesh receive the government's social pension.⁸ During the early years of South Africa's Child Support Grant, when targeting mechanisms were rigidly applied, only one in ten poor households with qualifying children was able to access the transfer.⁹

More intensive targeting can actually backfire and increase exclusion error, particularly when it aims to ration diminishing resources. If non-poor but well-connected individuals can more easily defend their share of the benefits, the residual remaining for the poor will shrink. This happened when the Malawi Starter Pack programme (free seeds and fertiliser) introduced community-based targeting, which caused the benefits to the poor to erode.¹⁰

Administrative costs

There are many ways of targeting benefits, but they all require people, skill, time and money. A means test, for example, will require the repeated verification of the income or assets of households in order to decide whether they should receive benefits. The dynamics of poverty in many countries significantly increases the cost of targeting. When people move in and out of poverty frequently, appropriate targeting requires regular assessment of the targeting criteria: "Targeting 'the poor' is an attempt to hit a moving target."¹¹

Private costs

Potential beneficiaries incur direct costs in order to demonstrate their eligibility. Private costs include expenses for transportation to apply for benefits, time expended in transit and in queues (with the associated loss of income or other foregone opportunities), and the fees for obtaining necessary documentation (including "informal" fees in some cases). Prospective workers in the Maharashtra Employment Guarantee Scheme sometimes need to provide cash payments to obtain and fill in appropriate forms, submitting them to the correct officials and enlist the attention of the social services committee.¹²

Indirect costs

Indirect costs may arise when beneficiaries change their behaviour in order to become eligible for the grant. There is evidence that poverty targeting can create disincentives to work, particularly if the targeting test is blunt, because an increase in income may disqualify beneficiaries from receiving the grant. If a person earning less than the equivalent of \$20 per month is entitled to a monthly transfer of \$80, the individual is unlikely to face incentives that

support increased work effort. Targeting transfers to those residing in specific areas may lead to increased migration – which can be costly for the beneficiary but is nevertheless preferable to destitution.

Social costs

Social costs from targeting can include stigma, the possible deterioration of community cohesiveness, and the potential erosion of informal support networks. While the provision of transfers can improve economic independence and reduce the impact of stigma, policy stances that reinforce negative stereotypes can exacerbate the psychological costs of the programmes. Policy makers in Armenia initiated a cash transfer programme by emphasising that it was only for the poor – aiming to employ stigma to promote self-targeting.¹³ In Jamaica, on the other hand, officials launched social transfers with television spots picturing the pregnant spouse of a cabinet minister registering for the programme, conveying a positive message about participation.¹⁴

“Self-targeting mechanisms that rely on social stigma, thereby reinforcing the social marginalisation of transfer recipients, are incompatible with current definitions of development that emphasise social objectives (e.g. empowerment and dignity) as well as economic objectives.¹⁵” The poor often depend on social networks that change when a beneficiary receives transfers. In some cases the beneficiary shares the added income, in other cases remittances to the grant recipient decline.¹⁶ Universal programmes, where benefits are regarded not as stigmatising but as entitlements to citizens, are more likely to build social cohesion. Further, poverty targeting can lead to increased social conflict, particularly within rural communities.

Political costs

Targeting the poor also imposes political costs – primarily by eliminating middle class beneficiaries who could lend their support to social transfers. Poverty-targeted programmes tend to have less support from the middle class and taxpayers – who are excluded from them – and as such tend to have smaller budgets. In contrast, universal programmes have broader-based support and larger budgets. As a result, universal programmes may be more likely to provide the poor with better and higher benefits. The greater the degree of marginalisation of the poor, the more likely that effective poverty targeting will actually reduce the total transfer of resources to the poor.¹⁷ When Sri Lanka began to more effectively target food subsidies using food stamps in the late 1970s, popular support for the social protection scheme deteriorated. In the face of steady inflation, policy makers neglected to adjust the nominal value of transfers for the relatively powerless poor beneficiaries. As a result, the real value of the benefit was cut in half, and poverty and malnutrition increased. The old, untargeted subsidy scheme had allied the middle classes with the poor – and provided more substantial social protection.¹⁸ Similarly, in Colombia, the shift of food subsidies to a poverty-targeted food stamp programme led to an erosion of political support, and the target programme was eliminated.¹⁹ As Sen has pointed out: “The beneficiaries

of thoroughly targeted poverty-alleviation programmes are often quite weak politically and may lack the clout to sustain the programmes and maintain the quality of services offered. Benefits meant exclusively for the poor often end up being poor benefits.”²⁰

The political vulnerability of well-targeted programmes is greatest when the pressure to cut budgets is fiercest – particularly in times of economic crisis. Argentina’s fiscal adjustment in the 1980s and 1990s placed pressure on all areas of the budget – but targeted social transfers and public works (Trabajar) were disproportionately hit. Expansions of the Trabajar programme more than proportionally benefited the poor, but economic contractions fell disproportionately on the most vulnerable while non-poor areas were protected. Targeting’s greater allocation to the poor is a mixed blessing – particularly when benefits are cut just at the time they are needed the most.²¹

Particularly in very low income countries, policy-makers and their constituencies are reluctant to support programmes that target the poorest, lending support instead to universal programmes.²² Namibia and Lesotho both operate universal pensions, with age as the only targeting mechanism. (However, Namibia has moved more towards a targeted approach in recent years.) In a Latin American poll that included upper income groups, more than 80% of the respondents supported increases in broadly targeted social transfer programmes.²³

The substantial nature of these political costs suggests that social transfer inclusion errors are not entirely wasteful. First of all, they potentially generate important support for social transfers. Second, they cost significantly less than their face value. The real cost of inclusion errors for grants paid to taxpayers is not the value of the transfer, but rather the social costs of the taxation process employed to raise the necessary funds. For many inclusion errors, taxpayers give with one hand (through tax) and receive with the other (from the universal benefits). The social cost of the tax process is often significantly less than the value of the transfers.

The targeting test

To formally answer the question of whether or not to target the poor, it is necessary to compare the net costs and benefits of targeting with the reference point of universal provision to the identified group. For example, for a pension programme categorically targeted to those 70 years and older, the universal reference point is the calculated cost of providing the transfer amount to all people in this age group.

The next step is to assess the benefit of targeting under ideal and realistic scenarios. Perfect targeting means the cost of the transfer is the poverty rate among the target group, multiplied times the cost of the universal transfer. For example, if the elderly face an 80% poverty rate, and a universal pension costs \$20 million, a transfer perfectly targeted to the poor would cost only \$16 million. The savings of \$4 million represents the amount of pensions

not paid to those who are not poor. The savings under the realistic scenario would be an amount less than the full \$4 million – perhaps estimated based on comparable international experiences or with the aid of a micro-simulation model. This analysis approximates the benefits of targeting.

The third step quantifies the costs of targeting. As discussed above, targeting imposes many costs on government, beneficiaries and society in general. Some of these can be relatively easily quantified, such as the direct administrative costs of implementing the targeting mechanism. Other costs can be estimated in theory, but precise measures are difficult in practice. For instance, the private costs of potential beneficiaries documenting their household income could be calculated with some accuracy, but in practice it's probably more appropriate to estimate this cost based on surveying a sample of older people and making a few reasonable assumptions. Other costs might not be possible to reliably estimate. For example, targeting can produce social stigma, which intensifies social exclusion. The cost is real, but quantifying it raises contentious issues that are difficult to resolve.

As a result, the application of the targeting test is subjective and normative, and requires policy maker participation. In spite of the complexities, through rigorous quantitative analysis of some costs, and objective qualitative analysis of other costs, the full information set can be provided to policy makers for application of the test. The result of the test is the policy decision of whether the benefits in terms of cost savings from targeting outweigh the net costs, or whether universal delivery would be a less expensive option. Key principles from the test are summarised in Box 8.2.

Targeting the beneficiary

Poverty targeting

In the cases where policy makers adopt a targeted approach in the programme identification phase, the technical design will require the specification of concrete targeting mechanisms. The following discussion provides background on some of the key options for targeting, with the associated benefits and costs.

Policy makers can target the poor individually or collectively, or they can delegate the task – to communities, or to the poor themselves.²⁴ Individual (or household) targeting involves evaluating the incomes, expenditures, assets or personal characteristics of individuals and households. Categorical targeting involves establishing easily-identifiable attributes that characterise poor households and providing benefits to those who share those traits, such as children, older people, or people who live in low-income areas. Self-targeting attempts to target the poor by making the resource provided relatively unattractive, so that only the poorest will want it, however this mechanism (self-targeting) is often costly and inefficient. Community targeting can involve any of these mechanisms, but the beneficiaries are determined at a community level, employing what may be seen as a more subjective approach that draws on local knowledge harvested by community representatives.

Box 8.2: The targeting test

The targeting test involves both subjective and political judgements and cannot be easily quantified. However, the framework provides some general principles for judging when targeting is more or less costly than universal provision. The following table discusses some of the main factors.

Country characteristic	Factors that suggest targeting will reduce the cost of social transfers	Factors that suggest targeting will increase the cost of social transfers
Government administrative capacity	If government capacity is strong , then it is more likely to succeed in implementing targeting.	If government capacity is weak , then targeting may over-tax the government's limited administrative resources – and may potentially prove counter-productive.
Poverty profile	If poverty rates are low , then targeting can potentially generate significant savings – and is more likely to reduce the cost of social transfers.	If poverty rates are high , then targeting has little potential to generate significant savings – and is less likely to reduce the cost of social transfers.
Social solidarity	If social solidarity is strong , the middle class is more likely to accept the need to allocate resources to the poor, and targeting will incur lower political costs.	If social solidarity is weak , the middle class is more likely to resent their exclusion from social transfer programmes – and the political backlash may compromise the success of the programme.
Formalisation	If the poor are well integrated into the formal economy , their economic status will be easier to verify – targeting will be less costly and more likely to succeed. In addition, the costs of complying with documentation requests and other private costs will likely be lower.	If the poor subsist in the informal economy , their economic status will be difficult to verify – targeting will be more costly and less likely to succeed. In particular, documentation to meet targeting requirements will likely be costlier.
Stigma	If the poor suffer little discrimination , stigma created by overt targeting mechanisms is likely to be less costly and targeting is more likely to reduce the costs of transfers.	If the poor suffer from significant social exclusion , targeting may highlight their plight and increase the psychological costs of poverty.

NOTE: The term “government” is used in the table to denote the agent responsible for the programme. This is not necessarily the national government – it may even be a private agent, development institution or donor.

Individual and household assessment

Individual (or household) assessment involves testing a person's or household's means for survival, a process often referred to as “means testing”. Usually this involves interviewing an applicant for a social transfer, and requesting information, and sometimes documentation, on income, assets and family relationships. Verifying this information is expensive, but so is the failure to do so: unverified means testing is susceptible to substantial under-reporting of incomes and assets. Proxy means testing provides an alternative form of individual assessment: instead of targeting based on income, it employs

more easily observed indicators associated with poverty, such as household demographics, characteristics of the household, durable goods and productive assets. Another type of individual (or household) assessment involves community participation in identifying beneficiaries, employing what may be seen as a more subjective approach that draws on the local knowledge harvested by community representatives.

Verified means testing provides a potentially accurate but often costly mechanism for targeting the poor. The prospective grant recipient must document individual or household income or assets, or any other variables the means test depends upon to demonstrate livelihoods. The formal evidence can be costly for the beneficiary – and can create significant costs as the government verifies the information. In particular, when the poor rely heavily on informal sector sources of income, the practicalities of verifying livelihoods substantially increase the cost of the means test. Because of these costs, this choice of targeting mechanism is rare in developing countries.²⁵

In 2000 in South Africa – two years after the implementation of the Child Support Grant – only 10% of eligible households received the social transfer. In the poorest provinces of the country, the take-up rates were the lowest. The poorest households were unable to navigate the bureaucracy and successfully qualify for the grant. Five years later – after the Department of Social Development effectively relaxed the means test and reduced the evidentiary burden (effectively moving closer to an unverified means test) – the take-up rate had increased by 500%. Take-up rates in the poorest provinces rose above the national average. Surprisingly, both inclusion and exclusion error rates fell.²⁶ Not surprisingly, unverified (simple) means tests are more common in developing countries.²⁷

The choice of means testing involves a decision about how important it is to the success of the programme to target accurately. The costs of improving targeting increase rapidly as one aspires to greater and greater accuracy. A programme can reach a point where it spends more on the administrative costs of excluding a beneficiary than it would spend on the benefit itself.²⁸

Proxy means tests

Instead of directly evaluating household earnings and expenditure, which is difficult and expensive to measure accurately, the proxy means test methodology uses assets or other variables such as “proxies” for income or wealth, to establish household well-being. Chile pioneered proxy means testing in 1980 with its Ficha CAS programme, and Colombia and Mexico subsequently adopted the technique for programmes involving public works and conditional cash transfers (as well as health insurance and skills training). Various forms of proxy means tests operate (in programmes or pilots) in Argentina, Armenia, Costa Rica, Ecuador, Egypt, Indonesia, Jamaica, Honduras, Nicaragua, Russia, Sri Lanka, Turkey, the West Bank and Gaza and Zimbabwe.²⁹

This targeting method generally uses national household surveys as its basis. To design a proxy means test, one must identify a manageable number of easily-observed or measured indicators associated with poverty – but

indicators that households cannot easily manipulate to qualify for the social transfer. Typical indicators can cover a range of areas, including demographic characteristics (such as age of household members and size of household), characteristics of the house (such as type of roof or floor), durable goods (such as refrigerators, televisions or cars) and productive assets (such as land or animals).³⁰ For example, World Bank consultants designing the Social Safety Net Reform Project in the West Bank and Gaza statistically analysed data from the Palestinian Housing Expenditure and Consumption Survey (PHECS) in order to construct a proxy means test formula involving a limited number of variables. They then conducted a targeting pilot (initiated in January 2004) to fine-tune the targeting instrument, testing and revising the formula to achieve a balance of inclusion and exclusion errors.³¹

While any one proxy may be relatively weakly correlated with welfare, correlations improve if multiple proxies are used. The proxy means test aims to find the set of proxies that best explain welfare, usually between ten and thirty proxies in total. Each variable is given a weight based on its estimated impact on household expenditure, and a statistical equation is used to calculate a “score” for each household using these weights. Households that score below a certain cut-off point are eligible for the programme.³² In constructing these formulas, national surveys can provide the breadth of coverage, while pilots provide the in depth information and dynamic feedback required to refine the formula in order to reduce inclusion and exclusion errors.

Proxy means tests pose some difficult practical challenges relating to the frequency of updating the formula, the degree of transparency, the requirements for strong administrative capacity and the importance of outreach. When household incomes fluctuate a great deal over time, the proxy means test tends to target poorly because it relies on static indicators.³³ Updating the formula and re-testing the population (referred to as “recertifying”) tends to be expensive, and is usually conducted on a three-year cycle or less frequently on an ad hoc basis.³⁴ Further, Paradoxically, transparency can undermine a proxy means test: by definition, the indicators are only proxies for living standards – so adequate knowledge of the formula can sometimes enable a household to manipulate their circumstances to qualify for the social transfer.³⁵ On the other hand, transparency can ensure that people are able to exercise their rights and can provide a greater sense of equity. Proxy means testing also calls for considerable administrative capacity: it requires both a technically proficient expert team to statistically determine the scoring formula and literate corps of enumerators who support the collection of data from poor applicants.³⁶ Effective outreach policies are critical to minimise exclusion errors but this management arrangement frequently fails to receive the necessary attention.³⁷

Proxy means tests work best when the easily observed proxy variables predict living standards well. One measure of the efficiency of a proxy means test formula is the variability in household expenditure levels that the formula explains. In theory, a “perfect” formula would explain 100% of the variability in expenditure across households – in other words, that formula would

accurately identify poor households. In practice, however, these formulas often perform very poorly. One formula used for a programme in Egypt explained only 43% of variability in consumption (compared to only 62% in the most data-intensive model tested), while proposed equations for a programme in Armenia explained only 25%. Frequently the formulae only explain about 50% of the variability in the identified measure of livelihoods.³⁸ Many factors help explain the large errors associated with proxy means testing. A key concern, however, is the use of the household survey as the basis of the targeting methodology. Household surveys themselves include inaccuracies and should never be regarded as more than an approximation of reality. While they are useful tools for measuring poverty, it is questionable whether they are reliable enough to be used as the foundation of a sophisticated targeting methodology.³⁹ Further, in-built into the proxy means test methodology is an error derived from the weakness of the correlation of the proxies with household income (or expenditure). The regressions used in proxy means tests rarely explain more than half of household income.⁴⁰ Consequently, even the initial weightings of the multiple proxies, when tested against the household surveys from which they were derived, have significant errors. Further, many households may never be surveyed and assessed; there is good evidence from Pakistan, Mexico and Nicaragua that even when all households are meant to be visited, some can be missed out.⁴¹

As such, the methodology incorporates significant errors even before it is implemented on the ground. International evidence indicates that further errors are incorporated during implementation. This can translate to high exclusion error, even among programmes that are highly regarded by advocates of proxy means testing: in Mexico's Oportunidades programme, around 70% of the poor are excluded, while 30% of the non-poor are beneficiaries,⁴² and in Jamaica's PATH programme, the exclusion and inclusion errors are around 50 percent, meaning that around half of intended beneficiaries are unable to access the programme.⁴³ There has also been an example of a proxy means test resulting in a regressive outcome – in Armenia – with more of the benefits going to the non-poor.⁴⁴

In sum, while proxy means testing generally reduces inclusion errors, it often leads to significant exclusion errors – particularly when the costs of application are high.⁴⁵ This pattern lowers their fiscal cost – but also substantially undermines their contribution to social protection.

Categorical targeting

Since obtaining information on income and assets can be costly, categorical targeting relies on easily observed traits that are associated with poverty. While this mechanism reduces the cost of implementing the targeting process, it can increase both inclusion and exclusion errors. Two kinds of traits are most commonly used: geographic and demographic. Geographic indicators aim to target the poor of a particular region, and are commonly used with conditional cash transfer programmes and in response to national disasters. Demographic indicators, such as age, sex, marital status, or disability, are more

Box 8.3: Depth of poverty across different types of households: income shortfall from the poverty threshold for different household types, as % of average income shortfall (poverty gap) for the country

	No elderly persons	Elderly persons only	Elderly & children only	Not headed by elderly	Headed by elderly	All persons
Burundi	100	154	143	100	100	100
Burkina Faso	100	113	116	100	99	100
Côte d'Ivoire	93	213	224	95	121	100
Cameroon	99	151	107	97	112	100
Ethiopia	98	168	120	97	117	100
Ghana	92	119	155	95	123	100
Guinea	88	181	208	92	123	100
Gambia	87	163	59	93	118	100
Kenya	96	128	136	96	124	100
Madagascar	101	96	99	101	93	100
Mozambique	101	92	122	99	105	100
Malawi	98	131	131	98	115	100
Nigeria	96	57	97	98	112	100
Uganda	99	185	151	98	109	100
Zambia	95	171	189	95	135	100

SOURCE: Schwarz (2003)

common with unconditional transfers – for example, age as the basis for child allowances and pensions.

In order to identify appropriate categories, it is useful to construct a typology of households and quantify the intensity of poverty for each type. The construction of these typologies depends on household survey data that is representative of a country's entire population, with information collected on income, expenditure, other measures of well-being, and household characteristics that identify the group categories. An example for fifteen African countries is presented in Box 8.3. Households are categorised by whether they include older people, children or others in various combinations. The reported statistics in the table reflect the relative poverty gap of each household type in each country, expressed as a percentage of the national poverty gap. For example, households in Burundi made up only of older persons have on average a poverty gap equal to 154% of the national average that is, 54% higher than the national average.

Box 8.3 shows that in most countries, households which include older people are poorer than households that exclude them, particularly when an older person is the household head. Households with only older people are even poorer on average, with the poorest types of households usually those that include just older people and children. While the targeting value of this

Box 8.4: Categorical targeting and the analysis of social impact

“In Burundi, Burkina Faso, Cote d’Ivoire, Ghana, and Guinea, the probability of female children not attending school increases when they shift into households headed by older people. The opposite is the case in Cameroon, Nigeria, Uganda and Zambia. In other words, in these four countries, female children living in older people’s homes do not suffer from schooling disadvantage. The policy conclusion of this very limited exercise is very simple and straightforward: it is important to be aware of gender differences in schooling outcomes

when children are looked after by households headed by older people. While our study does not offer conclusive proof, a social pension *targeted* to poor households headed by older people may have the potential for reducing the female disadvantage in schooling. More work is needed for understanding the gender impacts of a social pension program.”

SOURCE: Schwarz (2003)

information is crude, the data provides a guide if categorical targeting is to be pursued.

Further analysis of the poverty profile is required to better understand the nature of poverty and effectively design appropriate policy responses. For example, the high degree of poverty in households made up only of older people and children may reflect the AIDS pandemic – grandparents and other older people have become primary caregivers, often after a household has become impoverished coping with AIDS.⁴⁶

This type of data also enables one to calculate the cost of alternative social transfer programmes constructed on a categorical basis. For example, providing a pension equal to 70% of the poverty threshold to all individuals in Zambia aged 60 years and older would cost approximately 1.68% of national income, whereas perfectly targeting the benefit to the poor would cost 1.33% – a savings of only about 20%.⁴⁷ Does the cost of targeting the poorest rather than a category of individuals regardless of their poverty status – in terms of additional administration, incentive effects, private costs, stigmatisation, political and other costs – outweigh the fiscal savings from not paying the benefit to the non-poor? The benefit-cost calculation may vary from country to country. The policy implication may be different for a country like Ethiopia, where the cost of universal transfers is more than twice the value of benefits perfectly targeted to the poor. Box 8.4 illustrates how this categorical analysis can provide insight into the social impact of transfers.

Targeting older people

Targeting older people through a non-contributory pension is an important social transfer instrument in many developing countries. While universal pensions often benefit a large proportion of non-poor individuals, they are often more effective at reaching the poor. Box 8.3 shows that for a sample of fifteen low-income African countries, households that include older people are usually poorer than those that do not – and often much poorer. Some more specific categories – like households with only older people and children, households with only older people, or households headed by older people – were often much poorer.

A study of twelve social pension programmes around the world found they reached poor households, on average, significantly more than non-poor households – with some schemes performing extremely well but others yielding regressive outcomes.⁴⁸ Categorical targeting to older people can be combined with other mechanisms – Chile’s CAS-PASIS achieves a high degree of progressive incidence using a proxy means test, while Costa Rica’s non-contributory pension allows social workers wide discretion to make eligibility determinations during an office interview and effectively targets the poor.⁴⁹ However, these techniques undermine a rights-based approach and can erode transparency.

Universal pensions can appeal to taxpayers more than other approaches, because most taxpayers will eventually receive the inter-generational transfer. When cash transfers benefit taxpayers as well as the poor, the cost calculations are not directly comparable with those programmes that are effectively targeted to the poor. Some middle-income countries, such as South Africa and Brazil, combine individual assessment with categorical targeting in the form of a means- tested social pension. This model may pose significant challenges in low income countries, where the added complications and costs of means-testing may overwhelm the government’s administrative capacity. In addition to the public costs, the compliance requirements of means tests may be expensive for the targeted individuals and exclude many of the poor who cannot afford the private costs of qualifying for the pension.⁵⁰

The need for proof of age complicates categorical targeting to older people. Given the poor registry systems in many developing countries (particularly decades ago), many older people do not possess formal documentation of their age. If the costs of obtaining the documentation are high, many of the poorest are likely to be excluded. Publicising information about how to obtain appropriate information – and streamlining document access – can help improve targeting to the poor. South Africa’s experience documents the importance of government offices that respond to the needs of the poor. When people receive poor service at administrative offices and have to make multiple visits, high transportation costs can exclude the poorest⁵¹. The best practices balance flexibility with the need for fiduciary accountability. For example, Nepal allowed horoscopes and other widely-held documents as substitutes for birth certificates.

Targeting children

Approximately half of the world’s poor are children.⁵² One form of categorical targeting is the provision of benefits to households with children, particularly when the children are orphaned or otherwise vulnerable. Many of the issues relevant to targeting older people are also pertinent to children, including requirements for proof of age, the need for documentation that uniquely identifies the child, and motivations that stem from the demographic group’s association with poverty. In some countries, older people are increasingly responsible for grandchildren, especially in areas most affected by the HIV/AIDS pandemic. Other issues, such as linkages to education, are specific to

Box 8.5: Child poverty in developing countries

Country	Household/family characteristic	Poor	Non-poor	All households
Argentina	Number of children under 15	3.0	0.4	1.3
Bolivia	Number of children under 15	3.4	1.3	2.3
Brazil	Number of children under 15	3.6	0.8	1.8
Cameroon	% households with 6+ members	59	30	45
Chile	Number of children under 15	2.5	0.9	1.5
Costa Rica	Number of children under 15	3.3	1.0	2.0
Ecuador	Number of children under 15	3.4	1.4	2.9
El Salvador	Number of children under 15	3.7	1.1	2.4
Guyana	Number of children under 17	2.6	1.4	1.8
Honduras	Number of children under 15	4.2	1.7	3.1
Indonesia	Number of children under 9	1.7	N/A	1.2
Malawi	Household size	5.4	4.2	5.0
Mali	Household size	11.5	9.2	10.4
Mexico	Number of children under 15	4.0	1.1	2.3
Nepal	Number of children under 14	3.5	2.5	N/A
Nicaragua	Number of children under 15	4.9	1.8	3.3
Panama	Number of children under 15	3.2	0.8	1.9
Paraguay	Number of children under 15	4.3	1.3	2.8
Peru	Number of children under 15	3.7	1.1	2.4
Philippines	Household size	6.0	5.0	N/A
Tanzania	Household size	7.2	5.0	6.0
Uruguay	Number of children under 15	2.8	0.5	1.2

SOURCE: UNICEF (2000).

targeting children – particularly for programmes that impose conditions to ensure that children benefit from investment in human capital.

Social transfer programmes identify children as a vulnerable group for a number of reasons. In many countries, the number of children in poor households significantly exceeds the number of children in non-poor households. For example, in Mexico, Brazil and several other Latin American countries, the average poor household includes four children, while the average non-poor household has only one child. (Box 8.5 compares the composition of poor and non-poor households for a sample of developing countries.) In Africa, the number of orphans has steadily increased from 1990 to present – and forecasts project an increasing trend.⁵³ Politically, programmes that target children appeal to politicians and electorates because they support the principles of equal opportunity and support longer-term development objectives through their impact on nutrition, health and education.⁵⁴

Programmes that categorically targeting children frequently employ secondary mechanisms to further reduce the number of eligible beneficiaries. On paper South Africa applies a means test to its categorically targeted Child Support Grant programme, although in practice it takes a relaxed approach to implementing this since it found following the means test to the letter was a severe impediment to improving take-up rates.⁵⁵ A study of the United Kingdom’s categorical programmes for children found the administrative costs of the means-tested “Family Credit” were more than twice as high as the universal child benefit.⁵⁶ The relative administrative costs in developing countries are likely to be much higher, since reliable income documentation is more difficult to obtain. The International Labour Organisation has demonstrated the affordability of benefits universally targeted to children in low-income African countries. The cost of providing a benefit equal to approximately \$8 a month in purchasing power parity terms ranges from two to four percent of national income in most of the countries – and the estimated cost declines over time as the children’s projected share of the population falls.⁵⁷

Targeting people with disabilities

“Disability is both a cause and consequence of poverty.”⁵⁸ People with disabilities frequently face greater limits in access to education and employment opportunities.⁵⁹ Only a small fraction of people with disabilities in developing countries have access to assistance, rehabilitation and other appropriate services.⁶⁰ People with disabilities often face a greater likelihood of poverty for several reasons they generally incur greater costs due to medical expenses and coping mechanisms, their caregivers frequently forego alternative earnings, and they sometimes face work limitations.⁶¹ As a result, disability reinforces conditions of poverty, which in turn can exacerbate the debilitating impact of the disability and increase its incidence. For example, poverty fosters conditions like malnutrition that increase the likelihood of disability, and lack of adequate income hampers access to rehabilitative goods and services that mitigate disability.⁶² Social transfers can support part of a comprehensive approach to inclusion that addresses this vicious circle, potentially incorporating people with disabilities into mainstream social and economic activities and reducing the future incidence of disability. Figure 8.1 illustrates the forces that tend to perpetuate a vicious cycle of poverty, increased vulnerability, and a greater likelihood and increased debilitating impact of disability – which in turn reinforces poverty.

Different models of addressing disability yield alternative approaches to social protection.⁶³ The charity model sees people with disabilities as victims of circumstance which require welfare approaches. The medical model views disabilities as directly caused by diseases, injuries or other health impairments, requiring medical treatment and rehabilitation. The social model empowers people with disabilities in order to increase their social and economic participation and contribution, improving overall economic performance while reducing government spending on unnecessary care. This frees resources for

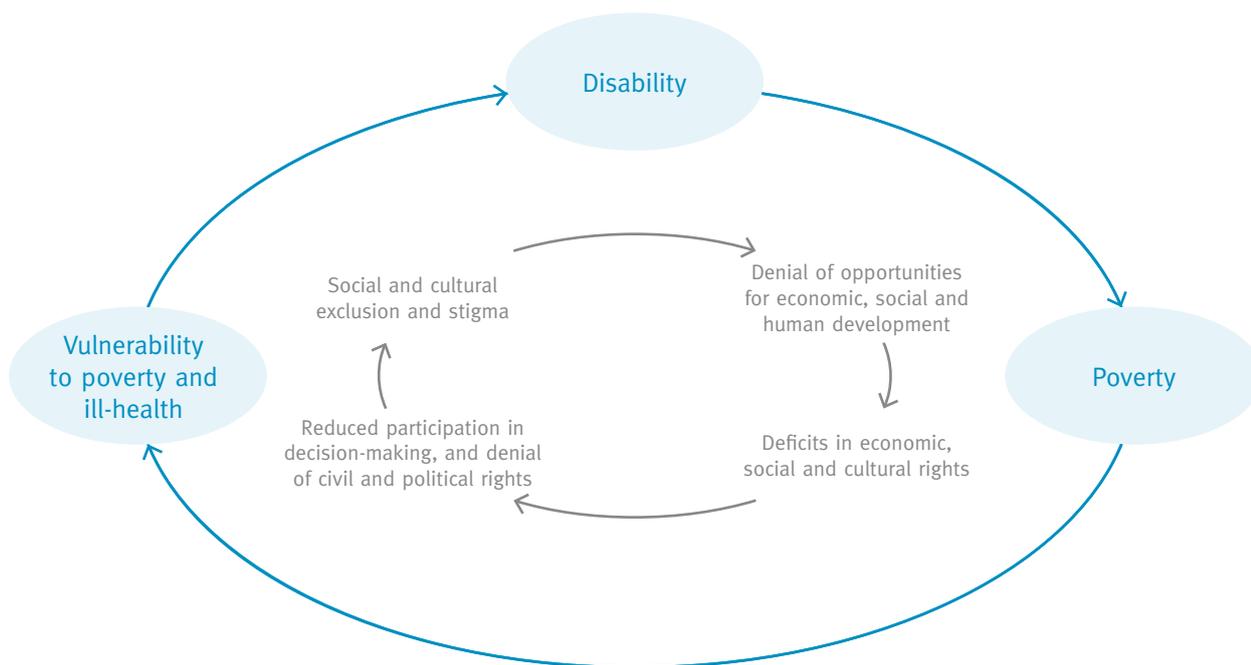


Figure 8.1 The vicious circle of poverty, vulnerability and disability

SOURCE: DFID (2000), page 4.

custodial care approaches that focus better on the minority of people with disabilities who need these services.

A key challenge in the categorical targeting of people with disabilities is the identification of the eligible impairment, limitation or participatory restriction. Most programmes identify beneficiaries through information provided by public entities, from censuses and other search activities, and through efforts to promote self-identification.⁶⁴ The typical screen is whether or not an individual faces a physical or mental impairment that leaves him or her unable to work. For example, in South Africa, “according to regulations issued by the national Minister of Social Development, a person is only eligible if the degree of his or her disability makes him or her incapable of entering a labour market.”⁶⁵ In Namibia, according to the Ministry of Health, a person with a 50 percent disability as determined by a medical doctor would be eligible for the disability grant. For example, “if an individual develops full-blown AIDS and is incapable of working, such a person would qualify for a disability grant by virtue of being incapable of adequately providing for their own maintenance, and not by virtue of being HIV-positive.”⁶⁶ However, in practice the determination of a person’s ability to work is difficult and subjective. Invisible and episodic impairments – such as severe lower back pain or certain mental illnesses – cannot be easily documented. The determination of disability is often a complex individual medical assessment, frequently requiring a visit to

Box 8.6: The disability pension in Namibia

Namibia is one of the few African countries to provide an extensive non-contributory pension to people with disabilities (together with Liberia, Mauritius, Mozambique and South Africa). Nevertheless, in 1995 only about a quarter of the estimated 44,000 people with disabilities received a social transfer – approximately ten thousand received a disability grant and one thousand receiving a pension for the blind. The 2001 Namibian census estimated a much higher incidence of disability, with over 85,000 people affected. In particular, women and those residing in certain rural regions receive disproportionately fewer pensions.

Many of the registered beneficiaries face difficulties in accessing the pension, particularly given the complexity of payment procedures. Persistent fraud poses a challenge, and many eligible people are unaware of registration and payment procedures. Registration procedures can be unduly burdensome for people with disabilities, requiring both birth and citizenship certificates, which many Namibians do not possess. The travel

required to obtain the documentation effectively excludes many of those with disabilities. The payment procedures require both an identification document and computer verification, and system failures periodically exclude eligible beneficiaries – who sometimes do not receive payments retroactively once the computer faults are corrected.

Since few rural Namibians have access to bank accounts, beneficiaries carry the pension home in the form of cash, creating vulnerability to theft and misappropriation. When the Department of Social Welfare privatised the pension payment service in the mid-1990s, the private contractor, Cashmaster Payment Service, reduced the number of distribution points, which exacerbates the problem and increased the transport costs incurred by people with disabilities. The privatisation also significantly increased administrative costs one estimate suggested an increase of 400%.

SOURCES: Subbarao (1998),
Barbro-Isabel Bruhns et al. (1995).

a health clinic or doctor. The subjectivity of the process inevitably gives rise to both inclusion and exclusion errors.⁶⁷

Ensuring inclusive social transfers requires providing physically and socially accessible advice centres and delivery mechanisms.⁶⁸ Physical accessibility requires appropriate architectural features and adapted mechanisms in the transportation system, but can also include social worker visits to the homes of people with disabilities, as well as their right to a personal family representative for programme application and enrolment, and to receive the social transfer.⁶⁹ Namibia's privatisation of social pension payment processes led to the closure of pay points, undermining physical accessibility for beneficiaries.⁷⁰ (Box 8.6 provides more details on Namibia's disability pension.)

Social accessibility requires appropriate outreach channels, including media and community publicity, and the management of administrative staff to ensure their attitudes encourage and facilitate access to benefits for people with disabilities.⁷¹ Appropriate programme design can foster the inclusion of people with disabilities.

Much of the debate over how to implement social transfers for people with disabilities revolves around the choice between mainstreaming and categorically targeting people with disabilities. "Mainstreaming" in this case refers to the policy of improving access for people with disabilities to the social transfers available within society. While targeting transfers to people

with disabilities provides critical resources and can improve the individual's bargaining power within the household, the targeting process requires substantial administrative capacity and cost and may lead to segregation rather than inclusion.⁷² Mainstreaming requires less administration, better promotes inclusion and more broadly reaches the poor, but may require a much longer term commitment to be effective.⁷³ In addition, mainstreaming social transfers for people with disabilities requires that the needed resources are available to the broader population.

Geographical targeting

Geographical targeting determines eligibility for benefits based, at least in part, on the location of the beneficiary's residence.⁷⁴ Disparities in living standards between regions and communities – caused by differences in climate, natural resources, geography and/or public policy – can be found in nearly every country.⁷⁵ Social transfer programmes, particularly those with limited resources, frequently adopt mechanisms that restrict their scope to those areas with the highest concentration of poverty. Conditional cash transfer programmes frequently employ poverty maps, surveys and administrative data in order to supplement other mechanisms with geographical targeting. For example, Brazil's Bolsa Familia targets specific poor communities within municipalities, Jamaica's PATH programme relies on collections of annual consumption data to target at a parish level, Mexico's Oportunidades programme employs a "marginality index" constructed from census data to identify poor communities, and Honduras relies on the "Height Census of First Grade School Children" in order to target poor communities based on the prevalence of malnutrition.⁷⁶ Similarly, public works programmes frequently rely on geographical targeting to identify communities that will most benefit from projects. For example, Argentina's Trabajar programme targeted projects based on the geographic distribution of unemployment by province.⁷⁷ On the other hand, unconditional transfer programmes – with more of a rights-based orientation – less frequently rely on narrow geographical targeting. For example, social pensions in Botswana, Lesotho, Mauritius, Namibia, and South Africa are available to all who qualify. Brazil, however, has separate social pension programmes for rural and urban areas.

One of the key advantages of geographical targeting is its potential simplicity. Particularly in acute emergencies, geographical targeting provides a mechanism for immediate delivery to the hardest hit areas. However, geographical targeting alone risks generating large errors of both exclusion and inclusion if poverty is not spatially concentrated. Particularly at a regional level where income disparities are usually large, geographical targeting includes many non-poor households while excluding regions which nonetheless contain many of the poor.⁷⁸

In both Nicaragua and Peru, social funds employing geographical mechanisms achieved relative success in identifying extremely poor communities, but proved less successful in targeting the poorest households.⁷⁹ A recent study of targeting identified 52 programmes that employed

geographical targeting – and all but one of these utilised an additional mechanism to further improve the identification of poor households.⁸⁰

The size of the geographical targeted area, however, can influence exclusion error, particularly in rural areas, where smaller target areas are more likely to host people of a similar poverty level. In many villages the size of landholdings determines income differences between households, but people share agro-climatic and spatial conditions and off-farm employment opportunities. At higher geographical levels, differences in distances to markets, road conditions, access to vital services and other factors usually contribute much greater variability. If detailed socio-economic information at a village level is available, targeting small areas can significantly improve the effectiveness of the geographic approach.⁸¹ However, targeting small areas on a national level requires reliable information in order to assess the prevalence of poverty at a disaggregated level. Many national household surveys include sample sizes large enough to distinguish poverty rates regionally, but lack the number of observations at a sub-regional level to allow significant inferences about differences in living standards at the level of finer geographical areas.⁸²

Self-targeting

Self-targeting refers to universal transfer programmes that are designed to be attractive primarily to the poor. The non-poor are supposed to voluntarily choose to forego the potential benefit – either because of the costs of participating, the resulting stigma, or the associated conditionalities (work requirements, access costs, or fulfillment of designated activities such as children attending school, household members visiting clinics or other conditionality requirements). Self-targeting was once considered less expensive than other mechanisms because the psycho-social costs of stigma were generally ignored.⁸³

However, in reality, self-targeting also involves significant inclusion and exclusion errors. For example, public works programmes often employ the combination of work requirements and low wages to promote self-targeting by the poor. In some very poor countries, however, members of less poor households may still seek employment at wages that are too low to even provide the very poorest with adequate social protection.⁸⁴ At the same time, the work requirement excludes those who are unable to supply labour to the programme – often the most vulnerable in the society.⁸⁵ In countries where the poor need transfers the most, the wage rate necessary to effectively self-target the poor is so low that such programmes could not honestly claim to be offering significant levels of social protection.⁸⁶

Community-based targeting

Community-based targeting is a relative newcomer to the tracts of social policy analysts, but its essential principles have supported informal mechanisms of social protection for centuries. Community-based targeting can be defined as “a state policy of contracting with community groups or intermediary agents to have them carry out one or more of the following activities: 1)

identify recipients for cash or in-kind benefits, 2) monitor the delivery of those benefits, and/or 3) engage in some part of the delivery process.”⁸⁷ Community-based targeting has most commonly been used within small-scale or pilot programmes, but there are some examples of its use at a national level, such as in Bangladesh and Indonesia.

The basis of community-based targeting is that communities themselves are best able to know and decide who is poor and deserving of support: community representatives are able to define poverty more appropriately in a local context, and they can more efficiently harvest information about individuals with whom they have personal connections.

Community responsibility for targeting also creates a participatory stake in the programme, providing a role for local ownership and control.⁸⁸ This model can also facilitate community mobilisation, empowering disadvantaged groups and legitimising the social transfers programme with positive political consequences.⁸⁹

Community targeting, however, faces its own risks. It is particularly prone to manipulation and capture by more powerful community members and local elites who may distribute resources in ways that deviate from targeting guidelines. Even when the process follows the recommended procedures, the close proximity of beneficiaries (the included), near-beneficiaries (the excluded) and the judges (community representatives) can foster costly social tension and resentment.⁹⁰ More frequently, decisions taken within the community tend to benefit as many people as possible, including the non-poor, regardless of targeting guidelines. Employing teachers to target poor children for a social transfer programme may seem technically efficient, but mandating teachers with this additional role may undermine their teaching effectiveness. Will a poor student excluded from a social transfer by a teacher’s decision feel resentment, and will this affect the student’s success in school?

In addition, since community-based targeting decentralises important policy elements of targeting, it may lead to varying benefit levels for the same groups in different regions. This undermines the objective of horizontal equity and in some cases may induce inefficient population movements. The costs may undermine political support for the programme.⁹¹

The empirical evidence on community targeting is mixed: the mechanism often yields very attractive outcomes, yet in other cases its performance is poor. Ethiopia’s experience with community targeting yielded both positive and negative results that varied from region to region, which apparently depended on the socio-economic and cultural circumstances of each locality.⁹² Evidence from Bangladesh suggests that the average impact is to improve targeting to the poor, but a great degree of variability exists across communities. In particular, in communities with the worst distributions of income, the poorest were the most excluded, probably because they tend to lack political power. Interestingly, as coverage increased, the proportion of benefits reaching the poor increased, suggesting that wider coverage will be more inclusive of the poor.⁹³ Malawi’s efforts to employ community targeting in its ‘Starter Packs’ programme in 2000 failed when community representatives refused to

categorise the finer layers of poverty, replying instead: “We are all poor”.⁹⁴

There is no consensus on how to best design a community targeting approach. One alternative provides a hybrid mechanism, where central authorities define and monitor the targeting categories, and community representatives implement the regulated process but with significant discretion. Public meetings, elected community representatives and external audits can improve transparency and accountability.⁹⁵

A similar approach identifies four ingredients that contribute to the effectiveness of community targeting:⁹⁶

1. Members of the community should understand the targeting process;
2. Information about rules and allocations should be available to the community;
3. Community representatives should be accountable, and those denied benefits should have recourse to appeal;
4. An impartial outside authority should audit the process.

Combinations of targeting mechanisms

Each of the mechanisms for poverty targeting has strengths and weaknesses. Appropriate combinations of instruments can provide complementarity, with the different strengths effectively offsetting the weaknesses. For example, Mexico’s Oportunidades conditional cash transfer programme combines geographical targeting, proxy means tests and community participation. Brazil’s Bolsa Escola employed a poverty line approach together with elements of community control. Many old-age pensions – like those in Brazil, India and South Africa – employ categorical targeting (age and sometimes gender) together with means testing. Other pension programmes, however – like those in Lesotho, Namibia, and Nepal – are universal within their categorical age targets.

Endnotes

- 1 Devereux (2002b), page 2.
- 2 For a thorough discussion of this framework, see Grosh (1994), chapter 8 (pages 131-149). Ravallion (1999) summarises the point: “One option that is probably feasible everywhere is a uniform distribution of the programme budget to every household (whether poor or not). If the transfer to the poor as a percentage of total spending on the programme is less than the percentage of households that are poor, then the uniform allocation is preferable” (pages 32–33).
- 3 Kakwani et al. (2005) find that “the values of PPP indices in conditions of perfect targeting show little difference from the values of indices resulted from universal transfers. This suggests that perfect targeting may not be necessary in cases such as these 15 African countries, where poverty is extremely high” (page 5).
- 4 Case et al. (2003), Duflo (2003), Samson et al. (2004).
- 5 Ravallion (2003) points out that economists who have tested household surveys with far more information than can effectively be harnessed into a targeting mechanism have usually been able to explain at best only half the variability in consumption or income (page 17).

- 6 The economic cost of rebating to taxpayers a lump sum amount is usually less than the value of the transfer. The costs (referred to as deadweight losses) include the administration costs and any distortions created by the tax system.
- 7 Devereux (2002b), page 4; Cornia and Stewart (1993).
- 8 Barrientos (2004), page 18.
- 9 Samson et al. (2006).
- 10 Levy and Barahona (2002), cited in Shepherd et al. (2005).
- 11 Devereux (2002b), page 14.
- 12 Pellissery (2005).
- 13 Coady et al. (2004), Lund (2002).
- 14 Coady et al. (2004), Grosh (1994).
- 15 Devereux (2002b).
- 16 Bush et al. (2001), page 2.
- 17 Gelbach and Pritchett (1995).
- 18 Ravallion (1999), page 47; Anand and Kanbur (1987, 1990); van der Walle (1998), page 240; Besley and Kanbur (1990), page 6.
- 19 Gelbach and Pritchett (1995), page 32; Grosh (1994).
- 20 Sen (1995), page 14.
- 21 Ravallion (2002), pages 118–119.
- 22 Subbarao (2003).
- 23 Rodrik (1999).
- 24 For a more detailed discussion, see Shepherd et al. (2005); Devereux (2002b), pages 7-15; Grosh (1994); Coady et al. (2004). For more on community targeting, see Conning and Kevane (2000).
- 25 Coady et al. (2004), page 13.
- 26 Samson, Mac Quene and van Niekerk (2005).
- 27 Coady, Grosh and Hoddinott (2004), page 13.
- 28 Besley and Kanbur (1993), page 71 and Devereux (2002b), page 7.
- 29 See Larrañaga (2003), Clert and Wodon (2001); Sancho (1992); Racynzski (1991); Castañeda (1990, 2003), Coady (2001) cited in Coady, Grosh and Hoddinott (2004), page 52.
- 30 Coady, Grosh and Hoddinott (2004), page 52.
- 31 World Bank (2004b), pages 2–4.
- 32 See Coady, Grosh and Hoddinott (2004), page 14, for further discussion of proxy means tests.
- 33 Coady, Grosh and Hoddinott (2004), page 14.
- 34 Ibid., page 53.
- 35 Subbarao et al. (1997), pages 20–21.
- 36 Coady et al. (2004), page 55.
- 37 Ibid.
- 38 Ahmed et al. (2001), Grosh and Glinskaya (1998), cited in Coady et al. (2004), page 54.
- 39 See Kidd and Wylde (2010) for a more in-depth discussion on the weaknesses of household surveys and their use in proxy means testing.
- 40 For examples of econometric studies that illustrate this, see Ahmed and Bouis (2002), Alatas et al. (2009), World Bank (2009b).
- 41 Adato et al. (2000); Adato and Roopnaraine (2004), page 16; Huber et al. (2009), page 46; GHK (2009), pages 107, 155.
- 42 Veras et al. (2007), page 2.

- 43 World Bank (2009a).
- 44 Coady et al. (2004).
- 45 Klugman (1999), page 11.
- 46 Subbarao and Coury, (2004).
- 47 Subbarao (1998), page 22.
- 48 Coady, Grosh and Hoddinott (2004), page 32.
- 49 Ibid., pages 69-70.
- 50 Schwarz (2003), page 19.
- 51 Samson (2002); Coady et al. (2004), page 70.
- 52 UNICEF (2000), page 41.
- 53 UNICEF (2005), page 73.
- 54 Coady, Grosh and Hoddinott (2004), page 7.
- 55 Samson et al., (2006).
- 56 Atkinson (1995) and Coady, Grosh and Hoddinott (2004), page 80.
- 57 Pal et al., (2005), pages 14–26.
- 58 DFID 2000, page 1.
- 59 Ibid.
- 60 Despouy (1993).
- 61 Mitra (2005), page 5; Mitra (2004), page 4; Elwan (1999), pages 21–24.
- 62 Mitra (2005), page 6; Mitra (2004), page 4; Elwan (1999), pages 24–25.
- 63 This paragraph summarises key ideas from Mitra (2004, 2005) which provide greater detail on these issues.
- 64 Mitra (2004), page 17.
- 65 Natrass (2006), page 3.
- 66 UN IRIN (2005).
- 67 Mitra (2004), page 9.
- 68 Ibid., page 15.
- 69 Mitra (2005), page 10.
- 70 Subbarao (1996), page 16.
- 71 Mitra (2005), page 11.
- 72 Ibid., page 15.
- 73 Ibid., page 16.
- 74 Coady et al. (2004), page 47.
- 75 Bigman and Fofack (2002), page 129.
- 76 Rawlings and Rubio (2003), page 3.
- 77 Coady et al. (2004), pages 22–23.
- 78 Bigman and Fofack (2002), page 136.
- 79 Coady et al. (2004), page 29; Paxson and Schady (2002).
- 80 Coady et al. (2004), page 62.
- 81 Bigman and Srinivasan (2001), page 3.
- 82 Bigman and Fofack (2002), page 136.
- 83 Stigma can be defined as “the feeling of shame that may come from an open admission that one is poor and in need of help” (Grosh 1994). Stigma can also be externally reinforced, increasing the social costs for those affected.
- 84 Chirwa et al. (2004).
- 85 Devereux (2002b), pages 8–9; Gebre-Medhin and Swinton (2001).

- 86 Devereux (2002b), pages 9–10.
- 87 Conning and Kevane (2000), page 2.
- 88 Devereux (2002b), pages 9–10.
- 89 Conning and Kevane (2000), pages 2–3.
- 90 Devereux (2002b), page 10.
- 91 Conning and Kevane (2000), page 3.
- 92 Devereux (2002b), page 10.
- 93 Subbarao (2003), page 25; Galasso and Ravallion (1999); see also Conning and Kevane (2000).
- 94 Devereux (2002b), page 11.
- 95 Conning and Kevane (2000), page 27.
- 96 Sharp (1998), page 92, in Devereux (2002b), page 11.