Development Economics 1

Lecture 2: Social protection

Stefano Caria

October 24, 2025

I will partly rely on materials developed by Clement Imbert.

Outline of the lecture

1. What is social protection?
2. Who should get social protection?
3. Is targeting on deprivation optimal?

What is social protection?

- Three main goals
  - *Redistribution*: raise the consumption of the poor
    $\rightarrow$ Social assistance programs
  - *Insurance*: support those individuals hit by a shock
    $\rightarrow$ Social insurance programs
  - *Graduation*: support individuals to increase earnings
    $\rightarrow$ Labour market programs

- Many programs pursue a mix of these goals.

Different types of programs (World Bank 2018)

TABLE 1.1 Social Protection and Labor Market Intervention Areas

Social protection and labor programs	Objectives	Types of programs
Social safety nets/social assistance (noncontributory)	Reduce poverty and inequality	• Unconditional cash transfers • Conditional cash transfers • Social pensions • Food and in-kind transfers • School feeding programs • Public works • Fee waivers and targeted subsidies • Other interventions (social services)
Social insurance (contributory)	Ensure adequate living standards in the face of shocks and life changes	• Contributory old-age, survivor, and disability pensions • Sick leave • Maternity/paternity benefits • Health insurance coverage • Other types of insurance
Labor market programs (contributory and noncontributory)	Improve chances of employment and earnings; smooth income during unemployment	• Active labor market programs (training, employment intermediation services, wage subsidies) • Passive labor market programs (unemployment insurance, early retirement incentives)

Source: World Bank 2012.
Note: ASPIRE = Atlas of Social Protection: Indicators of Resilience and Equity.

About half of the world receives some form of social protection

FIGURE 3.2 Share of Total Population and the Poorest Quintile That Receives Any Social Protection and Labor Programs, as Captured in Household Surveys, by Country Income Group

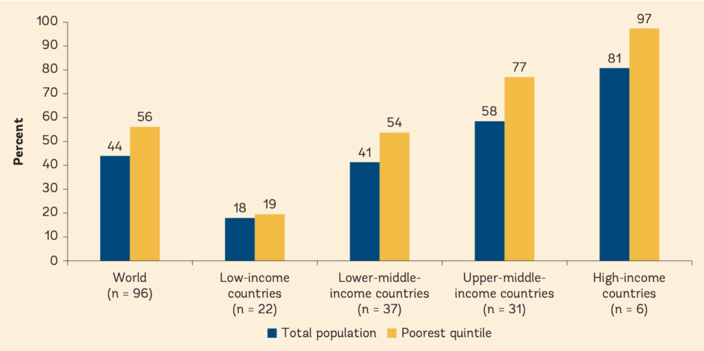

Source: ASPIRE database.
Note: The total number of countries per country income group included in the analysis appears in parentheses. Aggregated indicators are calculated using simple averages of country-level social protection and labor coverage rates across country income groups. Coverage is determined as follows: (number of individuals in the total population or poorest quintile who live in a household where at least one member receives the transfer)/(number of individuals in the total population). This figure underestimates total social protection and labor coverage because household surveys do not include all programs that exist in each country. The poorest quintile is calculated using per capita pre-transfer welfare (income or consumption). ASPIRE = Atlas of Social Protection: Indicators of Resilience and Equity.

LMICs have less social protection

FIGURE 3.3 Share of Poorest Quintile That Receives Any Social Protection and Labor Program, as Captured in Household Surveys, by Type of Social Protection and Labor Area and Country Income Group

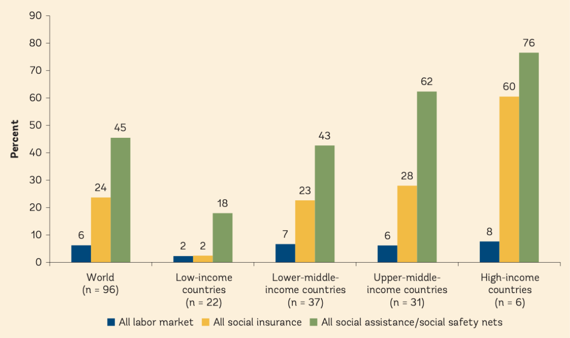

Source: ASPIRE database.
Note: The total number of countries per country income group included in the analysis appears in parentheses. Aggregated indicators are calculated using simple averages of country-level coverage rates for social insurance, social assistance, and labor market programs, across country income groups. Indicators do not count for overlap among programs types (people receiving more than one program); therefore, the sum of percentages by type of program may add up to more than 100 percent. Coverage is determined as follows: (number of individuals in the total population or poorest quintile who live in a household where at least one member receives the transfer)/(number of individuals in the total population). This figure underestimates total social protection and labor coverage because household surveys do not include all programs that exist in each country. The poorest quintile is calculated using per capita pretransfer welfare (income or consumption). ASPIRE = Atlas of Social Protection: Indicators of Resilience and Equity.

Fewer people in LMICs escape poverty due to social protection

FIGURE 3.27 Reductions in Poverty and Inequality from Social Safety Net Transfers, as Captured in Household Surveys, as a Share of Pretransfer Indicator Levels, by Country Income Group Using Relative Poverty Line

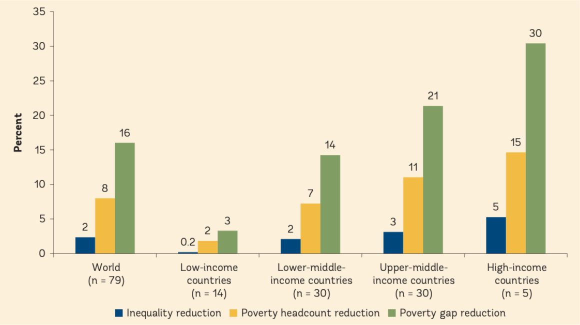

Source: ASPIRE database.
Note: The total number of countries per country income group included in the analysis appears in parentheses. This figure uses a relative measure of poverty defined as the poorest 20 percent of the welfare distribution (income or consumption). Impacts on poverty and inequality can be estimated only if monetary values are recorded in the household survey; for this reason, the sample of countries used in this figure is smaller than the one used to estimate coverage and beneficiary incidence. Percentages of poverty and inequality reduction are calculated as follows: (poverty headcount pretransfer – poverty headcount posttransfer)/(poverty headcount pretransfer). The same calculations apply for the Gini index and poverty gap percentage reductions. Aggregated indicators are calculated using simple averages of country-level percentage reduction of the indicator across country income groups. The reductions in poverty and inequality are underestimated because ASPIRE does not include data for every single country in the country income groups, and even for a given country the survey does not include all existing social safety net programs or provide monetary values for them. ASPIRE = Atlas of Social Protection: Indicators of Resilience and Equity.

In many LICs, social protection is entirely funded by foreign aid

FIGURE 2.2 Share of Donor-Funded Safety Nets in Sub-Saharan African Countries

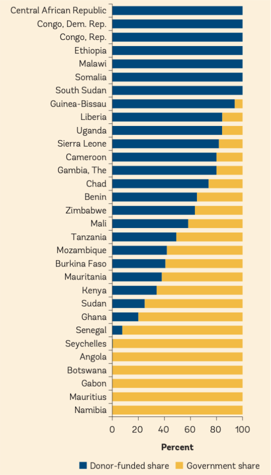

Source: Beegle, Coudouel, and Monsalve, forthcoming.

Key issues

- Who should get the transfer? *Information* problem + Normative problem
- Enjoys political support. *Acceptability* problem.
- Makes sure the poor get the transfer. *Enforcement* problem.
- Does not make the poor dependent. *Incentive* problem

In which way do these features look different in developing countries?

Specificities of LDCS

- Information: governments typically do not know households' income (informal sector, agriculture...).
  - Rely on surveys, community targeting or self-targeting.
- Acceptability: only rich minority with political power pays taxes + norms may differ in rural contexts.
  - Conditionality of transfers: CCT, public works.
  - Community targeting can help in rural contexts.
- Enforcement: governments rely on the bureaucracy to identify beneficiaries and transfer income.
  - Corruption + mismanagement $\Rightarrow$ leakages + mistargeting.
- Incentive problem: less of an issue in poorer settings (Crosta et al. 2025) but may take different forms.

Outline of the lecture

1. What is social protection?
2. Who should get social protection?
3. Is targeting on deprivation optimal?

The information problem: a simple model from Banerjee et al 2024

- A government wants to target a transfer to the poorest.

- The government observes:
  1. whether an individual is below the poverty line y^*. Income among the poor has density function $h(y)$
  2. a noisy, unbiased signal of their true income $y^e = y + \epsilon$;
  3. a self-report by each individual $\tilde{y}$.

- If individuals report an income different from y_e, they pay cost $F = \frac{a}{2}(y_e - \tilde{y})^2$

- The government forms a predicted income $y^p = \alpha\tilde{y} + (1-\alpha)y_e$

The information problem: a simple model from Banerjee et al 2024

- The government chooses:
  - A universal, untargeted transfer T to all poor households
  - A targeted transfer proportional to the individual poverty gap: $t(y^* - y^p)$
  - How much weight (α) to give to the self-report

- Social welfare is given by
  $\int_0^{y^*} g(y)h(y)E_\epsilon u(y + T + t(y^* - y^p))dy$
  - $g(y)$ is a welfare weight ($g'(y) < 0$).
  - $u()$ is the government concave utility function, which reflects concerns for equality

- The government has a budget of B.

Case 1: homogeneous misreporting costs a

The individual picks $\tilde{y}$ to maximise:

$$y + T + E_\epsilon[t(y^* - \alpha\tilde{y} - (1-\alpha)(y+\epsilon))] - \frac{a}{2}(y + \epsilon - \tilde{y})^2 \qquad (1)$$

$$\tilde{y}^* = y - \alpha t/a \qquad (2)$$

Trade-off between making the transfer more progressive and minimizing horizontal inequality

Suppose the government can only choose t, then

$$W'(t) = \int_0^{y^*} g(y)\, h(y)\, (\bar{y} - y)\, \mathbb{E}_\epsilon \big[u'\big(y + t(\bar{y} - y) + B - t(1-\alpha)\epsilon\big) \big]\; dy$$
$$- \int_0^{y^*} g(y)\, h(y)\, (1-\alpha)\, \mathbb{E}_\epsilon \big[\epsilon\, u'\big(y + t(\bar{y} - y) + B - t(1-\alpha)\epsilon\big) \big]\; dy. \tag{6}$$

First term: benefit from better targeting of the poor.
Second term: cost from greater horizontal inequality.

(Check discussion of the model in the Appendix. $\bar{y}$ is average income among the poor. (6) is obtained by substituting T by $B - t(y^* - \bar{y}) - \frac{\alpha^2 t^2}{a}$, which is implied by the budget constraint, and then differentiating by t).

Result 1

In the case with homogeneous costs, the optimal policy is

- $T = 0$
- $\alpha = 1$

Intuition:

- $t = 0$ is never optimal ($W'(0) > 0$).
- $\frac{\delta W}{\delta \alpha} > 0$: optimal to set $\alpha = 1$, so second term in (6) is 0
- If the second term is zero, $\frac{\delta W}{\delta t} > 0 \forall t$. So it is optimal to set $T = 0$ and to target all the transfer.

- This is the first best (a drops out from W, so the conclusion would hold for any value of a).

Heterogeneous costs, Result 2

Suppose now that there are two levels of a: $a_1 < a_2$.

If $\alpha > 0$, you transfer from those who are less willing (a_2) to those most willing to misreport (a_1) – horizontal inequality.

When you increase α, you reduce errors due to ϵ, but you increase the transfer from type a_2 to type a_1.

When you increase t, you increase the transfer from type a_2 to type 1 (and also errors due to ϵ if $\alpha < 1$).

Now it may be optimal to have

- $\alpha < 1$
- $T > 0$

How do $Var(\epsilon)$ and the difference between a_1 and a_2 shape the optimal α and T?

Figure A.1: Noise vs. heterogeneity

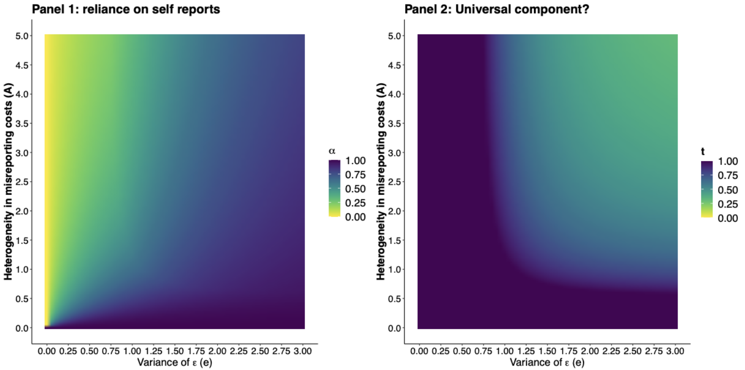

Note: This figure presents simulation results for equation 8 in Section A.2, over different values of A and $Var(\epsilon)$, and assuming a linear, CRRA utility function. To compute the values of α and t, we numerically optimise for alpha and t over each set of parameters. Optimal values of α and t are depicted by colours ranging from yellow (0) to dark blue (1).

How can the government obtain y_e?

- Three general methods:
  - Collect (direct or indirect) measures of poverty in a survey.
  - Ask community to decide who is poor.
  - Get poor households to *self-select* in the policy.
- These methods are often combined, e.g. initial survey + community updating.

Direct measures of poverty

Means testing.

- Developed world: easy to make transfers a function of income. LDCs: much harder.
- Surveys in LDCs ask individuals about consumption because income is variable/not well known.
- Individuals classified as poor if *consumption* below a threshold. Problems?
  - Costly: need long interviews with all households.
  - Frequency: do the survey every year?
  - Incentives: households may under-report consumption.

Governments in LDCs typically do not use *only* consumption data to determine eligibility.

Indirect measures of poverty

Proxy-means testing (PMT)

- Ownership of assets is highly correlated with poverty.
- Assets are easy to observe in a survey, and hard to hide.
- Indexes based on assets and household characteristics (*proxies* for poverty) often used to determine eligibility.

How does PMT perform in practice?

Figure 3

Predicted versus Actual per-capita Consumption for Households in Test Set Data

A: Indonesia

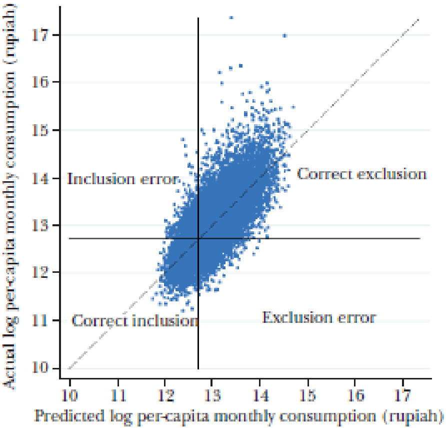

Source: Hanna and Olken (2018)

Community targeting

Idea: tell each community to determine which X households get a transfer.

- Advantages
  - Leverage local information: better measure (cf informal taxation)?
  - May capture local definitions of welfare better: more legitimacy?
  - Can adapt to change in circumstances: better insurance?
- Disadvantages
  - Elite capture/corruption as not rule based.
  - Ranking local households is complicated, demands local effort.
- Used in many contexts: Ethiopia's Food For Work program, Bangladesh's Food For Education program.

PMTs vs community targeting - Alatas et al (2012)

Indonesia, government wants to allocate a cash transfer to the poorest.

- RCT: some communities get a PMT, others community targeting, others a hybrid.
  - Community targeting: community members asked to rank households from richest to poorest during a meeting, poorest X households get the transfer.
- Survey to measure consumption and:
  - Households' relationship to community leaders - elite capture?
  - Households' own poverty ranking (self and others) - different preferences?
  - Household characteristics determining earning *potential*.

Alatas et al (2012) - community targeting in practice

PMT produces fewer errors than community or hybrid treatments

TABLE 3—RESULTS OF DIFFERENT TARGETING METHODS ON ERROR RATE BASED ON CONSUMPTION

		By income status		By detailed income status				
Sample:	Full population (1)	Inclusion error (2)	Exclusion error (3)	Rich (4)	Middle income (5)	Near poor (6)	Very poor (7)	Per capita consumption of beneficiaries (8)
Community treatment	0.031*	0.046**	0.022	0.028	0.067**	0.49	−0.013	9.933
	(0.017)	(0.018)	(0.028)	(0.021)	(0.027)	(0.038)	(0.039)	(18.742)
Hybrid treatment	0.029*	0.037**	0.009	0.020	0.052**	0.031	−0.008	−1.155
	(0.016)	(0.017)	(0.027)	(0.020)	(0.025)	(0.037)	(0.037)	(19.302)
Observations	5,753	3,725	2,028	1,843	1,882	1,074	954	1,719
Mean in PMT treatment	0.30	0.18	0.52	0.13	0.23	0.55	0.48	366

Notes: All regressions include stratum fixed effects. Robust standard errors in parentheses, clustered at the village level. All coefficients are interpretable relative to the PMT treatment, which is the omitted category. The mean of the dependent variable in the PMT treatment is shown in the bottom row. All specifications include stratum fixed effects.

***Significant at the 1 percent level.
**Significant at the 5 percent level.
*Significant at the 10 percent level.

But people are less satisfied with PMT allocations

TABLE 6—SATISFACTION

Panel A. Household endline survey

	Is the method applied to determine the targeted households appropriate? (1 = worst, 4 = best) (1)	Are you satisfied with the targeting activities in this subvillage in general? (1 = worst, 4 = best) (2)	Are there any poor HH that should be added to the list? (0 = no, 1 = yes) (3)	Number of HH that should be added to list (4)	Number of HH that should be subtracted from list (5)	*p*-value from joint test (6)
Community treatment	0.161***	0.245***	−0.189***	−0.578***	−0.554***	< 0.001
	(0.056)	(0.049)	(0.040)	(0.158)	(0.112)	
Hybrid treatment	0.018	0.063	0.020	0.078	−0.171	0.762
	(0.055)	(0.049)	(0.042)	(0.188)	(0.129)	
Observations	1,089	1,214	1,435	1,435	1,435	
Mean in PMT treatment	3.243	3.042	0.568	1.458	0.968	

Panel B. Subvillage head endline survey

Panel C. Comment forms and fund disbursement results

	Number of comments in the comment box	Number of complaints in the comment box	Number of complaints received by subvillage head	Did facilitator encounter any difficulty in distributing the funds? (0 = no, 1 = yes)	Fund distributed in a meeting (0 = no, 1 = yes)	
Community treatment	−0.944	−1.085***	−2.684***	−0.062***	0.082**	0.0014
	(0.822)	(0.286)	(0.530)	(0.023)	(0.038)	0.177
Hybrid treatment	−0.364	−0.554**	−2.010***	−0.045*	0.051	
	(0.821)	(0.285)	(0.529)	(0.026)	(0.038)	
Observations	640	640	640	621	614	
Mean in PMT treatment	11.392	1.694	4.34	0.135	0.579	

Notes: All estimation is by OLS with stratum fixed effects. Using ordered probit for multiple response and probit models for binary dependent variables produces the same signs and statistical significance as the results shown. These results are available from the authors upon request.

No evidence of elite capture

TABLE 7—ELITE TREATMENTS

	Attendance (survey data)	Full sample error rate	Full sample error rate		On beneficiary list	
	(1)	(2)	(3)	(4)	(5)	(6)
Community treatment	0.367***	0.029	0.033	0.048*	0.042*	0.054*
	(0.038)	(0.018)	(0.023)	(0.025)	(0.025)	(0.028)
Hybrid treatment	0.370***	0.027	0.024	0.008	0.025	0.012
	(0.037)	(0.018)	(0.022)	(0.024)	(0.022)	(0.023)
Elite subtreatment	−0.301***	0.004	0.016	−0.013	−0.015	−0.039
	(0.034)	(0.016)	(0.020)	(0.029)	(0.021)	(0.032)
Elite × hybrid				0.062		0.051
				(0.041)		(0.043)
Elite connectedness			−0.025	−0.025	−0.063***	−0.063***
			(0.021)	(0.021)	(0.021)	(0.021)
Elite connectedness × community treatment			−0.015	−0.013	−0.067**	−0.078**
			(0.035)	(0.038)	(0.033)	(0.036)
Elite connectedness × hybrid treatment			0.010	0.010	−0.013	−0.001
			(0.033)	(0.035)	(0.033)	(0.035)
Elite connectedness × elite treatment			−0.029	−0.034	0.041	0.064
			(0.031)	(0.047)	(0.030)	(0.042)
Elite connectedness × elite treatment × hybrid				0.003		−0.047
				(0.063)		(0.060)
Observations	287	5,753	5,753	5,753	5,756	5,756
Mean in PMT treatment	0.11	0.30	0.30	0.30	0.28	0.28

Notes: In column 1, an observation is a village and the dependent variable is the share of households surveyed in the endline survey where at least one household member attended a targeting meeting. The PMT mean in column 3 is not zero, because the question was worded generically to be about any targeting meeting, not just meetings associated with our project. The dependent variable in columns 2–4 is the dummy for error in targeting based on consumption, as in column 1 of Table 3. Dependent variable in columns 5 and 6 is a dummy for being a beneficiary of the program. All specifications in columns 3–6 include dummies for the community, hybrid, and elite treatment main effects, as well as stratum fixed effects; columns 4 and 6 also include a dummy for elite × hybrid. Robust standard errors are in parentheses, and standard errors are adjusted for clustering at the village level in columns 2–6. All specifications include stratum fixed effects.

***Significant at the 1 percent level.
**Significant at the 5 percent level.
*Significant at the 10 percent level.

Community treatment ranks less correlation with consumption ranks, but more correlated with self-perception ranks

TABLE 9—ASSESSING TARGETING TREATMENTS USING ALTERNATIVE WELFARE METRICS

	Consumption (r_g) (1)	Community survey ranks (r_c) (2)	Subvillage head survey ranks(r_e) (3)	Self-assessment (r_s) (4)
Community treatment	−0.065** (0.033)	0.246*** (0.029)	0.248*** (0.038)	0.102*** (0.033)
Hybrid treatment	−0.067** (0.033)	0.143*** (0.029)	0.128*** (0.038)	0.075** (0.033)
Observations	640	640	640	637
Mean in PMT treatment	0.451	0.506	0.456	0.343

Notes: The dependent variable is the rank correlation between the treatment outcome (i.e., the rank ordering of households generated by the PMT, community, or hybrid treatment) and the welfare metric shown in the column, where each observation is a village. Robust standard errors are shown in parentheses.

***Significant at the 1 percent level.
**Significant at the 5 percent level.
*Significant at the 10 percent level.

What is the community maximising? Earning capacity and ability to cope with shocks

Table 12—What is the Community Maximizing?

	Rank according to welfare metric			Targeting rank list in		
	Community survey ranks (r_c) (1)	Subvillage head survey ranks(r_e) (2)	Self-assessment (r_s) (3)	PMT villages (4)	Community villages (5)	Hybrid villages (6)
Log per capita consumption	0.176***	0.145***	0.087***	0.132***	0.197***	0.162***
	(0.008)	(0.008)	(0.004)	(0.013)	(0.014)	(0.014)
Panel A. Household demographics						
Log HH size	0.164***	0.134***	0.073***	−0.028	0.154***	0.078***
	(0.011)	(0.010)	(0.006)	(0.019)	(0.019)	(0.021)
Share kids	−0.125***	−0.094***	−0.037***	−0.296***	−0.068*	−0.141***
	(0.021)	(0.021)	(0.012)	(0.035)	(0.041)	(0.039)
Panel B. Ability to smooth shocks						
Elite connected	0.092***	0.044***	0.025***	0.062***	0.051***	0.043***
	(0.008)	(0.009)	(0.005)	(0.016)	(0.015)	(0.015)
Total connectedness	−0.039***	−0.021**	−0.015***	−0.016	−0.019	−0.054***
	(0.010)	(0.009)	(0.005)	(0.017)	(0.017)	(0.019)
Number of family members outside subvillage	0.012***	0.010***	0.006***	0.020***	0.001	0.001
	(0.004)	(0.003)	(0.002)	(0.006)	(0.006)	(0.006)
Participation through work to community projects	0.002	0.021**	0.005	0.000	0.010	0.003
	(0.011)	(0.010)	(0.006)	(0.018)	(0.019)	(0.019)
Participation through money to community projects	0.061***	0.041***	0.024***	0.056***	0.058***	0.034*
	(0.009)	(0.009)	(0.005)	(0.016)	(0.016)	(0.018)
Participation in religious groups	0.027***	0.033***	0.014**	0.033**	0.012	0.029
	(0.010)	(0.010)	(0.006)	(0.016)	(0.017)	(0.017)
Total savings	0.000	0.000	0.000	0.000	0.000	0.000
	(0.000)	(0.000)	(0.000)	(0.000)	(0.000)	(0.000)
Share of savings in a bank	0.096***	0.069***	0.052***	0.121***	0.103***	0.075***
	(0.011)	(0.010)	(0.006)	(0.018)	(0.021)	(0.020)
Debt as share of consumption	0.005***	0.004***	0.002***	0.002	0.007***	0.008***
	(0.001)	(0.001)	(0.000)	(0.002)	(0.001)	(0.001)
Panel C. Discrimination against minorities?						
Ethnic minority	−0.024*	−0.019	−0.003	0.012	−0.051**	−0.011
	(0.014)	(0.014)	(0.008)	(0.026)	(0.025)	(0.024)
Religious minority	0.012	−0.007	−0.014*	−0.018	0.025	0.012
	(0.018)	(0.017)	(0.008)	(0.030)	(0.032)	(0.033)
Panel D. Correcting for earnings ability						
HH head with primary education or less	−0.028***	−0.025***	−0.037***	−0.108***	−0.011	−0.066***
	(0.009)	(0.009)	(0.005)	(0.017)	(0.018)	(0.017)
Widow	−0.104***	−0.083***	−0.012	0.009	−0.108***	−0.026
	(0.014)	(0.014)	(0.008)	(0.027)	(0.024)	(0.028)
Disability	−0.045***	−0.037***	−0.026***	−0.079***	0.009	0.012
	(0.016)	(0.014)	(0.008)	(0.027)	(0.026)	(0.027)
Death	−0.041*	−0.031	−0.010	−0.111***	−0.013	−0.059
	(0.025)	(0.025)	(0.015)	(0.042)	(0.048)	(0.043)
Sick	−0.038***	−0.041***	−0.028***	0.007	−0.018	−0.044**
	(0.011)	(0.011)	(0.006)	(0.018)	(0.019)	(0.019)
Recent shock to income	−0.001	−0.005	−0.013**	−0.019	0.009	−0.012
	(0.009)	(0.009)	(0.005)	(0.016)	(0.016)	(0.017)
Tobacco and alcohol consumption	−0.0002***	−0.0002***	−0.0001***	−0.0002***	−0.0002***	−0.0001***
	(0.000)	(0.000)	(0.000)	(0.000)	(0.000)	(0.000)
Observations	5,337	4,680	5,724	1,814	1,876	1,889

Self-Targeting

Households are very good at identifying their own rank

- If households can tell whether or not they're above the poverty line, could ask them to **self-select**.
- Problem: all households want the transfer, will apply even they think it's unlikely they'll get it.
- Solution: offer a transfer that only the poor would want.
  - Something that the poor consume more of: basic food?
  - Something that uses something the poor value less than the rich: time (opportunity cost of time $=$ wage).
- 'Ordeals' (long administrative procedures) are typically used to target the poor/unemployed even in rich countries.
- Self targeting always distort households' behavior - imposes a cost on the poor (over-consumption of rice, or queuing).
- Is this an effective way to target the poor?

Self targeting: Alatas et al. (2016)

Test the efficacy of ordeals to target the poor:

- Same context as Alatas et al. (2012) but different cash transfer (conditional cash transfer PKH).
- Status-quo: government carries out an asset survey and automatically enrol beneficiaries who pass the PMT.
- Self-targeting intervention: beneficiaries need to travel a few km to enrol in the programme and take the test.

Conclude that as compared to automatic targeting, self-targeting:

- Reduces the number of non-poor whose eligibility is tested (cost-saving for government)
- Reduces inclusion errors, i.e. non-poor (based on consumption) who receive benefits.
- Does not increase exclusion errors, i.e. poor people not applying / not receiving benefits.

Outline of the lecture

1. What is social protection?
2. Who should get social protection?
3. Is targeting on deprivation optimal?

In the previous framework, households are passive recipients of transfers.

In practice, transfers are partly invested, resulting in different outcomes for different households.

$\rightarrow$ To maximise social welfare, we need to care both about the counterfactual outcomes of those receiving transfers ('deprivation'), and the effects of those transfers ('impact').

A recent paper by Haushofer et al. 2025 makes progress on this issue

Same setting in rural Kenya as Egger et al 2022.

Households receive a large cash transfer.

Has detailed baseline and endline data to study effects on household consumptions, income, assets.

A simple framework: The social welfare function

$$\sum_{h} \sum_{t=0}^{\bar{t}} \omega_h \delta^t \, W\big(Y_{h,t}(T_h)\big) \tag{3}$$

$$\sum_{h} \sum_{t=0}^{\bar{t}} W\big(Y_{h,t}(T_h)\big) \; = \; \sum_{h} \sum_{t=0}^{\bar{t}} W\big(Y^0_{h,t} + T_h \cdot \Delta_{h,t}\big) \tag{4}$$

$Y^0_{h,t}$ is the outcome of household h in the absence of treatment.
$\Delta_{h,t}$ is the effect of treatment T_h.

For simplicity, the second equation drops weights ω, and sets $\delta = 1$

What information does the government have?

$$\hat{Y}^0_t(X_h) \ \text{ of } \ \mathbb{E}\left[\, Y^0_{h,t} \mid X_h, t \,\right], \tag{5}$$

$$\hat{\Delta}_t(X_h) \ \text{ of } \ \mathbb{E}\left[\, Y^1_{h,t} - Y^0_{h,t} \mid X_h, t \,\right]. \tag{6}$$

Who should be targeted?

$$d\hat{W}(X_h) \equiv \sum_{t=0}^{\bar{t}} \Big[W\big(\hat{Y}_t^0(X_h) + \hat{\Delta}_t(X_h)\big) - W\big(\hat{Y}_t^0(X_h)\big)\Big], \tag{7}$$

$$\mathcal{R}^*(X_h) \equiv 1\Big\{ d\hat{W}(X_h) \geq q_{1-\phi}^{d\hat{W}} \Big\}. \tag{8}$$

Two alternative rules

$$\mathcal{R}^I(X_h) = 1\left\{\hat{\Delta}(X_h) \geq q_{1-\phi}^{\hat{\Delta}}\right\}. \tag{9}$$

$$\mathcal{R}^D(X_h) = 1\left\{\hat{Y}_0^0(X_h) \leq q_{\phi}^{\hat{Y}}\right\}. \tag{10}$$

Empirical approach

Assume CARA utility, and explore how conclusion changes as α parameter changes:

$$W(\hat{y}) = \begin{cases} \frac{1 - e^{-\alpha \hat{y}}}{\alpha}, & \alpha \neq 0, \\ \hat{y}, & \alpha = 0. \end{cases} \tag{11}$$

Obtain $E[Y_{h,t}^0 | X_h, t]$ with random forest. Obtain $E[Y_{h,t}^1 - Y_{h,t}^0 | X_h, t]$ with generalized random forests.

Simulate different rules, with $\phi = 50$.

Negative correlation between deprivation and impact

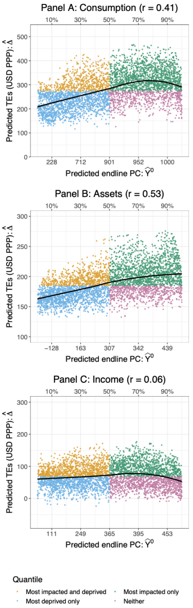

Optimal targeting under different preferences

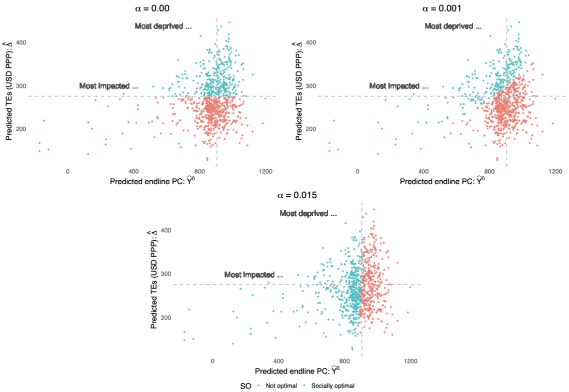

FIGURE 2. PLOTTING PREDICTED DEPRIVATION VERSUS IMPACT BY SOCIALLY OPTIMAL STATUS

Notes: This figure plots the predicted endlines and predicted treatment effects for households of 4 members (the median size). Socially optimal groups are highlighted for different curvature values using CARA. Dashed lines denote the thresholds for the most impacted and most deprived households. For exposition purposes, socially optimal households were selected without cross-fitted thresholds, using integrated predictions across quarters (static models), for households of the same size. A constant was added to the predicted endline outcomes so that the overall predicted mean matches the observed sample mean, since GRF models were trained with time-demeaned data. Monetary values are in USD PPP (2016).

Targeting solely on deprivation not optimal for reasonable curvatures

TABLE 1—MAIN SOCIAL WELFARE ANALYSIS: OVERLAP OF SOCIALLY OPTIMAL HOUSEHOLDS TO TARGET WITH MOST DEPRIVED AND MOST IMPACTED

	(1)	(2)	(3)	(4)	(5)
		Most deprived		Most impacted	
	CE	Share	p-val	Share	p-val
			$D>0.95$		$I>0.95$
Panel A: Consumption, CARA					
$\alpha = 0.0000$	\$50	0.26	0.00	1.00	1.00
$\alpha = 0.0005$	\$49	0.43	0.00	0.79	0.00
$\alpha = 0.0010$	\$49	0.54	0.00	0.68	0.00
$\alpha = 0.0075$	\$41	0.91	0.01	0.35	0.00
$\alpha = 0.0150$	\$33	0.96	0.31	0.29	0.00
Panel B: Other welfare measures, CARA					
Assets, $\alpha = 0.0010$	\$49	0.63	0.00	0.61	0.00
Income, $\alpha = 0.0010$	\$49	0.54	0.00	0.89	0.23
Panel C: Sensitivity checks on consumption					
CRRA, $\rho = 0.5$		0.36	0.00	0.87	0.04
CRRA, $\rho = 2$		0.53	0.00	0.68	0.00
Time discounting: $\beta = 15\%$, $\alpha = 0.0001$		0.54	0.00	0.67	0.00
Re-targeting dynamics, $\alpha = 0.001$		0.53	0.00	0.68	0.00
Pareto weights, $\alpha = 0.0005$		0.52	0.00	0.67	0.00
Saez-Stantcheva (2016), $\alpha = 0.0005$		0.52	0.00	0.67	0.00

Notes: Column 1 denotes the certainty equivalent (CE) of a 50-50 lottery over \$0 or \$100 under the specified CARA α parameter value. Column 2 (4) reports the share of households belonging to D (I) that are also "socially optimal" (those in the top 50% of households ranked by potential gains from treatment) for a planner to treat for a given utility function (CARA or CRRA) and parameter value (α or ρ). Reported shares are the mean of 150 5-fold GRF iterations; median ratios are similar (not shown). Columns (3) and (5) report p-values testing whether a planner would prefer to predominately target only the most deprived (D) or the most impacted (I). Panel C presents a variety of sensitivity analyses. For additional sensitivity checks, parameter values, and outcomes see Appendix tables A.1 (assets and income), A.2 (CRRA), A.3 (observable assets), C.1 (OLS and LASSO based prediction models), D.1 (additional robustness checks), and G.2 (pareto weights).

Reading

(*) Banerjee et al., 2024 Social Protection in the Developing World, Journal of Economic Literature.
→ Especially *Section 2.2* and Appendix A

(*) Alatas et al. 2012, Targeting the Poor: Evidence from a Field Experiment in Indonesia, Journal of Political Economy

(*) Haushofer et al. 2025, Targeting Impact or Deprivation American Economic Review

Alatas et al. 2016, Self-targeting: Evidence from a Field Experiment in Indonesia, Journal of Political Economy

World Bank, 2018, The State of Social Safety Nets