

Jones, Sam; Massingue, Telça; Muchabel, Hilário

Working Paper

When (not) to target social programs: The simple arithmetic of the effectiveness of scale-up

WIDER Working Paper, No. 80/25

Provided in Cooperation with:

United Nations University (UNU), World Institute for Development Economics Research (WIDER)

Suggested Citation: Jones, Sam; Massingue, Telça; Muchabel, Hilário (2025) : When (not) to target social programs: The simple arithmetic of the effectiveness of scale-up, WIDER Working Paper, No. 80/25, ISBN 978-92-9267-639-1, The United Nations University World Institute for Development Economics Research (UNU-WIDER), Helsinki, <https://doi.org/10.35188/UNU-WIDER/2025/639-1>

This Version is available at:

<https://hdl.handle.net/10419/331526>

Standard-Nutzungsbedingungen:

Die Dokumente auf EconStor dürfen zu eigenen wissenschaftlichen Zwecken und zum Privatgebrauch gespeichert und kopiert werden.

Sie dürfen die Dokumente nicht für öffentliche oder kommerzielle Zwecke vervielfältigen, öffentlich ausstellen, öffentlich zugänglich machen, vertreiben oder anderweitig nutzen.

Sofern die Verfasser die Dokumente unter Open-Content-Lizenzen (insbesondere CC-Lizenzen) zur Verfügung gestellt haben sollten, gelten abweichend von diesen Nutzungsbedingungen die in der dort genannten Lizenz gewährten Nutzungsrechte.

Terms of use:

Documents in EconStor may be saved and copied for your personal and scholarly purposes.

You are not to copy documents for public or commercial purposes, to exhibit the documents publicly, to make them publicly available on the internet, or to distribute or otherwise use the documents in public.

If the documents have been made available under an Open Content Licence (especially Creative Commons Licences), you may exercise further usage rights as specified in the indicated licence.

When (not) to target social programs

The simple arithmetic of the effectiveness of scale-up

Sam Jones,¹ Telça Massingue,² and Hilário Muchabel³

November 2025

Abstract: We develop a framework to assess when social programmes should prioritize expanding coverage versus improving targeting accuracy. Rather than treating universalism and targeting as a binary choice, we derive a practical rule-of-thumb based on observed eligibility rates and programme quotas to indicate when marginal welfare gains from expanding programme access may outweigh those from improved targeting. Applying the framework to Mozambique's national social pension program, we find that zero targeting effort may be optimal in over 60% of administrative areas, and reallocating existing resources using a binary RESbased rule could raise welfare by 13%. Our results support a differentiated approach to targeting: prioritize programme expansion in underserved high-poverty (rural) areas, and improve targeting effort where poverty is much lower (urban areas).

Key words: social protection, cash transfers, targeting, universalism, welfare optimisation, Mozambique

JEL classification: I38, H53, D61

Acknowledgements: The study could not have been conducted without the support of the Mozambican Ministry of Economy and Finance (MEF), the Ministry of Gender, Children and Social Action (MGCAS), as well as the National Institute for Social Action (INAS). Thanks to Kalle Hirvonen, Ravi Kanbur, Rachel Sabates-Wheeler, Veronika Wodsak and participants at the 2025 UNU-WIDER Annual Conference for constructive feedback.

¹ UNU-WIDER, Maputo, Mozambique, corresponding author: jones@wider.unu.edu; ² Centre for Economics and Management Studies, Eduardo Mondlane University, Maputo, Mozambique; ³ World Bank, Maputo, Mozambique.

This study is published within the project **Inclusive growth in Mozambique – scaling up research and capacity** implemented in collaboration between UNU-WIDER, University of Copenhagen, University Eduardo Mondlane, and the Mozambican Ministry of Economy and Finance. The project is financed through specific programme contributions by the governments of Sweden, and Switzerland.

Copyright © UNU-WIDER 2025

UNU-WIDER employs a fair use policy for reasonable reproduction of UNU-WIDER copyrighted content—such as the reproduction of a table or a figure, and/or text not exceeding 400 words—with due acknowledgement of the original source, without requiring explicit permission from the copyright holder.

Information and requests: publications@wider.unu.edu

ISSN 1798-7237 ISBN 978-92-9267-639-1

<https://doi.org/10.35188/UNU-WIDER/2025/639-1>

United Nations University World Institute for Development Economics Research provides economic analysis and policy advice with the aim of promoting sustainable and equitable development. The Institute began operations in 1985 in Helsinki, Finland, as the first research and training centre of the United Nations University. Today it is a unique blend of think tank, research institute, and UN agency—providing a range of services from policy advice to governments as well as freely available original research.

The Institute is funded through income from an endowment fund with additional contributions to its work programme from Finland and Sweden, as well as earmarked contributions for specific projects from a variety of donors.

Katajanokanlaituri 6 B, 00160 Helsinki, Finland

The views expressed in this paper are those of the author(s), and do not necessarily reflect the views of the Institute or the United Nations University, nor the programme/project donors.

1 Introduction

Social policy always operates under second-best conditions. The theoretical ideal of perfectly efficient, lump-sum transfers conditioned on individual ability is precluded by pervasive information and administrative constraints (Kanbur 1987; Mirrlees 1971). Such challenges are typically more severe in low- and middle-income contexts, where state capabilities and information infrastructures are comparatively weak. This has generated a long-standing debate between the merits of universal versus targeted assistance in the design of social programs, such as those to alleviate poverty.

The case for universal programmes draws on a range of ethical, practical and political economy considerations (Banerjee et al. 2019; Besley and Kanbur 1991; Mkandawire 2005; Schüring and Gassmann 2016; Sen 1995). Scholars have contended that targeted programmes are often politically fragile, lacking a broad base of social and political support necessary for effective performance. Targeting can be administratively complex and expensive to implement, potentially diverting resources from benefit payments to overheads. Also, evidence suggests that targeting often performs poorly in practice, generating material exclusion errors (e.g., Brown et al. 2018), while sometimes creating social divisions and stigma (e.g., Ellis 2012). Indeed, following Banerjee (1997), the finer the targeting, the greater the scope there may be for contestation and corruption in delivery.

On the other hand, a prevailing view holds that targeting is not only technically feasible but also cost effective (Banerjee et al. 2024; Hanna and Olken 2018). A key argument here is that universal programmes in resource-constrained settings are likely to involve (at best) moderate transfers to those in need, entailing low average impacts on welfare. At the same time, targeting is viewed as a fairly low-cost and sufficiently accurate tool, at least to exclude a considerable swathe of the non-eligible. Consequently, a large literature has investigated *how* best to target, developing and testing methods such as proxy means tests, ordeals and community targeting (see for instance Alatas et al. 2016; Devereux et al. 2017; Hanna and Karlan 2017).

Existing academic literature has generally framed the contrast between universalism and targeting in abstract binary terms.¹ However, from a practical standpoint, social programmes are rarely established at scale from scratch. Whilst universalism may be a long-run aspiration, social transfer programmes often evolve from (collections of) smaller programs, which expand and adjust gradually as resources, political priorities, and administrative capacities alter (see Slater 2023). For instance, as Lindert et al. (2007) describe, Brazil's flagship *Bolsa Família* programme was formed from the consolidation and harmonization of four separate cash trans-

¹ This is often implicit, as in Hoddinott (2007). Similarly, a conventional empirical framing contrasts a given targeting approach to a universal programme (e.g., Houssou and Zeller 2011).

fer schemes, the aim being to reduce what were ‘considerable gaps and duplications in coverage’. Thus, a key question for policymakers is not a binary one of whether or not to target, but a pragmatic one of resource allocation. Given existing funds and capacities, *when* is it more effective to invest in improving the accuracy of existing targeting schemes, versus simply scaling-up to reach more people?

This study seeks to answer this question. We begin by developing a simple analytical framework to clarify the trade-offs between scaling-up programme coverage (quantity expansion) and improving targeting accuracy (quality enhancement). To do so we present four propositions that link observed rates of poverty and programme quotas to the welfare effectiveness of alternative social programme designs. First, we show that targeting metrics, such as the share of eligible beneficiaries and coverage of the eligible, are interdependent, and algebraically constrained by observable parameters. Second, we derive analytical bounds on targeting performance metrics. These bounds tighten under minimal assumptions about non-regressive targeting and—critically—provide operational benchmarks even when detailed targeting data is not available (as is often the case).

Third, drawing on these preliminaries, we introduce a diagnostic rule-of-thumb, denoted the Relative Effectiveness of Scale-up (RES). This quantifies the marginal welfare gain from expanding coverage relative to improving targeting accuracy and can be computed directly from known values of poverty rates and programme quotas, thus serving as a guide for resource allocation decisions. Fourth, we extend the analysis to a stylized budget-constrained optimization problem, where targeting effort is costly and subject to diminishing returns. We show that optimal targeting effort is not only decreasing in the unit cost of effort but also in the poverty rate. As such, the model predicts that zero targeting effort, such as by random selection, can be optimal in high-poverty, low-capacity (high targeting cost) contexts.

We then apply this framework to Mozambique’s national social pension, the elderly component of the *Programa de Subsídio Social Básico* (PSSB). Combining administrative beneficiary data and 100% census microdata at the sub-district level, we compute poverty and quota rates for each administrative post. Our empirical results reveal substantial geographic variation. For approximately 60% of the relevant elderly population, the RES metric exceeds unity, implying that simple expansion of programme coverage is expected to yield higher welfare returns than improving targeting accuracy. Simulating optimal targeting effort under reasonable cost assumptions, we find that zero targeting effort is optimal in a majority of locations, especially in high-poverty rural areas. Counterfactual simulations show that reallocating existing programme budgets using a simple binary RES-based rule could raise welfare by 13%, without increasing total expenditure.

These results underscore the importance of going beyond a binary framing of the choice between universalism and targeting. Our contribution is to offer an approach to determine *when* investments in targeting are more or less likely to be a priority, providing a decision-making tool for social planners facing resource constraints and imperfect information. This builds on various strands of large existing literature, including models of targeting under imperfect information (e.g., Besley and Coate 1992; Lanjouw and Ravallion 1999; Ravallion and Chao 1989), debates on the political and operational feasibility of effective large-scale social programmes (e.g., Asri et al. 2022; Gelbach and Pritchett 2001), as well as estimates of the initial and recurrent costs of targeting (Coady et al. 2004; White et al. 2015).

More specifically, the paper connects to a general yet under-developed intuition that investments in targeting are unlikely to be economical in all circumstances. For example, White (2017) concludes that: “Whether to introduce a universal benefit or targeted scheme is ultimately a political decision. However, there are technical criteria to take into account Since targeting is not costless, the most obvious criterion is where poverty incidence is high so it is not worth the cost of targeting.” (p. 158). The main innovation of this study lies in formalizing this intuition, which recognizes the presence of a quantity-quality trade-off in social programme design, and presenting a simple yet tractable welfare-based analytical framework anchored on observable parameters. Operationalized via the RES metric, our approach introduces a criterion to evaluate the marginal effectiveness of targeting versus scale-up in real-world settings, the merit of which is demonstrated via the Mozambican case.

2 The arithmetic of targeting

2.1 Set-up

To fix ideas in a generic set-up, we model a poverty-targeted programme that allocates fixed per capita transfers of value t within a population of size N . Of these, Q^* individuals are deemed to be eligible (hereafter “the poor”), and Q are selected as actual beneficiaries by some administrative mechanism. The coverage rate or quota thus is $q = Q/N$, and the poverty rate is $q^* = Q^*/N$. Both are assumed to be known or estimated with reasonable precision at an aggregate spatial level (e.g., by administrative localities).

The probability of being selected as a beneficiary can be defined in two complementary ways. First we have:

$$\begin{aligned}
P(\text{Beneficiary}) &= P(\text{Beneficiary \& Eligible}) + P(\text{Beneficiary \& Non-eligible}) \\
&= P(\text{Eligible} \mid \text{Beneficiary}) \cdot P(\text{Beneficiary}) \\
&\quad + P(\text{Non-eligible} \mid \text{Beneficiary}) \cdot P(\text{Beneficiary}) \\
&= \alpha \cdot q + (1 - \alpha) \cdot q
\end{aligned} \tag{1}$$

Note this is just a weighted average of the unconditional selection probability, $P(\text{Beneficiary}) = q$, with weights given by the targeting success rate: $\alpha = P(\text{Eligible} \mid \text{Beneficiary})$, and the respective leakage or Type I error rate: $1 - \alpha$.

The second definition of the same selection probability is:

$$\begin{aligned}
P(\text{Beneficiary}) &= P(\text{Beneficiary} \mid \text{Eligible}) \cdot P(\text{Eligible}) \\
&\quad + P(\text{Beneficiary} \mid \text{Non-eligible}) \cdot P(\text{Non-Eligible}) \\
&= \lambda \cdot q^* + \gamma \cdot (1 - q^*)
\end{aligned} \tag{2}$$

Here, λ captures the probability that a given eligible individual is a beneficiary (responding to the question: ‘of all the poor, what share actually receives benefits?’), while γ captures the share of the non-eligible group incorrectly covered by the program.

Last, we consider an objective function facing policymakers. In an ideal world, a benign policy-maker may seek to maximize the total change in social welfare associated with the program, given by the sum of corresponding individual-specific utility gains. However, utility or close proxies, such as income or consumption, are rarely observed across the population. In many low-income contexts, for instance, informal economic activities are predominant, relevant data on these proxies are missing, outdated or subject to large measurement error, and incomes contain a large stochastic component (e.g., Beuermann et al. 2025).

Thus, we define a policy-relevant or operational maximand (Z) as the share of the population covered by the programme who are both eligible and selected as beneficiaries:

$$Z = \alpha \cdot q \tag{3}$$

This maximand focuses attention on the two primary levers available to policymakers: changes to the targeting success rate (α), which effectively proxies for the welfare impact of the programme given a fixed set of beneficiaries; and changes to the quota (q), which refers to the size of the set of beneficiaries. The relevant question is when one of these levers should be prioritized over the other.

While the above welfare metric is intuitive and operational, it is not innocuous. As noted by Coady and Skoufias (2004), α can be interpreted as a normalized measure of the welfare

impact of a fixed (uniform) transfer when all individuals below the poverty line take a welfare weight of one, and all the non-poor take a welfare weight of zero. Thus, our measure implicitly assumes each eligible individual contributes equally to social welfare, and that all transfers have the same marginal impact, regardless of the recipient's depth of poverty. Equally, the metric treats both leakage (benefits going to ineligible individuals) and exclusion errors (failure to reach eligible individuals) symmetrically in welfare terms, effectively ignoring the costs of type I errors. In reality, one might wish to assign different weights to these types of targeting errors. For instance, to the extent that social or political costs of including individuals above the eligibility threshold are material, this would require greater weight to be placed on targeting accuracy (see further below).

2.2 Propositions on targeting

We present three propositions speaking to the trade-off between quantity expansion and targeting quality enhancement in social programme design. In doing so, we maintain two assumptions: rates of programme coverage (q) and formal eligibility (q^*) in the population are known at meaningful levels of (dis)aggregation; but the policy-maker does not have accurate information on targeting performance. Our view is that it is precisely in such circumstances of imperfect information that robust operational guidance on the direction of social programme reform would be helpful.

Proposition 2.1. *The targeting identity*

In any targeting scheme, the number of beneficiaries who are eligible must equal the number of eligibles who are beneficiaries: $q^\lambda \equiv q\alpha$.*

Reflecting the equivalence between the first terms on the RHS of equations (1) and (2), this says that the product of the poverty rate and eligibility coverage rate is equal to the product of the quota and the targeting success rate.² An immediate corollary is that the ratio of targeting success to eligibility coverage (a ratio of two unknowns) equals the ratio of two knowns, namely that of the poverty rate and the quota:

$$\alpha/\lambda \equiv q^*/q \quad (4)$$

The implication is that targeting errors are mechanically interdependent: if coverage is below the poverty rate ($q < q^*$), under-coverage necessarily exceeds leakage.

² See Appendix Table B1 for illustration.

As an aside, note that the Coady-Grosh-Hoddinott (CDH) index (Coady et al. 2004) provides an alternative perspective on the targeting identity (equation 4). Proposed as a general measure of the extent to which a programme targets the poorest (progressiveness), it compares a metric of the share of benefits going to the poorest (the eligible) versus the share expected under a distribution-neutral benchmark. While this measure ideally defines coverage in terms of the total value of transfers, if programme benefits do not vary substantially across individuals (e.g., cash transfers are uniform) then the index can be approximated by the ratio of targeting success to the poverty rate:

$$\text{CDH} \approx \frac{\alpha}{q^*} = \frac{\lambda}{q} \quad (5)$$

Echoing analysis to come, this ratio equals unity under a benchmark assumption that selection is non-regressive and also suggests that absolute improvements in targeting performance will be valued in inverse proportion to the poverty rate.³

Proposition 2.2. *Bounds on targeting performance*

For any feasible allocation of programme benefits across eligible and non-eligible groups, the targeting performance metrics must satisfy:

$$\max \left(0, \frac{q + q^* - 1}{q} \right) \leq \alpha \leq \min \left(1, \frac{q^*}{q} \right) \quad (6)$$

$$\max \left(0, \frac{q + q^* - 1}{q^*} \right) \leq \lambda \leq \min \left(1, \frac{q}{q^*} \right) \quad (7)$$

These bounds, summarized in Table 1, are derived by fixing $\gamma \in (0, 1)$ at either of its extremes, and solving for the remaining unknown parameters in equations (1) or (2). The expressions show that upper bounds on the success rate (eligibility coverage) are inversely proportional to the quota (poverty rate)—i.e., for $q < q^*$ the higher the poverty rate, the lower the upper bound on the coverage of the eligible.

An important corollary is that if targeting is minimally non-regressive, meaning that selection of beneficiaries is no worse than random or that no priority is given to the non-eligible, then the targeting success metrics will satisfy tighter lower bounds:

$$\alpha \geq q^*; \lambda \geq q \quad (8)$$

These bounds are critical to move from a situation where targeting performance is unknown to one where it can be empirically bounded.

³ For instance: $\frac{\alpha + \Delta}{q_1^*} > \frac{\alpha + \Delta}{q_2^*}$ if $q_1^* < q_2^*$.

Table 1: Bounds on targeting success and error rates under alternative scenarios

	Targeting (α)	Elig. coverage (λ)
Upper	$\min(1, q^*/q)$	$\min(1, q/q^*)$
Neutral	q^*	q
Lower	$\max(0, (q + q^* - 1)/q)$	$\max(0, (q + q^* - 1)/q^*)$

Note: this table summarizes the theoretical bounds for two key targeting performance metrics—the targeting success rate (α) and the eligibility coverage rate (λ), as well as their corresponding error rates (leakage and under-coverage); the bounds are derived from algebraic identities that hold between observed programme coverage (q) and poverty rates (q^*); neutral values correspond to random assignment or non-regressive targeting.

Source: own derivations.

Proposition 2.3. *Relative Effectiveness of Scale-Up*

Given the welfare maximand Z (equation 3), the relative effectiveness of expanding programme coverage versus improving targeting accuracy is given by the marginal rate of technical substitution, denoted the *Relative Effectiveness of Scale-Up (RES)*:

$$RES = \frac{\partial Z / \partial q}{\partial Z / \partial \alpha} = \frac{\alpha}{q} \quad (9)$$

When $RES > 1$, marginal welfare gains from coverage expansion are expected to outweigh gains from improvements in targeting.

The RES is strictly bounded from below by: $\max(0, (q + q^* - 1)/q^2)$, and from above by: $\min(1/q, q^*/q^2)$. However, under the assumption of non-adverse targeting, the lower bound tightens to: q^*/q , which constitutes a practical rule-of-thumb for decision-makers.

By way of extension, consider a modified policy maximand that incorporates explicit linear costs of Type I errors, such that: $Z = \alpha \cdot q - \beta(1 - \alpha) \cdot q$, where β is a weight placed on errors of inclusion. Under this set-up, the neutral lower-bound RES becomes: $q^*/q - \beta/(1 + \beta)q$, which clearly lends to an increasingly conservative evaluation of pure scale-up as β rises.

2.3 Extension

The foregoing sought to formalize intuitions about targeting, offering practical tools to assess targeting quality in the absence of direct measures of leakage or under-coverage. Extending this analysis, the trade-off between programme expansion and selection is made explicit via a budget constraint:

$$B = tQ + cN \quad (10)$$

where B is the variable portion of the overall budget allocated to the program; and c is the cost of targeting per individual in the population. Defining $c = x \cdot (\mu B/N)$, re-expresses the targeting

cost as the product of two terms: $\mu B/N$, the maximum share $\mu \in (0, 1)$ of the total per population budget that can be feasibly allocated to targeting; and $x \in (0, 1)$, a decision parameter representing how much of this maximum cost the policymaker chooses to allocate. This constraint can be rearranged to express the quota as a function of the budget and other parameters: $Q = B(1 - x \cdot \mu)/t$.

To be clear, μ represents the slope of the trade-off between x and q , or the price of targeting effort. It is intended to reflect the full administrative and operational costs of improving targeting precision. This includes not only direct expenditures on data collection, verification, and eligibility scoring systems (e.g., implementing a proxy means test), but also indirect or hidden costs (see also Slater 2023). For example, the deployment of complex targeting systems may delay benefit delivery, divert staff time from routine programme administration, or create confusion among local implementers and beneficiaries. These disruptions are especially salient in low-capacity settings, where infrastructure for large-scale programme delivery is fragile or fragmented. Indeed, we reasonably expect μ to be substantially larger where existing administrative data such as population registries are deficient or not fully digitized.

In later simulations we adopt illustrative values for this parameter. However, we recognise there is a need for further empirical work to quantify (maximum) targeting costs more precisely and comprehensively. This includes understanding how costs scale with programme coverage, how they vary across targeting mechanisms, and how they interact with state capability constraints, a theme emphasized in recent work on the implementation capacity of social programmes in low-income settings (e.g., Andrews et al. 2017; Pritchett et al. 2010). While there is little empirical guidance to calibrate μ precisely, rough approximations are possible and can be refined via more detailed cost accounting (White et al. 2015).

Next, assume a behavioural equation linking costly effort to the targeting success rate, which for ease of analysis we assume is an increasing concave function in q^* and x :

$$\alpha = f(q^*, x) \tag{11}$$

meaning that $\partial\alpha/\partial x = f'_x \geq 0$, etc.. We also assume bounds on $0 \leq \alpha \leq 1$ are always satisfied and $f(q^*, x = 0) = q^*$, which says if no effort is allocated to targeting then $\alpha = q^*$, as per a benchmark neutral scenario discussed above.⁴

Pulling this together, the policy maximand of equation (3) is just a non-linear function of x :

$$Z(x) = \alpha \cdot q = f(q^*, x) \cdot \frac{b(1 - x\mu)}{t} \tag{12}$$

⁴ A more complete model would recognize that targeting is expected to be more difficult as the standard deviation of the underlying but unobserved target variable—say, permanent income—increases.

Solving the relevant first order condition and applying the constraints $0 \leq x \leq 1$ yields our final proposition.

Proposition 2.4. *Optimal targeting effort*

Optimal targeting effort, $x^ \in (0, 1)$, is decreasing in both the poverty rate and the cost of effort, satisfying:*

$$x^* = \begin{cases} 0, & \text{if } \mu > \frac{f'_x}{q^*} \\ 1, & \text{if } \mu < \frac{f'_x}{f'_x + \bar{\alpha}} \\ \frac{f'_x - \mu f(q^*, x^*)}{f'_x \mu}, & \text{otherwise} \end{cases} \quad (13)$$

For proof see Appendix A1.

While generic, this result raises the possibility that optimal targeting effort may be located at or close to its zero lower bound. This would occur where additional targeting effort induces a relatively large budgetary cost, where the poverty rate is already high, or where marginal gains in the success rate from higher effort are low (perhaps for other reasons). Minimally, it follows that in contexts with heterogeneous rates of poverty, a one-size-fits-all targeting effort is likely to be sub-optimal, and zero effort indeed may be optimal in some cases.

2.4 Numerical examples

Before considering a concrete case, some numerical examples shed light on the above propositions. Table 2 sets out a range of combinations for what are assumed to be two known inputs (q, q^*) , and derives corresponding bounds on the targeting success rate, the eligibility coverage rate and the RES. As might be expected, while the interval between the upper and lower bounds on the targeting parameters often can be large, they are far from uninformative and tighten considerably if we assume targeting is minimally no-regressive. For instance, in the first row, where the quota far exceeds the poverty rate, the upper bound on the share of beneficiaries who are eligible is at best 0.24 while a neutral (lower) bound is 0.20. Furthermore, the RES provides unambiguous guidance as to the expected relative gains from programme expansion versus improved targeting.

Table 2: Numerical examples of targeting performance bounds

Inputs		α		λ		RES	
q^*	q	Lower	Upper	Lower	Upper	Neutral	Upper
0.20	0.85	0.06	0.24	0.25	1.00	0.24	0.28
0.35	0.70	0.07	0.50	0.14	1.00	0.50	0.71
0.20	0.33	0.00	0.62	0.00	1.00	0.62	1.89
0.60	0.75	0.47	0.80	0.58	1.00	0.80	1.07
0.50	0.55	0.09	0.91	0.10	1.00	0.91	1.65
0.25	0.20	0.00	1.00	0.00	0.80	1.25	5.00
0.80	0.50	0.60	1.00	0.38	0.63	1.60	2.00
0.65	0.40	0.13	1.00	0.08	0.62	1.63	2.50
0.80	0.25	0.20	1.00	0.06	0.31	3.20	4.00

Note: table shows illustrative combinations taking poverty and coverage rates as inputs (q^* and q), then calculating resulting (lower, upper or neutral) bounds on the targeting success rate (α), eligibility coverage rate (λ), and the Relative Effectiveness of Scale-Up (RES) as per Table 1; when the RES exceeds unity then marginal gains from expanding programme coverage are expected to exceed gains from improving targeting accuracy.

Source: own derivations.

Turning to the final proposition, assume:

$$\alpha = q^* + x(1 - q^*)^\varphi \quad (14)$$


where φ governs how quickly marginal gains associated with a given additional targeting effort fall as q^* increases. This ensures that the response of targeting to additional funding (higher x) tends to decline with the poverty rate (i.e., $\partial^2 \alpha / \partial x \partial q^* < 0$); and larger values for φ imply greater inefficiencies or imperfections in the targeting process.

In this case, the solution for x^* becomes explicit (see Appendix A2 for derivations):

$$x^* = \begin{cases} 0, & \text{if } \mu > \frac{(1 - q^*)^\varphi}{q^*} \\ 1, & \text{if } \mu < \frac{(1 - q^*)^\varphi}{2(1 - q^*)^\varphi + q^*} \\ \frac{(1 - q^*)^\varphi - \mu q^*}{2\mu(1 - q^*)^\varphi}, & \text{otherwise} \end{cases} \quad (15)$$

Setting $\varphi = 3$ (see further below) allows us to simulate values for x^* in (μ, q^*) space, shown in Figure 1. A main insight is that at (very) high poverty rates, only extremely low values for μ would justify allocating *any* funds to improve targeting performance relative to a benchmark of random selection of beneficiaries (e.g., by lottery). For instance, when the poverty rate is 0.5, the maximum cost share must be below 0.2 to justify a non-zero spend on targeting. Intuitively, this simply reflects the trade-off between better targeting and greater programme coverage—i.e., for a given overall budget, any increase in the policy maximand associated with better targeting may be fully offset by a decline in coverage, especially when poverty is high.

Figure 1: Simulated optimal levels of targeting (x^*) given alternative poverty rates (q^*) and maximum costs (μ)



Note: the heatmap plots optimal values for the level of targeting, ranging from zero to one, obtained from equation (15) applying a systematic combination of different input values (for q^*, μ) and setting $\varphi = 3$.
Source: own elaboration.

3 Mozambique case

This previous section provided a practical framework for analysis of targeting performance, highlighting the critical relationship between the poverty rate, the price of targeting effort, and optimal targeting choices. Our results formalize the intuition that the optimal choice of targeting technology will vary across contexts in line with differences in various factors, including poverty rates. Naturally, this implies there is no single 'best' approach to targeting, and that even use of the same technology within the same country may be sub-optimal where contexts vary. The framework also points to the kind of empirical information that is critical to guide decisions in this domain. This includes data on actual programme coverage (q), target programme coverage (q^*), and targeting costs (e.g., μ). Furthermore, estimates of the extent to which targeting performance can be genuinely and consistently improved through allocating additional funds is likely to be helpful, akin to establishing empirically-valid versions of equation (11).

We now apply this framework to the elderly component of the *Programa de Subsídio Social Básico* (PSSB) in Mozambique. This is a social pension for the vulnerable elderly (aged 60+)

and is the country's largest social protection program. To conduct the analysis, we combine two main data sources: first, administrative data providing a complete register of all programme beneficiaries over the period 2020-2023 (yielding data on Q); and, second, granular population and poverty data based on the full records of the 2017 Population and Housing Census (yielding data on N, Q^*). For current purposes, poverty among the elderly is defined using a multi-dimensional index based on assets and living conditions, which closely mirrors the official national poverty analysis. All data is aggregated and then merged at the *posto administrativo* level, which is the third administrative tier, allowing for a highly localized analysis of targeting performance.⁵

A key input to the empirical application is the estimated poverty rate (q^*) at the subnational level, constructed from the 2017 national census. While this approach provides full spatial coverage and broad consistency with national poverty assessment criteria, it is not without limitations. First, measurement error in q^* is likely, particularly in remote or data-poor regions, where household-level indicators may be missing, outdated, or inconsistently reported. Such error may bias estimates of the RES metric as well as the inferred optimal effort level. Second, the choice of poverty definition matters. Multidimensional metrics, while normatively appealing, do not necessarily align with the implicit targeting criteria used in practice, such as community judgments or proxy means scores.⁶

With these caveats in mind, Table 3 summarises key information derived from these sources by province, showing bounds on the targeting metrics as per our earlier numerical examples.⁷ At the national level, we observe the poverty rate among the elderly ($q^* = 0.49$) and the programme quota ($q = 0.48$) are closely matched, suggesting the programme is appropriately scaled on aggregate. However, national averages mask profound sub-national differences, including relatively low programme coverage in the higher poverty regions of Zambézia and Nampula compared to over-coverage in the far south (e.g., Maputo Province and City). These patterns are illustrated in Appendix Figures B1 to B3, revealing significant dispersion in the neutral RES metric. Specifically, the cumulative distribution of the neutral RES suggests that

⁵ For further background on the history and structure of the PSSB see Almeida et al. (2025). The poverty measure applied here is an Alkire-Foster multidimensional poverty headcount index covering indicators across three main dimensions: education, housing conditions, and asset ownership. Households are classified as poor if deprived in at least five of six indicators. Aggregation of the registry and census data to lower levels is possible in theory, but faces challenges due to differing geographical codes and naming conventions in the datasets.

⁶ According to the program's operational manual (INAS 2022), eligibility for the elderly component of the PSSB is determined based on three criteria defined: (i) age — applicants must be 60 or older; (ii) residency — living for at least three months in the local community; (iii) poverty — household income is below a first-tier threshold, assessed via a Proxy Means Test (PMT). In practice, no PMT is performed and community-based targeting is used to identify the vulnerable persons without sources of income.

⁷ We focus here only on the *de jure* coverage of the program, abstracting from issues of non-payment which are material but nevertheless a subject of separate research. See MGCAS (2022) and IGM (2025) for further discussion.

around 60% of the elderly population reside in locations where this measure is greater than one (see Figure 2). In such locations, expanding coverage is generally expected to yield higher welfare returns than improvements in targeting.

Table 3: *De jure* coverage of PSSB-Elderly program, 2023

Province	q^*	q	α		λ		RES_n
			LB	UB	LB	UB	
Niassa	0.53	0.68	0.34	0.75	0.51	0.88	0.70
Cabo Delgado	0.61	0.51	0.29	0.88	0.26	0.73	1.35
Nampula	0.64	0.50	0.33	0.89	0.29	0.72	1.32
Zambezia	0.62	0.40	0.21	0.91	0.20	0.55	1.95
Tete	0.49	0.52	0.18	0.82	0.21	0.86	0.95
Manica	0.46	0.60	0.21	0.67	0.29	0.88	0.65
Sofala	0.43	0.46	0.17	0.69	0.19	0.81	1.01
Inhambane	0.49	0.45	0.14	0.87	0.16	0.76	1.28
Gaza	0.34	0.63	0.16	0.54	0.30	0.98	0.62
Maputo Prov.	0.08	0.34	0.02	0.18	0.05	1.00	0.10
Maputo City	0.01	0.27	0.00	0.03	0.03	1.00	0.02
National	0.49	0.48	0.21	0.74	0.23	0.78	1.05

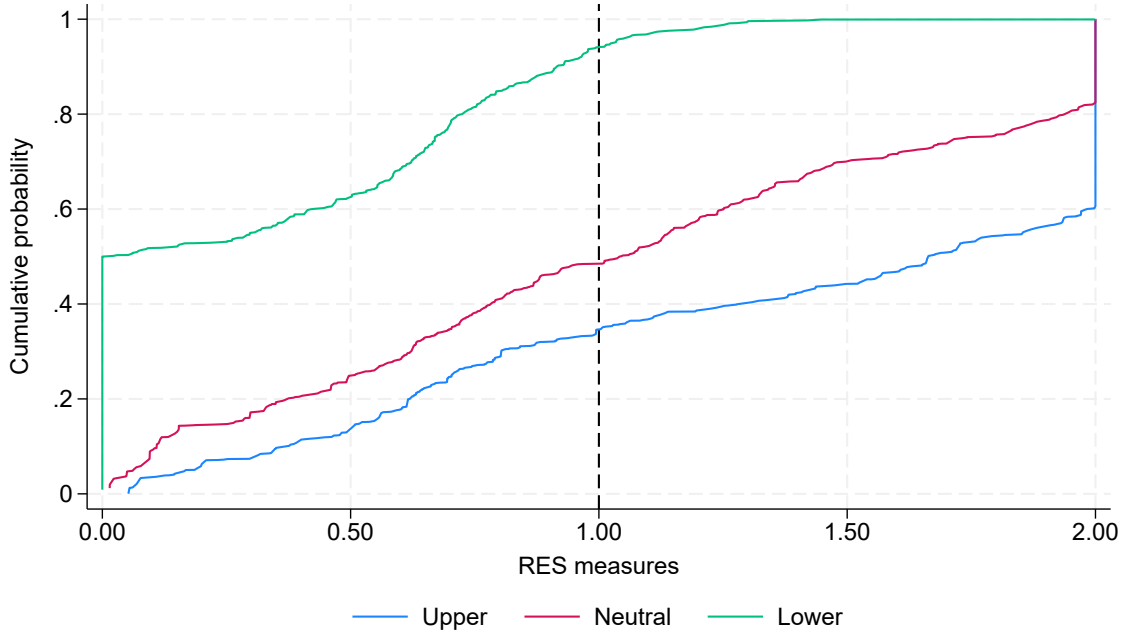
Note: table summarizes merged estimates from administrative records and census data, showing poverty rates (q^*), programme quotas (q), and derived bounds on targeting metrics (α , λ) for each province ordered North to South; UB, LB refer to upper and lower bounds respectively; the RES_n column reports the neutral Relative Effectiveness of Scale-up metric, highlighting where programme expansion is expected to yield greater welfare returns vs. improving targeting accuracy.

Source: own estimates.

As a further step we apply and calibrate the approximation in equation (14), which we use as a basis to model optimal targeting effort. Our main interest is to identify if and where zero targeting effort may be recommended, beyond the benchmark of random beneficiary selection. To do so, we use data from a bespoke survey of PSSB-Elderly programme members, described in detail below. We draw on the Vulnerable Lives Survey (VLS), conducted in 2023 as part of the Inclusive Growth in Mozambique (IGM) programme, a research collaboration between UNU-WIDER and the Government of Mozambique. The VLS was designed to generate robust empirical evidence on the performance and impacts of the PSSB-Elderly in shock-prone settings. Fieldwork was carried out in six purposively selected administrative posts across Mozambique's northern, central, and southern regions, chosen based on historical exposure to distinct types of shocks. In each location, approximately 400 households were surveyed, comprising both elderly beneficiaries and a relevant control group. For more details see IGM (2025).

This exercise provides an estimate $\varphi = 3$. Setting the maximum cost share at $\mu = 0.25$, meaning that maximal targeting effort (and success) would require one quarter of the variable component of the overall programme budget, we use equation (15) to calculate the optimal level of effort in each administrative post. Figure 3 plots these latter estimates against those of the neutral RES. The two measures are highly correlated, showing that zero effort is optimal for

Figure 2: Cumulative distributions of upper, neutral and lower bound RES estimates



Note: figure shows the cumulative distribution functions (CDFs) of upper, neutral, and lower bound estimates of the RES metric across all administrative posts; all RES metrics are windsorized at 2 for visual symmetry.
Source: own estimates.

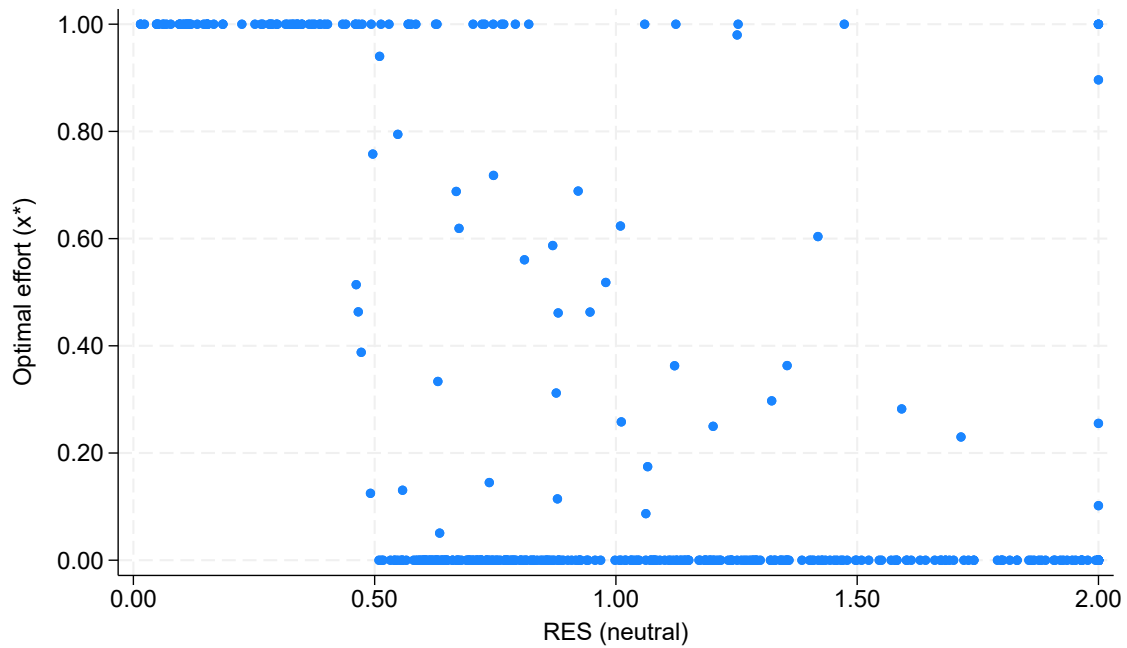
61% of locations. However, there is a clear bimodal distribution to the optimal effort estimates, with just 9% of urban locations in the zero effort regime. Appendix Figure B4 further examines the sensitivity of these estimates, showing the share of population residing in so-called zero optimum effort locations under alternative assumptions for φ and μ . This suggests one of our key insight remains robust—expanding programme scope without enhancing targeting is likely to be an optimal strategy in many locations across Mozambique.

Last, we employ the welfare metric Z to compare alternative allocations of the budget between targeting and transfers. To capture the status quo, we calculate baseline welfare as the sum across all locations of the product of the observed programme quota and an estimate of the target success rate, being the weighted average of the neutral and upper bounds on α . Thus, current total welfare is given by: $Z_0 = \sum_{\forall i} q_i \bar{\alpha}_i$, where $\bar{\alpha}_i = 0.7q_i^* + 0.3 \cdot \min(1, q_i^*/q_i)$ and i indexes administrative posts.

As a first counterfactual scenario, we maintain the same assumptions for targeting success ($\bar{\alpha}_i$) but reallocate the programme quota in proportion to observed poverty shares—i.e., if a location holds $y\%$ of all poor elderly persons, it is allocated $y\%$ of the total quota. This amounts to total welfare: $Z_1 = \sum_{\forall i} (Q_i^*/Q^*) \bar{\alpha}_i$. Third, retaining previous assumptions, we use our estimates for x^* to solve equation (14) and calculate the corresponding optimum quota.⁸ Last, we

⁸ For this, we require estimates of the ratio between the overall budget per head of population and the transfer value. These are derived by estimating x from observed data under the assumption $\alpha_i = \bar{\alpha}_i$. Due to noise in these

Figure 3: Scatter plot of optimal targeting effort versus the neutral RES metric, by administrative post



Note: figure shows the relationship between the neutral RES metric and the estimated optimal targeting effort (x^*) for each administrative post; RES values indicate the marginal rate of technical substitution between increasing programme coverage versus improvements in targeting effort.

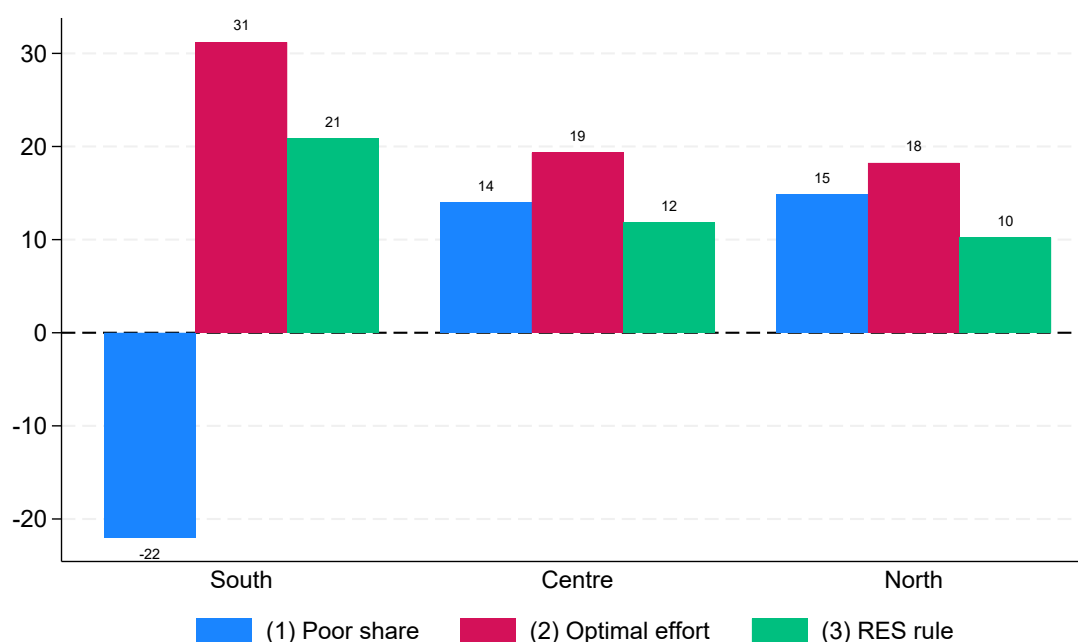
Source: own estimates.

re-run these latter calculations but now approximate x^* using a binary rule based on the observed neutral RES metric, namely we set: $\tilde{x}_i = 1[\text{RES}_{i,n} < 1]$.

Figure 4 plots these results, showing the percentage difference in total welfare per region relative to the baseline metric (Z_0). In scenario (1), we find that reallocating the proportional allocation of the quota to mirror the share of the elderly poor in each location (maintaining current targeting accuracy) would increase the national welfare metric by around 7%. An optimal allocation based on our formal model (scenario 2), which simultaneously adjusts both the quota and the targeting effort to their optimal levels in each location, achieves a 21% increase in welfare, including improvements in all locations. A more feasible approach, which applies a simple rule-of-thumb based on the RES (scenario 3) and involves eliminating additional targeting effort where RES is above unity and re-investing the savings into expanded coverage, yields a 13% welfare gain overall. The ability to achieve welfare gains across all locations, especially under scenario 3, demonstrates substantial gains are possible from simple, evidence-based, yet geographically-differentiated strategies, without requiring any increase to the overall budget.

estimates, we average x and the derived estimates for b/t by rural/urban locations in each of the three major geographical regions of the country.

Figure 4: Welfare changes under alternative targeting effort and quota allocation scenarios



Note: figure shows three counterfactual scenarios for reallocating programme budget across regions—(1) quota allocation is proportional to poverty shares, (2) applies the estimated optimal targeting effort (x^*), and (3) applies a binary RES-based rule, setting $x^* = 1$ if $RES_n < 1$ and 0 otherwise; the y-axis reports the percentage change in total welfare (Z) relative to the status quo (Z_0), not shown.

Source: own estimates.

4 Conclusion

With characteristic irony, Amartya Sen once remarked that: “I do not doubt that some expert in modern economics would find it helpful to say that targeting should be pushed exactly to the point at which the marginal benefit from it equals its marginal cost.” (Sen 1995: p.14). However, perhaps reflecting the practical difficulty of identifying an optimal ‘dosage’ of targeting, academic debates on social policy design have rarely focused on *how much* targeting is appropriate. Instead, they have tended to frame the issue in binary terms—i.e., as a choice between fully universal or targeted programs, which in turn has inspired a large literature assessing the relative performance of alternative targeting schemes.

This paper sought to (re)embrace the basic insight of marginalism, but also grasp its practical value. Rather than offering a general prescription for or against targeting, we provided an operational framework to assess when better targeting is likely to be a good use of scarce financial and human resources. We showed that the trade-off between targeting precision and programme coverage can be usefully summarized by the marginal rate of technical substitution, the Relative Effectiveness of Scale-up (RES). A neutral lower bound on this metric is easy to compute in settings with limited data. When it exceeds one, marginal welfare gains from expanding programme coverage are expected to outweigh gains from improving targeting accuracy. Extending this approach to take into account budget constraints, we showed that even

modest targeting costs may render efforts to enhance precision relatively inefficient, particularly in high-poverty contexts with tight quotas and low administrative capacity.

We applied this framework to Mozambique's national social pension programme (the elderly component of the PSSB). Reflecting large coverage inequalities, we found that for a majority of locations expanding coverage would yield greater welfare gains than investing in more precise beneficiary selection. Only in better-off urban areas would the converse be true. Counterfactual simulations suggested that reallocating resources according to either a simple RES rule or a model-derived optimum could generate welfare improvements of up to 21%, even without increases in programme expenditure.

These findings highlight a pathway beyond the polarized debate between universalism and targeting. Focussing on the administrative reality of contexts such as Mozambique, where many eligible individuals are not yet covered, a key challenge lies in determining when and where to focus on programme scale-up relative to making investments in targeting, which may be costly and disruptive. At the same time, for guidance taken from the (neutral) RES metric to be informative, low-cost selection schemes that are minimally non-regressive are essential. Thus, a relevant agenda is not how to achieve ever-better targeting, but rather how to avoid distortions associated with programme capture by the wealthy.

References

- Alatas, V., R. Purnamasari, M. Wai-Poi, A. Banerjee, B.A. Olken, and R. Hanna (2016). 'Self-Targeting: Evidence from a Field Experiment in Indonesia'. *Journal of Political Economy*, 124(2): 371–427. <https://doi.org/10.1086/685299>
- Almeida, S., H. Berkel, S. Jones, P. Justino, and T. Massingue (2025). 'Social Protection for Mozambique's Elderly: History, Structure, and Potential Effectiveness'. WIDER Working Paper 12/2025. Helsinki: UNU-WIDER. <https://doi.org/10.35188/UNU-WIDER/2025/569-1>
- Andrews, M., L. Pritchett, and M. Woolcock (2017). *Building State Capability: Evidence, Analysis, Action*. Oxford: Oxford University Press. <https://doi.org/10.1093/acprof:oso/9780198747482.001.0001>
- Asri, V., K. Michaelowa, S. Panda, and S.B. Paul (2022). 'The Pursuit of Simplicity: Can Simplifying Eligibility Criteria Improve Social Pension Targeting?' *Journal of Economic Behavior & Organization*, 200: 820–846. <https://doi.org/10.1016/j.jebo.2022.06.003>
- Banerjee, A., R. Hanna, B.A. Olken, and D.S. Lisker (2024). 'Social Protection in the Developing World'. *Journal of Economic Literature*, 62(4): 1349–1421. <https://doi.org/10.3386/10.1257/jel.20241646>
- Banerjee, A., P. Niehaus, and T. Suri (2019). 'Universal Basic Income in the Developing World'. *Annual Review of Economics*, 11(1): 959–983. <https://doi.org/10.1146/annurev-economics-080218-030229>

- Banerjee, A.V. (1997). 'A Theory of Misgovernance'. *The Quarterly Journal of Economics*, 112(4): 1289–1332. <https://doi.org/10.1162/003355300555484>
- Besley, T. and S. Coate (1992). 'Workfare Versus Welfare: Incentive Arguments for Work Requirements in Poverty-Alleviation Programs'. *The American Economic Review*, 82(1): 249–261.
- Besley, T. and R. Kanbur (1991). 'The Principles of Targeting'. In V.N. Balasubramanyam and S. Lall (eds), *Current Issues in Development Economics*, pages 69–90. London: Macmillan Education UK. https://doi.org/10.1007/978-1-349-21587-4_5
- Beuermann, D.W., B. Hoffmann, M. Stampini, D.L. Vargas, and D. Vera-Cossio (2025). 'Shooting a Moving Target: Evaluating Targeting Tools for Social Programs When Income Fluctuates'. *Journal of Development Economics*, 172: 103395. <https://doi.org/10.1016/j.jdeveco.2024.103395>
- Brown, C., M. Ravallion, and D. Van de Walle (2018). 'A Poor Means Test? Econometric Targeting in Africa'. *Journal of Development Economics*, 134: 109–124. <https://doi.org/10.1016/j.jdeveco.2018.05.004>
- Coady, D., M. Grosh, and J. Hoddinott (2004). 'Targeting Outcomes Redux'. *The World Bank Research Observer*, 19(1): 61–85. <https://doi.org/10.1093/wbro/lkh016>
- Coady, D. and E. Skoufias (2004). 'On the Targeting and Redistributive Efficiencies of Alternative Transfer Instruments'. *Review of Income and Wealth*, 50(1): 11–27. <https://doi.org/10.1111/j.0034-6586.2004.00109.x>
- Devereux, S., E. Masset, R. Sabates-Wheeler, M. Samson, A.-M. Rivas, and Te Lintelo, D. (2017). 'The Targeting Effectiveness of Social Transfers'. *Journal of Development Effectiveness*, 9(2): 162–211. <https://doi.org/10.1080/19439342.2017.1305981>
- Ellis, F. (2012). "We Are All Poor Here": Economic Difference, Social Divisiveness and Targeting Cash Transfers in Sub-Saharan Africa'. *Journal of Development Studies*, 48(2): 201–214. <https://doi.org/10.1080/00220388.2011.625408>
- Gelbach, J.B. and L.H. Pritchett (2001). 'Indicator Targeting in a Political Economy: Leakier Can be Better'. *The Journal of Policy Reform*, 4(2): 113–145. <https://doi.org/10.1080/13841280008523416>
- Hanna, R. and D. Karlan (2017). 'Designing Social Protection Programs: Using Theory and Experimentation to Understand How to Help Combat Poverty'. In *Handbook of Economic Field Experiments*, volume 2, pages 515–553. Elsevier. <https://doi.org/10.1016/bs.hefe.2016.07.002>
- Hanna, R. and B.A. Olken (2018). 'Universal Basic Incomes Versus Targeted Transfers: Anti-Poverty Programs in Developing Countries'. *Journal of Economic Perspectives*, 32(4): 201–226. <https://doi.org/10.1257/jep.32.4.201>
- Hoddinott, J. (2007). 'Social Protection: To Target or Not to Target'. *IDS Bulletin*, 38(3). <https://doi.org/10.1111/j.1759-5436.2007.tb00387.x>

- Houssou, N. and M. Zeller (2011). 'To Target or Not to Target? The Costs, Benefits, and Impacts of Indicator-Based Targeting'. *Food Policy*, 36(5): 627–637. <https://doi.org/10.1016/j.foodpol.2011.05.006>
- IGM (2025). Evaluating the Impact of the Basic Social Subsidy Programme on Elderly Well-Being in Mozambique. IGM report. Maputo: Inclusive Growth in Mozambique.
- INAS (2022). Manual De Procedimentos Do Programa Subsídio Social Básico. Internal procedures manual. Matola: Instituto Nacional de Acção Social, Ministério do Género, Criança e Acção Social.
- Kanbur, S.R. (1987). 'Measurement and Alleviation of Poverty: With an Application to the Effects of Macroeconomic Adjustment'. *Staff Papers-International Monetary Fund*, 34(1): 60–85. <https://doi.org/10.2307/3867024>
- Lanjouw, P. and M. Ravallion (1999). 'Benefit Incidence, Public Spending Reforms, and the Timing of Program Capture'. *The World Bank Economic Review*, 13(2): 257–273. <https://doi.org/10.1093/wber/13.2.257>
- Lindert, K., A. Linder, J. Hobbs, and B. De la Brière (2007). 'The Nuts and Bolts of Brazil's Bolsa Família Program: Implementing Conditional Cash Transfers in a Decentralized Context'. Social Protection Discussion Paper 709. Washington, DC: World Bank.
- MGCAS (2022). 'Relatório de Avaliação Intermédia da Estratégia Nacional de Segurança Social Básica 2016–2024'. Technical report. Maputo: MGCAS. Em colaboração com a OIT e parceiros.
- Mirrlees, J.A. (1971). 'An Exploration in the Theory of Optimum Income Taxation'. *The Review of Economic Studies*, 38(2): 175–208. <https://doi.org/10.2307/2296779>
- Mkandawire, T. (2005). 'Targeting and Universalism in Poverty Reduction'. Social Policy and Development Programme Paper 23. United Nations Research Institute for Social Development (UNRISD).
- Pritchett, L., M. Woolcock, and M. Andrews (2010). 'Capability Traps? The Mechanisms of Persistent Implementation Failure'. CGD Working Paper 234. Washington, DC: Center for Global Development. <https://doi.org/10.2139/ssrn.1824519>
- Ravallion, M. and K. Chao (1989). 'Targeted Policies for Poverty Alleviation Under Imperfect Information: Algorithms and Applications'. *Journal of Policy modeling*, 11(2): 213–224. [https://doi.org/10.1016/0161-8938\(89\)90015-X](https://doi.org/10.1016/0161-8938(89)90015-X)
- Schüring, E. and F. Gassmann (2016). 'The Political Economy of Targeting – a Critical Review'. *Development Policy Review*, 34(6): 809–829. <https://doi.org/10.1111/dpr.12185>
- Sen, A. (1995). 'The Political Economy of Targeting'. In D. Van De Walle and K. Nead (eds), *Public Spending and the Poor: Theory and Evidence*, pages 245–260. Baltimore: Johns Hopkins University Press.

- Slater, R. (2023). 'Universal, Categorical and Targeted Social Protection: Issues, Debates and Solutions'. In L. Patel, S. Plagerson, and I. Chinyoka (eds), *Handbook on Social Protection and Social Development in the Global South*, chapter 10, pages 180–204. Cheltenham: Edward Elgar Publishing Ltd. <https://doi.org/10.4337/9781800378421.00024>
- White, H. (2017). 'Effective Targeting of Social Programmes: an Overview of Issues'. *Journal of Development Effectiveness*, 9(2): 145–161. <https://doi.org/10.1080/19439342.2017.1317282>
- White, P., A. Hodges, and M. Greenslade (2015). 'Measuring and Maximising Value for Money in Social Protection Systems'. Technical report. London: Department for International Development (DFID).

A Appendix: Derivation of optimal targeting

A1 General case

We derive the optimal targeting effort $x^* \in [0, 1]$ by maximizing the welfare objective function:

$$Z(x) = \alpha \cdot q = f(q^*, x) \cdot \frac{b(1 - x\mu)}{t}, \quad (\text{A1})$$

where: $\alpha = f(q^*, x)$ is a concave function increasing in x , with $f(q^*, 0) = q^*$; $\mu \in (0, 1)$ is the maximum fraction of the total budget per capita of population that can be allocated to targeting effort; $x \in [0, 1]$ is the policymaker's decision on what share of that maximum is used; b is the per capita variable budget and t is the uniform transfer amount per beneficiary.

Step 1: First-Order condition

Differentiating $Z(x)$ with respect to x , and denoting the partial derivative $f'_x = \partial f(q^*, x) / \partial x$, we obtain:

$$\frac{\partial Z}{\partial x} = \frac{b}{t} [f'_x(1 - x\mu) - \mu f(q^*, x)]$$

Setting $\partial Z / \partial x = 0$ gives the interior first-order condition (FOC):

$$f'_x(1 - x\mu) = \mu f(q^*, x) \quad (\text{A2})$$

Solving equation (A2) for the optimum $x = x^*$, we isolate:

$$x^* = \frac{f'_x - \mu f(q^*, x^*)}{f'_x \mu}$$

Note that because f'_x and $f(q^*, x^*)$ are both functions of x , this expression is implicit and may require numerical solution methods. However, it is informative for establishing bounds and sufficiency conditions.

Step 2: Corner solutions and bounds

We now explore when the optimal x^* lies at the boundary of the feasible interval.

Case 1: Zero effort is optimal ($x^ = 0$).*

Substitute $x = 0$ into the first-order condition of equation (A2):

$$f'_x = \mu f(q^*, 0) = \mu q^*.$$

Hence, $x^* = 0$ is optimal if:

$$f'_x \leq \mu q^* \Rightarrow \mu \geq \frac{f'_x}{q^*}.$$

Case 2: Full effort is optimal ($x^ = 1$).*

Substitute $x = 1$ into the same first-order condition:

$$f'_x(1 - \mu) = \mu f(q^*, 1) = \mu \bar{\alpha},$$

where $\bar{\alpha} = f(q^*, 1)$. Rearranging gives:

$$f'_x = \mu(\bar{\alpha} + f'_x) \Rightarrow \mu \leq \frac{f'_x}{f'_x + \bar{\alpha}}.$$

Thus, $x^* = 1$ is optimal if the cost share is small relative to the marginal and maximum benefits of targeting.

Step 3: Interior dolution

If neither corner condition holds, the solution must lie in the interior $x^* \in (0, 1)$, and is defined implicitly by equation (A2), with:

$$x^* = \frac{f'_x - \mu f(q^*, x^*)}{f'_x \mu}.$$

Hence, the optimal effort level x^* satisfies equation (13) in the main text. This general formulation allows for application across a wide class of concave targeting functions, and provides a rigorous foundation for policy-relevant simulations.

A2 Specific example

We take the example from Section 2.4, defining: $\alpha = q^* + x(1 - q^*)^\varphi$. Differentiating $Z(x)$ with respect to x and setting to zero gives the first order conditions:

$$\frac{\partial Z}{\partial x} = \frac{B}{t} [(1 - q^*)^\varphi (1 - x\mu) + (q^* + x(1 - q^*)^\varphi)(-\mu)] = 0.$$

Expanding:

$$(1 - q^*)^\varphi - x\mu(1 - q^*)^\varphi - \mu q^* - x\mu(1 - q^*)^\varphi = 0.$$

Rearranging:

$$(1 - q^*)^\varphi - \mu q^* - 2x\mu(1 - q^*)^\varphi = 0.$$

Solving for x^* is the interior solution:

$$x^* = \frac{(1 - q^*)^\varphi - \mu q^*}{2\mu(1 - q^*)^\varphi}.$$

Applying constraints $0 \leq x \leq 1$, then a solution close to the lower bound ($x^* \geq 0$) obtains when:

$$\frac{(1 - q^*)^\varphi - \mu q^*}{2\mu(1 - q^*)^\varphi} \geq 0.$$

Since the denominator is always positive for $\mu > 0$, we require the numerator to be non-negative:

$$(1 - q^*)^\varphi - \mu q^* \geq 0.$$

In turn, solving for μ , yields:

$$\mu \leq \frac{(1 - q^*)^\varphi}{q^*}.$$

If $\mu > \frac{(1 - q^*)^\varphi}{q^*}$, then $x^* = 0$.

Similarly, the solutions close to the upper bound ($x^* \leq 1$) obtain when:

$$\frac{(1 - q^*)^\varphi - \mu q^*}{2\mu(1 - q^*)^\varphi} \leq 1.$$

Multiplying both sides by $2\mu(1 - q^*)^\varphi$:

$$(1 - q^*)^\varphi - \mu q^* \leq 2\mu(1 - q^*)^\varphi.$$

Rearranging:

$$(1 - 2\mu)(1 - q^*)^\varphi \leq \mu q^*.$$

Solving for μ , yields:

$$\mu \geq \frac{(1 - q^*)^\varphi}{2(1 - q^*)^\varphi + q^*}.$$

So, if $\mu < \frac{(1 - q^*)^\varphi}{2(1 - q^*)^\varphi + q^*}$, then $x^* = 1$.


B Appendix: Additional material

Table B1: Classification of population by beneficiary and eligibility status

		Beneficiary?		
		No	Yes	
Eligible?	No	a	b	$1 - q^*$
	Yes	c	d	q^*
		$1 - q$	q	1

Note: population is normalized to size one, such that $a + b + c + d = 1$; table indicates that the group of beneficiaries (q) can be expressed in two ways: $q = d + b = \alpha q + (1 - \alpha)q$, where $\alpha = d/(b+d)$; or $q = \lambda q^* + \gamma(1 - q^*)$, where $\lambda = d/(c+d)$ and $\gamma = b/(a+b)$.
Source: own elaboration.

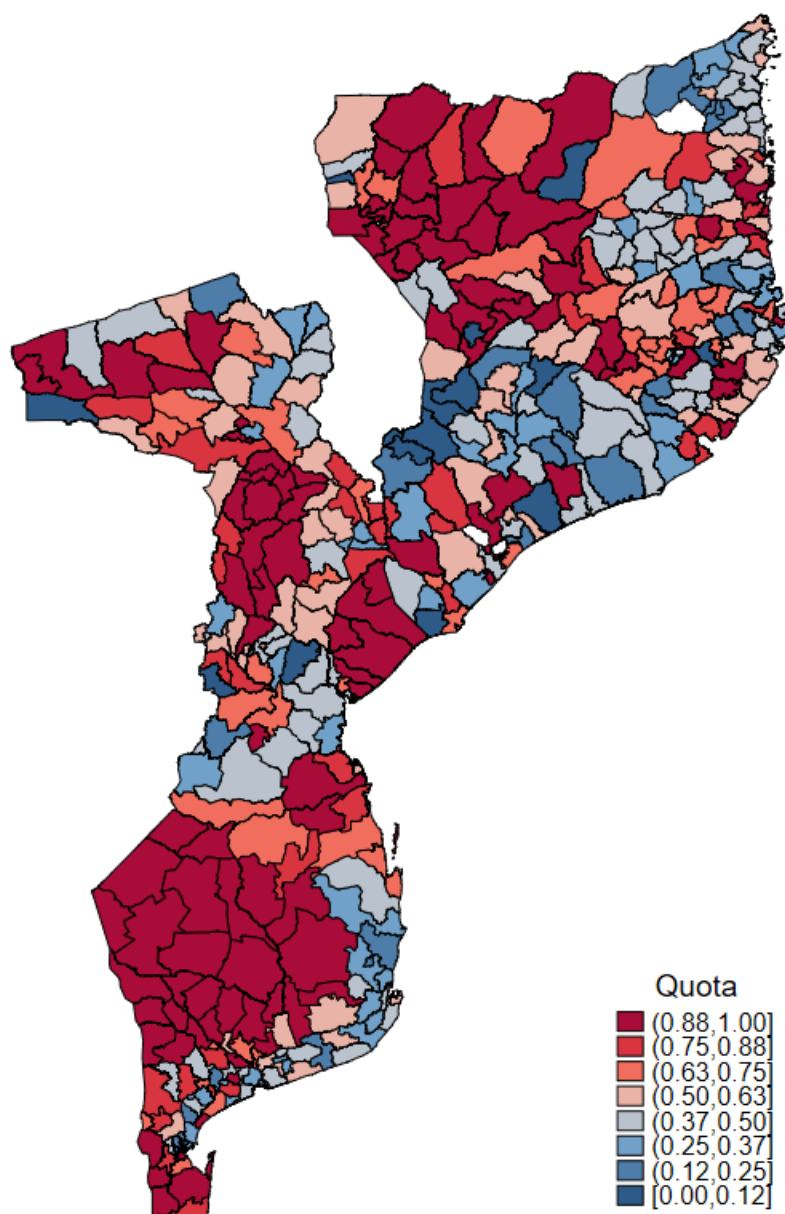
Figure B1: Estimated elderly poverty rates by administrative post



Note: the map shows estimated poverty rates among elderly individuals (aged 60+) across Mozambique's administrative posts; poverty is defined using a composite index incorporating indicators across health, education, and living standards dimensions, in line with official multidimensional poverty estimation methodologies.

Source: own elaboration.

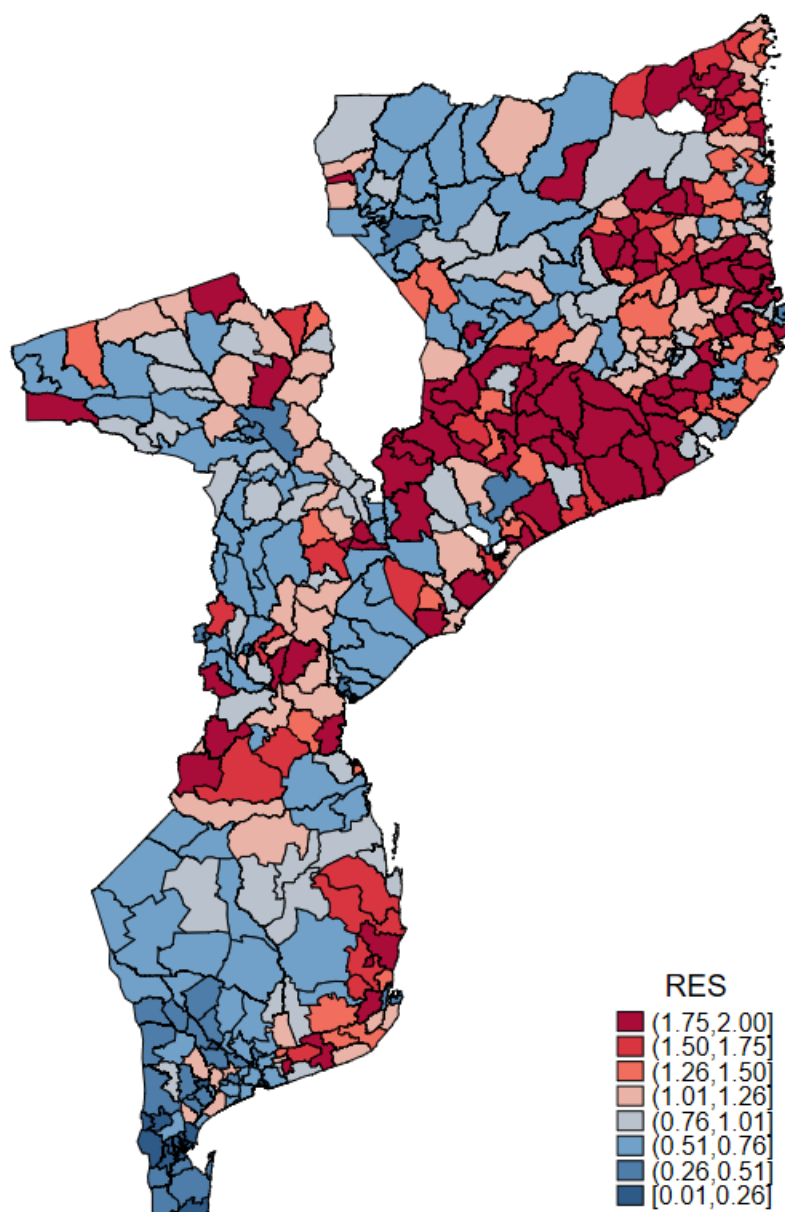
Figure B2: PSSB-Elderly quotas, by administrative post



Note: the map illustrates the *de jure* coverage of the PSSB-Elderly programme across administrative posts in Mozambique, measured as the ratio of enrolled programme beneficiaries to the estimated total elderly population (60+) in each location as of November 2023.


Source: own elaboration.

Figure B3: Estimates of the neutral RES, by administrative post



Note: the map shows estimates of the neutral Relative Effectiveness of Scale-up (RES) metric across administrative posts in Mozambique; the RES index compares the marginal welfare gain from expanding programme coverage relative to improving targeting accuracy, using a benchmark assumption of non-regressive targeting.
Source: own elaboration.

Figure B4: Heat map of population shares where zero targeting effort is estimated as optimal for alternative values of φ and μ



Note: figure displays a heat map of the share of the total total relevant population residing in locations where zero targeting effort is estimated to be optimal; estimates are based on simulations using alternative values for the targeting return parameter (φ) and the cost share parameter (μ); warmer colors (red/yellow) indicate higher population shares under zero-effort optimality, while cooler colors (blue) represent lower shares.
Source: own elaboration.