

ON THE TARGETING AND REDISTRIBUTIVE EFFICIENCIES OF ALTERNATIVE TRANSFER INSTRUMENTS

BY DAVID COADY*

International Food Policy Research Institute, Washington, D.C.

AND

EMMANUEL SKOUFIAS

Inter-American Development Bank, Washington, D.C.

The *distributional characteristic* provides an attractive alternative to conventional approaches used to evaluate the targeting performance of transfer programs. We decompose it into two components that are useful both conceptually and empirically; one capturing the *targeting efficiency* of the instrument, the other its *redistributive efficiency*. The redistributive index can also be generalized for the purposes of evaluating the degree of progressivity in tax-benefit structures. For illustrative purposes, we present an empirical application of the distributional characteristic and its decomposition using Mexican data. The welfare gains from using categorical targeting and means testing reflect improvements in redistributive and targeting efficiency respectively.

1. INTRODUCTION

In response to a tightening of government finances in developing countries, the emphasis on designing more efficient poverty-alleviation (or transfer) programs has become central to the ensuing policy debate. In terms of policy choices, this has involved a movement towards policy instruments which “target” the budget more efficiently to the “target group” (e.g. the “poor”).¹ In the literature, two approaches have become standard in evaluating the targeting performance of alternative transfer instruments, namely, a graphical approach and an index approach.

The graphical approach involves depicting the distributions of incomes and/or transfers (Case and Deaton, 1998; Schady, 2002). However, while this approach is useful for exploring rankings among alternative transfer schemes, in isolation it is of limited value especially if one wants to address issues of magnitude, i.e. how much better or worse one transfer system is compared to another. For this, one is forced to aggregate the distributional information into a single index, which can then be compared across alternative transfer schemes.

In this paper we are primarily interested in contributing to this latter aspect of the literature. We show that the *distributional characteristic*, already widely used in the commodity taxation literature (Newbery and Stern, 1987; Ahmad and Stern,

*Correspondence to: David Coady, Research Fellow, Food Consumption Nutrition Division, International Food Policy Research Institute, 2033 K St., N. W., Washington D.C. 20006, U.S.A. (d.coady@cgiar.org).

¹See Coady, Grosh, and Hoddinott (2002) for a recent review of targeted transfer programs in developing countries.

1991), provides an attractive alternative to various approaches found in the literature as a way of capturing the distributional power (or targeting performance) of transfer instruments. This approach has a number of advantages: (1) the value judgments underlying one's analysis are made more transparent and sensitivity analysis of results to these value judgments is straightforward; (2) it allows for a wider class of social welfare functions; and (3) it avoids the conceptual and practical controversy surrounding where to draw the poverty line which often unnecessarily distracts from important policy issues. The format of the paper is as follows. In Section 2 we provide a very brief overview of the most commonly used approaches. In Section 3 we present a very simple model to derive and motivate the distributional characteristic. In Section 4 we show how one can additively decompose this statistic into two terms, one which captures the *targeting efficiency* of the instrument, the other capturing the *redistributive efficiency* of the instrument. This decomposition is both conceptually and empirically very useful. We use this decomposition to provide an interpretation of the conventional leakage and undercoverage rates within standard welfare theory. Essentially, these can be seen as special cases of our targeting efficiency index. More generally, suitably redefined, the redistributive index provides a generalized index of progressivity-regressivity of tax-benefit systems. Finally, for the purpose of illustration, in Section 5 we present an empirical application of the decomposition using Mexican data and compare the distributional characteristic with the poverty-based alternatives. We show that the welfare gains associated with replacing universal food subsidies with categorically targeted transfers based on the demographic composition of households arise mainly from improved redistributive efficiency, while the gains from using means testing arise from improved targeting efficiency. Section 6 summarizes and concludes.

2. EXISTING APPROACHES

In this section we briefly describe the various approaches used in the literature to evaluate the relative targeting effectiveness of alternative transfer schemes. These can be usefully categorized into two approaches. The first of these involves graphically depicting the distributions of incomes and/or transfers. For example, one can use non-parametric techniques to compare the *probability density functions* of income including and excluding transfers. Progressive transfers will tend to shift the mass from the lower tail to the middle, with relatively little effect on the upper tail. While such a picture can give an initial feel for whether or not a transfer scheme is progressive, one often wants to have some idea of the magnitude of relative degrees of progression and this is not easily discernible from such a diagram.

Alternatively, one can compare non-parametric graphs of the *cumulative density functions* (cdf) for income including and excluding transfers. This approach orders households by income (from lowest to highest) and then plots their cumulative shares of total income. If the cdf including transfers lies everywhere above that excluding transfers, then one can unambiguously say that transfers are progressive, i.e. regardless of where in the income distribution one focuses. One can similarly compare across alternative transfer schemes. In practice, it may be suffi-

cient for policy making to restrict such comparisons to the lower parts of the distribution, i.e. to show dominance over a restricted range of lower incomes.

The existence of behavioral responses means that comparing incomes including and excluding transfers can be misleading since income excluding transfers is not then equal to pre-transfer income. For example, if progressive transfer schemes introduce disincentive labor supply effects at lower incomes then such comparisons will overstate the extent of progression. Alternatively, if transfers to low-income households generate positive multiplier effects then they will understate it. For this reason, it is useful to disentangle these issues by focusing on transfers in isolation. For example, one can construct *concentration curves*, which rank individuals by income and then plot their shares in total transfers. The curve for progressive transfers will lie everywhere above the 45-degree line with a positive and decreasing slope; the 45-degree line represents neutral transfers where low-income individuals receive their population shares in transfers. Similar dominance or restrictive dominance comparisons can be examined.

However, while the graphic approach is useful for exploring rankings among alternative transfer schemes, in isolation it is of limited value especially if one wants to address issues of magnitude, i.e. how much better or worse one transfer system is compared to another. For this one is forced to aggregate the distributional information into a single index, which can then be compared across alternative transfer schemes. Since this implicitly involves value judgments, it is important that these are made as transparent as possible and that one checks the robustness of one's general conclusions to alternative specifications.

One such index is the concentration coefficient, which is calculated as twice the covariance between individual transfers and income rankings all divided by total transfers (Milanovic, 1995). However, because this is based solely on a welfare *ranking* of households, it ignores important information regarding differences in incomes and thus also regarding the potential gains from targeting transfers. Two other sets of indices are commonly used, both of which have their foundations in the literature on poverty measurement.² The point of departure is usually taken as a situation where the government has a fixed budget allocated to an existing program aimed at reducing poverty, and wishes to determine how effective this program is at achieving this objective relative to feasible alternatives for distributing the program budget to households.

From the perspective of policy evaluation, two questions naturally arise: (1) How effective is the program at targeting "poor households," i.e. at identifying or promoting the participation of poor households in the program? (2) How effective is it in terms of reducing poverty? The two poverty approaches discussed below attempt to address each of these issues. Both start by identifying a *poverty line* based on some welfare index (here referred to as income). Households (or individuals) with incomes below the poverty line are classified as "poor" and those above as "non-poor."

One commonly used approach to evaluate the targeting performance of alternative transfer instruments is to compare leakage and undercoverage rates.

²For more detailed discussion of poverty measures, see Foster, Greer, and Thorbecke (1984), Atkinson (1987), Ravallion (1993b) and Deaton (1997).

Leakage is defined as the proportion of beneficiaries that are classified as non-poor, while undercoverage is defined as the proportion of poor households who are not beneficiaries. There are two obvious criticisms of this approach. Firstly, it ignores the seriousness of the targeting errors, e.g. it does not differentiate between the erroneous inclusion (exclusion) of non-poor (poor) households lying just above (below) the poverty line and those lying well above (below) the line. Secondly, it focuses only on who gets the transfers and not on how much households get (i.e. the size of the transfer budget).

The other commonly used poverty-based approach to evaluating the relative effectiveness of transfer instruments can be viewed as an attempt to incorporate the size of transfers and the budget explicitly into the analysis. Rather than asking how effective the program is at identifying the poor, it asks how effective it is at reducing poverty. It proceeds by comparing the relative impacts of the alternative instruments on the extent of poverty subject to a fixed common budget or, equivalently, the minimum cost of achieving a given reduction in poverty across instruments.

The extent of poverty is then usually calculated using the set of measures due to Foster, Greer, and Thorbecke (1984):

$$(1) \quad P_\alpha = N^{-1} \sum_{i=1}^N \left(1 - \frac{y_i}{z}\right)^\alpha I(y_i \leq z)$$

where y_i is income, z is the poverty line, N is the number of households, and $I(\cdot)$ is an indicator function taking the value 1 if its argument is true and zero otherwise. The parameter α captures the extent of our concern for the severity of poverty, with higher values corresponding to a greater aversion to severe poverty. The most commonly used poverty measures can be seen as special cases of this family of measures, namely, the poverty headcount index ($\alpha = 0$, or the percentage of households that are poor), the poverty gap ($\alpha = 1$, which captures the depth of poverty) and the severity index ($\alpha = 2$, which unlike the poverty gap is sensitive to redistribution among the poor). Only in the case of the severity index is there a higher value given to instruments that transfer more of the budget to the poorest households, and in all three poverty indices transfers to the non-poor are considered equally undesirable regardless of how close or otherwise they are to the poverty line. In addition, in both theory and practice, the identification of a poverty line is an extremely contentious issue, and in any case, as we will see below, the underlying concern for the poorest households can be adequately captured without such a precise categorization of households as being either “poor” or “non-poor.” Finally, in the case of the more acceptable severity index, value judgments (i.e. the relative weight given to income accruing to households in different parts of the income distribution) are not very transparent.

3. THE DISTRIBUTIONAL CHARACTERISTIC

In this section we develop a very simple model to help derive the so-called distributional characteristic more commonly used in the optimum commodity taxation and tax reform literature. Consider an economy with two groups, namely

households and the government. We assume a fixed transfer budget (B) and the objective of the “social planner” is to choose among alternative transfer programs with different levels of transfers across households, i.e. $dm = (dm^h)$. The planner’s problem can then be expressed as choosing the transfer program dm so as to:

$$(2) \quad \text{Max } \mathcal{L} \equiv W(\dots, V^h(p, m^h), \dots) + \lambda(B - \sum_h m^h)$$

where $W(\cdot)$ is the social welfare function, $V^h(\cdot)$ is the household’s indirect utility function, p is a vector of commodity and factor prices facing the household, λ is the social valuation of extra government revenue, and m^h is household non-factor income (which we assume comes solely from this transfer program, so that $dm^h = m^h$). The welfare impact of any given transfer program can be derived as:

$$d\mathcal{L} = \sum_h \frac{\partial W}{\partial V^h} \frac{\partial V^h}{\partial m^h} dm^h - \lambda dm^h = \sum \beta^h dm^h - \lambda dm^h$$

where β^h is the social valuation of extra income to household h (the so-called “welfare weight”). The first term captures the social welfare impact of the transfer program, while the second term represents the cost in terms of government revenue. Dividing the first term by the second term, and normalizing welfare weights such that $\lambda = 1$, we get a benefit-cost statistic for each program j defined as:

$$(3) \quad \lambda_j = \frac{\sum_h \beta^h dm^h}{\sum_h dm^h} = \sum_h \beta^h \theta^h$$

where θ^h is the share of the total budget received by household h . This statistic acts as a sufficient statistic for comparing the welfare impact of alternative transfer programs with a common fixed budget or, alternatively, for comparing budget reallocations between existing instruments.³

Equation (3) is analogous to the so-called distributional characteristic common in optimum commodity tax theory and, in the present context, can be interpreted as the social value of a unit of income distributed through the program. Therefore, λ_j will differ across transfer instruments both because welfare weights differ across households and because the structure of transfers differs across instruments. The greater the proportion of the budget ending up in the hands of the poorest households (i.e. those with relatively high β s), the higher the distributional characteristic. Note also that it is scale neutral in the sense that it does not change in response to a scaling up or down of transfer levels.

Underlying the concept of distributional power is a view that extra income to low-income (or poor) households is more socially valuable than extra income to high-income (or non-poor) households. The calculation of λ_j thus requires specifying the welfare weights, which are decreasing in household income. A very useful and common method for specifying these derives from Atkinson’s (1970) constant

³Strictly speaking, this is true only for “marginal transfers.” For non-marginal changes one simply replaces the first-order approximations with equivalent variations and makes welfare weights endogenous. We do not discuss such issues here since they just add unnecessary notational complications.

elasticity social welfare function where the (relative) welfare weight of household h is calculated as:

$$\beta^h = (y^k / y^h)^\varepsilon$$

where k is a reference household (e.g. the household on the poverty line, in which case $y^k = z$, the poverty line) and ε captures one's "aversion to inequality," with this aversion increasing in ε . For example, a value of $\varepsilon = 0$ implies no aversion to inequality (i.e. a dollar is a dollar no matter to whom it accrues) so that all welfare weights take on the value unity. A value of $\varepsilon = 1$ implies that if household h has twice (half) the income of household k then its welfare weight is 0.5 (2.0) as opposed to unity for k . A value of $\varepsilon = 2$ similarly implies a welfare weight of 0.25 (4.0) for h . Although there exists a qualitatively similar parameter for poverty indices, the underlying welfare weights are not as immediately transparent.

As ε approaches infinity, the welfare impact of transfers to the poorest households dominates the evaluation, consistent with a Rawlsian maxi-min social welfare perspective where one cares only about the welfare impact on the poorest households. For example, if we divide households into income quantiles and attach to them a welfare weight based on quantile mean income, then as ε increases the ranking of programs will be increasingly influenced by the share of transfers going to the poorest quantile. In fact, if we normalize welfare weights so that the welfare weight for the poorest quantile is unity, then from (3) it should be clear that as ε increases λ converges to the poorest quantile's transfer share. Also, in this case, λ ranges from unity to this quantile share as ε goes from zero to infinity. It should be obvious that the above approach to specifying welfare weights using higher values of ε can adequately incorporate our concerns for poverty without the need to introduce such a sharp distinction between "poor" and "non-poor" households as is the case for the poverty approaches outlined above.

4. TARGETING VS. DISTRIBUTIONAL EFFICIENCIES OF TRANSFER INSTRUMENTS

The distributional characteristic can be decomposed into two separate indices, which are both conceptually and empirically useful. Define dm^* as the average transfer to beneficiaries (i.e. total transfers divided by the number of beneficiaries, where beneficiaries are those with $dm^h > 0$). Then add and subtract dm^* across all beneficiaries (i.e. for all non-beneficiaries $dm^* = 0$) to get:

$$(4) \quad \lambda = \frac{\sum_h \beta^h \cdot dm^*}{\sum_h dm^h} + \frac{\sum_h \beta^h (dm^h - dm^*)}{\sum_h dm^h} = \lambda_T + \lambda_R$$

where λ_T is the *targeting efficiency* and λ_R is the *redistributive efficiency* of the transfer instrument. So λ_R captures the welfare impact, keeping fixed those who are receiving transfers (i.e. the targeting rule), of deviating from uniform transfers. One can interpret λ_T as the welfare impact of a program which transfers the poverty alleviation budget to the same beneficiary households but in equal amounts, and λ_R as the adjustment that needs to be made to allow for the differentiation of trans-

fers across households in a more progressive ($\lambda_R > 0$) or regressive ($\lambda_R < 0$) manner. Obviously, for programs with uniform transfers, $\lambda_R = 0$. The sense in which λ_R captures the redistributive efficiency of the policy instrument is made clearer by interpreting it as the welfare impact of a self-financing program which transfers dm^h to households and finances this by a lump-sum poll tax on all households with $dm^h > 0$ (i.e. all beneficiary households).

This decomposition of the welfare impact essentially defines progressivity (or regressivity) with respect to a distributionally “neutral” uniform transfer. Alternatively, one could define a neutral transfer program as one that is proportional with respect to household incomes.⁴ This perspective can also be easily accommodated within the above decomposition by subtracting a proportional transfer from beneficiaries, as opposed to a uniform transfer, with the factor of proportionality being determined both by the total incomes of beneficiary households and the total budget. As before, λ_T is still independent of the size of the budget. Note also that a uniform transfer financed by a proportional tax is progressive ($\lambda_R > 0$) but that a proportional transfer financed by a uniform poll tax is regressive ($\lambda_R < 0$). Therefore, defining neutrality with reference to deviations from a uniform transfer implicitly reflects a stronger concern for redistribution.

Strictly speaking λ_R is a *conditional* redistributive index since the program is assumed to be financed by a poll tax only on those receiving transfers. However, it is straightforward to construct a *generalized* (or *unconditional*) redistributive index by extending the poll tax across all households (in which case, of course, λ_T is constant across all programs). For example, if one is evaluating the degree of progressivity or regressivity of an individual tax instrument, or a set of taxes, then dm^h would represent actual taxes paid by each household and dm^* is the tax that would have been paid if the revenue was raised instead by a uniform poll tax (or proportional tax) applied to all households. This generalized redistributive index is essentially that derived in Duclos (1998) with his weights interpreted as welfare weights: it is a generalization in the sense that it can be applied to any quasi-concave social welfare function (and, thus, inequality indices consistent with such functions), for which a complete set of welfare weights across households can be derived.⁵

Later we use the above decomposition in an illustrative empirical analysis of the relative welfare impact of alternative transfer instruments. There it will be useful, for presentational purposes, to use the following transformation of the above decomposition. Consider some reference transfer scheme, j , e.g. the *status quo* or some optimal scheme. Then the welfare impact of moving to some alternative scheme, i , is:

⁴See Lambert (1993, pp. 164–7), Pfingsten (1986) and Besley and Preston (1988) for more detailed discussions of the concept of progression and the distributional “neutrality” of transfer schemes. For a discussion on the analysis of the progressivity of tax schedules and the reform of these schedules see, for example, Keen, Papapanagos, and Shorrocks (2000).

⁵Our derivation of this index makes the link between standard tax theory and measures of the progressivity or regressivity of taxes even more explicit. Essentially, these indices can be interpreted as simply the difference between the distributional characteristics of the actual tax payments and the hypothetical payments under some “reference” (or neutral) tax system. See also Pfahler (1987) on related issues, and Duclos (1997) for an interesting discussion on the relationship between social welfare, inequality and commonly used progressivity measures.

$$(5) \quad \frac{\lambda_i - \lambda_j}{\lambda_j} = \frac{\lambda_{Ti} - \lambda_{Tj}}{\lambda_j} + \frac{\lambda_{Ri} - \lambda_{Rj}}{\lambda_j}$$

where the first term on the r.h.s. captures the proportional change in welfare from moving from the reference scheme to the new scheme due to their different degrees of targeting efficiency, and the second term captures the proportional impact due to their different degrees of redistributive efficiency. Notice that a policy instrument which is poorly targeted may still have a relatively high welfare impact if the budget is allocated disproportionately to lower income households. For example, Case and Deaton (1998) find that pension transfers in South Africa were very progressive in spite of the absence of a means test because, although most households received some pension, the poorest received by far the most reflecting the fact that the elderly lived in poor households; in other words, their distributional power came from a high λ_R as opposed to a high λ_T .

We finish this section by interpreting undercoverage (U) and leakage (L) measures within the standard welfare framework described above. Consider a program which has a budget of \$1 for every poor household (N_i in number) and which distributes \$1 to N_i households but using an imperfect targeting rule. If, consistent with the poverty gap, everyone below the poverty line is given a welfare weight of unity and everyone above a welfare weight of zero, then for this program we have, using (3):

$$\lambda = \lambda_T = \frac{\text{Number of beneficiaries who are poor}}{\text{Total number of beneficiaries}}$$

which can also be interpreted as the proportion of the total budget (i.e. of N_i by construction) that reaches poor households, a measure suggested by Cornia and Stewart (1995). It is easily shown that this is just $(1 - L) = (1 - U)$, a measure of “coverage,” so that L and U both have some basis in welfare theory. However, these are valid measures of the welfare impact only for programs which transfer equal amounts to beneficiaries *and* only for a particular set of welfare weights (i.e. those consistent with the poverty gap indicator). More generally, using (4) it is clear that such a measure captures only the targeting efficiency (λ_T) of the program and ignores its redistributive efficiency (λ_R).

5. AN ILLUSTRATION FOR MEXICO

In this section we illustrate how the distributional characteristic can be used to evaluate the relative targeting performance of alternative transfer programs as well as in determining the factors influencing this performance. For the purposes of illustration we focus on the recent shift in Mexico’s poverty alleviation strategy towards better targeted transfer schemes. Our point of departure is one where universal food (i.e. cereal) subsidies (FSUB) constitute a central component of the overall poverty alleviation strategy. However, these are perceived as being poorly targeted with much leakage of benefits to non-poor households.

One can consider a number of alternative targeting strategies. Here we consider two broad approaches, *categorical* targeting and *means-tested* targeting:

TABLE 1

STRUCTURE OF DEMOGRAPHIC TRANSFERS (PESOS PER MONTH)		
Child's Age (years)	Male	Female
0–4	37.5	37.5
5–10	37.5	37.5
11–14	87.5	92.5
15–19	97.5	112.5

Notes: The transfer levels presented in this table are used to determine the transfer levels received by each household under each of the program alternatives considered. Therefore, under these programs, the transfer levels received by households will depend on their demographic structure.

(i) CHBEN: This program gives transfers to all households with children 0–4 years old; 37.5 pesos per child per month. This is akin to child benefit programs found in many countries.

(ii) SCHBEN: This program gives transfers to all households with children 5–19 years old, with transfers increasing by age and being higher for females over 10 years of age. This is akin to a program that subsidizes primary and secondary school enrollment and has a similar structure to a program recently introduced in Mexico (i.e. PROGRESA) as well as in other developing countries (Coady, 2001).

(iii) P-CHBEN: Same as (i) but only given to “poor” households with consumption per adult equivalent below 200 pesos per month.

(iv) P-SCHBEN: Same as (ii) but only given to “poor” households with consumption per adult equivalent below 200 pesos per month.

The P-SCHBEN program is used to determine the budget. The transfer levels resulting from the above schemes of payments are scaled up or down in order to exactly exhaust the program budget.

- (1) *Categorical targeting:* Under this form of targeting, whether or not a household receives transfers depends on such factors as its demographic composition, gender composition or where they live. We consider three alternatives. Under the first, which is similar to child benefit in developed countries, households receive a fixed lump-sum transfer for every child aged less than 5 years old (CHBEN). Under the second, only school-aged children aged 5–19 receive a lump-sum transfer (SCHBEN). The third alternative uses geographic targeting where only households in the poorest municipalities receive these SCHBEN transfers (GEOGM). The schedule of transfers is set out in Table 1.
- (2) *Means-tested targeting:* This approach involves basing eligibility for transfers on household income. For the purpose of illustration we assume that all households falling into the bottom 30 percent of the income distribution, using consumption per adult equivalent as our measure of income, are considered as being “poor” and thus eligible for transfers.⁶ Again we consider three alternative transfer schemes. Under the first, all poor

⁶Note that we are not addressing the issue of imperfect targeting, i.e. where one might have to use an imperfect indicator of income. For such an analysis, see Skoufias and Coady (2002).

TABLE 2
DECILE MEAN INCOME RATIOS AND CUMULATIVE TRANSFER SHARES

Income Decile	Mean Income Ratio	Program Cumulative Transfer Shares						
		FSUB	CHBEN	SCHBEN	GEOGM	P-UNIF	P-CHBEN	P-SCHBEN
Bottom	1.00	0.02	0.12	0.15	0.32	0.29	0.35	0.37
2	1.55	0.05	0.23	0.28	0.52	0.62	0.69	0.70
3	2.00	0.10	0.34	0.41	0.68	1.00	1.00	1.00
4	2.50	0.15	0.44	0.52	0.79	1.00	1.00	1.00
5	3.03	0.22	0.55	0.62	0.87	1.00	1.00	1.00
6	3.69	0.29	0.65	0.71	0.93	1.00	1.00	1.00
7	4.51	0.39	0.75	0.80	0.96	1.00	1.00	1.00
8	5.72	0.51	0.84	0.88	0.99	1.00	1.00	1.00
9	7.95	0.67	0.93	0.94	1.00	1.00	1.00	1.00
Top	17.73	1.00	1.00	1.00	1.00	1.00	1.00	1.00

Notes: Mean income ratio is the ratio of decile mean incomes to the mean for the bottom decile. The concept of income used is household consumption per adult equivalent. The programs are: universal food subsidies (FSUB), transfers to households with children aged less than 5 years (CHBEN), transfers to households with school-aged children aged 5–19 (SCHBEN), these to households in the poorest municipalities only (GEOGM), a uniform transfer to all poor households, and the above demographic transfers only to poor households (P-CHBEN and P-SCHBEN).

households receive a uniform lump-sum transfer (P-UNIF). Under the second and third alternatives, these poor households receive transfers based on household composition as above, denoted P-CHBEN and P-SCHBEN respectively.

The program P-SCHBEN is used to determine the budget, which is kept constant across all programs (see Table 1 for further discussion). The data source is the nationally representative Mexican household survey for 1996, which contains information on household expenditures as well as demographic composition. Mean income in our sample is \$445 per month and the poverty line comes out at around \$200 per month.⁷ The average poverty gap is 53 percent of the total income of all poor households, i.e. on average we would have to increase their incomes by this amount to eliminate poverty, and even then this would have to be distributed optimally (i.e., no leakage to the non-poor and transfers being just sufficient to bring the household to the poverty line). The budget comes out at nearly 12 percent of the poverty gap.

Table 2 presents the ratio of decile mean incomes to the mean income of the lowest decile as well as the cumulative transfer shares across deciles for each of the seven program alternatives. The first thing to note is that income inequality is quite high even among the poor. For example, the mean income in the second lowest decile is 1.55 times greater than that in the lowest decile. Comparing cumulative transfer shares, which is analogous to comparing concentration curves, we see a fairly robust ranking of programs with their welfare impacts increasing from left to right. The only departure from this ranking is for geographic targeting (GEOGM) and P-UNIF, with the former dominating when we focus on the share of transfers accruing to the bottom decile and the latter when we focus on this

⁷We use the \$ sign to denote Mexican pesos.

TABLE 3
EVALUATING WELFARE IMPACT OF PROGRAMS

Program	Distributional Characteristic			Poverty-based Measures			
	$\lambda(\varepsilon = 1)$	$\lambda(\varepsilon = 2)$	$\lambda(\varepsilon = 5)$	Poverty Gap Index	Severity of Poverty	Undercoverage	Leakage
	FSUB	0.24	0.10	0.04	0.090	0.040	0.00
CHBEN	0.45 (0.88)	0.29 (1.90)	0.19 (3.75)	0.077 (0.14)	0.031 (0.23)	0.48	0.62
SCHBEN	0.49 (1.04)	0.33 (2.30)	0.22 (4.50)	0.071 (0.21)	0.027 (0.33)	0.13	0.62
GEOGM	0.67 (1.79)	0.53 (4.30)	0.41 (9.25)	0.060 (0.33)	0.022 (0.45)	0.52	0.34
P-UNIF	0.71 (1.96)	0.54 (4.40)	0.38 (8.50)	0.046 (0.49)	0.017 (0.58)	0.00	0.00
P-CHBEN	0.74 (2.08)	0.60 (5.00)	0.44 (10.00)	0.057 (0.37)	0.022 (0.45)	0.49	0.00
P-SCHBEN	0.74 (2.08)	0.60 (5.00)	0.45 (10.25)	0.043 (0.52)	0.014 (0.65)	0.13	0.00

Note: Program details described in Tables 1 and 2. Prior to the program the poverty gap was 0.097 while the severity of poverty was 0.043. Numbers in brackets are welfare impacts relative to universal food subsidies.

share for the bottom two deciles. Note also that universal food subsidies are regressive with the poor receiving less than their population share.

Table 3 presents the welfare impacts of programs based on an analysis of their distributional characteristic (for alternative aversions to inequality) as well as on poverty-based indicators. The first column presents the welfare impact based on the distributional characteristic with $\varepsilon = 1$, i.e. with the underlying welfare weights given by the relative decile means presented in Table 2 above. As we take higher values for ε (and the welfare weights are raised to the power of two and five respectively) the distributional characteristic for each program converges towards the share of transfers received by the poorest decile. The resulting program ranking mirrors that emerging from a comparison of decile transfer shares and, consistent with this, the ranking of GEOGM and P-UNIF switch in favor of the former for $\varepsilon = 5$.

The poverty gap and severity of poverty indices give a very similar ranking of programs with the exception that the rank order of P-UNIF and P-CHBEN is reversed. Note also that both these poverty indices rank P-UNIF above GEOGM, consistent with the underlying welfare weights putting less weight on transfers to the poorest decile compared to the distributional characteristic with $\varepsilon = 5$.

An examination of program leakage and undercoverage rates helps to bring out a trade-off in moving to a more targeted program. Under universal food subsidies, undercoverage is zero and leakage is 70 percent, the highest it can be for these simulations. In moving to categorical targeting based on demographic composition (i.e. CHBEN and SCHBEN), leakage is reduced but at the expense of a higher undercoverage rate, especially for CHBEN. Adding geographic targeting similarly decreases leakage but again at the cost of higher undercoverage. Adding a means test criterion to CHBEN and SCHBEN (i.e. moving to P-CHBEN and P-SCHBEN respectively) helps to eliminate leakage totally (at least under our

naïve assumption of perfect information). Removing categorical targeting from these, i.e. by using a means test in isolation (P-UNIF), enables one to eliminate undercoverage.

The magnitude of the welfare gains from moving to better targeted programs is also presented in Table 3 (in brackets), which shows the welfare gains for each program relative to the welfare gains from universal food subsidies. We concentrate on comparing the distributional characteristics with the two poverty indices since, of the indices discussed earlier, the poverty indices are the only ones that use weights that reflect income differences and not solely income ranks. The striking difference is that both the poverty indices indicate substantially lower gains from targeting. For example, with the severity of poverty index, the maximum gain from targeting is a 65 percent increase for P-SCHBEN compared to universal food subsidies. This gain is substantially lower than those indicated by the distributional characteristic, which indicates that welfare increases by factors of around 2 and 10 times for the same program for $\varepsilon = 1$ and $\varepsilon = 5$ respectively. For increasing values of ε , the welfare gains will obviously converge to an even higher ratio, i.e. the ratio of transfer shares of the poorest decile across program comparisons. So, even for the lowest aversion to inequality we have considered (i.e. for $\varepsilon = 1$ where decile welfare weights are given by the ratio of decile mean incomes), the gains appear substantially higher than that suggested by the severity of poverty. This reflects two fundamental differences between poverty indices and the distributional characteristics considered: (1) the poverty indices value transfer amounts in excess of the poverty gap as having zero welfare value; and/or (2) the welfare weights underlying the poverty indices exhibit less aversion to inequality.⁸ In addition, the distributional characteristic attributes a positive weight to extra income accruing to households above the poverty line.

We finish by illustrating the application of the decomposition of the distributional characteristic presented earlier to identifying the source of the welfare gains associated with moving from universal food subsidies to the various targeted programs considered. In particular, we identify how much of these gains is due to improved targeting efficiency (i.e. which households are identified as beneficiaries) and how much is due to improved redistributive efficiency (i.e. the differentiation of transfer levels across households). Figure 1 presents the breakdown for the distributional characteristic with $\varepsilon = 2$ only, since the breakdowns for $\varepsilon = 1$ and $\varepsilon = 5$ are qualitatively identical. The contrast between categorical and means testing is clear. Most of the welfare gain from categorical targeting is due to improved redistributive efficiency, whereas most of the welfare gain from adding means testing is due to improved targeting efficiency. For CHBEN and SCHBEN improved targeting efficiency accounts for only 31 percent and 26 percent of the total welfare gain, whereas for the means-tested programs this accounts for between 72 percent and 78 percent of the gains. But note also that the magnitude of the gains associated with moving from universal food subsidies to categorically

⁸In the Appendix we derive an expression for calculating poverty weights (analogous to welfare weights) and compare these to welfare weights for various aversions to inequality. Our discussion there also highlights that, for the programs under review, virtually all of the difference between their relative welfare and poverty impacts reflects the lower aversion to inequality inherent in the severity of poverty index.

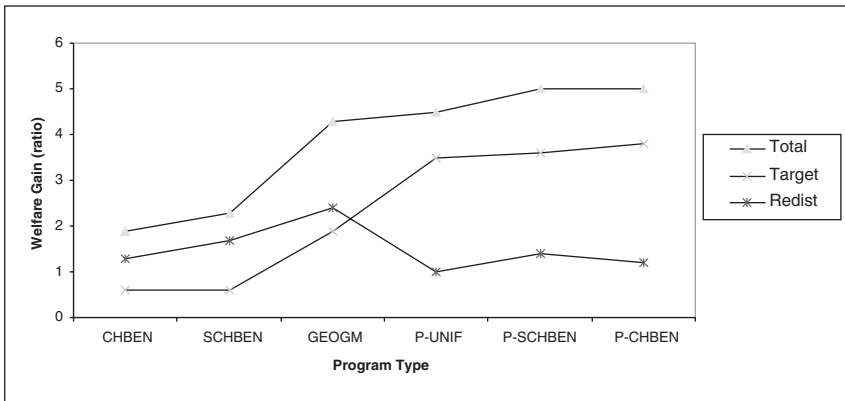


Figure 1. Total Welfare Gains Compared to Food Subsidies ($e = 2$)

targeted transfers based on household composition are only slightly lower compared to the additional gains from adding means testing to categorical targeting. This highlights the added distributional power that is potentially available from linking transfers to household composition, at least in the context of Mexico.

6. SUMMARY AND CONCLUSIONS

In this paper we argue that the distributional characteristic, which is already widely used in the literature on commodity taxation, provides an attractive alternative to conventional approaches used to evaluate the targeting performance of transfer instruments. It has a number of advantages: (1) the value judgments underlying one's analysis are made more transparent and sensitivity analysis of results to these value judgments is straightforward; (2) it allows for a wider class of social welfare functions; and (3) it avoids all of the controversy surrounding where to draw the poverty line which in practice often unnecessarily distracts from important policy issues. In a sense, poverty approaches can be interpreted as special cases. We also provide a simple, yet conceptually and empirically useful, additive decomposition of the distributional characteristic into a targeting efficiency index and a redistributive efficiency index. Using this decomposition, we show that the commonly used leakage and under-coverage rates are a special (and restrictive) case of the former; thus completely ignoring the latter, and providing a valid welfare measure only for programs with uniform transfers and even then for a set of welfare weights consistent only with the poverty gap. We also show that a simple extension of the redistributive efficiency index provides a useful generalization of standard indices of progressivity or regressivity.

Finally, we conclude by providing an illustration of the decomposition approach to a Mexican reform strategy which replaces universal food subsidies with categorically targeted (i.e. linking transfers to the demographic composition of households and geographic targeting) and means-tested transfer schemes. Our results suggest that linking transfers to the demographic composition of households can substantially improve the welfare impact of transfer programs through

improvements in redistributive efficiency. However, these gains may come at the expense of undesirably high levels of undercoverage of poor households. Means testing similarly increases the welfare impact through improvements in targeting efficiency.

APPENDIX: COMPARING WELFARE AND POVERTY WEIGHTS

In the text we indicated that there are two fundamental differences between poverty indices and the distributional characteristic: (i) the poverty indices value transfer amounts in excess of the poverty gap as having zero welfare value, and (ii) the welfare weights underlying the poverty indices exhibit less aversion to inequality. The discussion in this appendix develops these issues further.

One can derive an expression for household poverty weights by differentiating the poverty index in (1) with respect to an extra unit of income to household i to get:

$$\frac{\partial P(\alpha)}{\partial y_i} = -\frac{\alpha}{N \cdot z} \left(\frac{z - y_i}{z} \right)^{\alpha-1}$$

where y_i is household income, z is the poverty line, $(z - y_i)$ is the household poverty gap, N is the total number of households, and α captures our aversion to depth of poverty. This expression gives the contribution of a unit income to household i to reducing the overall poverty level.⁹ Consider $\alpha = 2$, i.e. the severity of poverty index. For all (non-poor) households with $y_i > z$, the poverty gap is set at zero so that poverty weights are set at zero. For poor households, their relative poverty weights are given by the ratio of their household poverty gaps, with those nearer the poverty line having lower poverty weights.

The above assumes that program transfers are “small” in the sense that the transfer received by a household is less than its poverty gap so that all of it contributes to decreasing the poverty index—we thus refer to these as “marginal poverty weights.” For “large” transfers one needs to adjust the above welfare weights to allow for the fact that transfers in excess of a household’s poverty gap do not count towards decreasing the poverty index. Denoting the transfer to household i as dy_i , the required adjustment for households with $dy_i > (z - y_i)$ is given by $(z - y_i)/dy_i$; otherwise no adjustment is required for poor households where the transfer is less than their poverty gap. We refer to these weights as “adjusted poverty weights.” Obviously, *ceteris paribus*, the larger the transfer levels (and the program budget) the lower the impact on poverty per unit expenditure.

In Figure A1 we present the marginal poverty weights for poor households in our data set, the adjusted poverty weights referring to the transfer levels under our targeted uniform transfer program (P-UNIF), as well as for another hypothetical transfer program with double these transfers (P-UNIF2). The marginal poverty weights (top line) decrease linearly with consumption. The adjusted poverty weights for our uniform transfer program, P-UNIF (middle line), coincide with these marginal weights for most poor households and are only slightly

⁹See Kanbur (1987) for a related discussion in the context of a uniform income transfer program.

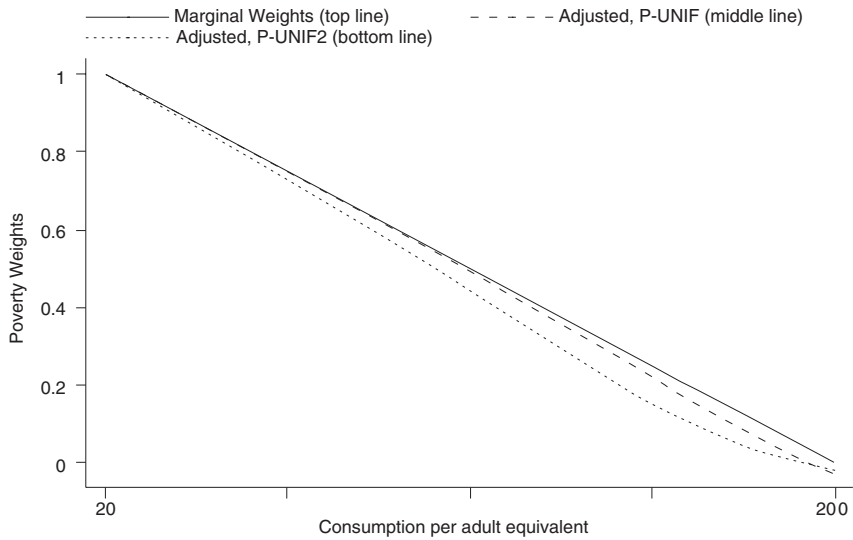


Figure A1. Poverty Weights for Severity of Poverty




Figure A2. Comparing Marginal Poverty and Welfare Weights

lower for those households very near the poverty line. For the other transfer programs considered in the text, the marginal and adjusted poverty weights are even more similar than for P-UNIF. One expects these weights to diverge more with higher transfers. This can be seen from the bottom line, which presents the adjusted weights for a program with twice the transfer levels (and thus budget) of P-UNIF,

denoted P-UNIF2. For these transfers the adjusted weights begin to diverge at much lower levels of consumption.

Of course, the higher transfer levels will still be expected to have a substantially higher impact on poverty. Before the programs, the severity of poverty index had the value 0.0432. The P-UNIF transfer program decreased this to 0.0166, a decrease of around 60 percent. Under P-UNIF2 the severity index decreases to 0.0064, a decrease of around 85 percent. Note that since the distributional characteristic is based on transfer shares, scaling up transfer levels will not affect its value.

Finally, in Figure A2 we compare the structure of poverty weights with welfare weights for different degrees of aversion to inequality; these have all been normalized by the weights attributed to the poorest household. The lines show clearly that the relative weights decrease more steeply for higher aversions to inequality. But, whereas the relative welfare weights decrease at a decreasing rate, the poverty rates decrease more or less linearly. The relative welfare weights for $\varepsilon = 2$ and $\varepsilon = 5$ are everywhere below the relative poverty weights for all poor households, whereas those for $\varepsilon = 1$ cross the relative poverty weights line left of the poverty line.

REFERENCES

- Ahmad, E. and N. Stern, *Tax Reform in Developing Countries: Theory and Practice*, Cambridge University Press, Cambridge, 1991.
- Atkinson, A., "On the Measurement of Inequality," *Journal of Economic Theory*, 2(3), 244–63, 1970.
- , "On the Measurement of Poverty," *Econometrica*, 55, 749–64, 1987.
- Baker, J. and M. Grosh, "Poverty Reduction Through Geographic Targeting: How Well Does It Work?" *World Development*, 22(7), 983–95, 1994.
- Besley, T. and I. Preston, "Invariance and the Axiomatics of Income Tax Progression: A Comment," *Bulletin of Economic Research*, 40, 159–63, 1988.
- Case, A. and A. Deaton, "Large Cash Transfers to the Elderly in South Africa", *Economic Journal*, 108, 1330–61, 1998.
- Coady, D., "An Evaluation of the Distributional Power of PROGRESA's Cash Transfers in Mexico," Discussion Paper, No. 117, Food Consumption Nutrition Division, International Food Policy Research Institute, Washington D.C., 2001.
- Coady, D., M. Grosh, and J. Hoddinott, "Targeting Outcomes Redux," Discussion Paper, No. 144, Food Consumption Nutrition Division, International Food Policy Research Institute, Washington D.C., 2002.
- Cornia, G. and F. Stewart, "Two Errors of Targeting," in D. van de Walle and K. Nead (eds), *Public Spending and the Poor: Theory and Evidence*, John Hopkins University Press, Baltimore, 1995.
- Deaton, A., *The Analysis of Household Surveys: A Microeconomic Approach To Development Policy*, John Hopkins University Press, Baltimore, 350–86, 1997.
- Duclos, J.-Y., "Measuring Progressivity and Inequality," *Research in Economic Inequality*, 7, 19–38, 1997.
- , "Social Evaluation Functions, Economic Isolation, and the Suits Index of Progressivity," *Journal of Public Economics*, 69(1), 103–21, 1988.
- Foster, J., J. Greer, and E. Thorbecke, "A Class of Decomposable Poverty Measures," *Econometrica*, 52, 761–6, 1984.
- Kanbur, R., "Measurement and Alleviation of Poverty," *IMF Staff Papers*, 34, 60–85, 1987.
- Keen, M., H. Papapanagos, and A. Shorrocks, "Tax Reform and Progressivity," *Economic Journal*, 110, 50–68, 2000.
- Lambert, P., *The Distribution and Redistribution of Income*, Manchester University Press, Manchester, 1993.
- Milanovic, B., "The Distributional Impact of Cash and In-Kind Transfers in Eastern Europe and Russia," in D. van de Walle and K. Nead (eds), *Public Spending and the Poor: Theory and Evidence*, John Hopkins University Press, Baltimore, 489–520, 1995.

- Newbery, D. and N. Stern (eds), *The Theory of Taxation for Developing Countries*, Oxford University Press, Oxford, 1987.
- Pfahler, W., "Redistributive Effects of Tax Progressivity: Evaluating a General Class of Aggregate Measures," *Public Finance*, 42(3), 1–31, 1987.
- Pfingsten, A., "The Measurement of Tax Progression," *Studies in Contemporary Economics*, 20, Springer-Verlag, Berlin, 1986.
- Ravallion, M., "Poverty Alleviation Through Regional Targeting: A Case Study for Indonesia," in K. Hoff, A. Braverman, and J. Stiglitz (eds), *The Economics of Rural Organization*, Oxford University Press, Oxford, 1993a.
- , *Poverty Comparisons: A Guide to Concepts and Measurement*, LSMS Working Paper 88, World Bank, Washington D.C., 1993b.
- Ravallion, M. and K. Chao, "Targeted Policies for Poverty Alleviation Under Imperfect Information: Algorithms and Applications," *Journal of Policy Modelling*, 11, 213–24, 1989.
- Schady, N., "Picking the Poor: Indicators for Geographic Targeting in Peru," *Review of Income and Wealth*, 34(3), 417–33, 2002.
- Skoufias, E. and D. Coady, "Are the Welfare Losses From Imperfect Targeting Important?" Discussion Paper, No. 125, Food Consumption Nutrition Division, International Food Policy Research Institute, Washington D.C., 2002.