Review of Income and Wealth
Series 48, Number 3, September 2002

PICKING THE POOR: INDICATORS FOR GEOGRAPHIC TARGETING IN PERU

BY NORBERT R. SCHADY*
The World Bank

Geographic targeting is perhaps the most popular mechanism used to direct social programs to the poor in Latin America. This paper empirically compares geographic targeting indicators available in Peru. To this effect, I combine household-level information from the 1997 Peru Living Standards Measurement Survey (LSMS) and district-level information from the 1993 Peru Population and Housing Census. I then conduct a series of simulations which estimate leakage rates, concentration curves, the impact of transfers on poverty as measured by the headcount index, poverty gap and P_2 measures of the FGT family, and non-parametric (kernel) densities when transfers are based on alternative indicators. I conclude that there is substantial potential for geographic targeting in Peru. However, the differences in outcomes across geographic targeting indicators are small, and are not statistically significant. These results are in keeping with earlier work which suggests that (among reasonable alternatives) the choice of geographic targeting indicator does not have an important bearing on poverty outcomes, and are at odds with more recent research which stresses the advantage of poverty maps which "impute" consumption or income.

1. INTRODUCTION

Lack of information is a serious constraint to targeting social programs effectively, especially in less developed countries (LDCs). Targeting social programs involves making distinctions between "deserving" (poor) and "undeserving" (non-poor) applicants. But this is no simple matter in countries where household characteristics such as income are rarely known. In such circumstances, policymakers intent on targeting are forced to choose among imperfect solutions. They can rely on observable household characteristics, such as land ownership, the ratio of working age-adults to dependants, or ownership of durable goods that seem likely to separate poor from non-poor households. They can "self-target" programs by designing them so that they appeal mainly to the poor—perhaps by offering employment at below-market wages, or subsidizing foodstuffs consumed primarily by the poor. Or they can use "geographic targeting" to direct resources to areas in which, on average, poverty appears to be greatest (Akerlof, 1978; Besley and Kanbur, 1990; Grosh, 1992).

Note: I thank particularly Christina Paxson for much helpful advice. I have also benefited from comments made by Greg Felker, Francisco Ferreira, Jesko Hentschel, Carlos Sobrado, John Waterbury, two anonymous referees, and participants at a seminar held at the World Bank. I am indebted to Marcos Robles for help obtaining the necessary data from INEI, and to Vajeera Dorabawila for help with the consumption aggregates for the 1994 and 1997 LSMS. An earlier version of this paper appeared as a World Bank Policy Research Working Paper 2477 (November 2000). The views and interpretations expressed in this paper are those of the author and do not necessarily represent the views and policies of the World Bank.

*Correspondence to: Norbert Schady, Poverty and Economic Management, Latin America and Caribbean, World Bank, 1818 H Street N.W., Washington, DC 20433, USA (nschady@worldbank.org).

Geographic targeting is appealing because it is comparatively simple to administer. Different parts of a country—regions, provinces, districts, even city blocks—are ranked by some measure of deprivation. This measure could be income-based poverty or, more commonly, an indicator of health, educational or nutritional status, or access to basic services, such as electricity or running water. Resources are then allocated in inverse proportion to average welfare, so that poor regions receive higher per capita transfers than rich ones. Alternatively, rich areas can be excluded from the program altogether. The simplicity of geographic targeting is an important advantage when lack of information or administrative capacity is a serious concern.

This paper compares a number of geographic targeting indicators that have been discussed by policy-makers in Peru. These include the infant mortality rate, which is used to target the Municipal Compensation Fund, the main block grant from the central government to local governments in Peru; the rate of chronic malnutrition; a composite "poverty" index developed by the Peruvian Social Fund (FONCODES), which FONCODES uses to target its projects; and an estimate of "imputed" poverty which combines census and survey data in an attempt to approximate money-based measures of welfare.

To test the potential impact on poverty of targeting with alternative indicators, I conduct a series of simulations which combine information on household expenditures from the 1997 Peru Living Standards Measurement Survey (LSMS) with district-level averages of various welfare measures. I use the results of the simulations to compare leakage rates, trace out concentration curves, estimate the impact on various poverty measures of the Foster–Greer–Thorbecke (FGT) family, and graph non-parametric (kernel) density estimates of the log of per capita expenditures (PCE) when transfers are made on the basis of alternative indicators.

The results show that geographic targeting with any of the indicators is a significant improvement over an untargeted regime in which resources are distributed equally across all districts. There is, therefore, substantial potential for geographic targeting in Peru. However, targeting outcomes are quite similar for all of the targeting indicators I consider, and differences in outcomes between indicators are not significant at conventional levels of significance. These findings are in keeping with earlier work which suggests that (among reasonable alternatives) the choice of geographic targeting indicator does not have an important bearing on poverty outcomes (Glewwe, 1992; Baker and Grosh, 1994), and are at odds with more recent research which stresses the advantages of poverty maps which "impute" consumption or income (Bigman and Fofack, 2000; Hentschel *et al.*, 2000).

The rest of the paper proceeds as follows. Section 2 briefly describes alternative geographic targeting indicators available in Peru. Section 3 describes the methodology I use to compare indicators. It also describes alternative performance measures, identifies formulas that can be used to allocate resources to districts, describes data requirements, and states the assumptions made for the simulation exercise. Section 4 summarizes the results of the analysis. Section 5 draws conclusions.

2. Geographic Targeting Indicators in Peru

In 1997 there were 13 administrative regions, 24 departments, 194 provinces, and 1812 districts in Peru (Webb and Fernández Baca, 1997, p. 112). Recent discussion about the geographic targeting of government programs has focused on the use of *district*-level averages. Districts can be quite small in Peru: according to the 1993 Population and Housing Census, the average district population was about 12,600 inhabitants, but some, predominantly rural districts had less than 200 inhabitants.

Infant Mortality

The infant mortality rate measures the fraction of children ever born who do not reach the first year of age. In Peru, as in many other developing countries, one would be loathe to estimate infant mortality on the basis of seriously incomplete death registries. To address this problem, demographers have developed indirect methods to estimate mortality at early ages (for example, Trussell and Menken, 1984). The Peruvian National Statistical Institute (INEI) used one such method—the preceding births technique—to estimate infant mortality in Peru.

The preceding births technique requires women to report the total number of children ever born and surviving at two different points in time. This information was available in Peru from the 1981 and 1993 population censuses. However, because the population of many districts is very small, applying the preceding births technique to individual districts would have produced highly uncertain results. INEI therefore followed a two-stage procedure. First, it estimated the infant mortality rate in every department, where the sample size was large enough. In addition, INEI regressed its estimate of the infant mortality rate in every department on some of its correlates for which information had been gathered in the 1993 Population and Housing Census—indicators such as women's average education level, household characteristics, and place of residence. In the second stage, the coefficients from these departmental regressions were applied to the 1993 district-level data to estimate district-level infant mortality rates (INEI, 1997).

Chronic Malnutrition

The district-level rate of chronic malnutrition measures the fraction of children whose height for age is at least two standard deviations below that of a reference population. In Peru, this information is available from a Census of Height and Age for all first-graders conducted by the Ministry of Education in 1993. (Enrolment rates in first grade are almost universal in Peru.)

The FONCODES Index

Peru has a long history of developing "poverty maps" based on composite indices of unmet basic needs. The first of these maps was constructed by Webb with information from the 1961 population census (Webb, 1977). This poverty map was updated with new censuses conducted in 1972, 1981, and 1993 (Amat y León, n.d.; Banco Central de Reserva, 1981; INEI, 1994). FONCODES, in turn, has developed composite poverty maps since its creation in 1991, often with technical assistance from the Inter-American Development Bank, the World Bank, and the German development agency GTZ.

The district-level poverty map analyzed in this paper was developed by FONCODES and the Ministry of the Presidency. It is based on eight indicators—the rates of chronic malnutrition, illiteracy, school-aged children not in school, overcrowded housing, inadequate roofing, and the proportion of the population without access to water, sewerage, and electricity. All of these indicators except the rate of chronic malnutrition were estimated with data from the 1993 Population and Housing Census (FONCODES, 1995, 1996).

Composite indices invariably involve some arbitrary weighting of individual indicators. The FONCODES index standardized each indicator by dividing it by its minimum value, multiplied the rate of chronic malnutrition by seven, and then added all of the individual indicators.[1] For ease of interpretation, FONCODES then standardized the index by dividing all values by the lowest value. The resulting index ranges from a value of 1 to 36.38.

Imputed Poverty and Income

In Peru, there are no survey-based estimates of income or expenditures at a level more disaggregated than the department: for example, household surveys conducted by INEI, which generally have sample sizes of 15,000 to 20,000 households, can only be used to compare income across "natural regions" and departments.[2] But INEI has combined variables which are common to both the 1993 census and one such survey conducted in 1995 to develop imputed district-level measures of income and poverty (INEI, 1996).

This procedure is conceptually similar to that used to estimate infant mortality at the district level. INEI estimated income and poverty levels in 1995 on the basis of the household survey, and then regressed income and an indicator variable for poverty status in every department on its correlates—household composition, education levels, access to basic services such as water, sewerage and electricity, ownership of durable goods, such as television, radio and refrigerator, and other variables included in both the census and the survey. The coefficients from the 24 department-level regressions were then used to impute average income in every district, and the fraction of the population in each district below the poverty line.

On the surface, the methodology applied by INEI for these imputations is similar in spirit to that proposed in Hentschel *et al.* (2000): both use information common to the census and a recent household survey to impute household income (or consumption). However, there appear to be some important differences. For

[1]This procedure had the unintended consequence that the greatest weight was given to those indicators with the greatest variance. Thus, while the intended weights were 50 percent for the rate of chronic malnutrition, and 7.14 percent each for the seven other measures, the actual weights in the index turned out to be 15.3 percent for the rate of chronic malnutrition, and 3.4, 2.2, 3.0, 38.3, 8.8, 7.4, and 21.6 percent for the measures of illiteracy, school attendance, overcrowding, inadequate roofing, and access to water, sewerage and electricity, respectively (World Bank, 1996, p. 7).

[2]These natural regions are Lima, and the urban and rural areas of the coast, sierra (highlands), and selva (jungle), respectively. Natural regions do not, in general, correspond to the administrative regions mentioned before.

one, the Hentschel *et al.* methodology estimates the probability of being poor for *individual* households based on their *individual* covariates. District estimates of poverty are then obtained as the mean of the households' probabilities of being poor in a given district. By contrast, INEI estimated mean *district* poverty levels directly, on the basis of the mean of the covariates at the *district* level. To see how this could make a difference, consider two districts with the same mean values for their covariates, but very different distributions of these covariates. These districts would have the same imputed poverty by the INEI calculations but could have different values of imputed poverty by the methodology proposed in Hentschel *et al.*[3]

3. The Analytic Framework

Measures of Performance

The simplest measure of targeting focuses on leakage and undercoverage rates (Grosh, 1992, pp. 16–17; Baker and Grosh, 1994). A poverty line is chosen to separate "poor" from "non-poor." Leakage rates are then defined as the fraction of total program resources which go to the non-poor, and undercoverage rates as the fraction of the poor who do *not* benefit from the program. By this measure, better geographic targeting indicators result in lower leakage and lower undercoverage rates.

A second approach simply ranks individuals by an indicator of welfare—say, per capita expenditures—and then cumulates the fraction of households and the fraction of resources transferred by different indicators. The results are often presented in terms of so-called "concentration curves" (see, for example, Milanovic, 1995), and I follow this practice below. By this measure, the best targeting indicator is that whose concentration curve is above all others at every point. The concentration curve method does not require use of a poverty line, which may be an advantage given the fact that setting poverty lines can be contentious.

Alternatively, one might want to compare the changes in various poverty measures which are likely to result when transfers are based on alternative targeting indicators (Chaudhuri and Ravallion, 1994). This is an exercise in comparative statics: what is total poverty before and immediately after the transfer? By this measure, the preferred geographic targeting indicator directs limited resources to areas where they would have the greatest short-term impact on poverty. More complex formulations, which might model the expected long-term returns from transfers to different districts, are beyond the scope of this paper.

In what follows I use three poverty measures from the Foster–Greer–Thorbecke (FGT) family—the headcount index, the poverty gap, and the P_2 measure (Foster, Greer, and Thorbecke, 1984). The FGT family of poverty measures follows the general formulation below:

$$P_\alpha = \frac{1}{N}\sum_{i=1}^{N}(1 - y_i/z)^\alpha \qquad (\text{for all } y_i < z) \tag{1}$$

[3]I thank Carlos Sobrado and an anonymous referee for this insight.

where y is income, z the poverty line, and α is a parameter which represents the aversion to inequality. When $\alpha = 0$, P_0 corresponds to the headcount index—the number of people below the poverty line; when $\alpha = 1$, P_1 corresponds to the poverty gap—a sum of the individual shortfalls in income for those below the poverty line, as a fraction of the poverty line itself. As α increases, the measure gives a greater weight to the poorest poor, and at very high values of α, P_α approaches a "Rawlsian" measure of welfare which gives weight only to the poorest household. The P_2 measure corresponds to a value of $\alpha = 2$.

Finally, one could look at changes in the entire distribution of log PCE, rather than just at changes for those below the poverty line. I use non-parametric (kernel) density estimates for this purpose.

Allocation Formulas

When there is no targeting, districts are simply allocated resources according to their share of the total population in the country, and everyone is assumed to receive the same per capita transfer. This no-targeting scenario serves as a benchmark to measure additional reductions in poverty that could be achieved when geographic targeting is conducted on the basis of some welfare indicator.

To compare targeting indicators, one must develop a formula which allocates resources across districts. I consider one such formula, which is relevant for Peru because it has been the basis of targeting by FONCODES, the program which made the most significant early advances developing targeting indicators in Peru. FONCODES ranks all districts by its poverty index. It then allocates resources to each district according to the following formula:

$$\text{Allocation}_i = \frac{(\text{Index}_i \times \text{Population}_i)}{\sum_{j=1}^{n} (\text{Index}_j \times \text{Population}_j)} \tag{2}$$

The "FONCODES method" thus makes all districts in the country eligible for benefits, but weights the population of each one by its poverty index. For example, Coronel Castañeda, the district with the highest value of the FONCODES index (36.38), and a population of 607 inhabitants, would be allocated $(36.38 \times 607)/346{,}201{,}217 = 0.0064$ percent of the total budget for that year. By contrast, Pacocha, the district with the lowest value of the FONCODES index (1.00), and a population of 6500, would be allocated $(1 \times 6{,}500)/346{,}201{,}217 = 0.0019$ percent of the total budget for that year. Per capita allocations to inhabitants of Coronel Castañeda would therefore be almost 37 times per capita allocations to inhabitants of Pacocha.[4] Note that if the index is a poverty rate, allocations to district i simply correspond to the fraction of the poor who live in district i. I adapt the "FONCODES method" to other indicators by substituting

[4]In actual fact, FONCODES' allocation mechanism has traditionally been a little more complicated than this. FONCODES first allocates 60 percent of resources to rural areas and 40 percent to urban areas. The final allocation to each district is then the sum of the rural and urban allocations—standardized to add up to 100 percent. Ad hoc adjustments are also made to privilege border areas, to coordinate investments with other public sector programs, and to ensure that every FONCODES regional office (which correspond roughly to individual departments) has a minimum operating budget. I do not take these "refinements" into account in the simulations below.

the infant mortality rate, the rate of chronic malnutrition, and the imputed poverty measure for the FONCODES index in equation (2) above, and compute the corresponding district-level allocations.[5]

The simulations in this paper make a number of assumptions. The most important assumption is that there is no targeting of program resources *within* a given district. This is clearly unrealistic: Paxson and Schady (2002) use non-parametric regressions to estimate the probability of benefiting from social programs as a function of the number of standard deviations a given household's income is above or below mean district income. Their results show that the within-district distribution of investments made by FONCODES is hump-shaped, peaking at about one and a half standard deviations above mean district income, while the within-district distribution of investments made by the school construction program INFES is mildly regressive, rising with household income. A more plausible assumption for the simulations in this paper is that the degree of within-district targeting is independent of the choice of welfare indicator which is used to distribute resources *across* districts. Under this assumption, which seems reasonable, the actual estimates of changes in poverty under different targeting regimes will be biased (up, if there is positive intra-district targeting, down if the converse is true), but the preferred rank-order of the indicators used to assign resources *across* districts should be unaffected.

The simulations assume that benefits from program investments in a district accrue entirely to the residents of that district. This might not hold, say, if beneficiaries of a food distribution program implemented in one district are residents of a different district. But inter-district spill-overs are unlikely to be systematic—that is, they should not consistently favor residents of one kind of district over residents of another. Inter-district spill-overs should therefore not affect the rank order of indicators either.

Some additional assumptions have to be made about the impact of transfers on various poverty measures. Poverty in Peru has generally been defined as an individual's inability to meet a specified level of expenditures—the poverty line—when individual expenditures are approximated by total household expenditures divided by the number of eligible household members.[6] We must therefore translate expenditures by social programs into household expenditures—by first translating program expenditures into changes in household income, and then estimating the proportion of additional disposable income that is spent. As a matter of convenience, I have assumed that all program expenditures translate

[5]Of course, this is only one of a potentially infinitely large number of formulas which could be used. For example, in the mid–1990s, the Technical Team of the Ministry of the Presidency proposed an allocation formula which made a distinction between districts with a high *proportion* of poor people (as measured by the FONCODES index) and districts with a large *number* of poor people (as measured by the product of the FONCODES index and population). Proposed investments would then be directed to 262 districts with the highest proportion of the population in poverty and 232 districts with the highest number of people in poverty.

[6]The 1995 and 1996 INEI household surveys, which used income to measure poverty, are exceptions.

into additional household income and that all of this additional income is spent.[7]

Finally, the simulations assume that the cost of administering programs is constant across regions, and ignore the effects of transfers on behavior such as migration towards districts which receive large per capita transfers, offsetting reductions in private intra-household transfers or employment, and the impact of taxes needed to finance poverty alleviation programs.

The Data Set

For all of the estimations below, I combine information from two sources: district-level averages of the infant mortality rate, the rate of chronic malnutrition, the FONCODES index, and the measure of imputed poverty, and household-level data on expenditures. District-level averages are available from INEI, and household-level data can be estimated from the 1997 LSMS.

District-level data can be used to estimate the proportion of total funds that would be allocated to every district under alternative targeting regimes. Further dividing this fraction by the total population of the district in question allows us to calculate the proportion of funds that would be allocated to every individual. Finally, multiplying this proportion by the total budget available for poverty alleviation programs, we can estimate per capita transfers.

The 1997 LSMS can be used to estimate the expenditures of the 3,840 households in the sample and, dividing total household expenditures by household size, for 19,562 household members. These estimates can be combined with information on poverty lines to calculate the headcount index, poverty gap, and P_2 measure at a national level before any transfers take place.[8]

The 1997 Peru LSMS drew households from 397 clusters, and the accompanying literature lists the districts from which each one of these clusters was drawn. Observations in the LSMS can be coded manually with district identifiers which match those used by INEI, and district-level and household-level data

[7]Two points are worth noting here. First, many social programs in Peru are involved with the construction of small-scale infrastructure. The wages paid to laborers in these projects are only a fraction of the total cost, and other benefits—say, of having an additional classroom—are unlikely to have a short-term impact on household income. We can easily relax the assumption of a one-for-one equivalence between changes in program expenditures and changes in household income, however, if we simply model a smaller budget. For example, if only 50 percent of the expenditures on poverty alleviation programs translate into short-term increases in income, the "relevant" budget would be half the actual budget. Second, households typically do not spend all of an increase in income. Assuming that all households spend a fixed fraction of additional income is not entirely satisfactory either, because the marginal propensity to save is likely to differ *systematically* across households. For example, if the fraction of income that is saved is higher in rich districts than in poor districts, simulations based on a constant marginal propensity to save might under-estimate the short-term impact of program investments in poor districts relative to rich districts, and under-estimate the relative performance of targeting indicators which assign a higher share of their resources to the poorest districts. One potential solution would be to estimate marginal propensities to save for households from the LSMS itself. The simplest way to do this would be to convert measures of income and expenditures into current prices, and then take the difference between them as a measure of savings (Paxson, 1992). By this measure, however, almost two-thirds (61 percent) of households in the 1994 LSMS *dissaved*. This seems unreasonable and suggests that income in these surveys is seriously under-estimated *vis-à-vis* expenditures.

[8]Note that regional price deflators are used throughout the paper to deflate *both* household expenditures and simulated transfers.

can then be merged. Having done this, I keep only those observations for which there are matching codes for place of residence in both data sets—in effect, discarding the district-level information for all but the 238 districts which were sampled in the 1997 LSMS.[9] The new, composite data set is representative in exactly the same way as the original LSMS data set, and can be used to make reasonable simulations about changes in expenditures and poverty at the national level. I also compute all of the results I report in this paper using the 1994 LSMS. These results are extremely similar to those calculated on the basis of the 1997 LSMS. I therefore report results only with the 1997 data (the 1994 results are available from the author upon request).

4. Results

What is the effect of geographic targeting with alternative indicators in Peru? As a first step towards answering this question, I present estimates of poverty and allocations of funds by region. About two-thirds of the population of Peru lives in urban areas, and well over a third of these urban residents live in the capital city, Lima. Table 1 decomposes poverty measures and allocations by alternative geographic targeting indicators into three categories: Lima, other urban, and rural. The values in each cell correspond to the proportion of total poverty or the proportion of total allocations by region, so that every row sums to 100 percent. Here, and for all of the results presented in the paper, individuals in the household surveys are weighted with the appropriate expansion factors.

Since the no targeting scenario makes an equal transfer to every Peruvian, regional allocations when there is no targeting correspond exactly to the fraction of the population living in each region. A comparison of these allocations with the various poverty measures shows that poverty in urban areas is below average: about 29 percent of the population lives in Lima, but only 13–20 percent of total poverty was found there in 1997. Other urban areas account for about 36 percent of the population, and about one-third of poverty. Poverty in rural areas, by contrast, is well above average: just over one-third of the population lives in rural areas, but between 47 and 57 percent of poverty was found there in 1997.

[9]In theory, the first step in the FGT approach implemented in the simulations in this paper would divide a given budget among the 1812 districts in Peru. How much gets allocated to each district would then depend on the targeting indicator and the allocation formula in question. But the fact that households in the 1997 LSMS were only drawn from 238 districts across the country raises a potential problem: since every targeting indicator allocates a different amount to each district, the total budget for this sample of districts would *not* be constant across indicators. For example, the proportion of the total budget allocated to the 238 districts in the 1997 LSMS would be smaller by the FONCODES index than by the measure of imputed poverty. As a result, the total amount transferred to the 3,840 households in the survey would be smaller when we use the FONCODES index than when we use the imputed poverty measure, even after each household in the survey is weighted by its expansion factor. At the heart of the problem is the fact that the LSMS draws a nationally-representative sample of households irrespective of the district in which these households live. If a number of samples were drawn, on average, the total budget would be the same across indicators, but this does not hold for any one sample. I have corrected for this problem by normalizing the budget—in effect, summing equation (2) above only over the sample of districts in the LSMS. Note that this is not an issue with the concentration curve approach because concentration curves graph out the *proportion* of a given budget that is allocated to each household in the survey—weighted, once again, by the appropriate expansion factors. Because concentration curves are mean-normalized in this way, the results are budget-independent.

TABLE 1

DISTRIBUTION OF POVERTY AND ALLOCATED EXPENDITURES, BY REGION AND TARGETING INDICATOR

	Year	Lima	Other Urban	Rural
Headcount index	97	20.24	32.66	47.11
Poverty gap	97	15.62	31.55	52.82
P_2 measure	97	12.87	30.61	56.52
Imputed poverty	97	15.76	36.53	47.71
Infant mortality	97	13.99	33.39	52.62
Chronic malnutrition	97	14.71	34.51	50.78
FONCODES index	97	16.59	32.62	50.78
No targeting	97	28.85	35.56	35.56

Note: The values in each cell correspond to the proportion of total poverty or the proportion of total allocations by region. Individuals in the household survey are weighted by the appropriate expansion factors.

Table 1 suggests that geographic targeting using any of the indicators under consideration appears to approximate the distribution of poverty reasonably well. When any of these indicators is used for geographic targeting, about one-half of all resources are transferred to rural areas, and about two-thirds of the remaining resources to urban areas outside Lima. Comparing the three indicators, the FONCODES index transfers more to Lima than the measures of infant mortality, chronic malnutrition, and imputed poverty, while the measure of infant mortality makes the largest transfers to rural areas.

In Table 2, I present Spearman correlation coefficients for the various targeting indicators. This table shows that all of the indicators are highly correlated with each other: to some extent, this is the case because all of the indicators draw on the same two "primary" data sets—the Population and Housing Census of 1993, and the Census of Height and Age of 1993. However, the Spearman correlation coefficients appear to be no lower for those indicators which draw on different data sources (for example, between the rate of chronic malnutrition, which uses the Census of Height and Age and the measure of imputed poverty, which uses the Population and Housing Census) than for those which draw on the same source (for example, the measures of imputed poverty and the infant mortality rate, both of which draw on the Population and Housing Census only). It appears that in Peru there is truly a very high degree of correlation between different

TABLE 2

SPEARMAN CORRELATION COEFFICIENTS BETWEEN ALTERNATIVE GEOGRAPHIC TARGETING INDICATORS

	Imputed Poverty	Infant Mortality	Chronic Malnutrition	FONCODES Index
Inmputed poverty	1.000			
Infant mortality	0.664	1.000		
Chronic malnutrition	0.672	0.781	1.000	
FONCODES index	0.736	0.647	0.702	1.000

Note: Spearman correlation coefficients calculated on the basis of the district-level data. All reported coefficients are significant at the 1% level or better.

TABLE 3

LEAKAGE RATES, BY TARGETING INDICATOR

Targeting Indicator	Leakage Rate
Imputed poverty	44.56
Infant mortality	44.45
Chronic malnutrition	45.08
FONCODES index	43.76
No targeting	51.15

Note: Leakage rates correspond to the fraction of the estimated transfers which are received by households above the poverty line.

district-level measures of welfare, a correlation which is not only a function of the fact that many indicators draw on similar data sets.

Leakage Rates

Table 3 presents leakage rates by targeting indicator. The leakage rate when there is no targeting (51.15) corresponds exactly to the fraction of the population which is not in poverty. Table 3 shows that leakage rates would be minimized with geographic targeting by the FONCODES index according to the 1997 LSMS. Bootstrapped standard errors (not reported, but available from the author upon request) suggest that geographic targeting with any of the indicators in question is a significant improvement on the no targeting scenario, while the differences in outcomes across indicators are not significant.

Concentration Curves

Initial simulations show that the concentration curves for the FONCODES index, the infant mortality rate, the rate of chronic malnutrition, and the imputed poverty rate are so close to each other as to be virtually indistinguishable from each other on a graph. For the sake of clarity, I therefore graph the *difference* between each of the concentration curves and the no targeting baseline case in Figure 1. (The "concentration curve" for the no targeting scenario is a straight, 45 degree line: every individual receives the same transfer, so the cumulative fraction of the transfer equals the cumulative fraction of the population at every point.) Figure 1 shows that this difference is positive throughout for all four curves: No matter where we take the cut-off between "poor" and "non-poor" to be, the poor receive a larger share of transfers when these are made on the basis of the infant mortality rate, the rate of chronic malnutrition, the FONCODES index, or the measure of imputed poverty than when transfers are not targeted. Figure 1 also shows that no single concentration curve lies everywhere above all others, although the FONCODES index curve always transfers more resources to poor households than the imputed poverty or chronic malnutrition measures.

Changes in Poverty

Figure 2 considers the impact of alternative geographic targeting regimes on poverty at various budget levels between 10 million soles and 5 billion soles (for

Figure 1. Difference in Concentration Curves, by Targeting Indicator

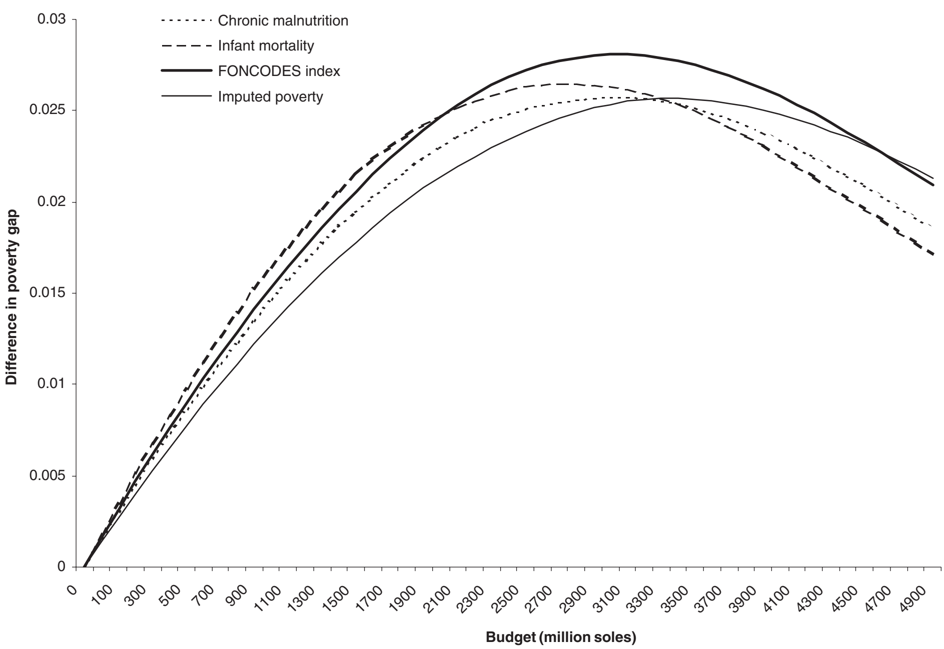

Figure 2. Differences in the Poverty Gap, by Targeting Indicator

a similar approach see Chaudhuri and Ravallion, 1994, and Jalan and Ravallion, 1998). As a point of reference, the Ministry of the Presidency, the single largest implementing agency of social programs in Peru, spent 2.2 billion soles on programs which could be targeted in 1995.[10] For the sake of parsimony, I present graphs only for the poverty gap measure. Results for the headcount index and the P_2 measure, which are available from the author upon request, are very similar.[11]

Once again, because targeting outcomes with different indicators are very similar to each other but are clearly superior to the no targeting regime, I graph the *differences* in poverty gaps. One way to understand the graphs is therefore as a double-difference: The first difference is the change in the poverty gap when a given budget is transferred—separately, for the no targeting scenario, and for geographic targeting using the FONCODES index, the infant mortality rate, the rate of chronic malnutrition, and the measure of imputed poverty. The second difference is the difference between the change in the poverty gap under the no targeting scenario, on the one hand, and the changes in the poverty gap when there is targeting by a given indicator. The value of the *y*-axis at any given budget and for any given curve is therefore the additional increase in the poverty gap which could be attained from switching from no targeting to targeting with the indicator in question.

Impact on Poverty

Like the concentration curve analysis, Figure 2 shows that targeting with any one of the indicators in question is clearly preferable to no geographic targeting: all of the curves are above zero throughout. Figure 2 also shows that none of the indicators clearly outperforms the others, although transfers based on imputed poverty appear to result in smaller decreases in poverty than the corresponding transfers made on the basis of infant mortality or the FONCODES index.[12]

Budget Savings

Table 4 presents the same information from a different angle: it considers the impact on the headcount index, poverty gap, and P_2 measures of spending the 2.2 billion sol reference budget without geographic targeting, and estimates the budget that would be necessary to achieve this same reduction in poverty

[10]These expenditures included programs in housing (Banco de Materiales, ENACE), potable water, sanitation, and electricity (UTE-FONAVI), educational infrastructure (INFES), food and nutrition (PRONAA), social services (INABIF), and numerous multi-sectoral programs (FONCODES, COOPOP, INADE) (World Bank, 1996).

[11]The curves for the headcount index tend to be much more choppy. The reason for this is that transfers will not reduce the headcount index unless at least one person in the country (or, in this case, one person in the survey) is bumped over the poverty line. This need not happen if resources are spent on those who are well below the poverty line. An increase in expenditures on the poor will always reduce poverty, however, if poverty is measured by the poverty gap or P_2 measures.

[12]Once again, bootstrapped standard errors suggest that the differences between, on the one hand, the no targeting case and, on the other hand, the measures of imputed poverty, infant mortality, chronic malnutrition, and FONCODES index are statistically significant at the 1 percent level or better. The differences between the infant mortality rate, the rate of chronic malnutrition, FONCODES index and imputed poverty are not significant (estimates available from the author upon request).

TABLE 4

ESTIMATED SAVINGS ASSOCIATED WITH GEOGRAPHIC TARGETING, BY INDICATOR

		Estimated Savings (million soles)		
	Year	Headcount Index	Poverty Gap	P_2 Measure
Imputed poverty	97	693	731	755
Infant mortality	97	770	838	845
Chronic malnutrition	97	766	781	789
FONCODES index	97	699	806	827

Note: Estimated savings are based on simulations with the reference budget of 2.2 billion soles.

when geographic targeting is conducted by the infant mortality rate, the rate of chronic malnutrition, the measure of imputed poverty, or the FONCODES index, respectively. The results show that the headcount index could be reduced by the same amount with a budget which is 693 million to 770 million soles smaller; the poverty gap could be reduced by the same amount with a budget which is between 731 million and 838 million soles smaller; while the P_2 measure, finally, could be reduced by the same amount with a budget which is 755 to 845 million soles smaller. Clearly, substantial savings can be achieved with geographic targeting in Peru.

Non-parametric Density Estimates

Figure 3, finally, graphs the non-parametric (kernel) density estimates of the log of PCE *after* hypothetical transfers totaling 2.2 billion soles have been made to households in the 1997 LSMS. To avoid cluttering the picture, and because

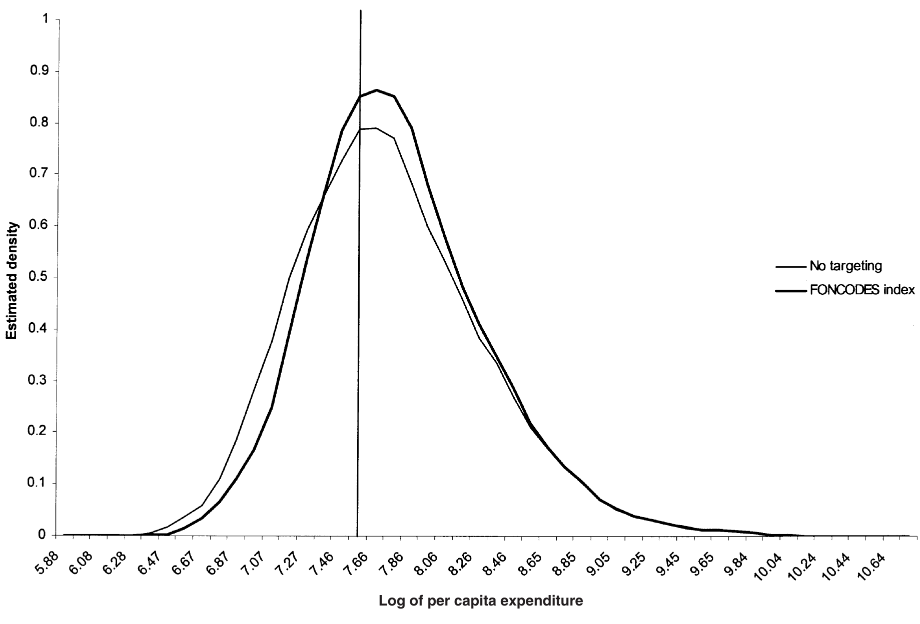

Figure 3. Kernel Density Functions of Log Per Capita Expenditure

outcomes are so similar across indicators, I present density estimates for the FONCODES index and the no targeting scenario only. A vertical line which corresponds to the log of PCE at the poverty line is also included. Figure 3 clearly shows that the left tail of the distribution for the no targeting scenario has more mass than the corresponding distribution for the FONCODES index: that is, the fraction of people with very low levels of per capita expenditure would be higher had transfers to districts been based simply on population rather than on the FONCODES index. By contrast, more mass is concentrated around the middle of the distribution when transfers are made using the FONCODES index. The *cumulative* density functions of log PCE (not presented, but available from the author upon request) show that the distribution corresponding to the FONCODES index first-order stochastically dominates the no targeting distribution up to the poverty line (the two distributions intersect at about the poverty line).

5. Conclusion

In this paper, I use a number of simulation techniques to empirically compare geographic targeting indicators available to policy-makers in Peru. The basic findings of the paper are two. First, all of the targeting indicators I consider are a significant improvement on the no targeting baseline scenario. This result is not surprising. When there is heterogeneity in the distribution of welfare across geographic jurisdictions, (aggregate) welfare gains can be attained with changes in the amount of transfers given to different jurisdictions. Second, all of the targeting indicators perform approximately as well as each other: the differences in outcomes are small, and are never statistically significant. This result is more surprising, as it suggests that in Peru the choice of indicator is not important—at least, among the alternatives considered in this paper.

The extent to which the Peruvian results are applicable elsewhere is not obvious: the joint distribution of various measures of welfare may vary a great deal from country to country. Also, the Peruvian results may be driven, in part, by a mis-application of the methodology proposed by Hentschel *et al.* (2000): imputations based on mean district characteristics, rather than on the individual household characteristics, could substantially bias the estimates. Still, the results do provide a word of caution for many of the more complex ways of estimating geographically-disaggregated measures of poverty which have become increasingly popular in the literature on targeting and poverty maps, as well as among policy makers in the developing world. There are likely to be inherent limitations in any methodology, no matter how sophisticated, which attempts to extract further information from a small number of variables from one or two primary data sets.

When the differences in poverty outcomes across indicators are small, other considerations may matter a good deal. Two are particularly worth considering: incentive effects associated with different indicators, and the degree to which indicators are perceived to be transparent or impartial. Incentive effects may matter if different indicators are more or less easy to manipulate: as is well known, if the costs of manipulating an indicator are small, and the potential benefits associated with transfers are large, households may attempt to manipulate their characteristics to ensure eligibility. In theory, composite measures of welfare are less subject to manipulation, especially where the list of variables and the weights are unknown by the household (as with the measures of imputed poverty, and the measures of the imputed infant mortality rate). I do not believe that this offers much guidance in the choice of *geographic* targeting indicator, however, as collective action problems would tend to make it very difficult for enough households in a district to manipulate their characteristics in order to ensure a higher level of aggregate transfers.

Transparency may be important to ensure the political acceptability of different transfer regimes. In Peru and elsewhere local leaders understandably try to get the largest level of transfers for their constituencies. Central government decision-makers setting allocation rules may have a harder time explaining and convincing local leaders of the fairness of transfer schemes when these are based on complicated statistical methods to "impute" welfare measures. More generally, such complication may make a transfer regime more susceptible to politically-motivated interference and corruption (as well as unintended mistakes). Under these circumstances, simpler allocation rules which are based on a single variable, or on a transparent (if arbitrary) aggregation of a handful of variables may be a preferable to more complicated schemes which are more appealing from a technocratic point of view.

References

Akerlof, George A., "The Economics of 'Tagging' as Applied to the Optimal Income Tax, Welfare Programs, and Manpower Planning," *American Economic Review* 68(1), 8–19, 1978.

Amat y León, Carlos, "Niveles de Vida Desiguales entre las Provincias del Peru," unpublished manuscript, n.d.

Baker, Judy L. and Margaret Grosh, "Measuring the Effects of Geographic Targeting on Poverty Reduction," World Bank LSMS Working Paper 99, Washington, D.C., 1994.

Banco Central de Reserva del Perú, "El Mapa de Pobreza," unpublished manuscript, 1981.

Besley, Timothy and Ravi Kanbur, "The Principles of Targeting," World Bank PRE Working Paper 385, 1990.

Bigman, David and Hippolyte Fofack, "Geographical Targeting for Poverty Alleviation: An Introduction to the Special Issue," *World Bank Economic Review*, 14(1), 129–45, 2000.

Chaudhuri, Shubham and Martin Ravallion, "How Well do Static Indicators Identify the Chronically Poor?," *Journal of Public Economics* 53(3), 367–94, 1994.

Datt, Gaurav and Martin Ravallion, "Regional Disparities, Targeting and Poverty in India," World Bank PRE Working Paper 385, 1990.

FONCODES (Fondo Nacional de Compensación y Desarrollo Social), "Resultado de las Acciones de Focalización" unpublished manuscript, 1995.

———, "Plan Operativo 1996," unpublished manuscript, 1996.

Foster, James, J. Greer, and Eric Thorbecke, "A Class of Decomposable Poverty Measures," *Econometrica*, 52(3), 761–5, 1984.

Glewwe, Paul, "Targeting Assistance to the Poor: Efficient Allocation of Transfers when Household Income is not Observed," *Journal of Development Economics*, 38(2), 297–321, 1992.

Glewwe, Paul and Jacques van der Gaag, "Identifying the Poor in Developing Countries: Do Different Definitions Matter?," *World Development* 18(6), 803–14, 1990.

Grosh, Margaret, *From Platitudes to Practice: Targeting Social Programs in Latin America*, The World Bank, Washington, D.C., 1992.

Hentschel, Jesko, Jean Olson Lanjouw, Peter Lanjouw, and Javier Poggi, "Combining Census and Survey Data to Trace the Spatial Dimensions of Poverty: A Case Study for Ecuador," *World Bank Economic Review* 14(1), 147–65, 2000.

(INEI) (Instituto Nacional de Estadística e Informática, *Perú*: *Mapa de Necesidades Básicas Insatisfechas de los Hogares a Nivel Distrital*, INEI, Lima, 1994.
————, *Metodología para Determinar el Ingreso y la Proporción de Hogares Pobres*, INEI, Lima, 1996.
———-, *Perú*: *Estimaciones de la Mortalidad Infantil en los Distritos*, INEI, Lima, 1997.
Jalan, Jyotsna and Martin Ravallion, "Transient Poverty in Post-reform Rural China," *Journal of Comparative Economics* 26(2), 338–57, 1998.
Kanbur, Ravi, Michael Keen, and Matti Tuomala, "Labor Supply and Targeting in Poverty Alleviation Programs," *World Bank Economic Review* 8(2), 191–211, 1994.
Milanovic, Branko, "The Distributional Impact of Cash and In-Kind Transfers in Eastern Europe and Russia," in Dominique Van de Walle and Kimberly Nead (eds), *Public Spending and the Poor: Theory and Evidence*, The Johns Hopkins University Press, Baltimore, MD, 1995.
Paxson, Christina H., "Using Weather Variability to Estimate the Response of Savings to Transitory Income in Thailand," *American Economic Review*, 82(1), 15–33, 1992.
Paxson, Christina H. and Norbert R. Schady, "The Allocation and Impact of Social Funds: Spending on School Infrastructure in Peru," *World Bank Economic Review*, 16(2), forthcoming, 2002.
Ravallion, Martin, "Poverty Alleviation through Regional Targeting: A Case Study for Indonesia," in N. Hoff, Karla Ruth, Avishay Braverman, and Joseph E. Stiglitz (eds), *The Economics of Rural Organization: Theory, Practice and Policy*, Oxford University Press, New York, 1993.
Republic of Peru, Ministry of the Presidency, *Elementos de la Estrategia Focalizada de Lucha Contra la Pobreza Extrema, 1996–2000*, Ministry of the Presidency, Lima, 1996.
Trussell, James and Jane Menken, "Estimating Levels, Trends, and Determinants of Child Mortality in Countries with Poor Statistics," *Population and Development Review*, 10 (Supplement), 1984.
Webb, Richard Charles, *Government Policy and the Distribution of Income in Peru, 1963–1973*, Harvard University Press, Cambridge, 1977.
Webb, Richard Charles and Gabriela Fernández Baca, *Perú 1997 en Números: Anuario Estadístico*, Cuanto, Lima, 1997.
World Bank, "Peru: Did the Ministry of the Presidency Reach the Poor in 1995?," unpublished manuscript, 1996.